8381742

THE UNITED STATES OF AMERICA

TO ALL TO WHOM THESE PRESENTS SHALL COME:

UNITED STATES DEPARTMENT OF COMMERCE

United States Patent and Trademark Office

June 06, 2023

THIS IS TO CERTIFY THAT ANNEXED IS A TRUE COPY FROM THE RECORDS OF THIS OFFICE OF THE FILE WRAPPER AND CONTENTS OF:

APPLICATION NUMBER: *17/714,965*
FILING DATE: *April 06, 2022*
PATENT NUMBER: *11564015*
ISSUE DATE: *January 24, 2023*

Certified by

Performing the Functions and Duties of the
Under Secretary of Commerce
for Intellectual Property
and Director of the United States
Patent and Trademark Office

PTO/AIA/15 (10-17)
Approved for use through 11/30/2020. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995 no persons are required to respond to a collection of information unless it displays a valid OMB control number.

UTILITY PATENT APPLICATION TRANSMITTAL

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Attorney Docket No.	14-TV1-CON7
First Named Inventor	Shkedi, Roy
Title	Targeted television advertisements based on online behavior
Priority Mail Express® Label No.	n/a - filed via EFS-Web

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

1. ☐ **Fee Transmittal Form** (PTO/SB/17 or equivalent)
2. ☐ **Applicant asserts small entity status.** See 37 CFR 1.27
3. ☐ **Applicant certifies micro entity status.** See 37 CFR 1.29. Applicant must attach form PTO/SB/15A or B or equivalent.
4. ☑ **Specification** [*Total Pages* 45] Both the claims and abstract must start on a new page. *(See MPEP § 608.01(a) for information on the preferred arrangement)*
5. ☑ **Drawing(s)** (35 U.S.C. 113) [*Total Sheets* 18]
6. **Inventor's Oath or Declaration** [*Total Pages* 1] *(including substitute statements under 37 CFR 1.64 and assignments serving as an oath or declaration under 37 CFR 1.63(e))*
   a. ☐ Newly executed (original or copy)
   b. ☑ A copy from a prior application (37 CFR 1.63(d))
7. ☑ **Application Data Sheet** ** See note below.* See 37 CFR 1.76 (PTO/AIA/14 or equivalent)
8. **CD-ROM or CD-R** in duplicate, large table, or Computer Program (*Appendix*)
   ☐ Landscape Table on CD
9. **Nucleotide and/or Amino Acid Sequence Submission** (*if applicable, items a. – c. are required*)
   a. ☐ Computer Readable Form (CRF)
   b. ☐ Specification Sequence Listing on:
      i. ☐ CD-ROM or CD-R (2 copies); or
      ii. ☐ Paper
   c. ☐ Statements verifying identity of above copies

ADDRESS TO:

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

ACCOMPANYING APPLICATION PAPERS

10. ☑ **Assignment Papers** (cover sheet & document(s))
    Name of Assignee INTENT IQ, LLC
11. ☐ **37 CFR 3.73(c) Statement** (*when there is an assignee*) ☑ **Power of Attorney**
12. ☐ **English Translation Document** (*if applicable*)
13. ☐ **Information Disclosure Statement** (PTO/SB/08 or PTO-1449)
    ☐ Copies of citations attached
14. ☐ **Preliminary Amendment**
15. ☐ **Return Receipt Postcard** (*MPEP § 503*) (*Should be specifically itemized*)
16. ☐ **Certified Copy of Priority Document(s)** (*if foreign priority is claimed*)
17. ☐ **Nonpublication Request** Under 35 U.S.C. 122(b)(2)(B)(i). Applicant must attach form PTO/SB/35 or equivalent.
18. ☐ **Other:** _____

***Note**: (1) Benefit claims under 37 CFR 1.78 and foreign priority claims under 1.55 **must** be included in an Application Data Sheet (ADS).
(2) For applications filed under 35 U.S.C. 111, the application must contain an ADS specifying the applicant if the applicant is an assignee, person to whom the inventor is under an obligation to assign, or person who otherwise shows sufficient proprietary interest in the matter. See 37 CFR 1.46(b).

19. CORRESPONDENCE ADDRESS

☑ The address associated with Customer Number: 26362 **OR** ☐ Correspondence address below

Name					
Address					
City		State		Zip Code	
Country		Telephone		Email	

Signature	/Louis J. Hoffman/	Date	04/06/2022
Name (Print/Type)	Louis J. Hoffman	Registration No. (Attorney/Agent)	38,918

This collection of information is required by 37 CFR 1.53(b). The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Privacy Act Statement

The **Privacy Act of 1974 (P.L. 93-579)** requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (*i.e.*, GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

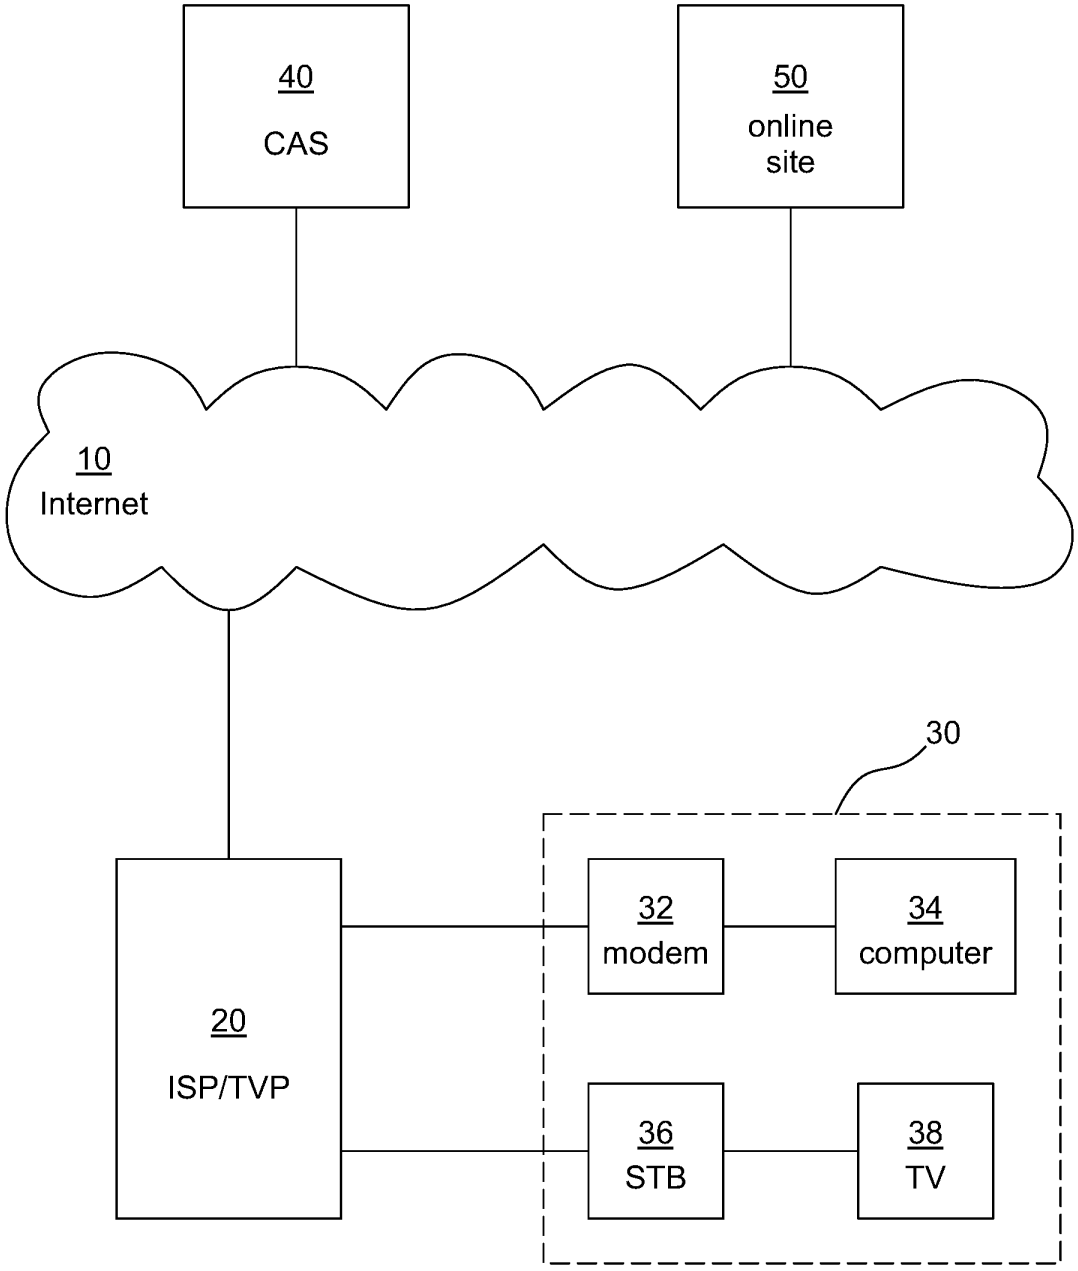

40
CAS
50
online
site
10
Internet
30
20
ISP/TVP
32
modem
34
computer
36
STB
38
TV

FIG. 1

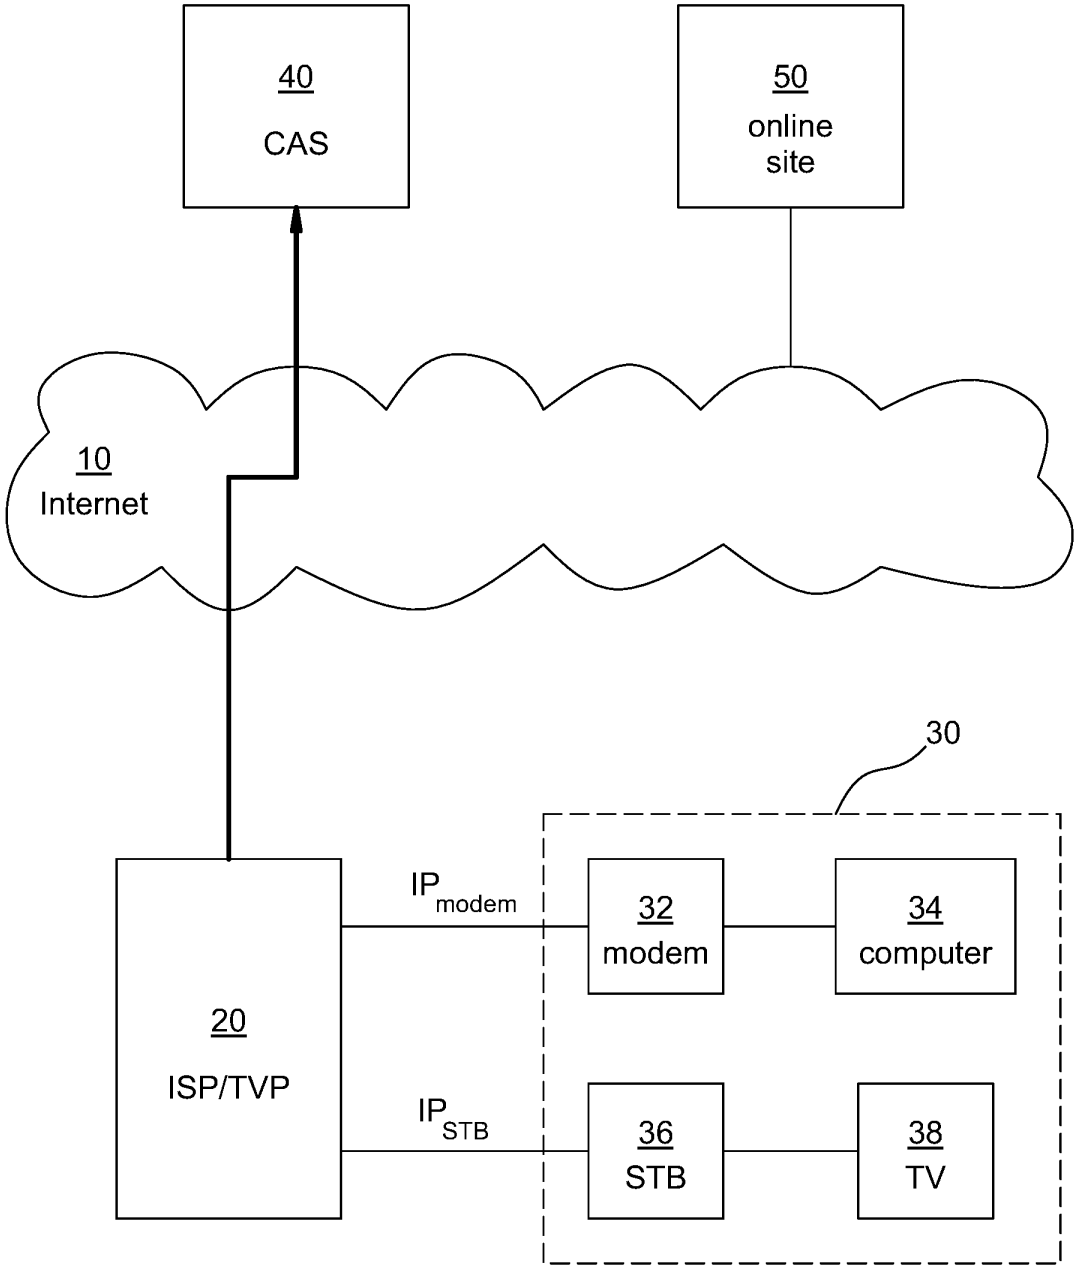

40
CAS
50
online
site
10
Internet
30
20
ISP/TVP
IP modem
IP STB
32
modem
34
computer
36
STB
38
TV

FIG. 2

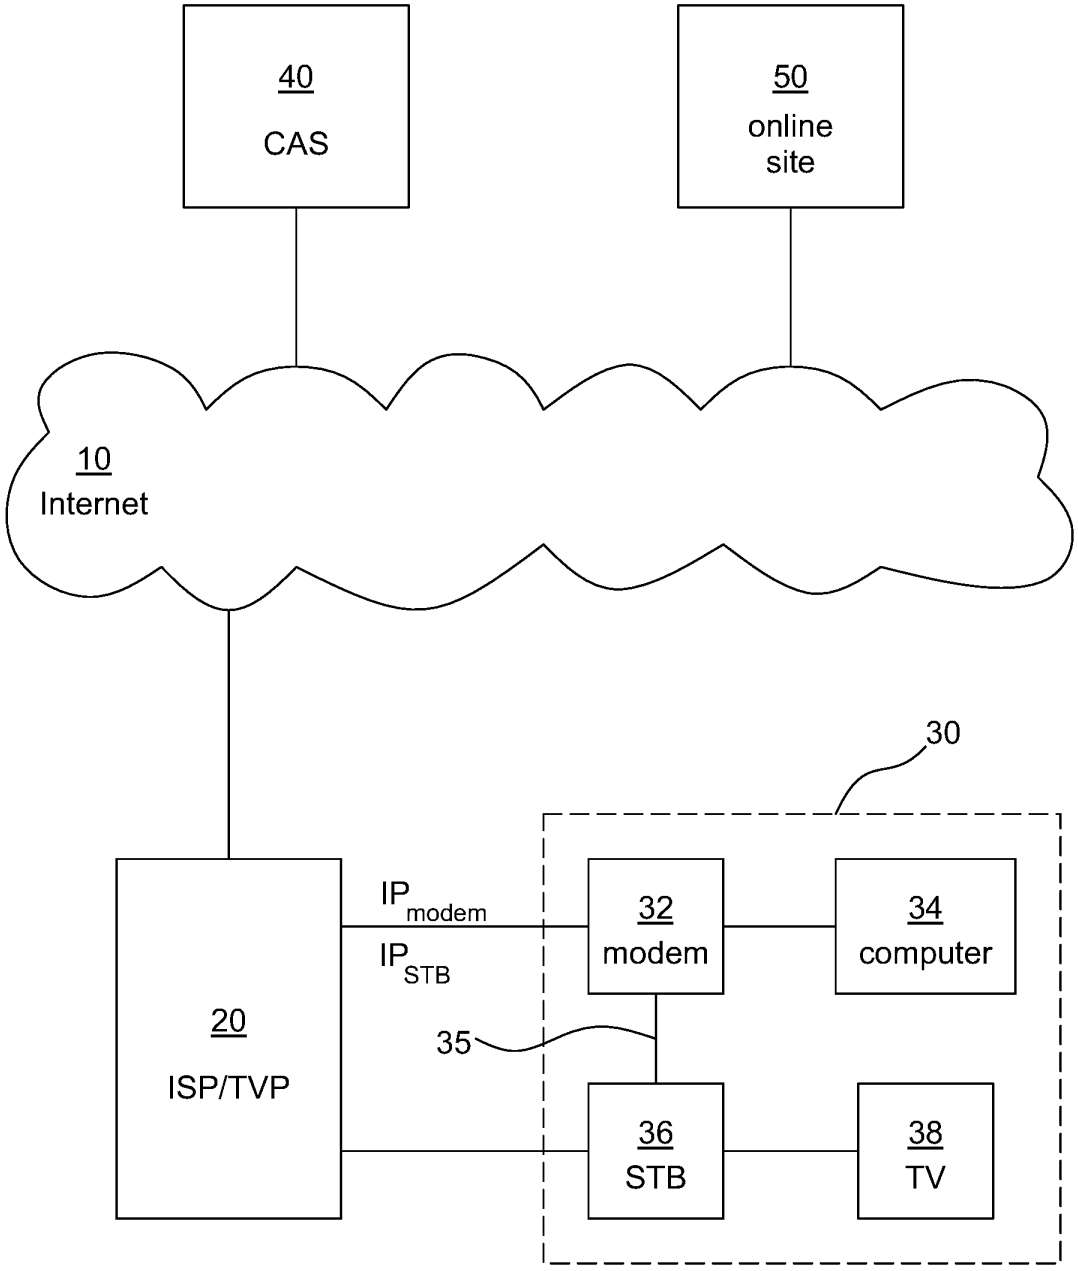

40
CAS
50
online
site
10
Internet
30
20
ISP/TVP
IP modem
IP STB
32
modem
34
computer
35
36
STB
38
TV

FIG. 3

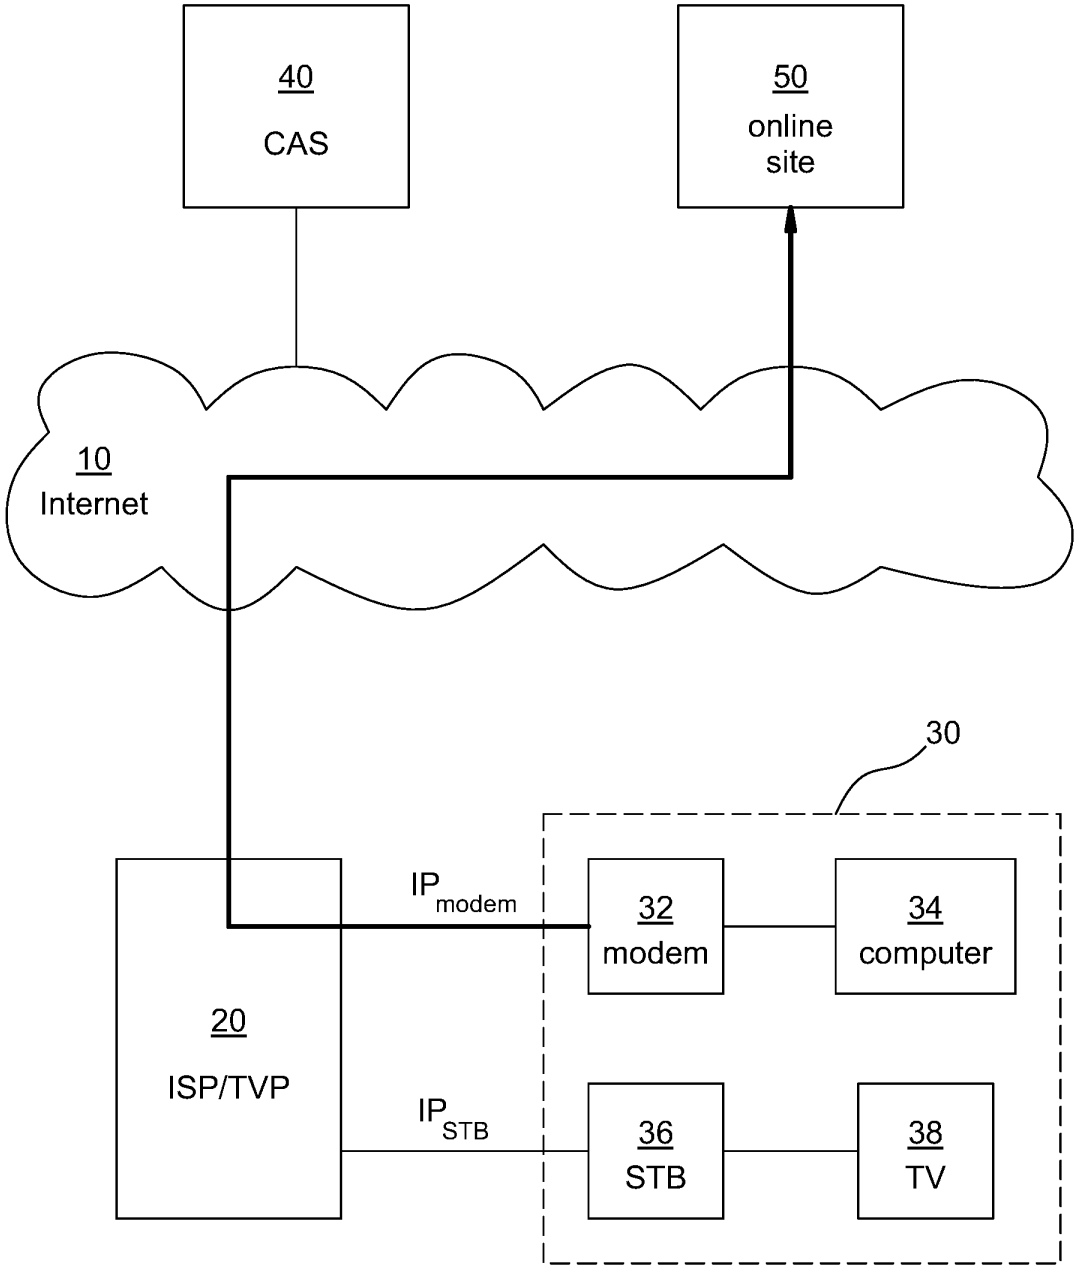

40
CAS
50
online
site
10
Internet
30
IP modem
32
modem
34
computer
20
ISP/TVP
IP STB
36
STB
38
TV

FIG. 4

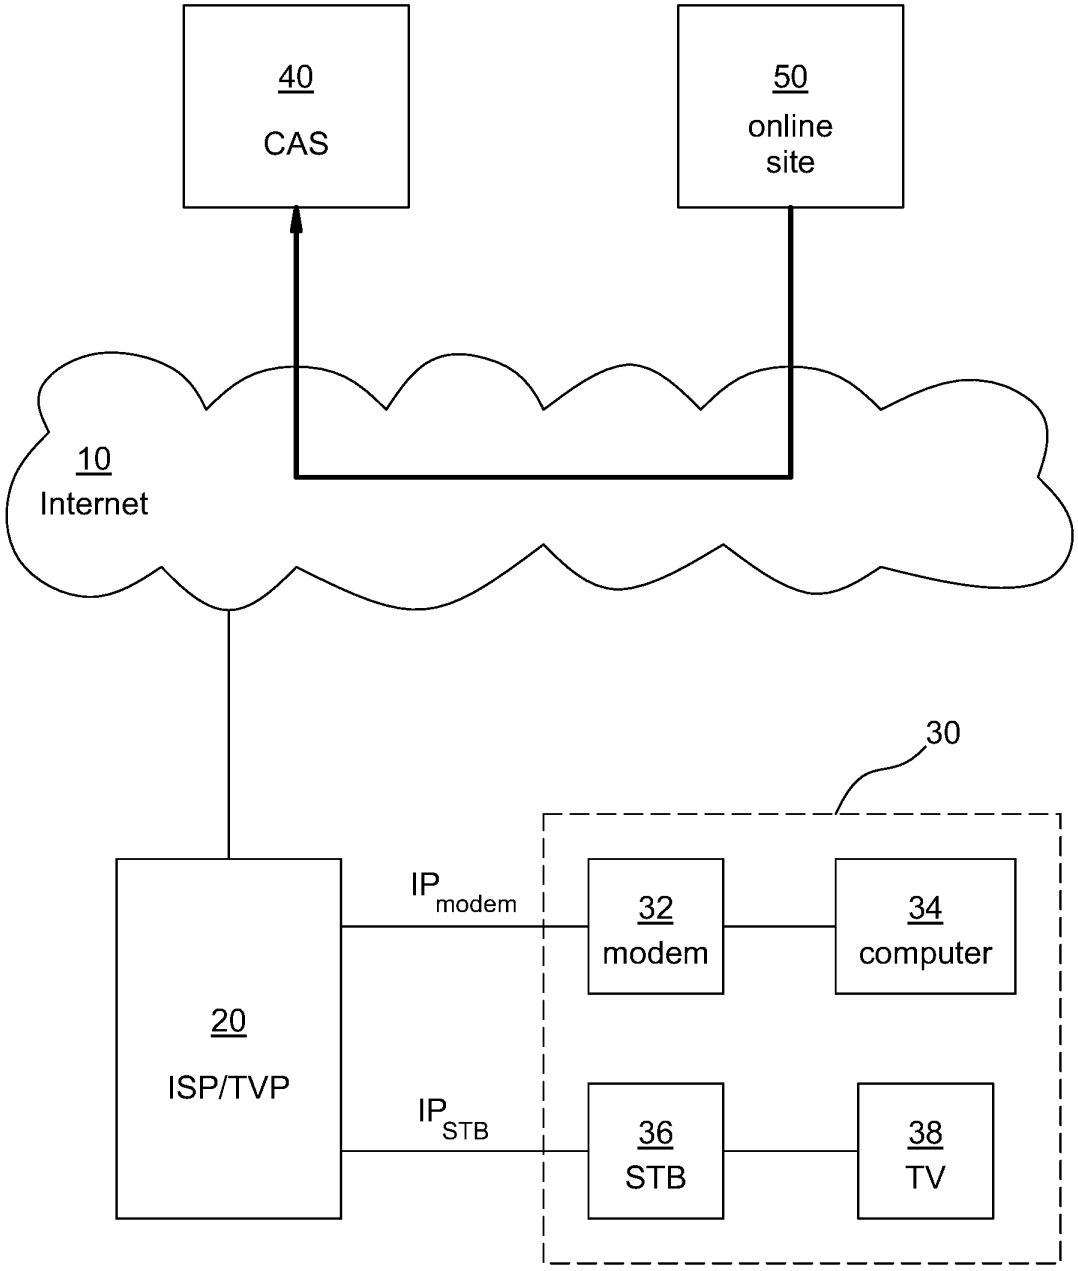

40
CAS
50
online site
10
Internet
30
20
ISP/TVP
IP modem
IP STB
32
modem
34
computer
36
STB
38
TV

FIG. 5

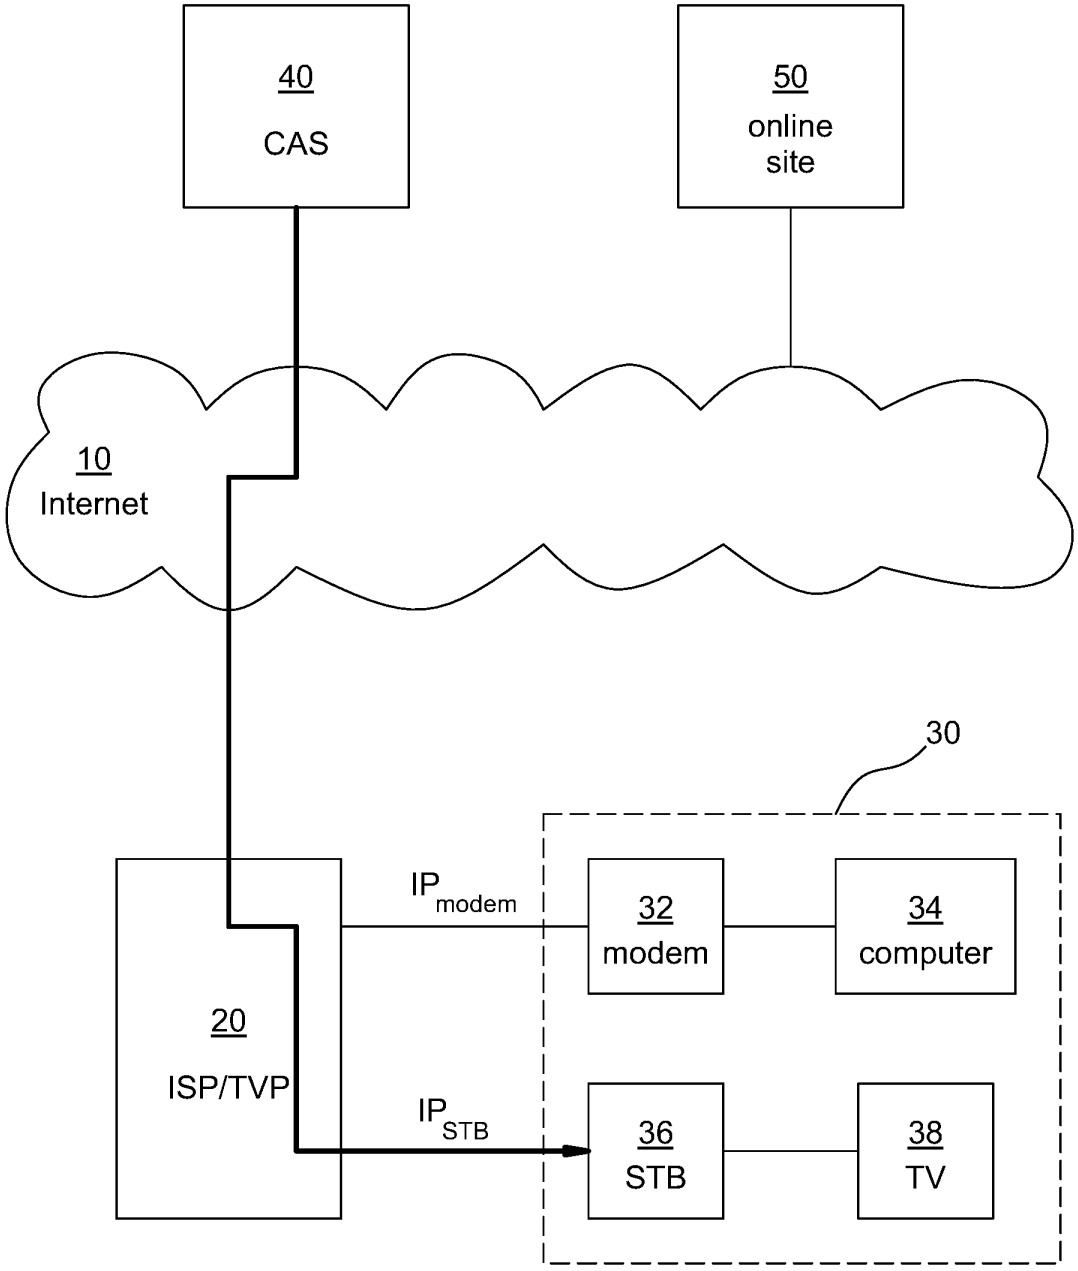

40
CAS
50
online
site
10
Internet
30
IP modem
32
modem
34
computer
20
ISP/TVP
IP STB
36
STB
38
TV

FIG. 6A

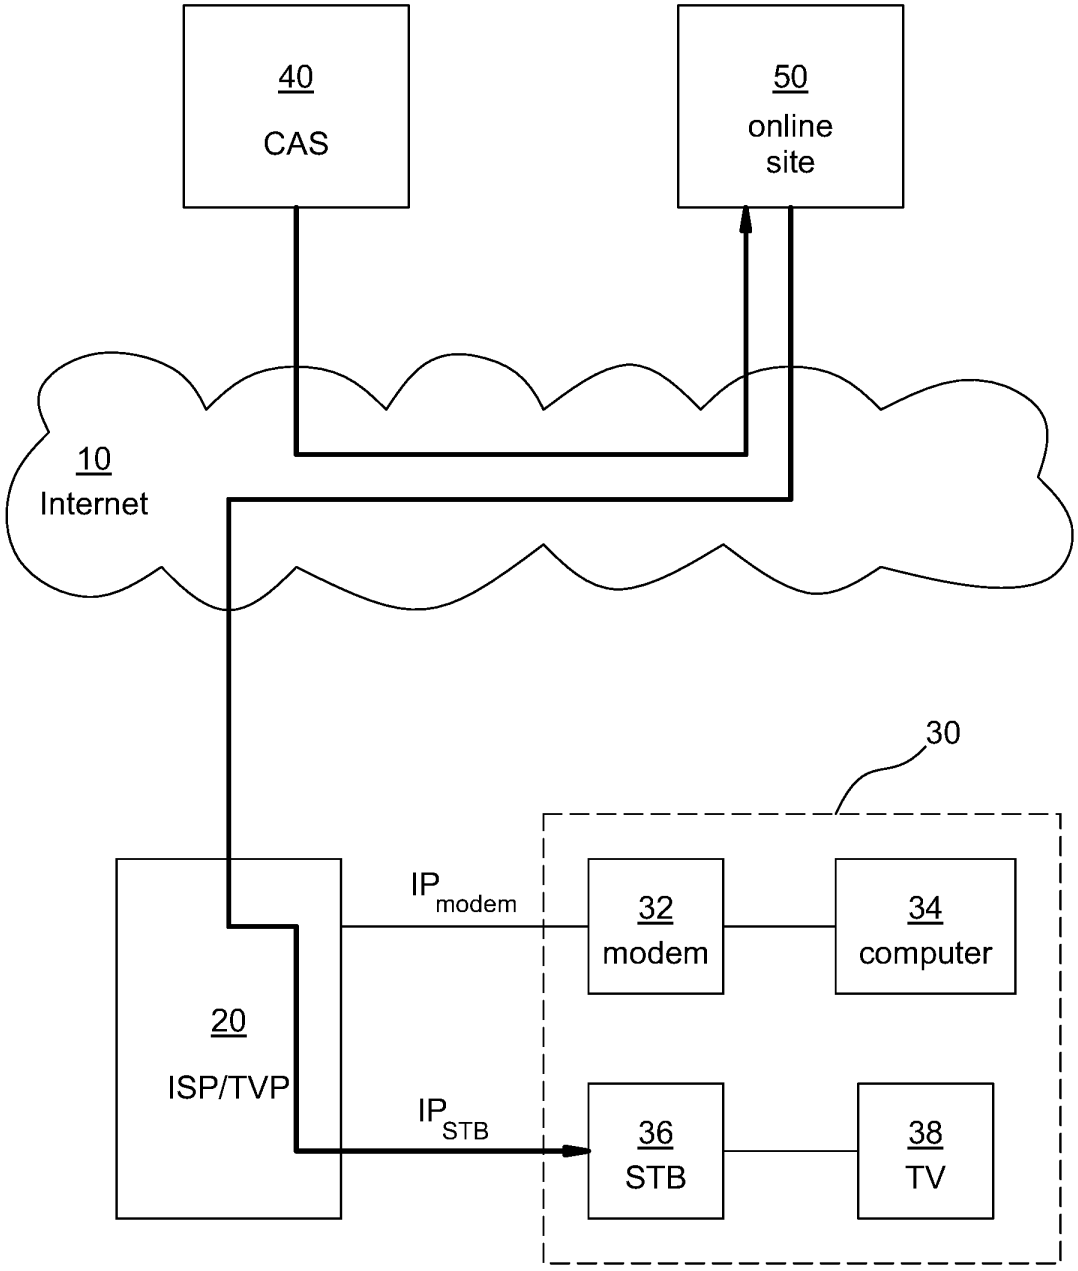

40
CAS
50
online
site
10
Internet
30
IP modem
32
modem
34
computer
20
ISP/TVP
IP STB
36
STB
38
TV

FIG. 6B

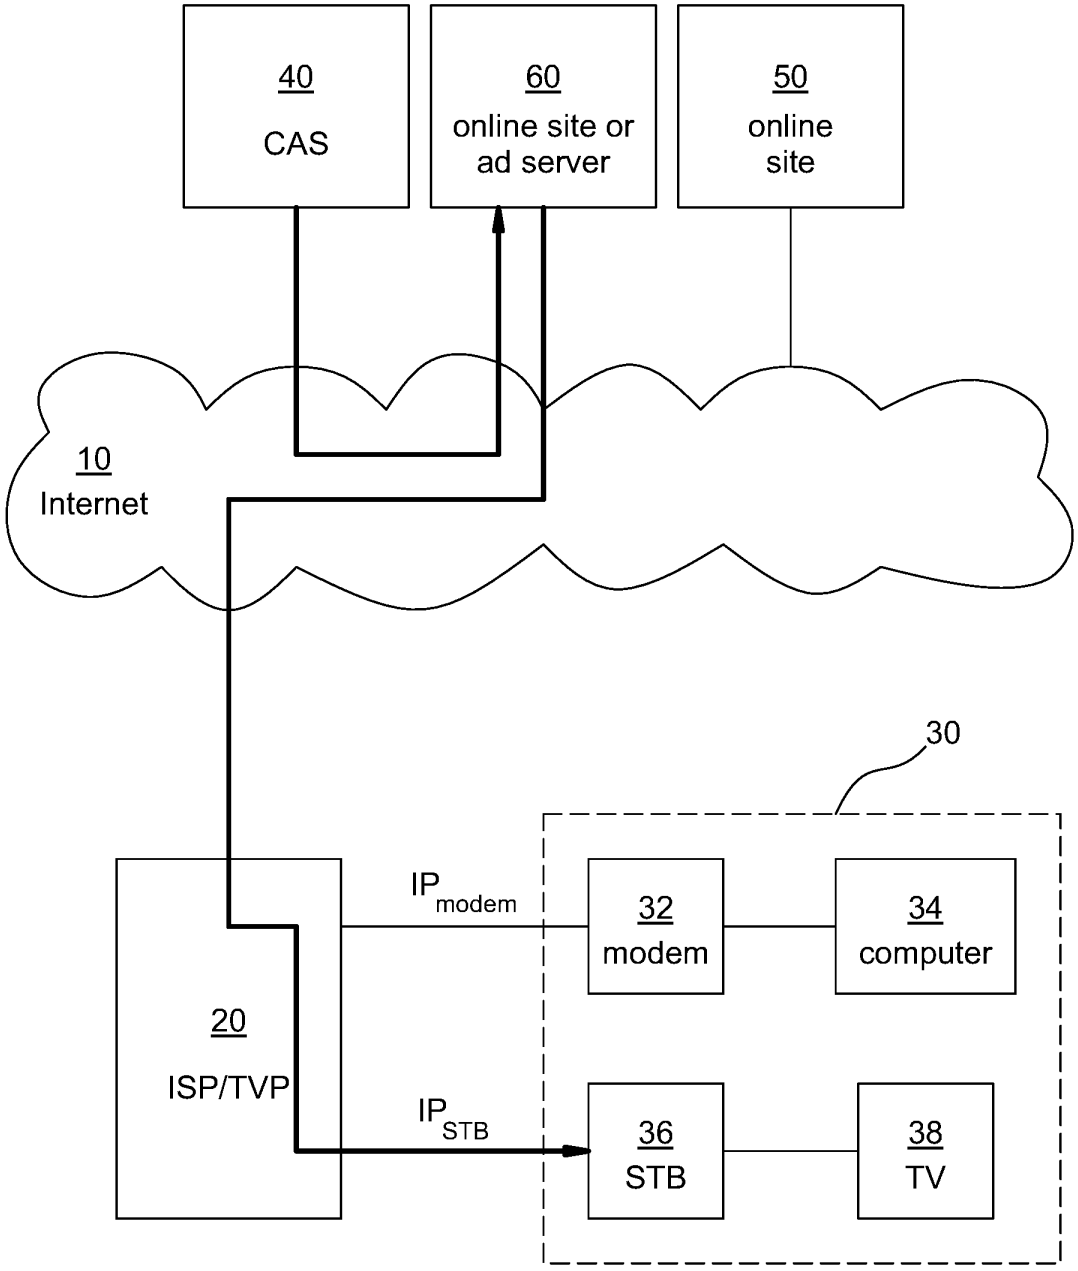

40
CAS
60
online site or
ad server
50
online
site
10
Internet
30
IP modem
32
modem
34
computer
20
ISP/TVP
IP STB
36
STB
38
TV

FIG. 6C

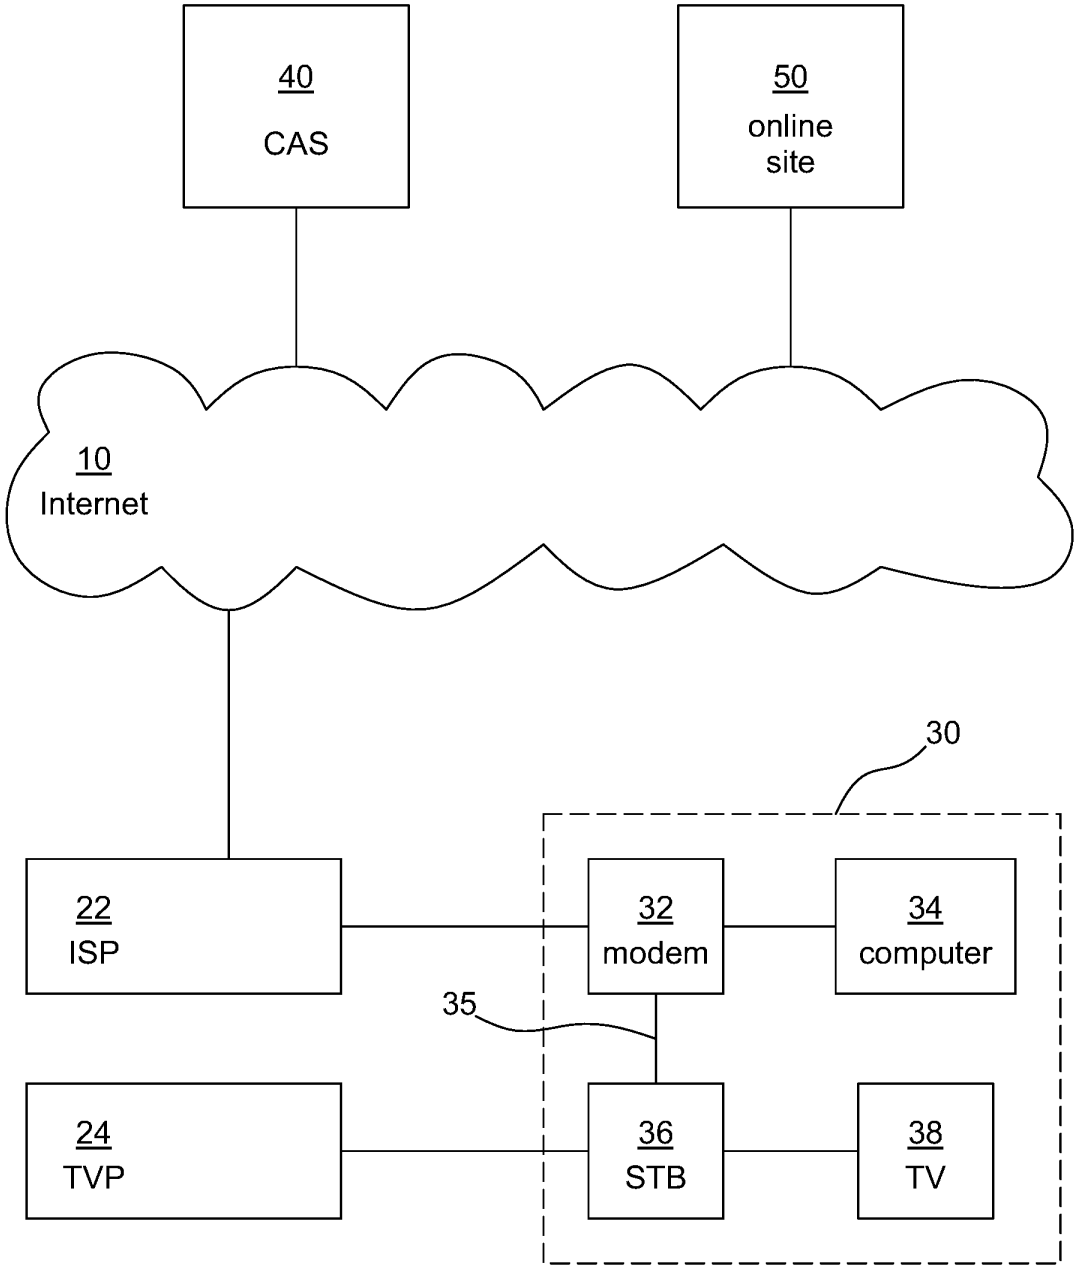

40
CAS
50
online
site
10
Internet
30
22
ISP
32
modem
34
computer
35
24
TVP
36
STB
38
TV

FIG. 7

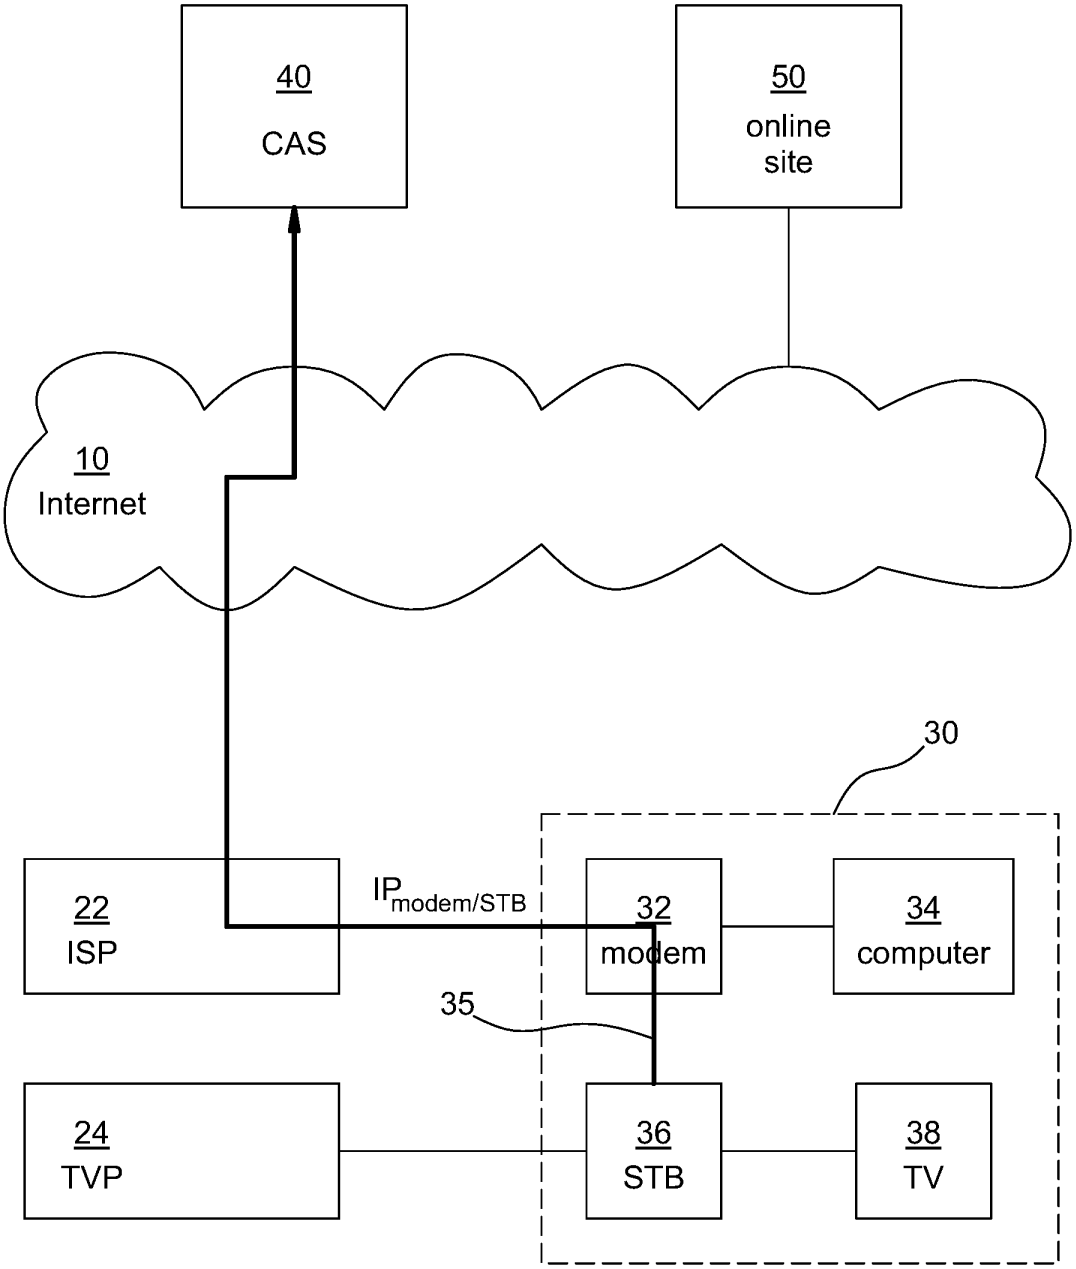

40
CAS
50
online site
10
Internet
30
22
ISP
IP modem/STB
32
modem
34
computer
35
24
TVP
36
STB
38
TV

FIG. 8

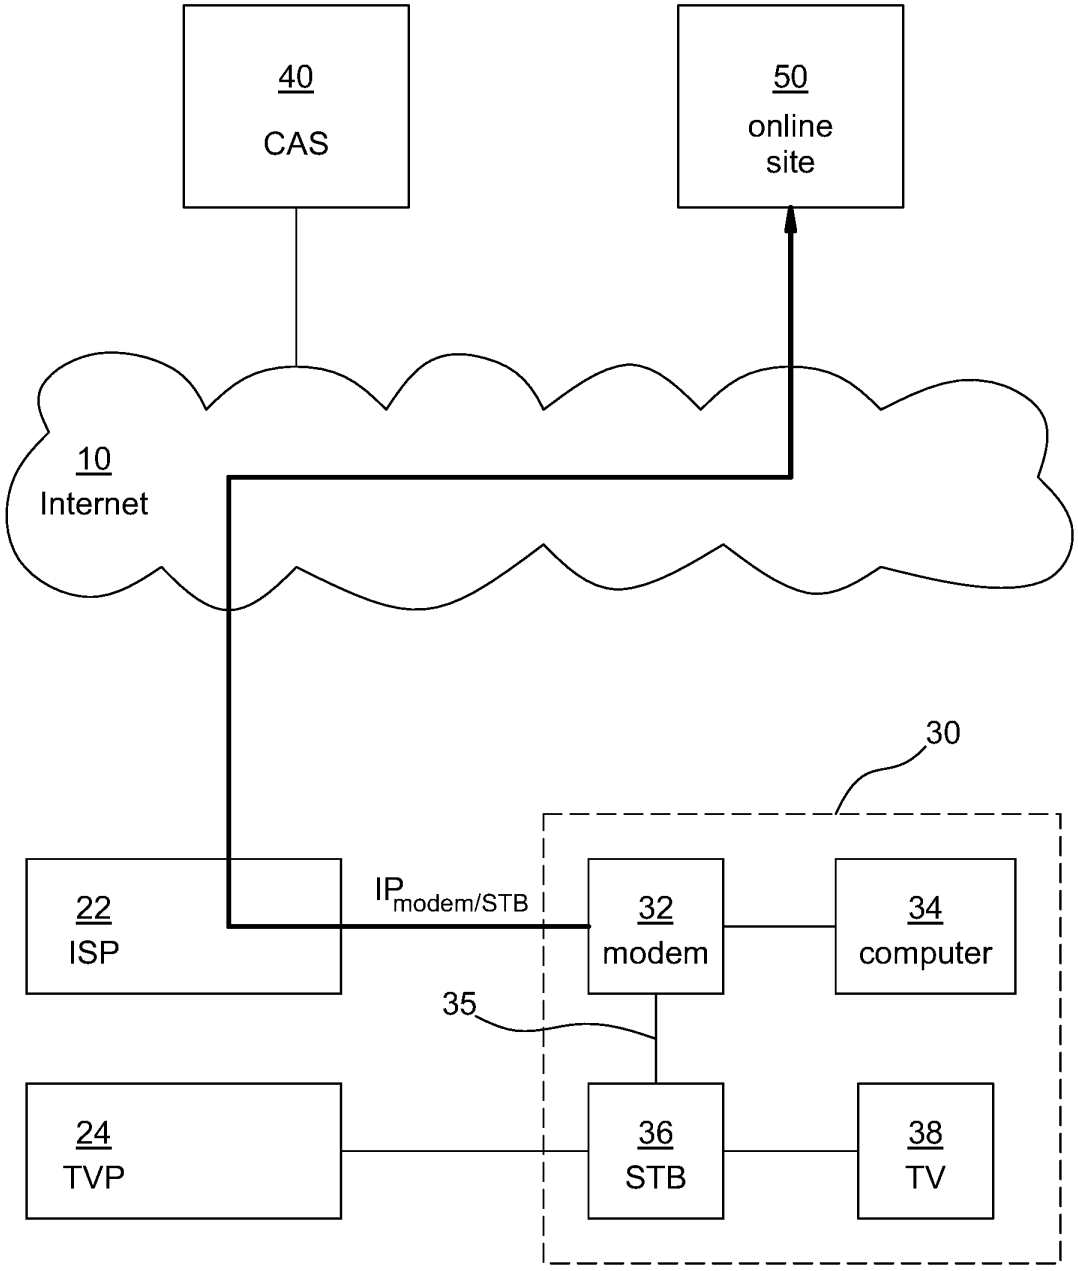

40
CAS
50
online
site
10
Internet
30
22
ISP
IP modem/STB
32
modem
34
computer
35
24
TVP
36
STB
38
TV

FIG. 9

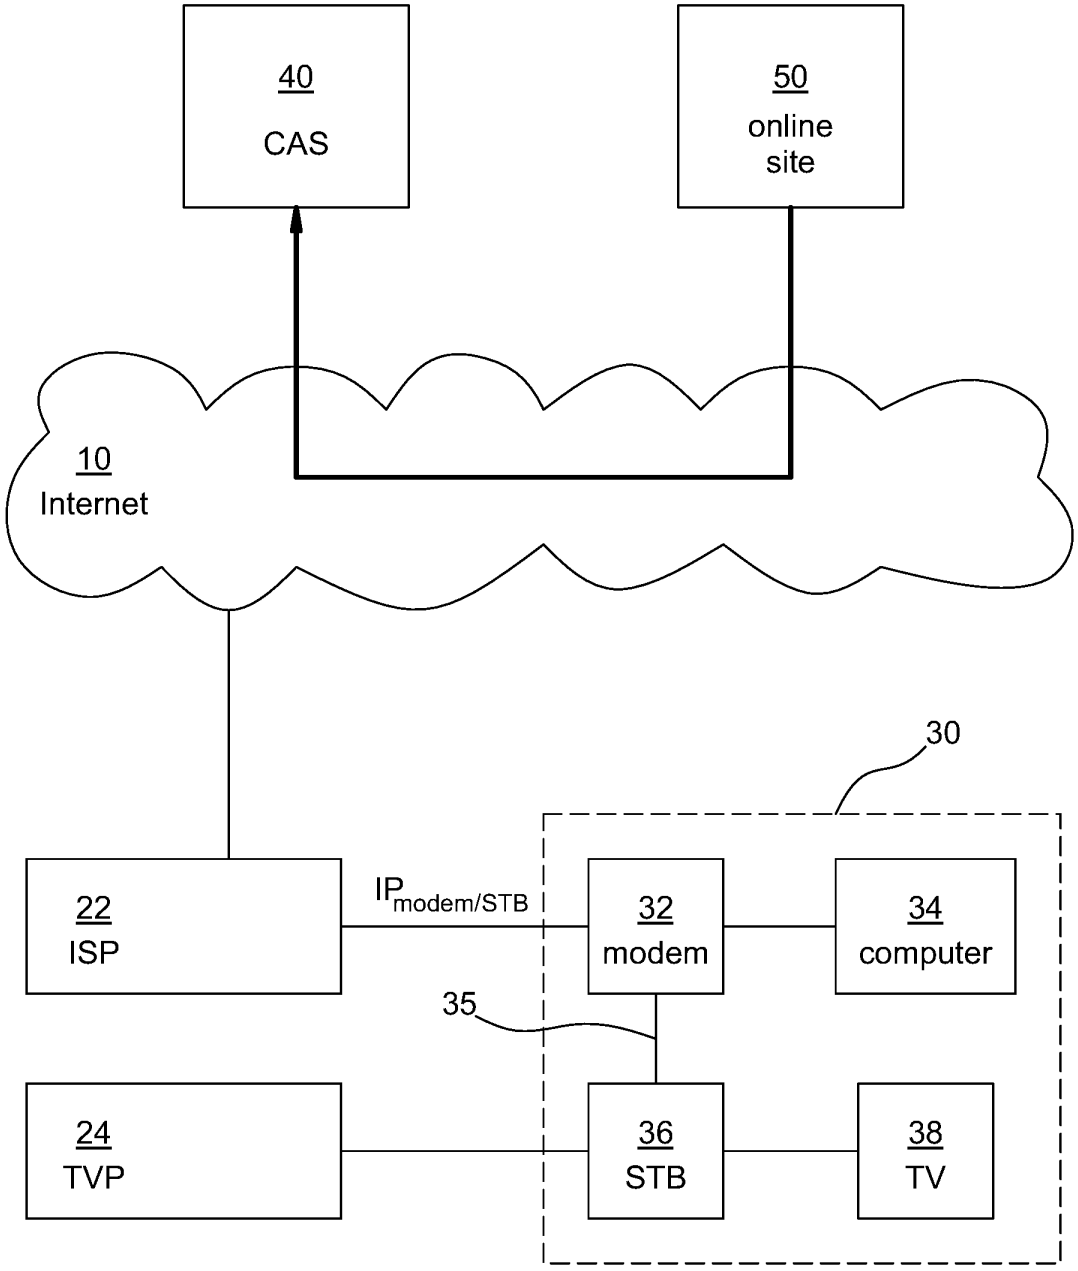

40
CAS
50
online site
10
Internet
30
22
ISP
IP modem/STB
32
modem
34
computer
35
24
TVP
36
STB
38
TV

FIG. 10

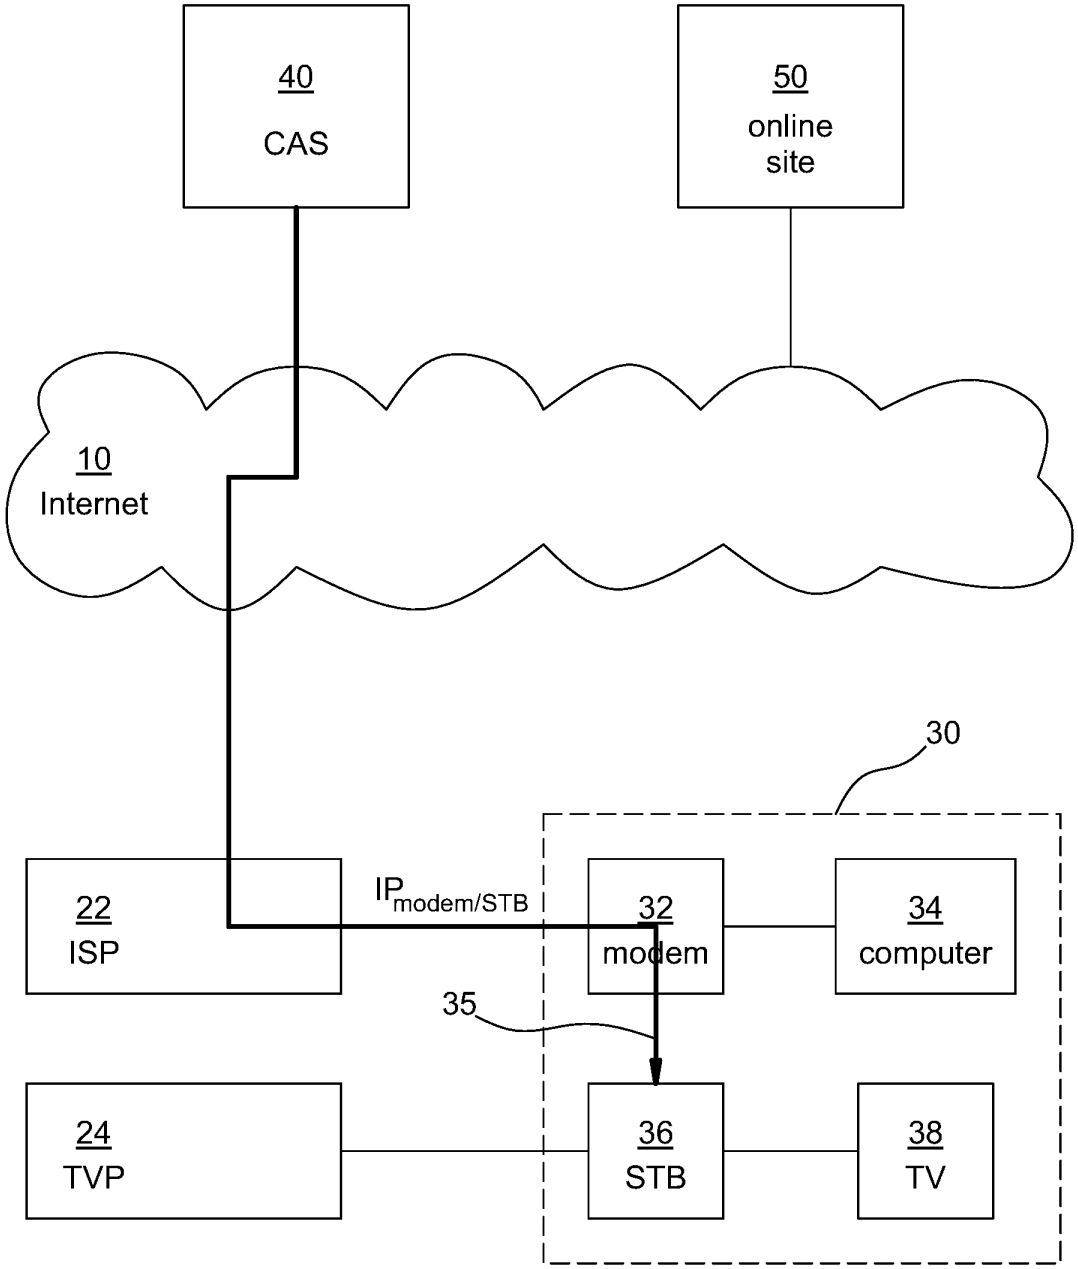

40
CAS
50
online site
10
Internet
30
22
ISP
IP modem/STB
32
modem
34
computer
35
24
TVP
36
STB
38
TV

FIG. 11A

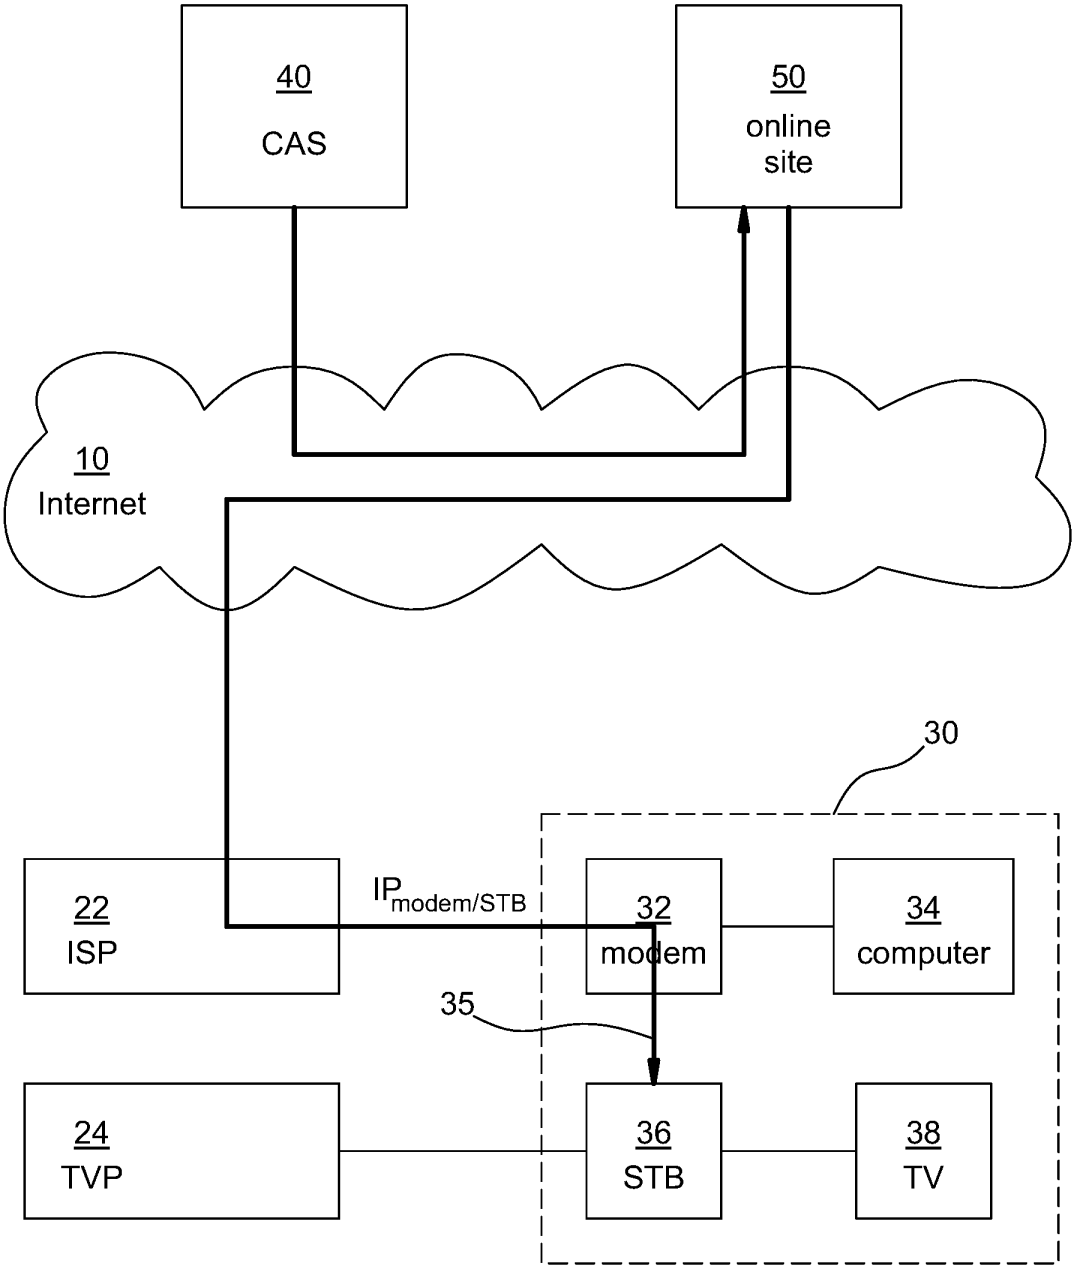

40
CAS
50
online
site
10
Internet
30
22
ISP
IP modem/STB
32
modem
34
computer
35
24
TVP
36
STB
38
TV

FIG. 11B

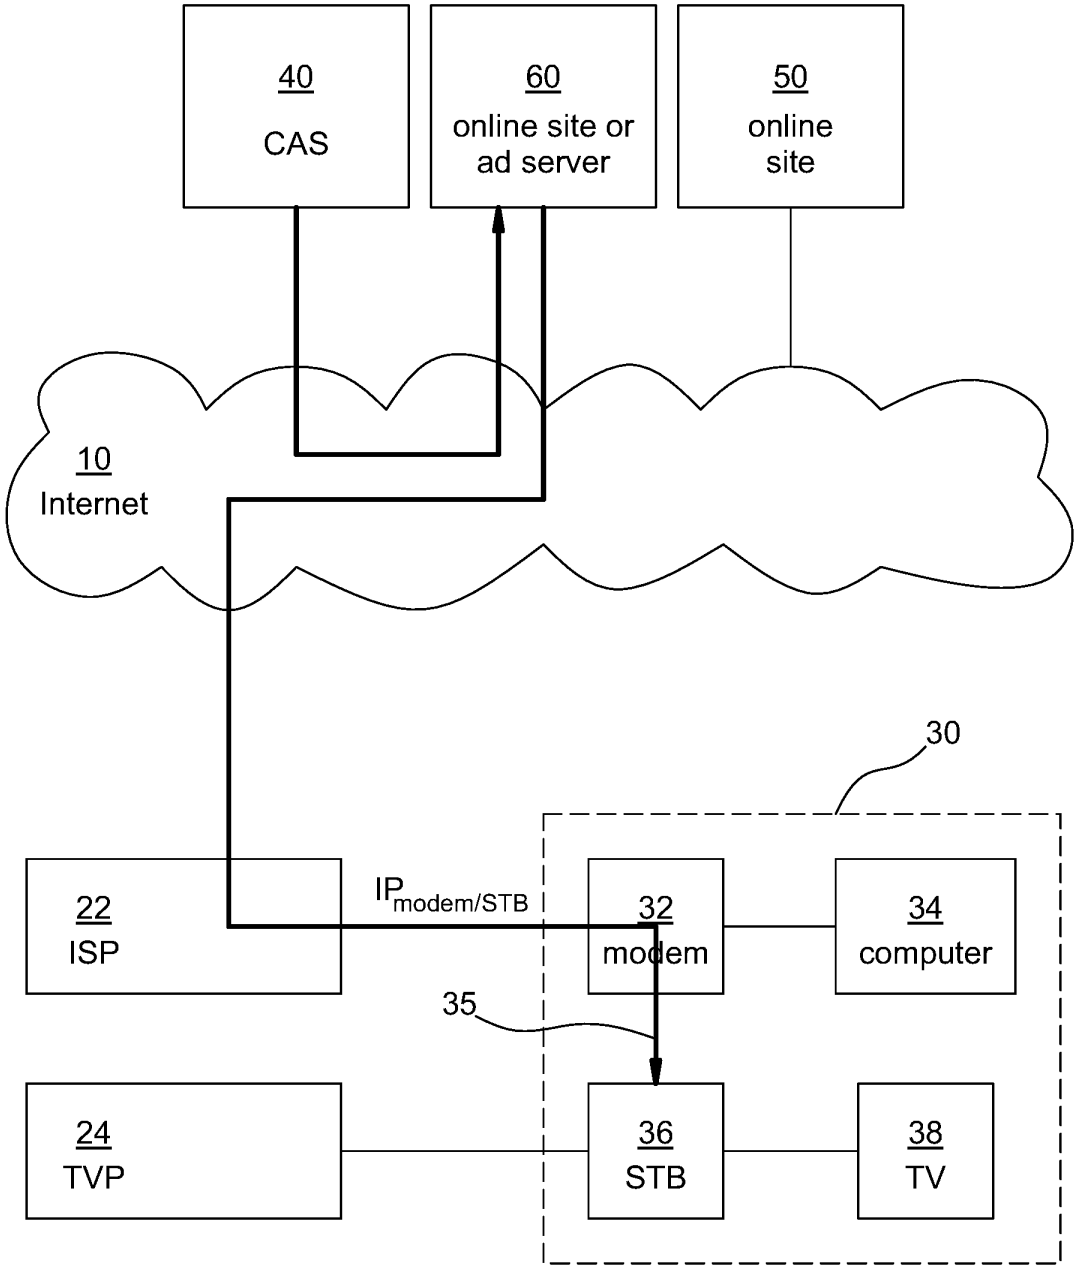

40
CAS
60
online site or
ad server
50
online
site
10
Internet
30
22
ISP
IP modem/STB
32
modem
34
computer
35
24
TVP
36
STB
38
TV

FIG. 11C

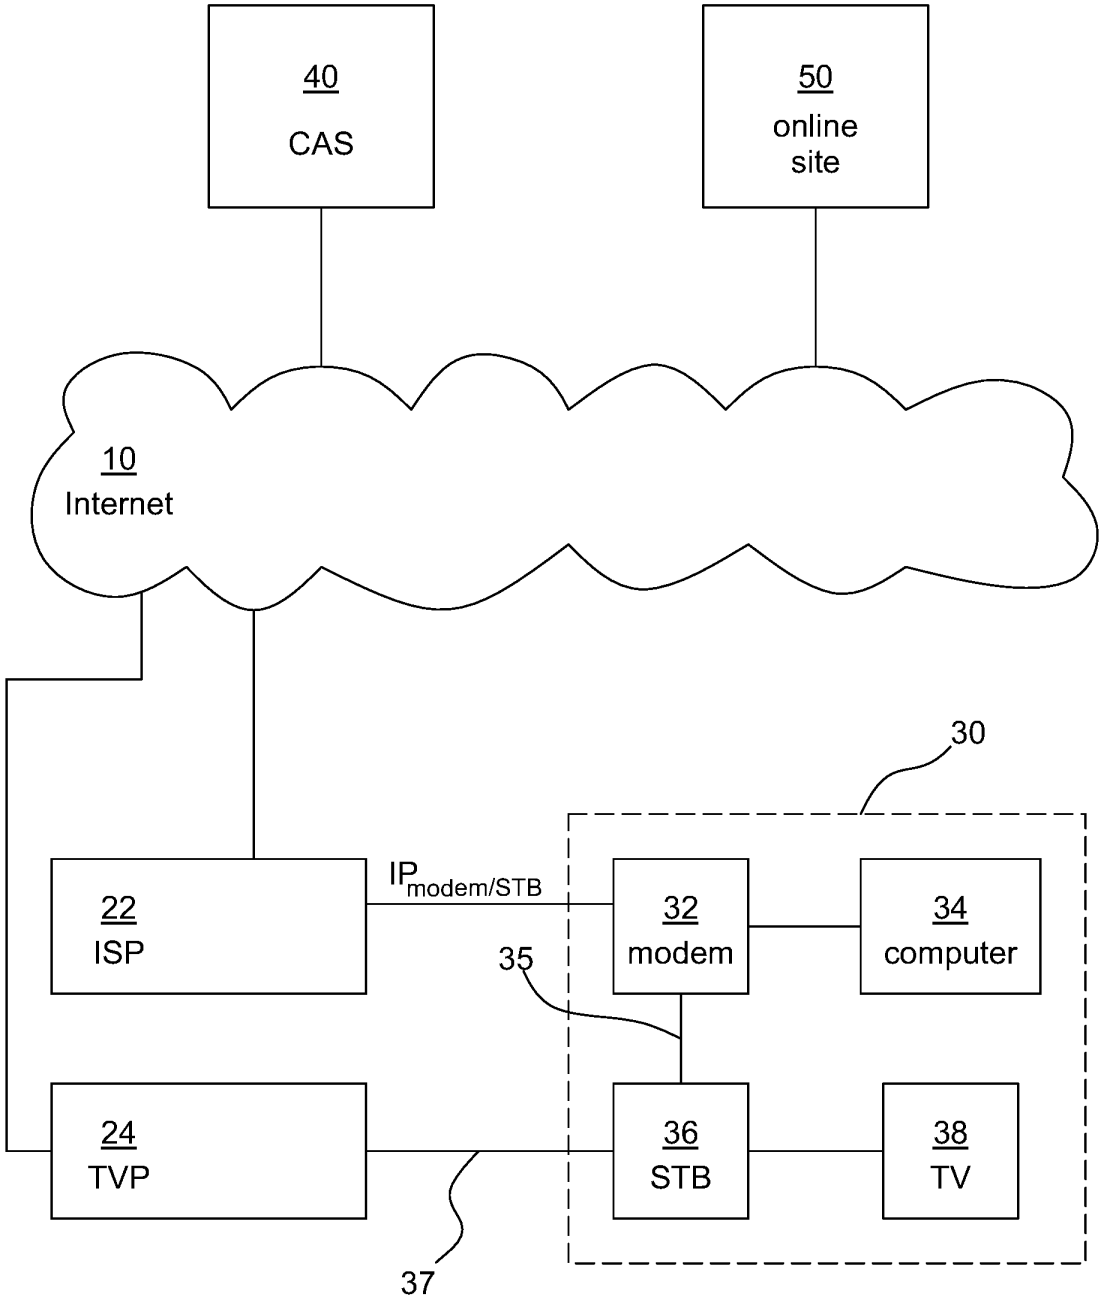

40
CAS
50
online site
10
Internet
30
22
ISP
IP modem/STB
32
modem
34
computer
35
24
TVP
36
STB
38
TV
37

FIG. 12

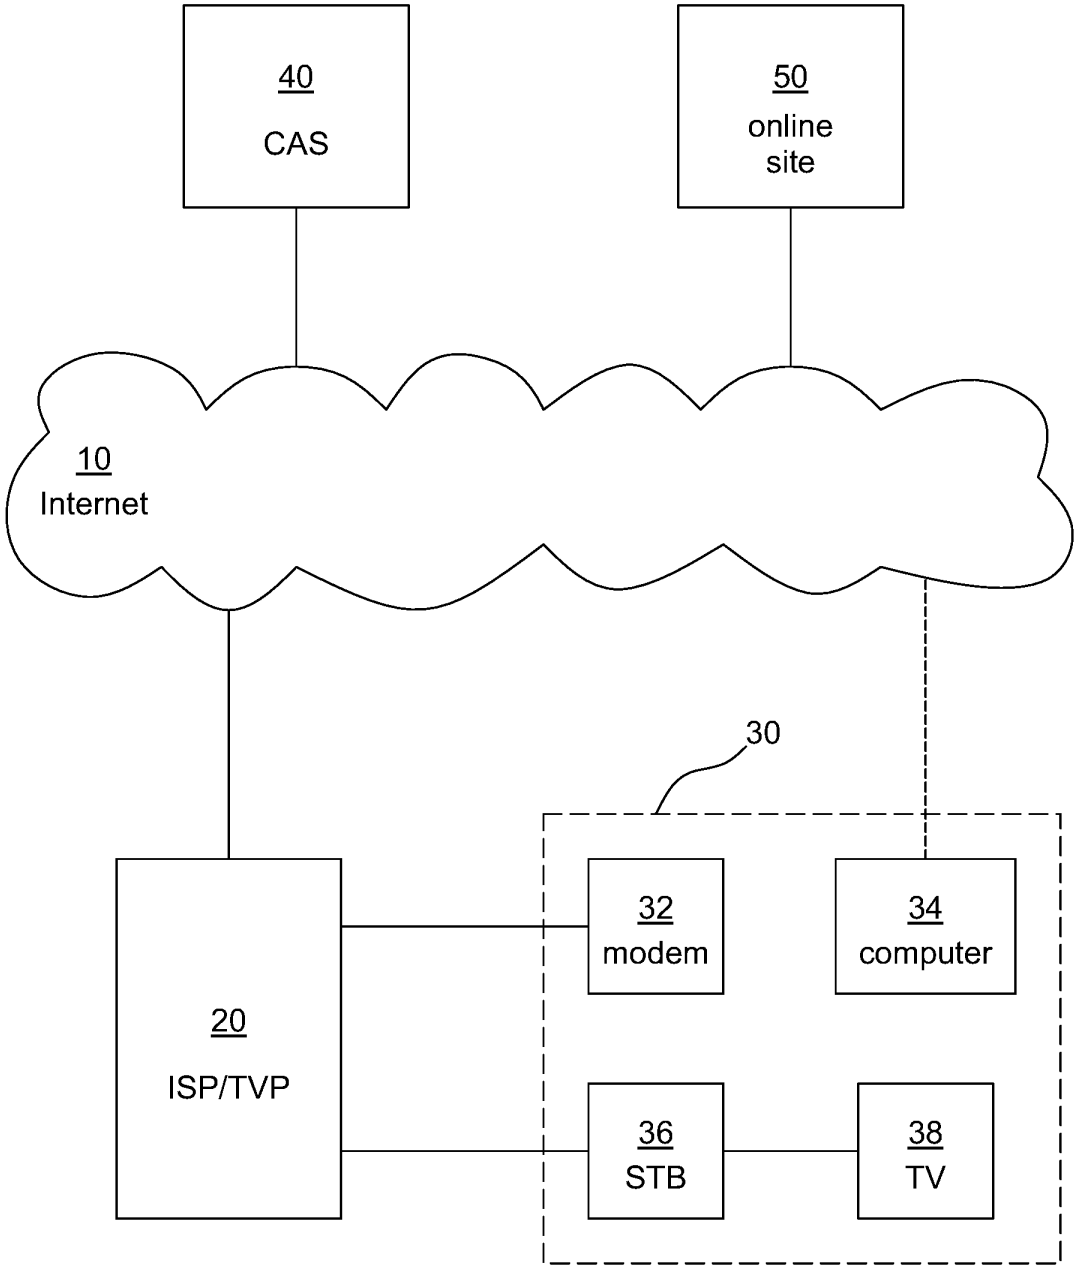

40
CAS
50
online
site
10
Internet
30
20
ISP/TVP
32
modem
34
computer
36
STB
38
TV

FIG. 13

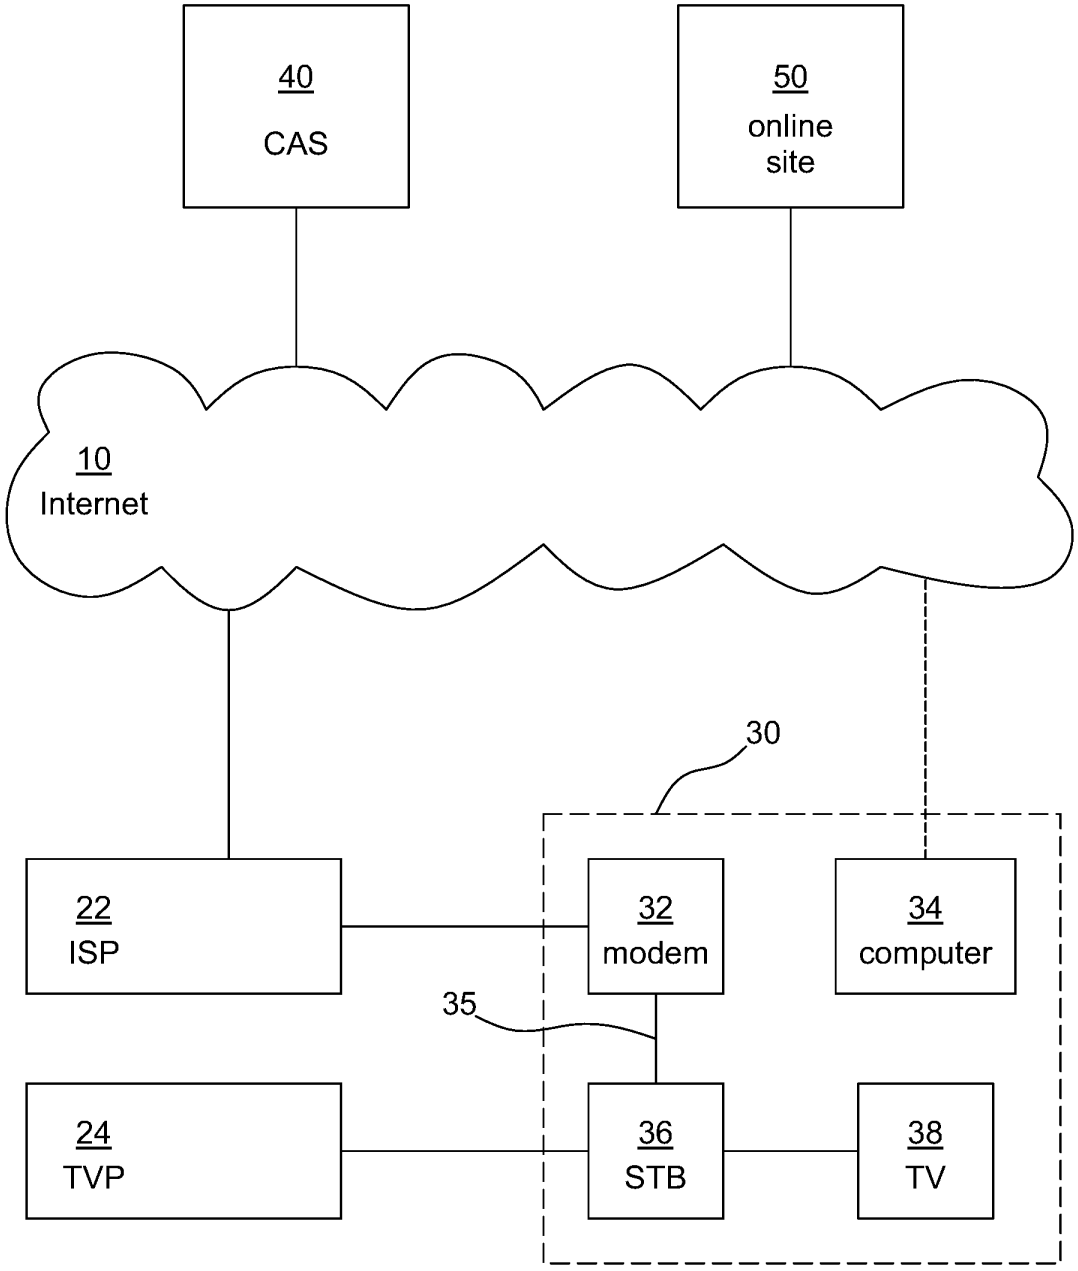

40
CAS
50
online
site
10
Internet
30
22
ISP
32
modem
34
computer
35
24
TVP
36
STB
38
TV

FIG. 14

PTO/AIA/01 (06-12)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

DECLARATION (37 CFR 1.63) FOR UTILITY OR DESIGN APPLICATION USING AN APPLICATION DATA SHEET (37 CFR 1.76)

Title of Invention	Targeted television advertisements based on online behavior

As the below named inventor, I hereby declare that:

This declaration is directed to:

- [■] The attached application, or
- [ ] United States application or PCT international application number _____

filed on _____.

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

WARNING:

Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.

LEGAL NAME OF INVENTOR

Inventor: Roy Shkedi Date (Optional): 04-04-2014

Signature: Roy Shkedi

Note: An application data sheet (PTO/SB/14 or equivalent), including naming the entire inventive entity, must accompany this form or must have been previously filed. Use an additional PTO/AIA/01 form for each additional inventor.

This collection of information is required by 35 U.S.C. 115 and 37 CFR 1.63. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 1 minute to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

TARGETED TELEVISION ADVERTISEMENTS BASED ON ONLINE BEHAVIOR

Applicant: Roy Shkedi

BACKGROUND

[0001] The field of the present invention relates to online access and targeted delivery of advertisements. In particular, television advertisements are targeted based on observed online (i.e., Internet) behavior of a television viewer without employing personally identifiable information.

[0002] Some of the terms used in the present disclosure or appended claims are defined as follows.

[0003] Television provider (TVP) – an entity that provides television service to a subscriber or user via any suitable transmission medium, including but not limited to coaxial cable, fiber-optical cable, network cable, phone line, satellite transmission, or VHF or UHF transmission.

[0004] Internet service provider (ISP; equivalently, an online access provider) – an entity that provides online access to a subscriber or user via any suitable transmission medium, including but not limited to coaxial cable, fiber-optical cable, network cable, phone line, satellite transmission, wireless transmission (e.g., WiMax, WiFi, other IEEE 802 wireless protocols, etc.), or VHF or UHF transmission. The online access enables the subscriber to access the Internet and its myriad online sites, or to access any future network successor to the Internet.

[0005] Set-top box (STB) - a device that connects a television and a signal source. The STB receives an incoming signal, extracts content from the received signal, and transmits the extracted content to the television to be presented to a viewer. The signal source can be a computer network cable (e.g., an Ethernet or other transmission-speed cable), a satellite dish, a coaxial cable connected to a cable television system, a telephone line or digital subscriber line (DSL), a wireless network connection, an antenna (VHF, UHF, digital, or other), or another suitable signal source. The content can include, but is not limited to, video (which often can include an audio portion),

audio, Internet web pages, interactive games, or other content. An STB may or may not include a dedicated television tuner. Despite its name, an STB need not be physically located on top of a television set literally. Under current technology, STBs often are located physically adjacent to the television set, such as in a media cabinet or the like, but it is not even necessary that the STB be located in proximity to the television. Nor is it necessary that the STB be a box, literally. Rather, a STB might be implemented, for example, as a circuit board, integrated circuit, set of integrated circuits, or software that is physically integrated with another "box," such as the television, a cable or other connection, a computer, or a building equipment or junction box, which also has other functions, or without being housed in any "box" at all.

[0006] Digital video recorder (DVR, alternatively personal video recorder or PVR) - a device that stores video content in a digitally encoded format on a digital storage medium, such as a hard drive, and enables playback of the stored content. A DVR can comprise a stand-alone unit connected to a television, an STB, or a signal source, or the DVR can comprise software that programs a computer to perform DVR storage and playback functions.

[0007] Video-on-Demand (VOD) - a system that allows users to select and view video content delivered from a signal source in response to a request from the user. Typically, the requested video content can be viewed at a time of the user's own choosing and can be paused, rewound, or fast-forwarded as desired by the user. A VOD system can "stream" the content (enabling viewing of portions of a requested item of video content while other portions are still being delivered from the signal source), or the VOD system can "download" the content and allow viewing only after a complete item is delivered from the signal source. Some VOD systems allow users to select and watch video content over a network as part of an interactive television system.

[0008] Interactive Television (interactive TV, iTV, idTV, or ITV) – any television system that enables a viewer to interact with video content delivered to a television. Interactive television can include, but is not limited to, access to Web sites through TV "crossover links," electronic mail and online chat, online commerce, or enhanced graphics (relative to standard television offerings).

[0009] Internet Protocol Television (IPTV) – a system wherein television content is delivered via a computer network using Internet Protocol (IP). For residential users, IPTV is often provided in conjunction with Video on Demand and can also be bundled with Internet services such as Internet access and Voice-over-IP (VoIP). Commercial bundling of IPTV, VoIP, and Internet access is sometimes referred to in the industry as a "triple play." Additional telecommunications services (e.g., mobile voice or data service) can be added, yielding a "quadruple play" and so forth. IPTV typically is supplied by a broadband service provider using a closed network infrastructure. IPTV also can be provided over the Internet or other publicly accessible computer network, in which case it might be referred to as Internet TV or TV-over-Internet. IPTV also can be used to deliver video or other content over a corporate LAN or other business network.

[0010] Online user interface device – any user interface device used to access a remote network such as the Internet, including but not limited to a cell phone or mobile handset, a personal digital assistant (PDA), or a networked computer (desktop, workstation, notebook, laptop, or other).

[0011] Online access device – any device used to connect an online user interface device to a remote network such as the Internet, including but not limited to a modem, a wired or wireless router, a wireless access point, a wired network adapter (e.g., Ethernet adapter), a wireless network adapter (e.g., IEEE 802.11, ED-VO, EDGE, HSPA, CDMA, GSM, or other), or an optical fiber based network adapter (e.g., a network interface unit or optical network terminal). Different types of online access devices can and sometimes are combined into a single unit (e.g., a modem that also functions as a router for a LAN). An online user interface device and an online access device can be, and sometimes are, combined into a single unit (e.g., a computer with a built-in Ethernet adapter, wireless adapter, or modem).

[0012] Router – any device that acts as a junction between networks, to buffer and transfer data between or among them. For example, a router can be employed to connect a local area network (LAN) to the Internet, thereby enabling online user interface devices connected to the LAN to share a connection to the Internet through the router. The router receives data from devices on the LAN and transmits them to the Internet directed to their corresponding destinations, and receives data from the Internet

Internet and directs them to the corresponding devices on the LAN.

[0013] Modem – a device that enables online access by a user by acting as an interface between the online access provider's network transmission system and the user's computer or other online user interface device. Modems vary according to the type of provider network transmission system. Unless a specific type of modem is specified, the term "modem" shall encompass telephone modems, cable modems, DSL modems, wireless modems, satellite modems, or modems for providing online access to any other suitable network transmission system.

[0014] Cable modem - a type of modem that enables digital data transmission over cable television infrastructure. Cable modems are primarily used to deliver broadband Internet access using modulation frequencies that are not used for traditional television transmission.

[0015] Asymmetric Digital Subscriber Line (ADSL, or often DSL) - a data communications technology that enables faster data transmission over a copper telephone line than a conventional modem can provide. ADSL data transmission employs frequencies that typically are not used for voice transmission, usually frequencies beyond the range of normal human hearing (higher than about 20kHz). Propagation of such high-frequency signals over standard copper telephone lines typically is poor, typically limiting the use of DSL to distances less than about 5 km. Once the signal reaches the telephone company's nearest central office (CO), the ADSL signal is stripped off and routed into a conventional data network, while any voice-frequency signal is routed into the conventional telephone network. That arrangement enables a single telephone line to be used for both data transmission and telephone calls simultaneously.

[0016] DSL modem – an ADSL transceiver, also known as an ADSL modem, used to connect one or more computers to a phone line to use an ADSL service. A DSL modem also can be referred to as a remote ADSL termination unit (ATU-R). An ADSL modem can also be configured to act as a router, managing the connection and sharing of the ADSL service with multiple computer or other networked devices. Such a combined device can be referred to as a DSL modem/router or similar terminology.

[0017] Digital Subscriber Line Access Multiplexer (DSLAM) – a device enabling high-speed connections between telephone lines and a computer network such as the Internet. Typically it is located at a telephone company central office (CO) and connects multiple Digital Subscriber Lines (DSLs) to a computer network, typically the Internet, using a suitable multiplexing technique.

[0018] Cable Modem Termination System (CMTS) – equipment typically found in a cable company's headend that is used to provide data communication over the cable television infrastructure, thereby enabling the cable company to offer services such as broadband Internet access or VoIP to its subscribers. To provide high-speed data services (i.e., broadband access), a cable company typically connects its headend to the Internet using high capacity data links, directly or through a network service provider. On the subscriber side of the headend, the CMTS enables data communication with each subscriber's cable modem. Various CMTSs are capable of serving different cable modem population sizes, which can range from 4,000 cable modems to 150,000 or more. A given cable company headend may have only a few CMTSs, or a dozen or more, depending on the size of the cable modem population serviced by that headend.

[0019] Dynamic Host Configuration Protocol (DHCP) – a system that automatically assigns Internet Protocol addresses (IP addresses), subnet masks, default routers, and other IP parameters, which are required for proper routing of data transmissions to or from a particular device connected to the network. The assignment usually occurs when a DHCP-configured computer, modem, router, or other device boots up or regains connectivity to the network. The DHCP client (i.e., DHCP software resident in the computer) sends out a query requesting a response from a DHCP server on the network. The query is typically initiated immediately after booting up and before the client initiates any IP-based communication with other hosts. The DHCP server then replies to the client with its assigned IP address, subnet mask, domain name server (DNS), and default gateway information (referred to as "stateful" assignment). The assignment of the IP address usually expires after a predetermined period of time, at which point the DHCP client and server renegotiate a new IP address from the DHCP server's predefined pool of IP addresses. Because, under DHCP, the IP address of a

given computer varies over time, various network-related functions are more difficult. For example, configuring firewall rules to allow access to or from a machine that receives its IP address via DHCP is more complicated because the IP address varies from time to time. Network administrators typically must enable access to an entire remote DHCP subnet for a particular TCP/UDP port. Such complications arise in other instances as well. Many residential routers and firewalls are configured in the factory to act as DHCP servers for home networks. A computer can also be used as a DHCP server. Internet Service Providers (ISPs) generally use DHCP to assign individual IP addresses to subscribers. DHCPv6, which is the Dynamic Host Configuration Protocol for Internet Protocol Version 6 (IPv6), enables local generation of IP addresses (i.e., "stateless" assignment). Although such stateless address auto-configuration of IPv6 substantially eliminates a prime motivation for DHCP in IPv4, DHCPv6 can still be used to statefully assign addresses if needed or desired by a network administrator. DHCPv6 can also be used to distribute information that is not otherwise discoverable, e.g., the domain name server.

[0020] Behavioral targeting – the delivery of specific advertisements to a subscriber, the advertisements being selected on the basis of activity of the subscriber, typically recent activity, including but not limited to: online searches conducted by the subscriber; content accessed by the subscriber online or on television; online advertisements viewed, clicked on, or otherwise accessed by the subscriber; online shopping or purchases made by the subscriber; and any other form of previous subscriber activity.

[0021] Central Ad Server (CAS) – a computer server that manages delivery of online advertisements to visitors of online sites. A local ad server can be typically run by a single online publisher to serve ads to visitors of websites of that publisher's Internet domains, or by a single advertiser to serve ads into ad space acquired by the advertiser on various other websites. A third-party or remote ad server typically is remotely located and delivers advertisements of various advertisers to visitors of websites of multiple domains owned by multiple publishers. The remote ad server acts as a central conduit for delivering advertisements, enabling advertisers and publishers to track the distribution of their online advertisements and to control the rotation and distribution of their advertisements across the Internet from one location. The advertisements can be

stored on the CAS for later delivery, can be transmitted to the CAS and then delivered from the CAS upon receiving an ad request, or can be delivered from another source in response to an ad request received and routed by the CAS. Examples of third-party ad servers include DoubleClick's DART for Publishers central ad server (also known as DFP) and DoubleClick's DART for Advertisers central ad server (also known as DFA).

[0022] Profile provider – An entity that collects profile information that is used to target advertisements. In the context here, the profile provider cooperates with a CAS, which receives all or part of the collected profile information from the profile provider for use in targeting TV advertisements. User profile information derived from online activity can include observed online behavior of a user accessing the Internet or demographic information collected from a user accessing the Internet. Examples of profile providers can include, but are not limited to, any entity that owns or uses: (1) a visited Internet site server; (2) a server delivering content, images, audio, video, text, or any combination directed to an online user interface device (such as a computer or other online interface device) via an online access device (such as a modem or router), either directly or indirectly (e.g., via a redirect); (3) a server delivering an ad to an online user interface device via an online access device on behalf of an advertiser or an ad network; (4) a server recording an activity conducted from an online user interface device such as a click on an ad or a link to an ad, a viewing of an ad, a click on a link to particular content, a search, a request for product information, receipt of particular content, a product purchase, a telephone call made, or any other selected and definable user activity; or (5) a server facilitating instant messages or any other kind of communication on behalf of the user. Another example of a profile provider is: (6) a company sponsoring and having access to a computer program located on the user's computer or other online user interface device that can observe the user's online activity (with the user's permission), such as a browser toolbar or desktop search software. A profile provider, broadly, can be: (7) any entity able to collect behavioral profiles (observed online activity) or demographic profiles (provided by the user), preferably for purposes here including the IP address used when the profile was observed or collected and the date and time the profile was observed or collected, regardless of whether or not the entity collected a given profile directly through contact with the user's computer or indirectly from another entity such as those listed in this paragraph. In some cases, a

user's online activity will result in direct contact between the online user interface device via an online access device and the profile provider, e.g., if the profile provider is an online commerce site, the user makes a purchase at the site, and the online commerce site generates a profile for that user. In other instances there may be no direct contact between the profile provider and the user, e.g., if the user makes a purchase at an online commerce site that in turn reports information pertaining to the user to the profile provider. In some situations, also, a profile provider might also own or otherwise control a CAS, in which case user profiles can be immediately available to the CAS without need for transmission between separate entities.

[0023] Profiles or partial profiles provided by a profile provider to a CAS can contain any quantity of profile information, such as, in one example, just an online access IP address used by a person at the time his profile was collected and the identity of the profile provider. The IP address can be provided by the profile provider itself or might be obtained by the CAS when a user engages in any online activity or provides an item of demographic information and is redirected by a profile provider to the CAS. In another example, a profile can be more extensive and can include demographic or behavioral information, such as an extensive browsing history, shipping or purchase histories, content viewed, and other information concerning the user's characteristics or the user's activities. Although the profile provider is an entity, many or most of the actions attributed to the profile provider are actually performed by equipment under the administrative control of the profile provider, such as computers, servers, software running on those computers or servers, network connection hardware or software, or other equipment. Such actions may still be characterized as being performed "by the profile provider," whether performed automatically, semi-automatically, or manually.

[0024] Personally Identifiable Information (PII) – information that can be used to identify a specific person, including but not limited to: name, Social Security number (SSN), date of birth, street address, email address, static IP address (if any), phone number (home, work, wireless), financial account numbers (bank accounts, credit accounts, or any other financial data), driver's license number, vehicle registration number, vehicle license number, facial photographs, fingerprints, handwriting or signature, or any other information that can assist in identifying a specific person.

[0025] Non-Personally-Identifiable Information (non-PII) – information about a person that typically cannot be used to specifically identify that person, including but not limited to: city, state, or country of residence, age, gender, race, ethnicity, school or workplace (if sufficiently large), salary or income, hobbies, dynamically assigned IP addresses, online sites visited, online searches conducted, or other information that is useful to know about a person but done not by itself allow one knowing the information to identify the particular person.

[0026] Cookie – a text file placed on a user's computer by a server that also serves content to the user's computer using browser software. The cookie typically can be read or altered only by a server operating under the same Internet domain as the server that originally placed the cookie. The cookie file can be used to identify a computer that has already been in contact with the same domain and can also be used to store PII or non-PII pertaining to a user of that computer. In a first example, a cookie can store non-PII such as previous searches conducted at the site, or pages viewed or visited at the site, by the computer user. In a second example, a cookie can be used to store a username used by the user to access a site, customized preferences of the user, or various pieces of PII. It should be noted that a cookie file can also be created, altered, or deleted by software located on the user's computer.

[0027] Television advertisement (TV ad) – a full screen video ad, a partial screen video ad, a banner ad, a text ad, an audio ad, or any other form of advertisement suitable for delivery to and visual or audible presentation by a television set.

[0028] Various systems are used currently for targeting advertisements based on user/viewer/customer behavior. Many of these rely on the collection of personally identifiable information (PII) to enable correlation of the person exhibiting the behavior and advertisements targeted at that person. There are some examples wherein advertisements can be targeted without collecting PII, but in such examples it is typically the case that the medium of the behavior and that of the advertisement are the same. For example, many grocery stores hand out so-called "club cards" that need not be linked to PII. A shopper presents the card at checkout to receive various discounts, thereby allowing the store to link the list of purchased items to the card. As the system "learns" the shopper's purchasing habits, the system begins issuing coupons targeted at

at purchases that the shopper has made previously or that the system predicts the shopper may wish to make based on past purchases. In another example, online advertisements are readily targeted based on an Internet user's online activities without using PII. The use of cookies enables an ad server to recognize an Internet site visitor who has been previously presented with ads by the ad server or who has conducted searches or accessed content at sites linked to the ad server. The ad server can target future advertising to the site visitor based on that previous activity. A user who has searched for airline tickets to southern California on an online travel site, for instance, can later receive targeted online advertisements for Disneyland, delivered perhaps while visiting some other online site, to the user's computer from an ad server that collected the user's search information from the online travel site.

[0029] It becomes more difficult to avoid the use of PII when it is desired to target advertising in one medium based on activity in another. Various schemes currently are implemented, under development, or being considered wherein PII is used, e.g., to target television advertisements based on a viewer's online behavior. Some of those schemes involve agreements or alliances among television providers, online access providers, online search portals, or online sites. The PII has typically been required to make the connection between the different media, because different devices are typically employed to access each one.

[0030] A significant shortcoming of many such cross-media ad-targeting systems is the need to use PII to target advertisements delivered in one medium based on user behavior in another medium. Battelle (John Battelle, The Search, Portfolio, New York, 2005) describes the delivery of personalized ads to a DVR based on observed online behavior of the DVR user (conducted searches, sites visited, etc., including PII) when he or she was using a personal computer for online access and further describes financially rewarding the user for watching the ad (e.g., by reducing the fee for the TV service or by offering "free" TV service). However, many consumers object to the merging or correlation of personally identifiable information and online behavior. Financial incentives such as free or reduced-fee games, screen-savers, content, or Internet access in return for information useful for targeting ads requiring collection and use of PII and data pertaining to online behavior generally have not proven to be

attractive to the public in the past. Such combining of PII with online behavior data has proven particularly unattractive to consumers having larger incomes, who advertisers are especially interested in reaching.

[0031] An example of a cross-media ad-targeting system, planned to be implemented as a result of an announced alliance between Google and BSkyB, is a system wherein BSkyB can use searches conducted by their subscribers on the BSkyB Internet search portal to deliver targeted ads to the DVR's of those same subscribers. BSkyB already has PII pertaining to its television subscribers, so linking search results from its own search portal site may not seem too alarming to its subscribers. Nevertheless, PII is still required to target the television advertisements. However, many and perhaps most search sites and content sites accessed online are not owned or controlled by television service providers. In addition, many and perhaps most search sites and content sites do not typically collect PII to identify those users who access them online.

[0032] It is therefore desirable to provide systems and methods for delivering targeted television advertisements to users or viewers based on their online behavior (searches performed, sites visited, online ads viewed, an so on), but without using PII to link the television viewer to his or her online activity.

SUMMARY

[0033] In a method for delivering targeted television advertisements based on online behavior, IP addresses indicating online access devices and IP addresses indicating television set-top boxes are electronically associated, for a multitude of users. Using user profile information derived from online activity from one of the online access IP addresses, a television advertisement is selected, such as by using behavioral targeting or demographic information, and automatically directed to the set-top box indicated by the set-top IP address associated with that online access IP address. Preferably neither the user profile information nor the electronic association of online access and set-top box IP addresses includes personally identifiable information.

[0034] Objects and advantages pertaining to delivery of targeted television advertisements based on online behavior may become apparent upon referring to the exemplary embodiments illustrated in the drawings and disclosed in the following written description and/or claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0035] Figs. 1-5 and 6A-6C illustrate schematically a system for targeting television advertisements based on online behavior in which online access and television service are provided by a common provider.

[0036] Figs. 7-10, 11A-11C, and 12 illustrate schematically a system for targeting television advertisements based on online behavior in which online access and television service are provided by different providers.

[0037] Figs. 13 and 14 illustrate schematically a system for targeting television advertisements based on online behavior in which online access is provided by a mobile device.

[0038] The embodiments shown in the figures are exemplary and should not be construed as limiting the scope of the present disclosure and/or appended claims.

DETAILED DESCRIPTION OF EXEMPLARY EMBODIMENTS

[0039] In the following disclosure and appended claims, the term "user" shall be construed as one or more persons receiving online access or television service at a delivery end-point within a household, office, business, or other site or establishment served by an online access provider or television service provider. Thus, delivery of a television ad to "the user" based on access of an online site by "the user" might involve a single person, or one person might access the online site while the corresponding television ad might be delivered to another person in the household, office, business, or other site or establishment. The delivery end-point may be one television set or multiple sets of a given user; likewise, online access may involve one or more computers or other online interface devices of a given user. In some instances, a business having multiple physical locations may be served by separate online or television services, but in other instances, a business may have an internal LAN or WAN that extends service provided to multiple physical locations. Also, some computers and televisions are portable, and can access the service provided from remote locations. Accordingly, the term "user" is intended to refer to that person or those persons who receive the provided services and does not require a fixed or single location in such instances.

[0040] In the disclosed exemplary systems and methods for delivering targeted television advertisements based on online behavior, IP addresses indicating online access devices and IP addresses indicating television set-top boxes are electronically associated, for a multitude of users. Using user profile information derived from online activity from one of the online access IP addresses, a television advertisement is selected, such as by using behavioral targeting or demographic information, and automatically directed to the set-top box indicated by the set-top IP address associated with that online access IP address. Preferably neither the user profile information nor the electronic association of online access and set-top box IP addresses includes personally identifiable information. The online access and set-top box IP addresses can be electronically associated for many or all of the users of the multitude before directing television ads to any of the multitude of users. Alternatively, online access and set-top box IP addresses can be electronically associated for one or a few of a

multitude of users at any given time as needed for directing television ads to users, with additional IP addresses electronically associated at later times for directing other television ads to additional users.

[0041] In an exemplary system for serving targeted television advertisements based on a user's online behavior, the user receives both television service and Internet access from a common service provider that acts as both an internet service provider and a television provider. Accordingly, such a provider entity can be referred to as ISP/TVP. The basic layout of this scenario is illustrated schematically in Fig. 1, wherein online access device 32, online user interface device 34, set-top box 36, and television 38 are associated with a user 30. In some instances online access device 32 and online user interface device 34 are integrated in a single unit, while in other instances they comprise separate units. Likewise, in some instances set-top box 36 and television 38 are integrated in a single unit, while in other instances they comprise separate units. The common service provider 20 can offer Internet access via any suitable online access device 32 (a modem, router, or network adapter suitable for connecting to cable, DSL, wireless, satellite, Ethernet, or any other transmission system) along with the television service provided through set-top box 36 (STB 36). For purposes of the subsequent discussion, online access device 32 shall be described and shown in the figures as a modem and online user interface device 34 shall be described and shown in the figures as a computer separate from modem 32. These descriptions shall not be construed as limiting the scope of the present disclosure or appended claims.

[0042] Many users of online access service (i.e., subscribers) are not provided with a static IP address; instead a dynamic IP address (designated as IP_{modem} in the figures) is allocated to the user's modem 32 for online access, typically through DHCP. The online access IP address is allocated by the Cable Modem Termination System (CMTS) if ISP/TVP 20 is a cable company, whereas the online access IP address is allocated by the Digital Subscriber Line Access Multiplexer (DSLAM) if ISP/TVP 20 is a phone company. Other types of online access providers allocate dynamic IP addresses in a similar manner using analogous equipment. It is to the modem IP address (i.e., the online access IP address) that traffic from a remote network is routed to the modem 32

and hence to other devices sharing the modem connection to the remote network; the modem IP address is the only address "seen" by the remote network. The dynamically assigned online access IP address is replaced after some time interval by another IP address allocated to modem 32. ISP/TVP 20 transmits (as in Fig. 2) to a Central Ad Server (CAS 40), or perhaps to multiple CASs, the online access IP addresses of a multitude of users along with corresponding pseudonyms or aliases associated with the IP addresses, e.g., user XY123 is connected to the Internet through a modem 32 that is assigned IP address aaa.bbb.ccc.ddd. Each online access IP address and its associated pseudonym can be transmitted to the CAS 40 immediately upon its allocation or reallocation by ISP/TV 20, which could be advantageous in sending TV ads to the user's STB quickly. Alternatively, online access IP addresses and the corresponding pseudonyms can be transmitted periodically to CAS 40 at any necessary or desirable time interval. The pseudonyms and online access IP addresses are non-PII. The multitude of users can be any set made up of a large number of users and does not necessarily include every user having online access or television service through a particular company or provider. For example, the system described can operate only with respect to a subset of users, such as those users who have capabilities needed to implement this system (e.g., suitable hardware, software, operating system, etc.), those who have done some sort of subscription, or those selected based on criteria as to which operation of the system is considered desirable, less expensive to implement, or profitable. For example, those users might be chosen that have STBs with an operating system, software, or hardware capable of receiving the software required to accept TV ads from a CAS. In any event, the fact that other users may exist as to which the system does not operate is not intended to negate the advantages of the system as to those users for which the system does operate.

[0043] Because the user pseudonym and online access IP address are associated with the modem 32, multiple computers or other devices sharing the modem connection for online access will share the online access IP address and the user pseudonym. When the user's STB 36 is used to access a web-based server (for accessing an online programming guide, for example, through the cable television transmission system or through a connection to the Internet), ISP/TVP 20 additionally transmits (as in Fig. 2) to the CAS 40 the IP address allocated to the STB (designated as IP_{STB} in the figures) and

associates it with the pseudonym previously associated with the user's modem 32, e.g., user XY123 has an STB at IP address eee.fff.ggg.hhh in addition to a modem for online access at IP address aaa.bbb.ccc.ddd. The association of the user's online access IP address (i.e., modem IP address) and the user's STB IP address at CAS 40 constitutes non-PII. ISP/TVP 20 knows the IP address of STB 36 because the web based server accessed by STB 36 is typically operated by ISP/TVP 20. If STB 36 has its own modem (an increasingly common scenario as the transition is made from IPv4 to IPv6), the IP address of STB 36 is allocated by ISP/TVP 20 via the CMTS or DSLAM or other similar system, typically using DHCP. Alternatively, in the event that STB 36 accesses the web-based server via a connection 35 to the same modem 32 that provides the user's online access (as in Fig. 3), STB 36 will have reported to CAS 40 the same IP address that was reported for the user's online access modem 32. In all of those cases, the CAS 40 can now associate the IP address for the user's online access (modem 32) with an IP address for the user's STB 36, and the association is not PII.

[0044] ISP/TVP 20 reports to CAS 40 that user XY123 is accessing the Internet using online access IP address aaa.bbb.ccc.ddd and has STB 36 using STB IP address eee.fff.ggg.hhh. These associations can be stored on CAS 40, or only by ISP/TVP 20, as desired. If user XY123 engages in online activity from IP address aaa.bbb.ccc.ddd (for example, accessing an online site 50, as in Fig. 4), then online site 50 can electronically communicate (as in Fig. 5) an indication of the user's online activity to CAS 40, which in turn can cause a television advertisement to be directed to the STB IP address that is associated with online IP address aaa.bbb.ccc.ddd. In this example the online site 50 acts as a profile provider. Without relying on PII, CAS 40 knows that IP address aaa.bbb.ccc.ddd is associated with user XY123 and that user XY123 has a set-top box with IP address eee.fff.ggg.hhh, and CAS 40 causes the television advertisement to be directed to IP address eee.fff.ggg.hhh that indicates STB 36 (as in Figs. 6A-6C). In some implementations, the television advertisements are delivered to STB 36 using an ad server controlled by ISP/TVP 20, using either an IP-based protocol (an example of which is shown in the referenced figures) or a television delivery system (e.g., a dedicated cable advertisement channel). In one such implementation, that ad server report to CAS 40 the user pseudonym (XY123) and CAS 40 provides that ad server with TV ads targeting user XY123 based on profile information collected from IP

addresses reported by ISP/TVP 20 as belonging to XY123. In another such implementation, that ad server reports the IP addresses used by each of a multitude of users of Internet services provided by ISP/TVP 20 along with the times and dates that those addresses were used, and CAS 40 provides that ad server with TV ads targeting each of the multitude of Internet users based on profile information received from the reported IP addresses at the corresponding reported times and dates. Personal privacy is maintained in these implementations because (1) no PII is used by CAS 40 in these implementations, and (2) although ISP/TVP 20 has the capacity to identify the users, the ad server of ISP/TVP 20 does not learn the profile information collected for and received by CAS 40.

[0045] The particular television advertisement thus directed typically is selected for delivery to the user's STB 36 based on the user's online behavior (e.g., merely accessing the online site 50, viewing a particular web page, performing a particular search, clicking on a particular link or ad, online shopping or purchases, or other activity). Information about observed online behavior (i.e., indications of online activity) or collected demographic information can be used by the CAS or by the profile provider to direct a targeted TV ad. In the latter case, the CAS has just the IP address of the profiled user and the identity of the profile provider but not the observed online behavior or demographic information collected by the profile provider, so the CAS can direct a TV ad to the user merely by enabling the profile provider to arrange for the delivery of a targeted TV ad based on the observed online behavior or demographic information available to the profile provider. Examples of profile providers are given above.

[0046] The profile provider (such as the proprietor of online site 50) can send an electronic communication to the CAS, preferably automatically such as by having the profile provider's server pre-programmed to send such electronic communication, which electronic communication contains information that the CAS will use in causing a selected TV ad to be transferred to the STB at the proper time. The electronic communication received by the CAS contains information derived from the interaction of the user with an online site (i.e., online activity), which can be online site 50 or another online site for which the profile provider has access to data about user interactions. The electronic communication can take a variety of forms, including: (1) an indicator of

an indicator of the TV ad that should be shown, such as an URL or ad title given by the profile provider to retrieve the correct TV ad from the profile provider's ad server, or even the actual TV ad; (2) an indicator of information about the user (likely not PII), such as (a) an IP address pointing to the user, preferably with time and date of visits; (b) a description of profile information about the user, or (c) a code or keyword used by the profile provider to fetch a visitor's profile from a database; or (3) an indicator describing user activity, such as a code or keyword indicating (a) certain categories of users, such as users likely interested in mortgages or users likely interested in travel, etc., or (b) more simply, what the visitor searched for on the online site, such as "car insurance." The electronic communication can be sent for each instance of a user interacting with the online site, or a list can be created and transmitted containing information about a number of user contacts. The electronic communication can include any combination of items 1-3 above; for example, a list of visitor IP addresses (preferably with date/time of visits) can be usefully paired with an indicator, in the form of a URL or code, identifying specific TV ads that the profile provider would like to deliver to each of the listed visitors.

[0047] Once requested, the television advertisement can be transmitted to the user's STB in a variety of ways. The advertisement can be previously stored on the CAS and transmitted directly from the CAS to the user's STB (Fig. 6A). The advertisement stored on the CAS can be transmitted from the accessed online site, from another online site, or from another ad server and stored on the CAS as part of the transmission request from the accessed online site, and then transmitted to the user's STB. The advertisement can be transmitted directly to the user's STB from the accessed online site (Fig. 6B) or from another online site or ad server 60 (Fig. 6C). The user's STB can be instructed by the CAS to retrieve the advertisement directly from the accessed site, from the CAS, from another online site, or from another ad server, using for example the URL given to the CAS by the profile provider. In any of those alternatives, the television advertisement need not be transmitted directly to the STB; instead it can be transmitted through one or more intermediate servers (i.e., a server hosted by an ISP, cable company, or telecommunications company, for example). The television advertisement can be streamed or otherwise delivered in real time, or it can be delivered to the user's STB or an associated DVR for later viewing.

[0048] As IP address are dynamically assigned and reassigned, ISP/TVP 20 reports updated online access and STB IP addresses to CAS 40 (as in Fig. 2). The IP addresses and their associated pseudonyms can be transmitted to the CAS 40 immediately upon allocation or reallocation by ISP/TV 20, or can be transmitted periodically to CAS 40 at any necessary or desirable time interval (fixed or variable). Although assignment of a pseudonym associated with both the online access IP address and the STB IP address can facilitate delivery of television ads to STB 36 based on the user's online behavior (via modem 32), it is not necessary. All that is required is an association between the online access IP address and the STB IP address to enable delivery of television ads based on the user's online activities without relying on PII.

[0049] In another exemplary system for serving targeted television advertisements based on a user's online behavior, the user need not receive both television service and Internet access from a common service provider. The basic layout for such a system is illustrated schematically in Fig. 7, with modem 32 connected to Internet service provider (ISP) 22 for providing online access, and STB 36 connected to television provider (TVP) 24 for providing television service. Again, modem 32 and computer 34 are described as examples of online access and interface devices, respectively; the use of these specific examples shall not be construed as limiting the scope of the present disclosure or appended claims. An online access IP address is assigned to modem 32 as described above. In cases where the user's STB 36 can also obtain online access via the same IP address as the user's online access (i.e., via modem 32), that common IP address comprises the electronic association that can be used to perform targeted television advertising without requiring PII. That case can occur by connecting STB 36 to the user's modem 32 via connection 35, in addition to connecting STB 36 to TVP 24. For example, STB 36 can be connected to a local area network (LAN) that shares device 32 for online access. In such an arrangement, network traffic is routed to a common IP address (i.e., the IP address of modem 32) for STB 36 and other devices connected to the LAN. The user can receive online access from any online service provider via any suitable modem 32 (for cable, DSL, wireless, satellite, or any other transmission system) or other suitable online access device and can receive television service from any suitable television service provider through STB 36. The term

"common IP address" denotes the IP address of a modem or other online access device to which remote network traffic is routed for each of multiple devices sharing a common connection to the remote network through that modem, e.g., the IP address of a modem connecting to the Internet a router in a local area network (LAN) would constitute the common IP address for multiple devices connected to the LAN for online access, even though those multiple devices typically also have individual IP addresses assigned by the router for routing traffic within the LAN. An online site or server accessed by any one of the multiple devices only "sees" the common IP address (the modem IP address).

[0050] In an exemplary implementation, a user's STB 36 can periodically contact a CAS 40 (or multiple CASs) via the IP address of the user's modem 32 that is common to STB 36 (as in Fig. 8). The common IP address is dynamically assigned by ISP 22 to the user's modem 32 or other online access equipment. Modem 32 can include a router to route network traffic properly among the user's computer (or computers or other devices) 34 and the user's STB 36. Alternatively, a separate router can be employed, or computer 34 or STB 36 or television 38 can include routing functions, for properly routing network traffic passing through modem 32, with STB 36 providing routing functions or being connected to the separate router, to computer 34, or to television 38 as appropriate. STB 36 reports the common IP address to CAS 40 with which it is in contact. However, STB 36 itself might not know the common IP address it is using, because it communicates through modem 32. In such circumstances, STB 36 nevertheless can report the common IP address to CAS 40 simply by initiating communication through modem 32 to CAS 40, thereby enabling CAS 40 to extract the common IP address from the communication. The common IP address enables delivery of television ads to STB 36 based on access of an online site via the common IP address (or based on other user profile information derived from online activity observed as originating from the online access IP address of modem 32, i.e., from the common IP address). The existence of the common IP address associates the online access IP address and the set-top box IP address. Again, this does not constitute PII. The time interval for the periodic contact of the STB and the CAS can be set to any appropriate time period, e.g., every five minutes, every hour, or at some other necessary or desirable time interval. The time interval can be selected to be

commensurate with a typical or average time interval for dynamic reassignment of the online access IP address by the user's Internet Service Provider. The selected time interval can be made adjustable if needed or desired. Alternatively or additionally, the contact between the STB and the CAS may not be set to occur at a fixed time interval but rather irregularly, such as instances where it is triggered by dynamic reassignment of the online access IP address by the ISP.

[0051] When a user engages in online activity (for example, accessing online site 50, as in Fig. 9), user profile information derived from the user's online activity from the online access IP address can be delivered electronically to CAS 40 (from the accessed site 50 as in Fig. 10, or from another profile provider), which in turn can cause a targeted television advertisement to be directed to the user's STB, if there is an STB IP address associated with the online access IP address. In the example shown in Fig. 9, online site 50 acts as a profile provider, and the user's online access IP address (the common IP address), from which the online site was accessed by the user, is included with the electronic delivery. CAS 40 causes the television advertisement to be directed to STB 36 at the common IP address (as in Figs. 11A-11C), in accordance with the alternative methods for delivery of the TV ad discussed herein in connection with Figs. 6A-6C. The particular television advertisement thus transmitted typically is selected for delivery to the user's STB 36 based on the user's online behavior (examples of various sorts of online behavior, or demographic information provided or determined online, are discussed above). The accessed online site 50 transmitting the electronic communication (examples of which are listed above) can be any profile provider, or any profile provider having access to information about user activity with respect to online site 50 can transmit the electronic communication (examples of alternative types of profile providers are also given above).

[0052] The television ad can be directed to the STB that most recently reported to CAS 40 as the common IP address the online access IP address from which online site 50 was accessed. In that case, the IP addresses can already be associated when the user profile information is used or the association can be created when the user profile is received or afterwards. Or the television ad can be directed to the STB that next reports that IP address to CAS 40 as the common IP address. In that case, the IP

addresses might not be associated until the STB and CAS 40 are next in communication with each other and the STB IP address is received by CAS 40. Associations between STBs and online access IP addresses could be stored for future usage or could be created in real-time when a STB communicates with a CAS to receive TV ads and not stored for future usage (i.e., the association is dropped as soon as the TV ads are delivered to the STB). If the periodic contact interval is sufficiently small relative to the typical or average time interval between reallocation of IP addresses, then it is likely that those two IP addresses will be the same and that the television ad will be directed to the correct STB. There will be a small fraction of television ads that are misdirected, if the online access IP address is reassigned between the accessing of the online site (i.e., deriving the user profile information) and the relied-upon contact between the user's STB and the CAS (i.e., associating the IP addresses). That fraction can be reduced by reducing the periodic time interval between STB reports of its common IP address to the CAS or configuring the STB to report changes to the common IP address (i.e., to enable updating of the association between the IP addresses). Associations between online access and STB IP addresses can be stored for future use, or the IP addresses can be associated in real-time only when an STB communicates with a CAS to receive TV ads and not stored for future use (e.g, the association can be purged upon directing the TV ads to the STB).

[0053] Instead of relying on or waiting for contact from the STB reporting its common IP address to one or more CASs, the online site accessed by the user could instead query one or more CASs to find an STB with a common IP address matching the online access IP address from which the online site was accessed. If a matching STB IP address is found, a request is transmitted to the corresponding CAS to transmit a television advertisement to the STB. If the query, match, and request are completed sufficiently quickly, then the probability can be kept low that a television ad will be directed to the wrong STB.

[0054] In another exemplary implementation, the user's STB 36 can be assigned an identifier, which need not include PII. Such an identifier can be permanently or temporarily assigned to STB 36. Examples of such identifiers could include an equipment serial number or a portion thereof, a device MAC address or a portion

thereof, a username, a pseudonym, a confirmation number, a tag, or any other identifying code or data element. The CAS can place an identifier on the STB the first time it is in contact with the STB and read or amend the identifier every time afterwards the STB and the CAS are in contact. Such an identifier is an analogy of a “cookie” used by online servers to recognize visitors’ computers. The identifier can include the IP address the STB is using when in contact with the CAS, so the next time the CAS and the STB are in contact, the CAS could compare the IP address then used by the STB with the previous IP address used by the STB without looking in a database. From a privacy perspective, there is an advantage to using an identifier placed by the CAS on the STB, because it enables the user to control the identifier from user’s TV screen and set privacy settings to match user preference. A user may, for example, erase the identifier, alter the identifier to inform the CAS the next time it is in communication with the STB that the user is not interested in targeted ads, interested only in ads targeting online behavior from the last seven days, interested in ads in certain categories only, or interested in all types of ads except for certain excluded categories (e.g., adult content). In addition to reporting the common IP address to CAS 40, STB 36 can also report the identifier. When an online site 50 is accessed from the common IP address, online site 50 requests that a television advertisement be transmitted to the STB having the common IP address. The identifier of the STB most recently reporting the common IP address before the request is compared to the identifier of the STB that next reports that IP address to CAS 40 after the request. If the compared identifiers match, then it is nearly certain that the STB corresponds to the same user online access IP address from which online site 50 was accessed and the television advertisement is correctly directed. If a unique identifier is used (such as a complete MAC address or a complete device serial number) then misdirecting the television ad is virtually impossible. Even if the identifier is not unique (e.g., a partial MAC address or serial number, or a simple four-digit number), it is still highly unlikely that a given IP address would be reallocated between STB’s that happen to have the same identifier. Use of an identifier can enable accurate directing of television ads to the correct STB while requiring the STB to report its IP address to the CAS less frequently, reducing network traffic and other resource requirements.

[0055] In another exemplary implementation, instead of reporting the common IP address of STB 36 to CAS 40 relatively frequently, STB 36 can be configured to store its common IP addresses and corresponding times and dates that each of the addresses is assigned to modem 32 (and hence also to the LAN to which the STB 36 is connected). The common IP address can be periodically sampled and the sampling stored (e.g., every 1, 5, or 10 minutes, or some other suitable interval), or the specific times and dates of IP address reassignments can be stored with the corresponding IP addresses. STB 36 can be configured to periodically (e.g., every 4, 6, 12, or 24 hours, or some other suitable time interval) report to CAS 40 all of the common IP addresses assigned to modem 32 and the corresponding dates and times they were used for some preceding span of time (e.g., 1, 7, 14, 30, or 90 days, or some other suitable interval). Alternatively, the common IP address can be stored on the STB as part of a cookie or tag set and subsequently read or reset by CAS 40. CAS 40 also can be configured to receive IP addresses and corresponding access times and dates for all visitors to an online site 50 that wishes to deliver television advertisements based on online user behavior. Online site 50 periodically reports to CAS 40 (as in Fig. 10) all IP addresses from which the online site was accessed during some preceding period, along with the corresponding times and dates. The frequency and time periods of the online site's reporting can be similar to those of the STB reporting, or they can differ substantially, as needed or desired. The two lists (one from the online site, one from the STB) of IP addresses are compared, and delivery of a television advertisement is requested if any matches are found, i.e., whenever it is found that the STB was connected to a modem that was assigned an IP address at a time and date when the online site was accessed from that IP address. Such cross-matching ensures accurate delivery of the television advertisement and requires substantially less communication between the STB and the CAS.

[0056] The common IP addresses (and corresponding times and dates) can be stored on the CAS. Alternatively, privacy can be enhanced in this implementation by storing the common IP addresses (and corresponding times and dates) locally on the STB, instead of on the CAS, so the only device that monitors the different IP addresses assigned to the same access device is the STB. If stored, whether on the STB, on the CAS, or elsewhere, the list of common IP addresses can be purged. Purging of the IP

addresses can be done periodically, at any suitable time interval, and can be manual or automated. Similarly, the list of online access IP address (and corresponding times and dates) from which an online site has been access can be purged. The STB can be employed to set privacy options, thereby enabling the user to control at least partly if and how TV advertisements are targeted. For example, the user might choose whether or not, or for how long, the CAS or STB will record and retain the IP addresses assigned to the user's modem, or can choose to erase all presently stored IP address data.

[0057] Typically a router does not provide the IP address of the cable, DSL, or other modem to the computer(s) or device(s) connected to it or to a LAN managed by it. STB 36 requires this information to store it for later reporting to CAS 40. One method for enabling this functionality includes providing an Internet server (perhaps as part of CAS 40, or perhaps an independent server) that receives periodic requests from STB 36 and replies to STB 36 with the common IP address (i.e., the IP address of modem 32, through which STB 36 is provided with online access). Any other method for monitoring the common IP address can be employed, including configuring the modem 32 to directly report its IP address to STB 36. STB 36 periodically monitors the IP address of the modem 32 at a selected time interval (described above), and stores the modem IP address (i.e., the common IP address) and the corresponding times and dates each common IP address was allocated to the user's modem.

[0058] In any embodiment or implementation in which STB 36 reports IP address data to CAS 40 (either stored on STB 36 or for storage at CAS 40), that communication can be achieved in several ways. STB 36 can report the common IP address through the online access modem 32. Alternatively, STB 36 can report to CAS 40 through a dedicated upstream data channel 37 that is part of the television system (as in Fig. 12). The reported IP address data need not be transmitted directly to CAS 40; instead it can be transmitted to an intermediate server (e.g., hosted by an ISP, cable company, or telecom company).

[0059] It should be noted that some STBs have a dedicated modem and typically would not need to obtain online access through the user's online access modem to deliver television service to the user. However, other functionality of an STB includes,

or is expected to include, delivery of audio, video, images, and other content from a user's computer 34 or other networked devices to the user's television 38. To provide that functionality, STB 36 must be connected to the same LAN or router as the user's computer(s), even if the STB has its own modem. Although using its own modem for most of its functions, such an STB would still have additional online access through the user's online access modem 32, enabling the user's online access and STB to have a dynamically assigned common IP address (as described above). The STB can report the common online access IP address through the online access modem 32 at the common IP address or through the STB's dedicated modem at its own IP address. The television advertisement transmitted to the STB can be received through the online access modem 32 at the common IP address or can be received through the STB's dedicated modem at its own IP address (if, for example, the CAS is provided with the STB modem IP address along with the common IP address).

[0060] However the user's online access and STB IP addresses are electronically associated, the association can be maintained even if the online user interface device 34 is a mobile device (e.g., an Internet-enabled cell phone, handset, PDA, or laptop computer) that is intermittently disconnected from modem 32 and is connected to the Internet through another connection (e.g., another LAN, a wireless hotspot, or a cell phone network, as in Figs. 13 or 14). When connected to the Internet through modem 32 (as in Figs. 1 or 7), online sites visited using such a mobile device 34 can request delivery of television advertisements to STB 36 in any of the ways described above. Mobile device 34 can be programmed to store, in much the same manner as described above for STB 36, its online access IP addresses and the corresponding dates and times when disconnected from modem 32 and connected to the Internet by another connection (i.e., when "roaming"). When mobile device 34 is later reconnected to modem 32, those stored roaming online access IP addresses, along with the corresponding dates and times, can be transmitted to CAS 40 as described previously. Television advertisements can be directed to STB 36 based on matches found between the roaming online access IP addresses (and dates and times) from mobile device 34 and online access IP addresses (and dates and times) reported by online site 50 to CAS 40 for visits to online site 50.

[0061] Furthermore, the selected television advertisements can be directed to STB 36 even after a mobile device 34 is disconnected from modem 32 and begins accessing the Internet through roaming online access IP addresses. A tag or cookie placed on mobile device 34 (or stored in a database elsewhere, such as on CAS 40, by CAS 40, or somewhere accessible on STB 36 or its associated LAN), which tag or cookie can include online access IP addresses (and corresponding dates and times) of modem 32 or STB 36. Those IP addresses (and dates and times) can be reported to CAS 40 by mobile device 34 while roaming (following a redirect from visited online site 50 for example), and a television advertisement can be directed to STB 36 (using the IP address reported by the roaming mobile device 34) at the request of online site 50. Because a new IP address might be assigned to modem 32 and STB 36 after roaming device 34 is disconnected from modem 32, an identifier for STB 36 (as described earlier) can be employed to verify that the television advertisement is directed to the correct STB. The identifier can be stored on mobile device 34 as part of the cookie or tag.

[0062] However the user's online access and STB IP addresses are electronically associated, the STB can be configured to filter the targeted television advertisements based on user preferences. For example, a user might program the STB to reject certain ad categories (e.g., women's clothing, or health ads) and accept certain other ad categories (e.g., automobile ads, electronic equipment ads, or travel ads). The filtering can be configured as "opt in" (i.e., receive ads only in selected categories) or "opt out" (i.e., receive ads in all but selected categories). In another example, a user might filter ads based on how long ago an online site was visited; e.g., the user might choose to accept ads only from Internet sites visited in the last seven days. Such filtering or preference data can be applied before or after an ad is selected to be transmitted. The STB can be configured to notify the CAS or the online site to prevent selection, or transmission, of the television advertisement, or the STB can accept the advertisement but then not display it to the user. In an exemplary implementation, such user preference or filtering information can be stored in the STB as part of a cookie or tag set and subsequently read or amended by the CAS. User preference or filtering information can be stored in any suitable location or format.

[0063] In another implementation, no software is required to be installed on the mobile device, only cookies are placed on the device, which makes implementation in the market much simpler. Once a STB is confirmed to have been using a certain IP address at a certain date and time (for example by having a uniquely tagged STB consecutively communicate with the CAS using the same IP address or by the STB reporting the common IP address recorded by it as used by the online access device), the CAS database can be searched to find all the computers, mobile devices, and other STBs that used the same IP address at the same date and time, and those will be assumed to be devices that are part of the same household and using the same LAN. Then, an association or link is created among the devices, by linking the device identifiers or tags with the STB's identifier in the CAS database or by placing linking codes or information identifying the linked devices within the identifiers (cookies) of each of the devices along with the date and time the association was created. The association is kept for no more then a pre-determined period of time (such as seven days), to ensure that the delivered TV ads will be based only on online activity of those people that live in the house and not on online activity of people that happened to use the house's LAN only once, such as houseguests.

[0064] Once the association is created, profile information (both behavioral and demographic) collected about the computers and mobile devices before and after the association was created could be used to target ads to the STB, regardless of the IP addresses used by the computers and mobile devices (such as cell phones) to access the web at the time the profiles were collected.

[0065] For example, consider a common scenario where a laptop, used both at home and in the office, is linked to a home STB as a result of it being used at home. The day after the association was created, the laptop is taken to work, where it is used to visit a home mortgage site (using a different LAN or online access device than the one used at home). The home mortgage site redirects the laptop to the CAS along with a request to transmit a mortgage-related TV ad from one of the site's advertisers to the laptop user. The CAS, reading its own cookie that was placed on the laptop, checks its database and finds that the laptop is linked to a certain STB. In the evening, when the laptop user returns home and turns on the TV, the user will see the mortgage-related TV ad, even if

TV ad, even if the user left the laptop at work that day.

[0066] In another example, suppose a laptop user using a laptop at work one morning views a home mortgage advertiser's ad, before any association between the laptop and a home STB is created. When the mortgage advertiser ad server redirects the laptop to the CAS with a request to transmit a TV advertisement to the user, the CAS places a cookie on the laptop, or reads its own cookie if it encountered the laptop before. The CAS, finding no association between the laptop and a STB, stores the IP address used to view the advertiser's online ad, the CAS cookie used to mark the laptop, the date and time the ad was viewed, and a link (e.g., a URL) to the advertiser's TV ad. In the afternoon, suppose the laptop user returns home with the laptop. When the user accesses the Internet via the home's LAN with the laptop, suppose the user visits a travel site interested in delivering travel-related TV ads to the user from the site's advertisers. When the travel site redirects the laptop to the CAS, the CAS reads its own cookie and records the IP address used by the laptop to access the travel site, the CAS cookie used to mark the laptop, the date and time the site was visited, and links to TV ads (sold by the site to its travel advertisers). Next, the home STB is in contact with the CAS in the evening to check whether there are any TV ads available for its user. The CAS then uniquely identifies the STB using a tag it previously placed on the STB and recognizes that the STB is using the same IP address as the STB did several days ago. The CAS checks which profile providers were visited by computers and mobile devices using the same IP address since the STB's last contact. The CAS finds the laptop visit to the travel site using the common IP address. The CAS further finds the laptop's view of the online mortgage advertiser's ad, using the cookie used by the CAS to mark the laptop. The CAS can then redirect the STB to fetch the mortgage TV ad, using the link provided by the mortgage advertiser (despite the fact that the online mortgage ad was viewed by the laptop before making the association with the home STB) and redirect the STB to fetch the travel TV ads from the links provided by the travel site.

[0067] An Internet site can determine in one of several ways which CAS to contact to arrange for targeted delivery of television advertisements to a visitor of the online site:

[0068] (i) Each ISP (or ISP/TVP) has a block of IP addresses from which it dynamically allocates IP addresses to users using DHCP. Based on the online access

IP address of a visitor to the site, an Internet site can determine which ISP the visitor is using to access the site. The site can then transmit a request to the correct CAS for that ISP. The request could be transmitted by redirecting the user to the correct CAS, for example.

[0069] (ii) The CAS or the CAS owner can transmit to participating Internet sites a file that includes a list of online access IP addresses that the CAS currently recognizes. Updated files can be transmitted at suitable intervals. An Internet site can receive such files from multiple CASs, and according to the information in the files can determine the proper CAS from which to request that an ad be served to the Internet site visitor.

[0070] (iii) The Internet site can request ads from all or most major CASs, relying on the probability that at least one of them may have registered the IP address allocated to the site visitor by the ISP that provides the site visitor's online access. If that is the case, that CAS can transmit to the Internet site a confirmation that the CAS recognizes the site visitor's online access IP address. The Internet site can deliver its request by redirecting a visitor to all or most major ISPs that operate a CAS, for example, or by sending all or most of the major CAS servers a file with the IP addresses and corresponding times and dates of visitors to the site.

[0071] The CAS that receives a request to deliver a television ad to an Internet site visitor might already have electronically associated an STB IP address with the visitor's online access IP address. If not, the CAS can store the site visitor's online IP address, the identity of the requesting Internet site, and the time and date of the visit to the Internet site. If the site visitor's STB later accesses the Internet server and the ISP/TVP transmits the site visitor's STB IP address to the CAS, the CAS can then arrange delivery of requested television ads. Or if the site visitor's STB reports a common IP address that matches the site visitor's IP address, the CAS can then arrange delivery of requested television ads.

[0072] Although the CAS is described as an entity separate from the profile provider, it could well be owned by or otherwise controlled by the profile provider. For example, if a profile provider is an online site (a large site such as Google could be a good example), in that implementation no redirect (or other inter-party communication) to the CAS will be required because the CAS is the server collecting the profile on behalf of the online

site. For example, Google's CAS could collect the IP addresses of users who search on Google, what those users search for, at what times and dates the searches took place, and possibly the Google cookies used to mark those users' computers. Due to the very large amount of profile information about so many people accumulated by Google, the Google CAS could be used to serve TV ads to STBs. An STB communicates with the Google CAS to check whether a STB user has used another computer or mobile device to search on Google (the Google CAS uses the STB's common IP address to find those other computers or mobile devices, as explained above), and if Google finds such searches were made, it arranges for delivery of targeted TV ads, sold by Google to its advertisers, to the STB of users who made such searches.

[0073] As a security precaution, many residential or home office network routers are set by default to block access from the Internet to devices connected to router. In this default mode, television advertisements can be transmitted to the STB through the modem/router only upon if an initiated request is transmitted by the STB. The STB can be configured to query the CAS periodically for television ads awaiting transmission to the STB, and in response the CAS can either transmit the ad or direct the STB to retrieve the ad from the accessed online site or from another online site or ad server. Alternatively, enabling "Port Forwarding" on the router would allow direct access from the Internet to a STB connected to the Internet through the modem/router, thereby enabling the CAS or online site to transmit, at its own initiative, an advertisement to the STB.

[0074] However the transmission of the advertisement is initiated, and whatever the source of the transmitted advertisement, the STB can save the advertisement (in its own memory or storage or in a DVR) for viewing at a later time by the user (perhaps repeatedly), or the advertisement can be presented in "real-time" or "streaming" format. The advertisement (stored or real-time) can be presented independently, or can be integrated into other television content. Any suitable procedures for transmitting the targeted television advertisement to the online user's STB can be employed.

[0075] An online site or other profile provider can provide to the CAS a profile of a visitor, in return for a fixed payment or a commission based on revenue generated by a

television advertisement transmitted to the profiled visitor. In such a case the transmission request received by the CAS from the profile provider includes the visitor's profile (searches performed, pages viewed, provided demographic attributes, etc.; typically not PII). The CAS can match the profile with one or more other television advertisements to be delivered to the visitor's STB (in addition to those requested by the online site). Such additional ads can be sold by the CAS or by another party, and can be transmitted directly from the CAS or pulled by the STB from another site or ad server to which the CAS redirects the STB.

[0076] As described previously herein, it is often the case that more than one computer or other online user interface device is connected to the Internet through a single modem or router, and all devices sharing that connection will appear to an online site to be connected through the same IP address at any given time. For example, in a home scenario, a parent and a child in the same household might each use their own computer to access the Internet, with the parent visiting sites pertaining to news or finance and the child visiting sites pertaining to pop music or television. (The same is true in a business setting, where many computers may access websites through a common modem or router.) Each such site visited by any person on the LAN reports to the CAS the IP addresses and corresponding times and dates of each of those visits (which are the same at any given time for both computers). The STB in the same household reports to the CAS the common IP address used for online access. The CAS can now match the STB to the various sites visited. Television advertisements can be selected based not only on requests made by the visited sites (as already described) but also on content being viewed on the television when the TV advertisement is delivered. For example, if programming is being viewed on the television that is typically targeted at adults (e.g., news, financial reports, or adult drama) then television advertisements requested by the news or finance online sites are presented. Alternatively, if programming is being viewed on the television that is more typically targeted at children or teenagers, then television advertisements requested by the pop music or fashion online sites are presented. Both types of ads can be directed to the STB and stored for later viewing (during viewing of suitable television content) or can be directed to the STB and viewed as the suitable television content is viewed. The same discrimination between types of sites visited using multiple online access

multiple online access devices can be employed regardless of whether the computers in the household are fixed or mobile devices, through the techniques discussed previously herein.

[0077] In another implementation, the CAS arranges for delivery of TV ads to the STB along with associated television program targeting criteria, such as targeted age (kid, teen, adult, or senior), targeted gender, or type of content (travel, finance, news, or other subjects). The STB chooses the TV program within which to serve each of the delivered ads based on comparing the television program targeting criteria to actual program characteristics. Program characteristics can be ascertained by virtue of their being included in the electronic TV guide available to the STB or sent with the program as accompanying embedded data, for example, using known or later-discovered techniques. In a variation, when the STB delivers a television program to the user, or when the television program is about ready to show an advertisement, the STB can connect with the CAS and report to the CAS the identity of the viewed program, and the CAS can respond with an instruction identifying which of the ads stored in the STB (previously delivered from the CAS to the STB) should be delivered during the viewed program, based on comparing the characteristics of the program being viewed with the targeting criteria.

[0078] The television advertisements directed to users' set-top boxes at the request of a profile provider can come from a variety of sources. In one example, a given profile provider can request that the CAS direct an advertisement to the user's STB from (or on behalf of) the profile provider itself (if, for example, the profile provider is an online advertiser site interested in delivering TV ads to people that visited its online site) or from (or on behalf of) a related site (perhaps a site in the same Internet domain as the given profile provider). In that example, the proprietor of the CAS, as the facilitator of the targeted TV ad delivery based on the online site visit, can receive a payment from the profile provider. The CAS proprietor can keep a share of revenue as a commission and pay the TVP for the ad space (except in the case where the CAS proprietor owns the ad space). The TVP can either retain the entire amount paid by the CAS proprietor, if the TVP owns the TV ad space within which the TV ad was delivered, or the TVP can pay some or all of the amount paid by the CAS proprietor to another entity such as a TV

broadcasting network, if the TVP does not own the ad space. In the case where the STB is provided by a third party other than the TVP (such as a service provider like Tivo), the STB provider can arrange with content owners (such as a TV broadcasting network) to have the service provider's STBs replace broadcasted TV ads with targeted TV ads, where a targeted TV ad is available, and the service provider and the TV broadcasting network can share the revenue with the CAS proprietor as discussed above. In some cases, the CAS can be owned by a TVP, an ISP, an STB provider, an online site, a TV ad space owner, or a TV content provider, in which case the payments will be divided in accordance to the respective roles. In general, a payment from a TV advertiser is used to pay the TV ad space owner and the entities facilitating the delivery of the targeted TV ad.

[0079] In another example, the online site or other profile provider can request that an ad from another advertiser of its own choosing be directed to the user's STB to whom the online site has sold an amount of television advertising space. The online site generates a new revenue stream by selling TV ads to its advertisers, where the TV ads are delivered to members of the online site's audience while they watch TV. In that example the online site can collect a revenue amount from the site-chosen advertiser while paying the CAS, ISP, TVP, ISP/TVP, or TV ad space owner in return for assisting to facilitate delivery to the STB of the ad provided by the site-chosen advertiser.

[0080] In another example, the profile provider provides a profile to the CAS in return for a payment triggered (1) every time the provided profile (whether including behavioral or demographic data) is used to deliver a TV ad, (2) when a TV ad that has been delivered using the provided profile generates revenues, or (3) simply on account of delivery of the profile without regard to whether, how often, or how effectively it is used.

[0081] The TV advertisement itself can be sold by the CAS, ISP, TVP, ISP/TVP, STB provider, TV ad space owner, or any third party such as a re-seller or a firm that represents ad space owners or profile providers and sells to advertisers. In that example the profile provider can collect a revenue amount from any of the entities benefiting from the delivery of the TV ad based on the profile provided by the profile provider, including the TV ad space owner, STB provider, CAS, ISP, TVP, ISP/TVP, or a third party selling the targeted ad space. The profile provider, re-seller, or the other

advertisers can also pay a revenue amount to the CAS, ISP, TVP, or ISP/TVP in return for directing the ad to the STB.

[0082] The systems and methods disclosed herein can be used to generate revenue in a variety of ways for various of the involved entities, not limited to the examples given here, that fall within the scope of the present disclosure or appended claims. It should be noted that the terms “pay,” “collect,” “receive,” and so forth, when referring to revenue amounts, can denote actual exchanges of funds or can denote credits or debits to electronic accounts, especially automatic payment implemented with computer tracking. The terms can apply whether the payments are characterized as commissions, royalties, referral fees, holdbacks, overrides, purchase-resales, or any other compensation arrangements giving net results of split advertising revenues as stated above. Payment can occur either immediately, such as through micro-payment transfers, periodically, such as daily, weekly, or monthly, or upon accumulation of payments from multiple events totaling above a threshold amount. The systems and methods disclosed herein can be implemented with any suitable accounting modules or subsystems for tracking such payments or receipts of funds.

[0083] The systems and methods disclosed herein can be implemented as general or special purpose computers or other programmable hardware devices programmed through software or as hardware or equipment “programmed” through hard wiring, or a combination of the two. Computer programs or other software, if used, can be implemented in temporary or permanent storage or in replaceable media, such as by including programming in microcode, object-oriented code, web-based or distributed software modules that operate together, RAM, ROM, CD-ROM, DVD-ROM, hard drives, thumb drives, flash memory, optical media, magnetic media, semiconductor media, or any future storage alternatives.

[0084] It is intended that equivalents of the disclosed exemplary embodiments and methods shall fall within the scope of the present disclosure and/or appended claims. It is intended that the disclosed exemplary embodiments and methods, and equivalents thereof, may be modified while remaining within the scope of the present disclosure or appended claims. For example, where the term “product” is used in connection with a

shopping or purchase request, the product can be a physical product, intangible product, or a service.

[0085] For purposes of the present disclosure and appended claims, the conjunction "or" is to be construed inclusively (e.g., "a dog or a cat" would be interpreted as "a dog, or a cat, or both"; e.g., "a dog, a cat, or a mouse" would be interpreted as "a dog, or a cat, or a mouse, or any two, or all three"), unless: (i) it is explicitly stated otherwise, e.g., by use of "either…or," "only one of…," or similar language; or (ii) two or more of the listed alternatives are mutually exclusive within the particular context, in which case "or" would encompass only those combinations involving non-mutually-exclusive alternatives. For purposes of the present disclosure or appended claims, the words "comprising," "including," and "having," and variants thereof, shall be construed as open ended terminology, with the same meaning as if the phrase "at least" were appended after each instance thereof.

CLAIMS

What is claimed is:

1. A computer-implemented method comprising:
   (a) electronically associating online access IP addresses and received set-top box IP addresses for each of a multitude of users; and
   (b) using user profile information derived from online activity from a first one of the online access IP addresses to automatically cause a first television advertisement to be directed to the set-top box indicated by the set-top box IP address associated with the first online access IP address.

2. The method of Claim 1 wherein neither the user profile information in part (b) nor the association between the first online access IP address and its associated set-top box IP address includes personally identifiable information.

3. The method of Claim 1
   (i) further comprising using user profile information derived from online activity from a second one of the online access IP addresses to automatically cause a second television advertisement to be directed to the set-top box indicated by the set-top box IP address associated with the second online access IP address.
   (ii) wherein which television advertisement is delivered to which set-top box depends on differences between the online activities from the first and second online access IP addresses, respectively.

4. The method of Claim 3 wherein neither the user profile information in part (i) nor the association between the second online access IP address and its associated set-top box IP address includes personally identifiable information.

5. The method of Claim 1 further comprising storing electronic associations between the online access IP addresses and the received set-top box IP addresses.

6. The method of Claim 1 wherein the online activity comprises requesting a search.

7. The method of Claim 1 wherein the first television advertisement is caused to be directed to the set-top box upon first being contacted by the set-top box after the user profile information is derived.

8. The method of Claim 1 wherein the first television advertisement is selected from among a plurality of possible television advertisements based on the user profile information derived from the online activity from the first online access IP address.

9. The method of Claim 8 wherein the first television advertisement is selected based at least partly on demographic or behavioral attributes included in the user profile information.

10. The method of Claim 1 wherein the first television advertisement is selected based at least partly on the basis of television content viewed at the time the first television advertisement is to be presented by the set-top box.

11. The method of Claim 1 wherein the online access and set-top box IP addresses are both received from a service provider that provides both online access and television service to each of the multitude of users.

12. The method of Claim 11 wherein the association of the online access and set-top box IP addresses of each user is based on at least a pseudonym assigned to the associated IP addresses of that user.

13. The method of Claim 11 wherein part (a) comprises the act of storing the online access IP addresses and the set-top box IP addresses, and further comprising updating at least one of the received online access IP addresses or at least one of the received set-top box addresses.

14. The method of Claim 1 wherein the association of the online access and set-top box IP addresses of each user is based on a common IP address to which network traffic is routed for both the online access device and set-top box of that user.

15. The method of Claim 14 wherein the common IP address is an IP address dynamically assigned to a modem or router used by the user for online access and connected to the user's set-top box.

16. The method of Claim 14 further comprising:
    (i) receiving an identifier corresponding to the set-top box of the user; and
    (ii) causing the television advertisement to be transmitted only if directed to the set-top box corresponding to the identifier.

17. The method of Claim 14 wherein part (a) comprises the act of storing the common IP address, and further comprising updating the common IP address of the user.

18. The method of Claim 14 wherein part (a) comprises the act of storing the common IP address, and further comprising periodically updating the common online access and set-top box IP address of the user.

19. The method of Claim 14 wherein:
    (i) the user profile information in part (b) includes a list of online access IP addresses corresponding to specific online activities and times and dates corresponding to each of the online activities; and
    (ii) the television advertisement is caused to be directed to the set-top box of one of the users only if any of the listed online access IP addresses corresponding to the specific online activities and the corresponding dates and times of the online activities match the common IP address assigned to that user and the date and time corresponding thereto.

20. The method of Claim 19 wherein part (a) comprises the act of storing the common IP address and dates and times corresponding thereto, and further comprising periodically updating (i) the common IP address and times and dates corresponding thereto or (ii) the list of online access IP addresses corresponding to specific online activities and time and dates corresponding to each of the online activities.

21. The method of Claim 1 wherein the first television advertisement is received at a central ad server and directed to the set-top box.

22. The method of Claim 1 wherein part (b) comprises directing instructions to a profile provider, to an online site, or to an ad server to direct the first television advertisement to the set-top box.

23. The method of Claim 1 wherein part (b) comprises directing instructions to the set-top box to retrieve the first television advertisement from a profile provider, from an online site, or from a central ad server.

24. The method of Claim 1 wherein the first television advertisement is caused to be directed only in response to a request for advertisements originating from the user's set-top box.

25. The method of Claim 1 wherein the first television advertisement can be directed without a request for advertisements originating from the user's set-top box.

26. The method of Claim 1 wherein the online access IP address of each user is an IP address associated with a modem or router used by that user for online access.

27. The method of Claim 1 wherein only advertisements satisfying criteria selected by each user are directed to the set-top box of that user.

28. The method of Claim 1 further comprising receiving a revenue amount as a result of directing the advertisement.

29. The method of Claim 1 further comprising paying a revenue amount to a profile provider as a result of using the user profile information in part (b).

30. A computer-implemented method comprising electronically delivering, to a server that electronically associates a multitude of online access IP addresses and IP addresses indicating corresponding television set-top boxes, user profile information derived from online activity from a first online access IP address, as a result of which a first selected television advertisement can be directed to the set-

top box indicated by the set-top box IP address associated with the first online access IP address.

31. The method of Claim 30 wherein neither the user profile information nor the association between the first online access IP address and its associated set-top box IP address includes personally identifiable information.

32. The method of Claim 30 further comprising electronically delivering, to the server, user profile information derived from online activity from a second one of the online access IP addresses, as a result of which a second television advertisement, different from the first television advertisement, can be directed to the set-top box indicated by the set-top box IP address associated with the second online access IP address, wherein which television advertisement is to be delivered to which set-top box depends on differences between the online activities from the first and second IP addresses, respectively.

33. The method of Claim 32 wherein neither the user profile information derived from the online activity from the second online access IP address nor the association between the second online access IP address and its associated set-top box IP address includes personally identifiable information.

34. The method of Claim 30 further comprising receiving a revenue amount from a proprietor of the server as a result of the electronic delivering act.

35. The method of claim 30 further comprising receiving a revenue amount as a result of the electronic delivering act, wherein the revenue amount is received in response to each instance in which revenue is generated as a result of displaying the selected television advertisement on the set-top box associated with the first online access IP address.

36. The method of Claim 30 further comprising electronically identifying the first television advertisement for subsequent direction to the set-top box.

37. The method of Claim 36 wherein the first television advertisement is selected based at least partly on a specific search request, which request has originated from the online access IP address.

38. The method of Claim 36 wherein the first television advertisement is selected based at least partly on demographic or behavioral attributes included in the derived user profile information.

39. The method of Claim 36 further comprising receiving a revenue amount from a provider of the television advertisement as a result of electronically delivering the user profile information.

40. The method of claim 39 wherein the revenue amount is received from an advertiser that has contracted to display a TV ad to persons who have characteristics contained in the user profile information.

41. The method of Claim 30 further comprising electronically delivering the first television advertisement to the server for subsequent direction to the set-top box.

42. The method of Claim 30 further comprising directing the television advertisement to the set-top box in response to instructions received from the server.

43. The method of claim 30 wherein (i) the step of electronically delivering user profile information comprises delivering the first user online access IP address.

44. A computer-implemented method of operating an Internet site comprising electronically delivering user profile information derived from online activity on the Internet site from a first online access IP address, as a result of which: (a) a server, which electronically associates a multitude of online access IP addresses and IP addresses indicating corresponding television set-top boxes, causes direction of a select television advertisement to the set-top box indicated by the set-top box IP address associated with the first online access IP address, and (b) a revenue amount is received.

45. The method of Claim 44 wherein the user profile information derived from the online activity from the first online access IP address does not include personally identifiable information.

46. The method of Claim 45 wherein the user profile information derived from the online activity comprises an indication of a search request originating from the first online access IP address.

47. The method of Claim 45 wherein the user profile information derived from the online activity comprises demographic attributes collected through the Internet site from the first online access IP address.

48. The method of claim 44 wherein (i) the step of electronically delivering user profile information derived from the online activity comprises delivering a first user online access IP address; (ii) the selection of a television advertisement to be delivered to the STB associated with the first user online access IP address is performed by the online site; and (iii) the selection of a television advertisement is based on user profile information not delivered by the site.

49. The method of claim 44 wherein the step of electronically delivering user profile information derived from the online activity comprises redirecting a user accessing the site from the first online access IP address.

50. The method of claim 44 wherein the revenue amount of part (b) is received from an advertiser that has contracted with the proprietor of the Internet site to display a TV ad to an audience of the Internet site.

51. The method of claim 44 wherein the revenue amount of part (b) is received in response to each instance in which revenue is generated as a result of displaying the select television advertisement on the set-top box associated with the first online access IP address.

ABSTRACT

In a method for delivering targeted television advertisements based on online behavior, IP addresses indicating online access devices and IP addresses indicating television set-top boxes are electronically associated for a multitude of users. Using user profile information derived from online activity from one of the online access IP addresses, a television advertisement is selected, such as by using behavioral targeting or demographic information, and automatically directed to the set-top box indicated by the set-top IP address associated with that online access IP address. Preferably neither the user profile information nor the electronic association of online access and set-top box IP addresses includes personally identifiable information.

PTO/AIA/14 (02-18)
Approved for use through 11/30/2020. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	14-TV1-CON7
	Application Number	
Title of Invention	Targeted television advertisements based on online behavior	

The application data sheet is part of the provisional or nonprovisional application for which it is being submitted. The following form contains the bibliographic data arranged in a format specified by the United States Patent and Trademark Office as outlined in 37 CFR 1.76.
This document may be completed electronically and submitted to the Office in electronic format using the Electronic Filing System (EFS) or the document may be printed and included in a paper filed application.

Secrecy Order 37 CFR 5.2:

☐ Portions or all of the application associated with this Application Data Sheet may fall under a Secrecy Order pursuant to 37 CFR 5.2 (Paper filers only. Applications that fall under Secrecy Order may not be filed electronically.)

Inventor Information:

Inventor 1 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
Mr.	Roy		Shkedi	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	Forest Hills	State/Province	NY	Country of Residence	US

Mailing Address of Inventor:

Address 1	37-18 Northern Boulevard, Suite 404		
Address 2			
City	Long Island City	State/Province	NY
Postal Code	11101	Country	US

All Inventors Must Be Listed - Additional Inventor Information blocks may be generated within this form by selecting the **Add** button. Add

Correspondence Information:

Enter either Customer Number or complete the Correspondence Information section below.
For further information see 37 CFR 1.33(a).

☐ **An Address is being provided for the correspondence Information of this application.**

Customer Number	26362		
Email Address	louis@valuablepatents.com	Add Email	Remove Email
Email Address	donald@valuablepatents.com	Add Email	Remove Email

Application Information:

Title of the Invention	Targeted television advertisements based on online behavior		
Attorney Docket Number	14-TV1-CON7	Small Entity Status Claimed	☐
Application Type	Nonprovisional		
Subject Matter	Utility		
Total Number of Drawing Sheets (if any)	18	Suggested Figure for Publication (if any)	

PTO/AIA/14 (02-18)
Approved for use through 11/30/2020. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	14-TV1-CON7
	Application Number	
Title of Invention	Targeted television advertisements based on online behavior	

Filing By Reference:

Only complete this section when filing an application by reference under 35 U.S.C. 111(c) and 37 CFR 1.57(a). Do not complete this section if application papers including a specification and any drawings are being filed. Any domestic benefit or foreign priority information must be provided in the appropriate section(s) below (i.e., "Domestic Benefit/National Stage Information" and "Foreign Priority Information").

For the purposes of a filing date under 37 CFR 1.53(b), the description and any drawings of the present application are replaced by this reference to the previously filed application, subject to conditions and requirements of 37 CFR 1.57(a).

Application number of the previously filed application	Filing date (YYYY-MM-DD)	Intellectual Property Authority or Country

Publication Information:

☐ Request Early Publication (Fee required at time of Request 37 CFR 1.219)

☐ **Request Not to Publish.** I hereby request that the attached application not be published under 35 U.S.C. 122(b) and certify that the invention disclosed in the attached application **has not and will not** be the subject of an application filed in another country, or under a multilateral international agreement, that requires publication at eighteen months after filing.

Representative Information:

Representative information should be provided for all practitioners having a power of attorney in the application. Providing this information in the Application Data Sheet does not constitute a power of attorney in the application (see 37 CFR 1.32).
Either enter Customer Number or complete the Representative Name section below. If both sections are completed the customer Number will be used for the Representative Information during processing.

Please Select One:	◉ Customer Number	○ US Patent Practitioner	○ Limited Recognition (37 CFR 11.9)
Customer Number	26362		

Domestic Benefit/National Stage Information:

This section allows for the applicant to either claim benefit under 35 U.S.C. 119(e), 120, 121, 365(c), or 386(c) or indicate National Stage entry from a PCT application. Providing benefit claim information in the Application Data Sheet constitutes the specific reference required by 35 U.S.C. 119(e) or 120, and 37 CFR 1.78.
When referring to the current application, please leave the "Application Number" field blank.

Prior Application Status	Patented	Remove

Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)	Patent Number	Issue Date (YYYY-MM-DD)
	Continuation of	16925655	2020-07-10	11303973	2022-04-12

PTO/AIA/14 (02-18)
Approved for use through 11/30/2020. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	14-TV1-CON7
	Application Number	
Title of Invention	Targeted television advertisements based on online behavior	

Prior Application Status	Patented			Remove	
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)	Patent Number	Issue Date (YYYY-MM-DD)
16925655	Continuation of	16228479	2018-12-20	10715878	2020-07-14
Prior Application Status	Patented			Remove	
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)	Patent Number	Issue Date (YYYY-MM-DD)
16228479	Continuation of	15803631	2017-11-03	10178442	2019-01-08
Prior Application Status	Patented			Remove	
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)	Patent Number	Issue Date (YYYY-MM-DD)
15803631	Continuation of	15179356	2016-06-10	9813778	2017-11-07
Prior Application Status	Patented			Remove	
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)	Patent Number	Issue Date (YYYY-MM-DD)
15179356	Continuation of	14246525	2014-04-07	9369779	2016-06-14
Prior Application Status	Patented			Remove	
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)	Patent Number	Issue Date (YYYY-MM-DD)
14246525	Division of	13097885	2011-04-29	8695032	2014-04-08
Prior Application Status	Patented			Remove	
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)	Patent Number	Issue Date (YYYY-MM-DD)
13097885	Division of	12860666	2010-08-20	8281336	2010-12-28
Prior Application Status	Patented			Remove	
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)	Patent Number	Issue Date (YYYY-MM-DD)
12860666	Continuation of	11736544	2007-04-17	7861260	2010-12-28

Additional Domestic Benefit/National Stage Data may be generated within this form by selecting the **Add** button.

Foreign Priority Information:

This section allows for the applicant to claim priority to a foreign application. Providing this information in the application data sheet constitutes the claim for priority as required by 35 U.S.C. 119(b) and 37 CFR 1.55. When priority is claimed to a foreign application that is eligible for retrieval under the priority document exchange program (PDX)[i] the information will be used by the Office to automatically attempt retrieval pursuant to 37 CFR 1.55(i)(1) and (2). Under the PDX program, applicant bears the ultimate responsibility for ensuring that a copy of the foreign application is received by the Office from the participating foreign intellectual property office, or a certified copy of the foreign priority application is filed, within the time period specified in 37 CFR 1.55(g)(1).

PTO/AIA/14 (02-18)
Approved for use through 11/30/2020. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	14-TV1-CON7
	Application Number	
Title of Invention	Targeted television advertisements based on online behavior	

Application Number	Country[i]	Filing Date (YYYY-MM-DD)	Access Code[i] (if applicable)

Remove

Additional Foreign Priority Data may be generated within this form by selecting the **Add** button.

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications

☐ This application (1) claims priority to or the benefit of an application filed before March 16, 2013 and (2) also contains, or contained at any time, a claim to a claimed invention that has an effective filing date on or after March 16, 2013.
NOTE: By providing this statement under 37 CFR 1.55 or 1.78, this application, with a filing date on or after March 16, 2013, will be examined under the first inventor to file provisions of the AIA.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	14-TV1-CON7
	Application Number	
Title of Invention	Targeted television advertisements based on online behavior	

Authorization or Opt-Out of Authorization to Permit Access:

When this Application Data Sheet is properly signed and filed with the application, applicant has provided written authority to permit a participating foreign intellectual property (IP) office access to the instant application-as-filed (see paragraph A in subsection 1 below) and the European Patent Office (EPO) access to any search results from the instant application (see paragraph B in subsection 1 below).

Should applicant choose not to provide an authorization identified in subsection 1 below, applicant **must opt-out** of the authorization by checking the corresponding box A or B or both in subsection 2 below.

NOTE: This section of the Application Data Sheet is **ONLY** reviewed and processed with the **INITIAL** filing of an application. After the initial filing of an application, an Application Data Sheet cannot be used to provide or rescind authorization for access by a foreign IP office(s). Instead, Form PTO/SB/39 or PTO/SB/69 must be used as appropriate.

1. Authorization to Permit Access by a Foreign Intellectual Property Office(s)

A. Priority Document Exchange (PDX) - Unless box A in subsection 2 (opt-out of authorization) is checked, the undersigned hereby **grants the USPTO authority** to provide the European Patent Office (EPO), the Japan Patent Office (JPO), the Korean Intellectual Property Office (KIPO), the State Intellectual Property Office of the People's Republic of China (SIPO), the World Intellectual Property Organization (WIPO), and any other foreign intellectual property office participating with the USPTO in a bilateral or multilateral priority document exchange agreement in which a foreign application claiming priority to the instant patent application is filed, access to: (1) the instant patent application-as-filed and its related bibliographic data, (2) any foreign or domestic application to which priority or benefit is claimed by the instant application and its related bibliographic data, and (3) the date of filing of this Authorization. See 37 CFR 1.14(h)(1).

B. Search Results from U.S. Application to EPO - Unless box B in subsection 2 (opt-out of authorization) is checked, the undersigned hereby **grants the USPTO authority** to provide the EPO access to the bibliographic data and search results from the instant patent application when a European patent application claiming priority to the instant patent application is filed. See 37 CFR 1.14(h)(2).

The applicant is reminded that the EPO's Rule 141(1) EPC (European Patent Convention) requires applicants to submit a copy of search results from the instant application without delay in a European patent application that claims priority to the instant application.

2. Opt-Out of Authorizations to Permit Access by a Foreign Intellectual Property Office(s)

☐ A. Applicant **DOES NOT** authorize the USPTO to permit a participating foreign IP office access to the instant application-as-filed. If this box is checked, the USPTO will not be providing a participating foreign IP office with any documents and information identified in subsection 1A above.

☐ B. Applicant **DOES NOT** authorize the USPTO to transmit to the EPO any search results from the instant patent application. If this box is checked, the USPTO will not be providing the EPO with search results from the instant application.

NOTE: Once the application has published or is otherwise publicly available, the USPTO may provide access to the application in accordance with 37 CFR 1.14.

PTO/AIA/14 (02-18)
Approved for use through 11/30/2020. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	14-TV1-CON7
	Application Number	
Title of Invention	Targeted television advertisements based on online behavior	

Applicant Information:

Providing assignment information in this section does not substitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Applicant 1

If the applicant is the inventor (or the remaining joint inventor or inventors under 37 CFR 1.45), this section should not be completed. The information to be provided in this section is the name and address of the legal representative who is the applicant under 37 CFR 1.43; or the name and address of the assignee, person to whom the inventor is under an obligation to assign the invention, or person who otherwise shows sufficient proprietary interest in the matter who is the applicant under 37 CFR 1.46. If the applicant is an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest) together with one or more joint inventors, then the joint inventor or inventors who are also the applicant should be identified in this section.

Clear

◉ Assignee	○ Legal Representative under 35 U.S.C. 117	○ Joint Inventor
○ Person to whom the inventor is obligated to assign.	○ Person who shows sufficient proprietary interest	

If applicant is the legal representative, indicate the authority to file the patent application, the inventor is:

Name of the Deceased or Legally Incapacitated Inventor:

If the Applicant is an Organization check here. ☒

Organization Name	INTENT IQ, LLC

Mailing Address Information For Applicant:

Address 1	37-18 Northern Boulevard, Suite 404		
Address 2			
City	Long Island City	**State/Province**	NY
Country	US	Postal Code	11101
Phone Number	646-217-0930	Fax Number	
Email Address			

Additional Applicant Data may be generated within this form by selecting the Add button.

Assignee Information including Non-Applicant Assignee Information:

Providing assignment information in this section does not substitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

PTO/AIA/14 (02-18)
Approved for use through 11/30/2020. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	14-TV1-CON7
	Application Number	
Title of Invention	Targeted television advertisements based on online behavior	

Assignee 1

Complete this section if assignee information, including non-applicant assignee information, is desired to be included on the patent application publication. An assignee-applicant identified in the "Applicant Information" section will appear on the patent application publication as an applicant. For an assignee-applicant, complete this section only if identification as an assignee is also desired on the patent application publication.

If the Assignee or Non-Applicant Assignee is an Organization check here. ☐

Prefix	**Given Name**	Middle Name	**Family Name**	Suffix

Mailing Address Information For Assignee including Non-Applicant Assignee:

Address 1			
Address 2			
City		**State/Province**	
Country		Postal Code	
Phone Number		Fax Number	
Email Address			

Additional Assignee or Non-Applicant Assignee Data may be generated within this form by selecting the Add button.

Signature:

NOTE: This Application Data Sheet must be signed in accordance with 37 CFR 1.33(b). **However, if this Application Data Sheet is submitted with the INITIAL filing of the application and either box A or B is not checked in subsection 2 of the "Authorization or Opt-Out of Authorization to Permit Access" section, then this form must also be signed in accordance with 37 CFR 1.14(c).**

This Application Data Sheet **must** be signed by a patent practitioner if one or more of the applicants is a **juristic entity** (e.g., corporation or association). If the applicant is two or more joint inventors, this form must be signed by a patent practitioner, **all** joint inventors who are the applicant, or one or more joint inventor-applicants who have been given power of attorney (e.g., see USPTO Form PTO/AIA/81) on behalf of **all** joint inventor-applicants.

See 37 CFR 1.4(d) for the manner of making signatures and certifications.

Signature	/Louis J. Hoffman/			Date (YYYY-MM-DD)	2022-04-06
First Name	Louis	Last Name	Hoffman	Registration Number	38,918

Additional Signature may be generated within this form by selecting the Add button.

PTO/AIA/14 (02-18)
Approved for use through 11/30/2020. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	14-TV1-CON7
	Application Number	
Title of Invention	Targeted television advertisements based on online behavior	

This collection of information is required by 37 CFR 1.76. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 23 minutes to complete, including gathering, preparing, and submitting the completed application data sheet form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1 The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3 A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent CooperationTreaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	45414490
Application Number:	17714965
International Application Number:	
Confirmation Number:	7004
Title of Invention:	Targeted television advertisements based on online behavior
First Named Inventor/Applicant Name:	Roy Shkedi
Customer Number:	26362
Filer:	Louis J. Hoffman/Michelle Lisa
Filer Authorized By:	Louis J. Hoffman
Attorney Docket Number:	14-TV1-CON7
Receipt Date:	06-APR-2022
Filing Date:	
Time Stamp:	20:29:50
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Transmittal of New Application	22-04-06-IIQ-TV1_CON7-App_Trans.pdf	286210 e5a2a615a45fd60a4ca98e518950d7397e6a1b97	no	2

Warnings:

Information:					
2	Drawings-only black and white line drawings	TV1-Figures.pdf	126789 6087085a487506a1c304b5d2fee0a189d5f9408a	no	18
Warnings:					
Information:					
3	Oath or Declaration filed	IIQ-TV1-Declaration-parent_copy.pdf	144649 e5ffe97bf4423cd8efcab355dcf2f719bf54bf99	no	1
Warnings:					
Information:					
4		TV1-Specification.pdf	184224 f4ca8650698e2831f2ebb4910bbdbe13af7c9d3c	yes	45

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Specification	1	37
	Claims	38	44
	Abstract	45	45

Warnings:					
Information:					
5	Application Data Sheet	22-04-06-IIQ-TV1_CON7-ADS.pdf	497290 24e67b88799678aba6d6a986e4bf46e4e6b5052b	no	9
Warnings:					
Information:					
This is not an USPTO supplied ADS fillable form					
		Total Files Size (in bytes):	1239162		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

Doc Code: PA..
Document Description: Power of Attorney

PTO/AIA/82A (07-13)
Approved for use through 03/31/2021. OMB 0651-0035
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

TRANSMITTAL FOR POWER OF ATTORNEY TO ONE OR MORE REGISTERED PRACTITIONERS

NOTE: This form is to be submitted with the Power of Attorney by Applicant form (PTO/AIA/82B) to identify the application to which the Power of Attorney is directed, in accordance with 37 CFR 1.5, unless the application number and filing date are identified in the Power of Attorney by Applicant form. If neither form PTO/AIA/82A nor form PTO/AIA82B identifies the application to which the Power of Attorney is directed, the Power of Attorney will not be recognized in the application.

Application Number	17/714,965
Filing Date	4/6/2022
First Named Inventor	Shkedi, Roy
Title	Targeted television advertisements based on online behavior
Art Unit	2425 (in parent app.)
Examiner Name	Ekpo, Nnenna Ngozi (in parent app.)
Attorney Docket Number	14-TV1-CON7

SIGNATURE of Applicant or Patent Practitioner

Signature	/Louis J. Hoffman/	Date (Optional)	04/08/2022
Name	Louis J. Hoffman	Registration Number	38,918
Title (if Applicant is a juristic entity)			
Applicant Name (if Applicant is a juristic entity)			

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4(d) for signature requirements and certifications. If more than one applicant, use multiple forms.

☐ *Total of _____ forms are submitted.

This collection of information is required by 37 CFR 1.131, 1.32, and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Doc Code: PA..
Document Description: Power of Attorney

PTO/AIA/82B (07-13)
Approved for use through 11/30/2014. OMB 0651-0051
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number

POWER OF ATTORNEY BY APPLICANT

I hereby revoke all previous powers of attorney given in the application identified in either the attached transmittal letter or the boxes below.

Application Number	Filing Date
14/246,525	04/07/2014

(Note: The boxes above may be left blank if information is provided on form PTO/AIA/82A.)

☑ I hereby appoint the Patent Practitioner(s) associated with the following Customer Number as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the application referenced in the attached transmittal letter (form PTO/AIA/82A) or identified above: 26362

OR

☐ I hereby appoint Practitioner(s) named in the attached list (form PTO/AIA/82C) as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the patent application referenced in the attached transmittal letter (form PTO/AIA/82A) or identified above. (Note: Complete form PTO/AIA/82C.)

Please recognize or change the correspondence address for the application identified in the attached transmittal letter or the boxes above to:

☑ The address associated with the above-mentioned Customer Number

OR

☐ The address associated with Customer Number:

OR

Firm or Individual Name					
Address					
City		State		Zip	
Country					
Telephone		Email			

I am the Applicant (if the Applicant is a juristic entity, list the Applicant name in the box):

Intent IQ, LLC

☐ Inventor or Joint Inventor (title not required below)

☐ Legal Representative of a Deceased or Legally Incapacitated Inventor (title not required below)

☑ Assignee or Person to Whom the Inventor is Under an Obligation to Assign (provide signer's title if applicant is a juristic entity)

☐ Person Who Otherwise Shows Sufficient Proprietary Interest (e.g., a petition under 37 CFR 1.46(b)(2) was granted in the application or is concurrently being filed with this document) (provide signer's title if applicant is a juristic entity)

SIGNATURE of Applicant for Patent

The undersigned (whose title is supplied below) is authorized to act on behalf of the applicant (e.g., where the applicant is a juristic entity).

Signature	Roy Shkedi	Date (Optional)	04-07-2014
Name	Roy Shkedi		
Title	Chairman		

NOTE: Signature - This form must be signed by the applicant in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications. If more than one applicant, use multiple forms.

☐ Total of forms are submitted.

This collection of information is required by 37 CFR 1.131, 1.32, and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor	:	Shkedi, Roy	Art Unit	:	2425 (in parent)
Serial No.	:	17/714,965	Examiner	:	Ekpo, Nnenna (in parent)
Filing Date	:	04/06/2022	Conf. No.	:	7004
Title	:	Targeted television advertisements based on online behavior			

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Filed via EFS-Web on April 8, 2022

INFORMATION DISCLOSURE STATEMENT

Dear Sir:

Assignee discloses the references listed on form PTO-1449 attached to this document. All foreign and non-patent literature references listed were cited in parent application Serial No. 16/925,655, therefore, pursuant to Office rules, copies of those references are not included here.

Assignee also wishes to direct Examiner to the applications listed in the table below (which share the same first inventor) as well as the references cited therein:

Serial No.	Filing Date	Earliest Parent Filing Date
17/401,078	8/12/2021	12/31/2007
17/097,867	11/13/2020	06/16/2006
17/092,651	11/09/2020	06/13/2013
17/390,211	7/30/2021	08/03/2011 (prov.)
17/401,052	08/12/2021	06/19/2006 (prov.)
17/699,297	3/1/2022	08/28/2020 (prov.)

If the Office has any questions, please feel free to contact assignee's undersigned attorney of record.

Respectfully submitted,

INTENT IQ, LLC.

by its attorney

Dated: April 8, 2022

/Louis J. Hoffman/
Louis J. Hoffman
Reg. No. 38,918

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, Arizona 85260
(480) 948-3295

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Page 1 of 21

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	1	5220655	06/15/1993	Tsutsui
	2	5359508	10/25/1994	Rossides
	3	5442633	08/15/1995	Perkins et al.
	4	5572643	11/05/1996	Judson
	5	5636346	06/03/1997	Saxe
	6	5710884	01/20/1998	Dedrick
	7	5717923	02/10/1998	Dedrick
	8	5724424	03/03/1998	Gifford
	9	5724521	03/03/1998	Dedrick
	10	5751956	05/12/1998	Kirsch
	11	5752238	05/12/1998	Dedrick
	12	5774670	06/30/1998	Montulli
	13	5794210	08/11/1998	Goldhaber, et al.
	14	5819092	10/06/1998	Ferguson et al.
	15	5819285	10/06/1998	Damico et al.
	16	5835718	11/10/1998	Blewett
	17	5848396	12/08/1998	Gerace
	18	5855008	12/29/1998	Goldhaber, et al.
	19	5870546	02/09/1999	Kirsch
	20	5870550	02/09/1999	Wesinger Jr. et al.
	21	5901287	05/04/1999	Bull et al.
	22	5918014	06/29/1999	Robinson
	23	5933811	08/03/1999	Angles et al.
	24	5937390	08/10/1999	Hyodo
	25	5937392	08/10/1999	Alberts
	26	5948061	09/07/1999	Merriman, et al.
	27	5959623	09/28/1999	van Hoff et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	28	5963915	10/05/1999	Kirsch
	29	5999912	12/07/1999	Wodarz et al.
	30	6014638	01/11/2000	Burge et al.
	31	6016504	01/18/2000	Arnold et al.
	32	6026368	02/15/2000	Brown et al.
	33	6026369	02/15/2000	Capek
	34	6026374	02/15/2000	Chess
	35	6044376	03/28/2000	Kurtzman, II
	36	6047310	04/04/2000	Kamakura et al.
	37	6049835	04/11/2000	Gagnon
	38	6055513	04/25/2000	Katz, et al.
	39	6073241	06/06/2000	Rosenberg et al.
	40	6128663	10/03/2000	Thomas
	41	6133912	10/17/2000	Montero
	42	6134532	10/17/2000	Lazarus et al.
	43	6141010	10/31/2000	Hoyle
	44	6177931	01/23/2001	Alexander et al.
	45	6188398	02/13/2001	Collins-Rector et al.
	46	6205432	03/20/2001	Gabbard et al.
	47	6216112	04/10/2001	Fuller et al.
	48	6216129	04/10/2001	Eldering
	49	6269361	07/31/2001	Davis et al.
	50	6285987	09/04/2001	Roth et al.
	51	6285991	09/04/2001	Powar
	52	6289318	09/11/2001	Barber
	53	6308202	10/23/2001	Cohn et al.
	54	6324519	11/27/2001	Eldering

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	55	6366298	04/02/2002	Haitsuka et al.
	56	6377986	04/23/2002	Philyaw et al.
	57	6385592	05/07/2002	Angles, et al.
	58	6401075	06/04/2002	Mason et al.
	59	6408278	06/18/2002	Carney et al.
	60	6415270	07/02/2002	Rackson et al.
	61	6415322	07/02/2002	Jaye
	62	6449657	09/10/2002	Stanbach Jr. et al.
	63	6484492	11/26/2002	Meholic, et al.
	64	6487538	11/26/2002	Gupta, et al.
	65	6505169	01/07/2003	Bhagavath et al.
	66	6615247	09/02/2003	Murphy
	67	6633850	10/14/2003	Gabbard et al.
	68	6523041	02/18/2003	Morgan et al.
	69	6654725	11/25/2003	Langheinrich et al.
	70	6766327	07/20/2004	Morgan et al.
	71	6810527	10/26/2004	Conrad et al.
	72	6850988	02/01/2005	Reed
	73	6868392	03/15/2005	Ogasawara
	74	6868525	03/15/2005	Szabo
	75	6907566	06/14/2005	McElfresh et al.
	76	6983379	01/03/2006	Spalink et al.
	77	6985882	01/10/2006	Del Sesto
	78	7007074	02/28/2006	Radwin
	79	7038637	05/02/2006	Eller et al.
	80	7039594	05/02/2006	Gersting
	81	7039599	05/02/2006	Merriman et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	82	7043531	05/09/2006	Seibel et al.
	83	7058593	06/06/2006	Merritt
	84	7065550	06/20/2006	Raghunandan
	85	7085682	08/01/2006	Heller et al.
	86	7155508	12/26/2006	Sankuratripati, et al.
	87	7165105	01/16/2007	Reiner et al.
	88	7181438	02/20/2007	Szabo
	89	7188076	03/06/2007	Bensemana
	90	7194424	03/20/2007	Greer et al.
	91	7246150	07/17/2007	Donoho et al.
	92	7299195	11/20/2007	Tawakol et al.
	93	7337133	02/26/2008	Bezos et al.
	94	7360251	04/15/2008	Spalink et al.
	95	7386599	06/10/2008	Piersol et al.
	96	7530020	05/05/2009	Szabo
	97	7552069	06/23/2009	Kepecs
	98	7558795	07/07/2009	Malik et al.
	99	7580855	08/25/2009	Fernandes
	100	7606810	10/20/2009	Jeavons
	101	7698422	04/13/2010	Vanderhook et al.
	102	7739594	06/15/2010	Vasilik
	103	7756744	07/13/2010	Schiff et al.
	104	7797274	09/14/2010	Strathearn et al.
	105	7822636	10/26/2010	Ferber et al.
	106	7899801	03/01/2011	Bronson et al.
	107	7930252	04/19/2011	Bender et al.
	108	7949553	05/24/2011	Wisniewski et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	109	8015184	09/06/2011	Zito et al.
	110	8019881	09/13/2011	Sandhu et al.
	111	8036929	10/11/2011	Reisman
	112	8050970	11/01/2011	Barry et al.
	113	8082332	12/20/2011	Gilbert
	114	8095450	01/10/2012	Walker et al.
	115	8117068	02/14/2012	Shuster et al.
	116	8135799	03/13/2012	Mekikian
	117	8175908	05/08/2012	Anderson
	118	8364540	01/29/2013	Soroca et al.
	119	8510174	08/13/2013	Walker et al.
	120	8527352	09/03/2013	Chatwin et al.
	121	8554940	10/08/2013	Jennings et al.
	122	8751310	06/10/2014	van Datta et al.
	123	8768766	07/01/2014	Ellis et al.
	124	8880677	11/04/2014	Julia et al.
	125	8930204	01/06/2015	Igoe et al.
	126	8996406	03/31/2015	Jain
	127	9117217	08/25/2015	Wilson et al.
	128	20010014868	08/16/2001	Herz et al.
	129	20010020242	09/06/2001	Gupta, et al.
	130	20010034637	10/25/2001	Lin et al.
	131	20010034646	10/25/2001	Hoyt et al.
	132	20010039515	11/08/2001	Mayadas
	133	20010042002	11/15/2001	Koopersmitch
	134	20010044743	11/22/2001	McKinley et al.
	135	20020002541	01/03/2002	Williams

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	136	20020026360	02/28/2002	McGregor et al.
	137	20020038350	03/28/2002	Lambert et al.
	138	20020046286	04/18/2002	Caldwell et al.
	139	20020072965	06/13/2002	Merriman et al.
	140	20020072971	06/13/2002	DeBusk, et al.
	141	20020077891	06/20/2002	Castle et al.
	142	20020082910	06/27/2002	Kontogouris
	143	20020082923	06/27/2002	Merriman, et al.
	144	20020091568	07/11/2002	Kraft et al.
	145	20020099603	07/25/2002	Bandura
	146	20020112013	08/15/2002	Walsh
	147	20020112171	08/15/2002	Ginter et al.
	148	20020123928	09/05/2002	Eldering et al.
	149	20020133399	09/19/2002	Main
	150	20020143632	10/03/2002	Walter et al.
	151	20020161633	10/31/2002	Jacob et al.
	152	20020161636	10/31/2002	Takahashi et al.
	153	20020184088	12/05/2002	Rosenberg
	154	20020194058	12/19/2002	Eldering
	155	20030023481	01/30/2003	Calvert et al.
	156	20030028451	02/06/2003	Ananian
	157	20030036944	02/20/2003	Lesandrini et al.
	158	20030046156	03/06/2003	Cromer, et al.
	159	20030046689	03/06/2003	Gaos
	160	20030050827	03/13/2003	Hennessey
	161	20030050863	03/13/2003	Radwin
	162	20030065558	04/03/2003	Shaw et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	163	20030065805	04/03/2003	Barnes Jr.
	164	20030083931	05/01/2003	Lang
	165	20030083938	05/01/2003	Smith et al.
	166	20030101126	05/29/2003	Cheung et al.
	167	20030110497	06/12/2003	Yassin, et al.
	168	20030135413	07/17/2003	Nishi, et al.
	169	20030135853	07/17/2003	Goldman, et al.
	170	20030149938	08/07/2003	McElfresh, et al.
	171	20030154126	08/14/2003	Gehlot et al.
	172	20030171990	09/11/2003	Rao, et al.
	173	20030208404	11/06/2003	Michie
	174	20030220837	11/27/2003	Asayama
	175	20030229542	12/11/2003	Morrisroe
	176	20040024632	02/05/2004	Perry
	177	20040044566	03/04/2004	Bostelmann et al.
	178	20040059632	03/25/2004	Kang et al.
	179	20040068436	04/08/2004	Boubek et al.
	180	20040117322	06/17/2004	Bjorksten et al.
	181	20040176995	09/09/2004	Fusz
	182	20040186776	09/23/2004	Llach
	183	20040193488	09/30/2004	Khoo et al.
	184	20040260605	12/23/2004	McIntyre et al.
	185	20050010472	01/13/2005	Quatse et al.
	186	20050021397	01/27/2005	Cui, et al.
	187	20050028188	02/03/2005	Latona, et al.
	188	20050038698	02/17/2005	Lukose et al.
	189	20050038702	02/17/2005	Merriman et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	190	20050080772	04/14/2005	Bem
	191	20050120045	06/02/2005	Klawon
	192	20050120091	06/02/2005	Casais
	193	20050131762	06/16/2005	Bharat et al.
	194	20050131762	06/16/2005	Bharat et al.
	195	20050165644	07/28/2005	Beyda, et al.
	196	20050177613	08/11/2005	Dresden
	197	20050187818	08/25/2005	Zito et al.
	198	20050192863	09/01/2005	Mohan
	199	20050216336	09/29/2005	Roberts et al.
	200	20050216823	09/29/2005	Petersen, et al.
	201	20050222905	10/06/2005	Wills
	202	20050278443	12/15/2005	Winner, et al.
	203	20050289590	12/29/2005	Cheok et al.
	204	20060020508	01/26/2006	Gorti et al.
	205	20060020510	01/26/2006	Vest
	206	20060026061	02/02/2006	Collins
	207	20060026061	02/02/2006	Collins
	208	20060026064	02/02/2006	Collins
	209	20060029296	02/09/2006	King et al.
	210	20060053047	03/09/2006	Garcia et al.
	211	20060064351	03/23/2006	Hamor, et al.
	212	20060069614	03/30/2006	Agarwal, et al.
	213	20060095322	05/04/2006	Dierks
	214	20060098899	05/11/2006	King et al.
	215	20060106710	05/18/2006	Meek, et al.
	216	20060116924	06/01/2006	Angles, et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	217	20060122879	06/08/2006	O'Kelley
	218	20060129455	06/15/2006	Shah
	219	20060161553	07/20/2006	Woo
	220	20060183467	08/17/2006	Stewart
	221	20060200377	09/07/2006	Wolfe et al.
	222	20060212349	09/21/2006	Brady
	223	20060212350	09/21/2006	Ellis et al.
	224	20060224496	10/05/2006	Sandholm et al.
	225	20060253432	11/09/2006	Eagle et al.
	226	20060271690	11/30/2006	Banga et al.
	227	20060282327	12/14/2006	Neal et al.
	228	20060287916	12/21/2006	Starr et al.
	229	20070027850	02/01/2007	Chan et al.
	230	20070027901	02/01/2007	Chan et al.
	231	20070038516	02/15/2007	Apple et al.
	232	20070067297	03/22/2007	Kublickis
	233	20070079331	04/05/2007	van Datta et al.
	234	20070088603	04/19/2007	Jouppi et al.
	235	20070121843	05/31/2007	Atazky et al.
	236	20070130015	06/07/2007	Starr et al.
	237	20070130070	06/07/2007	Williams
	238	20070150234	06/28/2007	Wicks
	239	20070150457	06/28/2007	Goldman et al.
	240	20070150468	06/28/2007	Goldman et al.
	241	20070150483	06/28/2007	Goldman et al.
	242	20070150721	06/28/2007	Goldman et al.
	243	20070157228	07/05/2007	Bayer et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	244	20070157245	07/05/2007	Collins
	245	20070162396	07/12/2007	Goldman et al.
	246	20070179845	08/02/2007	Jain
	247	20070179846	08/02/2007	Jain et al.
	248	20070179848	08/02/2007	Jain
	249	20070179849	08/02/2007	Jain
	250	20070179853	08/02/2007	Feige et al.
	251	20070192166	08/16/2007	Van Luchene
	252	20070192168	08/16/2007	Van Luchene
	253	20070198385	08/23/2007	McGill et al.
	254	20070203794	08/30/2007	Hajdukiewicz et al.
	255	20070214037	09/13/2007	Shubert et al.
	256	20070214048	09/13/2007	Chan et al.
	257	20070239517	10/11/2007	Chung et al.
	258	20070239518	10/11/2007	Chung et al.
	259	20070260624	11/08/2007	Chung et al.
	260	20080052392	02/28/2008	Webster et al.
	261	20080109306	05/08/2008	Maigret et al.
	262	20080262901	10/23/2008	Banga et al.
	263	20080263633	10/23/2008	Banga et al.
	264	20080319850	12/25/2008	Shaul et al.
	265	20090106100	04/23/2009	Mashinsky
	266	20090222329	09/03/2009	Ramer et al.
	267	20090234711	09/17/2009	Ramer et al.
	268	20090234861	09/17/2009	Ramer et al.
	269	20090254971	10/08/2009	Herz et al.
	270	20090281954	11/12/2009	Waelbroeck et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	271	20090299857	12/03/2009	Brubaker
	272	20100257053	10/07/2010	Ferber et al.
	273	20100268605	10/21/2010	Waelbroeck et al.
	274	20100293057	11/18/2010	Haveliwala et al.
	275	20100293057	11/18/2010	Haveliwala et al.
	276	20100332583	12/30/2010	Szabo
	277	20110112903	05/12/2011	Goldman et al.
	278	20110191140	08/04/2011	Newman et al.
	279	20110258049	10/20/2011	Ramer et al.
	280	20120209705	08/16/2012	Ramer et al.
	281	20120209706	08/16/2012	Ramer et al.
	282	20120209707	08/16/2012	Ramer et al.
	283	20120209708	08/16/2012	Ramer et al.
	284	20120209709	08/16/2012	Ramer et al.
	285	20120209710	08/16/2012	Ramer et al.
	286	20120215635	08/23/2012	Ramer et al.
	287	20120215639	08/23/2012	Ramer et al.
	288	20130097019	04/18/2013	Ferber et al.
	289	20130097026	04/18/2013	Ferber et al.
	290	20130232001	09/05/2013	van Datta et al.
	291	20130254035	09/26/2013	Ramer et al.
	292	20140324576	10/30/2014	van Datta et al.
	293	6832207	12/14/2004	Shkedi
	294	6925440	08/02/2005	Shkedi
	295	6973436	12/06/2005	Shkedi
	296	7072853	07/04/2006	Shkedi
	297	7428493	09/23/2008	Shkedi

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	298	7454364	11/18/2008	Shkedi
	299	7822637	10/26/2010	Shkedi
	300	7822639	10/26/2010	Shkedi
	301	7861260	12/28/2010	Shkedi
	302	7890609	02/15/2011	Shkedi
	303	7979307	07/12/2011	Shkedi
	304	8051444	11/01/2011	Shkedi
	305	8239264	08/07/2012	Shkedi
	306	8244574	08/14/2012	Shkedi
	307	8244582	08/14/2012	Shkedi
	308	8244583	08/14/2012	Shkedi
	309	8244586	08/14/2012	Shkedi
	310	8280758	10/02/2012	Shkedi
	311	8281336	10/02/2012	Shkedi
	312	8494904	07/23/2013	Shkedi
	313	8566164	10/22/2013	Shkedi et al.
	314	8589210	11/19/2013	Shkedi
	315	8595069	11/26/2013	Shkedi et al.
	316	8600815	12/03/2013	Shkedi
	317	8607267	12/10/2013	Shkedi
	318	8677398	03/18/2014	Shkedi
	319	8683502	03/25/2014	Shkedi et al.
	320	8695032	04/08/2014	Shkedi
	321	8713600	04/29/2014	Shkedi
	322	8775249	07/08/2014	Shkedi
	323	8997138	03/31/2015	Shkedi
	324	9071886	06/30/2015	Shkedi

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	325	9078035	07/07/2015	Shkedi et al.
	326	9083853	07/14/2015	Shkedi
	327	9131282	09/08/2015	Shkedi
	328	9226019	12/29/2015	Shkedi
	329	9271024	02/23/2016	Shkedi et al.
	330	9351053	05/24/2016	Shkedi
	331	9369779	06/14/2016	Shkedi
	332	9571865	02/14/2017	Shkedi
	333	9591380	03/07/2017	Shkedi et al.
	334	9756372	09/05/2017	Shkedi
	335	9800917	10/24/2017	Shkedi
	336	9813778	11/07/2017	Shkedi
	337	10009636	06/26/2018	Shkedi
	338	10026100	07/17/2018	Shkedi
	339	10070200	09/04/2018	Shkedi
	340	10178442	01/08/2019	Shkedi
	341	10306282	05/28/2019	Shkedi
	342	10321198	06/11/2019	Shkedi et al.
	343	10405058	09/03/2019	Shkedi et al.
	344	10628857	04/21/2020	Shkedi
	345	10645438	05/05/2020	Shkedi
	346	10715878	07/14/2020	Shkedi
	347	10771860	09/08/2020	Shkedi et al.
	348	10839423	11/17/2020	Shkedi
	349	10984445	4/20/2021	Shkedi
	350	11082753	4/8/2021	Shkedi et al.
	351	11093970	9/12/2019	Shkedi

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	352	11095952	8/17/2021	Shkedi et al.
	353	11282105	03/03/2022	Shkedi et al.
	354	20050086112	04/21/2005	Shkedi
	355	20090049468	02/19/2009	Shkedi
	356	201100035256	2/10/2011	Shkedi et al.
	357	20110040607	2/17/2011	Shkedi
	358	20090172723	07/02/2009	Shkedi et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

FOREIGN PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee	Translation Yes / No
_____	359	WO 2005/046237	05/19/2005	Gutta et al.	No
_____	360	WO 2004/084109	09/30/2004	NHN Corp.	No
_____	361	KR 10-2004-0099222	11/26/2004	Echo Marketing Co., Ltd.	Yes
_____	362	KR 10-2001-0096343	11/07/2001	MyOne Card Corp.	Yes
_____	363	JP 2003-331183	11/21/2003	Masuda et al	Yes
_____	364	KR 20020037411	05/21/2002	Sin	Yes
_____	365	JP 2012-073990	04/12/2012	Takada	Yes
_____	366	WO1997021183A1	06/12/1997	Bell Communications Research, Inc.	No
_____	367	WO1999057660A1	11/11/1999	IWEB LTD.	No
_____	368	CA 2246746	03/12/1998	Blumenau	No
_____	369	WO 99/44159	09/02/1999	Jaye	No
_____	370	WO 01/54034	07/26/2001	Phillips et al.	No
	371	WO2005022887A2	03/10/2005	Gottfurcht	No

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	372	Engage Technologies, Inc., Form S-1A, filed with the SEC, pp. 36-50, July 19, 1999
_____	373	The e-Marketing Report, Morgan Stanley Dean Whitter, May, 2000, portions
_____	374	ClickOver, Inc. introduces ClickWise, the first inteligent, high-performance ad management system for electronic advertising, Business Wire, March 24, 1997
_____	375	Advertiser Solutions, https://web.archive.org/web/20050325042542/http://www.premiumnetwork.com/advertisers.shtml, archived March 25, 2005
_____	376	baseview.com website (ClassManagerPro) from www.archive.org, archived January 16, 1997
_____	377	Media Kit – Marketers, RevenueScience, https://web.archive.org/web/20051124105455/http://www.revenuescience.com/datafile/Mediakit_marketers_02.pdf, archived November 24, 2005

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
		Publisher Solutions, https://web.archive.org/web/20050325061154/http://www.premiumnetwork.com/publishers.shtml, archived March 25, 2005
	378	
	379	S-1/A SEC Filing, filed by FlyCast Communications Corporation on April 30, 1999
	380	RevenueScience (for Publishers), Behavioral Targeting, Delivering on the Promise of One-to-One Marketing (2005)
	381	24/7 RealMedia, The Science of Digital Marketing, Company Overview (2007)
	382	aaddzz.com webpages from www.archive.org, archived January 30, 1998
	383	AdForce, Inc., User Guide, A Complete Guide to AdForce, Version 2.6 (1998)
	384	FAWCETT, Adrienne W., Publishers thankful but nervous about boom, Magazinedotcom, Hot spending boosts …, adage.com, November 1, 1999
	385	Ad-Star.com website archive from www.archive.org (archived on April 12, 1997 and February 1, 1997)
	386	CHANG et al., "Goodies" in exchange for Consumer Information on the Internet, The Economics and Issues, Proceedings of the Thirty-First Annual Hawaii International Conference on System Sciences, 1998
	387	avenuea.com webpage, FAQ, www.avenuea.com/corpinfo/faq.html, archived October 2, 1999
	388	Advertising-A Unified Approach to Adverstising Management, Baseview Products, Inc. brochure, October 14, 1999
	389	RIBEIRO-NETO et al., Extracting Semi-Structured Data Through Examples, Proceedings of the Eighth international Conference on Information and Knowledge Management (1999)
	390	RASKUTTI and BEITZ, Acquiring Customer Preferences for Information Filtering, a Heuristic-Statistical Approach, Fifth Annual Conference on User Modeling (1996)
	391	AdKnowledge, Campaign Manager-Quick Reference Card (1996)
	392	EMERT, 2 Web Ad Firms Decide to Marry, Palo Alto's Focalink, ClickOver Merging, San Francisco Chronicle, November 17, 1997
	393	ELDERBROCK & BORWANKAR, Building Successful Internet Businesses, IDG Worldwide (1996)
	394	SCHUMANN, David W. & THORSON, Ester, Advertising and the World Wide Web, Lawrence Erlbaum Associates (1999)
	395	WILLIAMSON, Debra Aho, Web ads mark 2nd birthday with decisive issues ahead, adage.com, October 21, 1996

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	396	WILLIAMSON, Debra Aho, Targeting distinguishes AdForce from the pack, adage.com, March 1, 1999
_____	397	Doubleclick.com webpages from www.archived.org, archived February 5, 1998
_____	398	TITTEL et al., Foundations of World Wide Web Programming with HTML & CGI, IDG Books Worldwide (1995)
_____	399	HARRISON et al., An Emerging Marketplace for Digital Advertising Based on Amalgamated Digital Signage Networks, IEEE International Conference on E-Commerce, June 24-27, 2003
_____	400	GRAHAM, Ian S., HTML Sourcebook, A Complete Guide to HTML, John Wiley & Sons (1995)
_____	401	Aaddzz Brochure, Information Access Technologies, Inc. (1997)
_____	402	GILBERT, SPECIAL REPORT, Agencies centralize Web ad serving, adage.com, Mar 1, 1999
_____	403	PALMER & ERIKSEN, Digital Newspapers Explore Marketing on the Internet, Communications of the ACM, Vol. 42, No. 9, September 1999
_____	404	ANGWIN, Julia, Focalink Will Monitor World of Online Ads, San Francisco Chronicle, February 7, 1996
_____	405	GALLAGHER & PARSONS, A Framework for Targeting Banner Advertising on the Internet, Proceedings of the Thirtieth Hawaii International Conference on System Sciences, January 1997
_____	406	KAZIENKO & ADAMSKI, Personalized Web Advertising Methods, Third International Conference, AH 2004, August 23-26, 2004
_____	407	CLELAND, Kim, Accipiter latest entrant in race to manage ads, company sets sights on helping content providers target audience, Adage.com, September 9, 1996
_____	408	CLELAND, Kim, New Player joins fray of Web ad targeters, adage.com, February 24, 1997
_____	409	BAYNE, Kim M., AdKnowledge rolls out Web ad evaluation tool, adage.com, July 8, 1999
_____	410	LANGHEINRICH et al., Unintrusive Customization Techniques for Web Advertising, COMPUTER NETWORKS 31, 1259–1272, May 1999
_____	411	MAEKAWA et al., A Java-Based Information Browsing System in a Remote Display Environment, IEEE (2004)
_____	412	MEEKER, Mary, The Internet Advertising Report, Harper Business (1997)
_____	413	YEAGER & MCGRATH, Web Server Technology, The Advance Guide for World Wide Web Information Providers, Morgan Kaufmann Publishers (1996)

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	414	net.Genesis & Devra Hall, Build a Website, The Programmer's Guide to Creating, Building, and Maintaining a Web Presence, Prima Publishing (1995)
_____	415	NetGravity Launches AdServer 3.5 with Beefed Up Targeting, clickz.com, September 10, 1998
_____	416	NOVAK & HOFFMAN, New Metrics for New Media, Toward the Development of Web Measurement Standards, Project 2000, elab.vanderbilt.edu, September 26, 1996
_____	417	RealMedia.com webpages, www.archive.org, archived on May 15, 19977 and June 7, 1997
_____	418	BRUNER, Rick E., Interactive, NetGravity to launch new ad service bureau, adage.com, December 1, 1997
_____	419	BRUNER, Rick E., Small networks chase per-click ad business, Aaddzz, ValueClick compete, AdvertisingAge, September 8, 1997
_____	420	ZEFF & ARONSON, Advertising on the Internet, 1st Edition, John Wiley & Sons (1997)
_____	421	ZEFF & ARONSON, Advertising on the Internet, 2nd Edition, John Wiley & Sons (1999)
_____	422	SINGH & DALAL, Web Home Pages as Advertisements, Communications of the ACM, August 1999
_____	423	The ClickZ Network, Netscape Selects Real Media's Open AdStream, January 13, 1998
_____	424	The ClickZ Network, Real Media Signs More Media Companies, December 30, 1998
_____	425	www.Spinbox.com webpages, www.archive.org, archived beginning Nov. 11, 1998
_____	426	BAKOS, The Emerging Role of Electronic Marketplaces on the Internet, Communications of the ACM, Vol. 41, No. 8, August 1998
_____	427	KOHDA, Youiji & ENDO, Susumu, Ubiquitous Advertising on the WWW, Merging Advertisement on the Browser, Fifth International World Wide Web Conference, May 1996
_____	428	Baseview.com webpages from www.archive.org, archived January 16, 1997
_____	429	AdForce, Inc., Quick Reference Guide, Version 3.0 (1999)
_____	430	AdManagerPro, Ad Entry Manual, User's Manual V.1.05, 1997
_____	431	AdManagerPro, Ad Entry Manual, User's Manual V.2.0, 1997
_____	432	AdManagerPro, Administration and Receivables Manual, Administrator's Manual Version 1.0.5, November 1997
_____	433	AdManagerPro, Administration and Receivables Manual, Administrator's Manual Version 2.0, June 1998

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
		Commercialization of the World Wide Web, The Role of Cookies, vanderbilt.edu
_____	434	web pages, archive.org, February 1, 1997
_____	435	CRABTREE et al., Adaptive Personal Agents, Personal Technologies (1998)
_____	436	DALTON, Gregory, OPS, Answer to Cookies?, InformationWeek, October 13, 1997
_____	437	DUNN et al., Privacy and Profiling on the Web, Submitted to W3C on June 2, 1997
_____	438	firefly.com web pages, archive.org, archived 1998-2000
		GILBERT, Jennifer, CMGI's Web dealmaking hits big-time, Advertising Age, October
_____	439	4, 1999
		HERSLEY et al., Implementation of OPS Over HTTP, Submitted to W3C on June 2,
_____	440	1997
		HERSLEY et al., Proposal for an Open Profiling Standard, Submitted to W3C on June
_____	441	2, 1997
		HERSLEY et al., Standard Practices of OPS Systems, Submitted to W3C on June 2,
_____	442	1997
		junkbusters.com web pages, archive.org, archived April 1997 through November
_____	443	2000
		KARPINSKI, Richard, ICE application to revolutionize e-commerce, Business
_____	444	Marketing, March 1998
_____	445	matchlogic.com web pages, archive.org, April 13, 2000
		ROSCHEISEN, R. Martin, A Network-Centric Design for Relationship-Based Rights
_____	446	Management (Dissertation), December 1997
		KRISTOL, D. & MONTULLI, L., HTTP State Management Mechanism, Internet
_____	447	Engineering Task Force, available at https://tools.ietf.org/html/rfc2109, February 1997
		BERNERS-LEE, TIM, Information Management: A Proposal, CERN, available at
_____	448	http://www.w3.org/History/1989/proposal.html, May 1990
		BERNERS-LEE, TIM, Uniform Resource Locators (URL), A Syntax for the Expression of Access Information of Objects on the Network, Internet Engineering Task Force,
_____	449	available at http://www.w3.org/Addressing/URL/url-spec.txt, March 21, 1994
		BERNERS-LEE, TIM, The Original HTTP as defined in 1991, W3C, available at
_____	450	http://www.w3.org/Protocols/HTTP/AsImplemented.html (1991)
		BERNERS-LEE, TIM, et al., Hypertext Transfer Protocol – HTTP/1.0, Network Working Group Request for Comments: 1945, Internet Engineering Task Force,
_____	451	available at https://tools.ietf.org/html/rfc1945, May 1996
		LEINER, BARRY M. ET AL., Brief History of the Internet, Internet Society, October
_____	452	15, 2012

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
		BERNERS-LEE, T. & CONNOLLY, D., Hypertext Markup Language – 2.0, Network Working Group Request for Comments: 1866, Internet Engineering Task Force,
_____	453	available at https://tools.ietf.org/html/rfc1866, November 1995
_____	454	TUNGATE, MARK, Adland-A Global History of Advertising, 2d Ed. (2013)
_____	455	JOHNSON, BRADLEY, Interactive, Planet Oasis Gives Web Sites Promotion Clout, Advertising Age, available at http://adage.com/article/news/interactive-planet-oasis-web-sites-promotion-clout/78467, July 8, 1996
_____	456	JONES, DUNCAN, The History of Pay Per Click (PPC) Marketing, Online Marketing Bog (www.duncanjonesnz.com) available at https://www.duncanjonesnz.com/history-of-ppc-marketing, September 23, 2014
_____	457	MOUKHEIBER, Z., DoubleClick is watching you, Forbes Magazine, November 4, 1996
_____	458	FIELDING, R. et al., Hypertext Transfer Protocol – HTTP/1.1, Network Working Group, RFC 2068, January 1997
_____	459	GLENN, A. and MILLMAN, D., Access Management of Web-based Services, Ariel Glenn, D-Lib Magazine, September 1998
_____	460	SAMAR, V., Single Sign-on Using Cookies for Web Applications, IEEE 8th Intl. Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises, June 16-18, 1999
_____	461	KIYOMITSU, H., An E-Commerce Technology: Web Personalization on the Patron Web, Kobe University Economic Review 46 (2000)
_____	462	WARKENTIN, M. and BAPNA, R., Intelligent Agent-Based Data Mining in Electronic Markets, AMCIS 2000 Proceedings (2000)
_____	463	BONCELLA, R., Internet Privacy – At Home and At Work, Communications of AIS, Vol. 7, Article 14, September 2001
_____	464	New Engage ECHO Remarketing Solution Aimed at Decreasing Customer Acquisition Costs, Improving Customer Retention, BusinessWire, September 25, 2000
_____	465	BILCHEV, G. and MARSTON, D., Personalized advertising – exploiting the distributed user profile, BT Technology Journal, Vol. 21, No. 1, January 2003
_____	466	PAN, Y. and ZINKHAN, G.M., Netvertising characteristics, challenges and opportunities: a research agenda, Int. J. Internet Marketing and Advertising, Vol. 1, No. 3, 2004

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	467	Google Launches Self-Service Advertising Program, http://googlepress.blogspot.com/2000/10/google-launches-self-service.html, October 23, 2000
_____	468	Google Introduces New Pricing For Popular Self-Service Online Advertising Program, http://googlepress.blogspot.com/2002/02/google-introduces-new-pricing-for.html, February 20, 2002
_____	469	AdWords Separate Content Bidding, http://googlepress.blogspot.com/2005/11/adwords-separate-content-bidding_22.html, November 22, 2005
_____	470	WashingtonPost.Newsweek Interactive Selects Revenue Science Audience SearchTM for Behavioral Targeting Services, https://www.audiencescience.com/washingtonpost-newsweek-interactive-selects-revenue-science-audience-search-for-behavioral-targeting-services/, September 20, 2004
_____	471	Revenue Science Hosts Behavioral Targeting Workshop at Interactive Advertising World, https://www.audiencescience.com/revenue-science-hosts-behavioral-targeting-workshop-at-interactive-advertising-world/, September 17, 2004
_____	472	Online Advertising and Privacy Survey Shows Consumers Hold Strong Preference for Targeted Advertising, https://www.audiencescience.com/online-advertising-and-privacy-survey-shows-consumers-hold-strong-preference-for-targeted-advertising/, September 9, 2004
_____	473	Yahoo's Invalidity Contentions (pages related to parent patents in this family); AlmondNet, Inc. et al. v. Yahoo! Inc., Civil Action No. 16-cv-1557-ILG-SMG, U.S. District Court for the Eastern District of New York; June 30, 2017
_____	474	Exhibits C-1 through C-12 (as revised) of Yahoo's Invalidity Contentions cited above; June 30, 2017 and August 22, 2017

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Electronic Acknowledgement Receipt

EFS ID:	45433134
Application Number:	17714965
International Application Number:	
Confirmation Number:	7004
Title of Invention:	Targeted television advertisements based on online behavior
First Named Inventor/Applicant Name:	Roy Shkedi
Customer Number:	26362
Filer:	Louis J. Hoffman/Michelle Lisa
Filer Authorized By:	Louis J. Hoffman
Attorney Docket Number:	14-TV1-CON7
Receipt Date:	08-APR-2022
Filing Date:	
Time Stamp:	16:07:49
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Power of Attorney	22-04-08-IIQ-TV1_CON7-POA.pdf	627682 22ddbcb840b2e9a5ee2bcaace63422b843f31b5b	no	2

Warnings:

Information:					
2	Transmittal Letter	22-04-08-IIQ-TV1_CON7-IDS_Cover.pdf	452782 f5f4a7a1ea87427ab5081658aebdd9198db7592b	no	2
Warnings:					
Information:					
3	Information Disclosure Statement (IDS) Form (SB08)	22-04-08-IIQ-TV_CON7-FORM_1449.pdf	605076 f76d629d74a333e401bd85072d9e8647c8d67a12	no	21
Warnings:					
Information:					
This is not an USPTO supplied IDS fillable form					
			Total Files Size (in bytes):	1685540	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING or 371(c) DATE	GRP ART UNIT	FIL FEE REC'D	ATTY.DOCKET.NO	TOT CLAIMS	IND CLAIMS
17/714,965	04/06/2022		0.00	14-TV1-CON7	51	3

CONFIRMATION NO. 7004

FILING RECEIPT

26362
LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, AZ 85260

OC000000133104240

Date Mailed: 04/12/2022

Receipt is acknowledged of this non-provisional utility patent application. The application will be taken up for examination in due course. Applicant will be notified as to the results of the examination. Any correspondence concerning the application must include the following identification information: the U.S. APPLICATION NUMBER, FILING DATE, NAME OF FIRST INVENTOR, and TITLE OF INVENTION. Fees transmitted by check or draft are subject to collection.

Please verify the accuracy of the data presented on this receipt. If an error is noted on this Filing Receipt, please submit a written request for a corrected Filing Receipt, including a properly marked-up ADS showing the changes with strike-through for deletions and underlining for additions. If you received a "Notice to File Missing Parts" or other Notice requiring a response for this application, please submit any request for correction to this Filing Receipt with your reply to the Notice. When the USPTO processes the reply to the Notice, the USPTO will generate another Filing Receipt incorporating the requested corrections provided that the request is grantable.

Inventor(s)

Roy Shkedi, Forest Hills, NY;

Applicant(s)

Intent IQ, LLC, Long Island City, NY;

Power of Attorney: The patent practitioners associated with Customer Number 26362

Domestic Priority data as claimed by applicant

This application is a CON of 16/925,655 07/10/2020 PAT 11303973
which is a CON of 16/228,479 12/20/2018 PAT 10715878
which is a CON of 15/803,631 11/03/2017 PAT 10178442
which is a CON of 15/179,356 06/10/2016 PAT 9813778
which is a CON of 14/246,525 04/07/2014 PAT 9369779
which is a DIV of 13/097,885 04/29/2011 PAT 8695032
which is a DIV of 12/860,666 08/20/2010 PAT 8281336
which is a CON of 11/736,544 04/17/2007 PAT 7861260

Foreign Applications for which priority is claimed (You may be eligible to benefit from the **Patent Prosecution Highway** program at the USPTO. Please see http://www.uspto.gov for more information.) - None.
Foreign application information must be provided in an Application Data Sheet in order to constitute a claim to foreign priority. See 37 CFR 1.55 and 1.76.

Permission to Access Application via Priority Document Exchange: Yes

Permission to Access Search Results: Yes

Applicant may provide or rescind an authorization for access using Form PTO/SB/39 or Form PTO/SB/69 as appropriate.

If Required, Foreign Filing License Granted: 04/11/2022

The country code and number of your priority application, to be used for filing abroad under the Paris Convention, is **US 17/714,965**

Projected Publication Date: To Be Determined - pending completion of Missing Parts

Non-Publication Request: No

Early Publication Request: No

Title

Targeted television advertisements based on online behavior

Preliminary Class

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications: No

PROTECTING YOUR INVENTION OUTSIDE THE UNITED STATES

Since the rights granted by a U.S. patent extend only throughout the territory of the United States and have no effect in a foreign country, an inventor who wishes patent protection in another country must apply for a patent in a specific country or in regional patent offices. Applicants may wish to consider the filing of an international application under the Patent Cooperation Treaty (PCT). An international (PCT) application generally has the same effect as a regular national patent application in each PCT-member country. The PCT process **simplifies** the filing of patent applications on the same invention in member countries, but **does not result** in a grant of "an international patent" and does not eliminate the need of applicants to file additional documents and fees in countries where patent protection is desired.

Almost every country has its own patent law, and a person desiring a patent in a particular country must make an application for patent in that country in accordance with its particular laws. Since the laws of many countries differ in various respects from the patent law of the United States, applicants are advised to seek guidance from specific foreign countries to ensure that patent rights are not lost prematurely.

Applicants also are advised that in the case of inventions made in the United States, the Director of the USPTO must issue a license before applicants can apply for a patent in a foreign country. The filing of a U.S. patent application serves as a request for a foreign filing license. The application's filing receipt contains further information and guidance as to the status of applicant's license for foreign filing.

Applicants may wish to consult the USPTO booklet, "General Information Concerning Patents" (specifically, the section entitled "Treaties and Foreign Patents") for more information on timeframes and deadlines for filing foreign patent applications. The guide is available either by contacting the USPTO Contact Center at 800-786-9199, or it can be viewed on the USPTO website at http://www.uspto.gov/web/offices/pac/doc/general/index.html.

For information on preventing theft of your intellectual property (patents, trademarks and copyrights), you may wish to consult the U.S. Government website, http://www.stopfakes.gov. Part of a Department of Commerce initiative, this website includes self-help "toolkits" giving innovators guidance on how to protect intellectual property in specific

countries such as China, Korea and Mexico. For questions regarding patent enforcement issues, applicants may call the U.S. Government hotline at 1-866-999-HALT (1-866-999-4258).

LICENSE FOR FOREIGN FILING UNDER

Title 35, United States Code, Section 184

Title 37, Code of Federal Regulations, 5.11 & 5.15

GRANTED

The applicant has been granted a license under 35 U.S.C. 184, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" followed by a date appears on this form. Such licenses are issued in all applications where the conditions for issuance of a license have been met, regardless of whether or not a license may be required as set forth in 37 CFR 5.15. The scope and limitations of this license are set forth in 37 CFR 5.15(a) unless an earlier license has been issued under 37 CFR 5.15(b). The license is subject to revocation upon written notification. The date indicated is the effective date of the license, unless an earlier license of similar scope has been granted under 37 CFR 5.13 or 5.14.

This license is to be retained by the licensee and may be used at any time on or after the effective date thereof unless it is revoked. This license is automatically transferred to any related applications(s) filed under 37 CFR 1.53(d). This license is not retroactive.

The grant of a license does not in any way lessen the responsibility of a licensee for the security of the subject matter as imposed by any Government contract or the provisions of existing laws relating to espionage and the national security or the export of technical data. Licensees should apprise themselves of current regulations especially with respect to certain countries, of other agencies, particularly the Office of Defense Trade Controls, Department of State (with respect to Arms, Munitions and Implements of War (22 CFR 121-128)); the Bureau of Industry and Security, Department of Commerce (15 CFR parts 730-774); the Office of Foreign AssetsControl, Department of Treasury (31 CFR Parts 500+) and the Department of Energy.

NOT GRANTED

No license under 35 U.S.C. 184 has been granted at this time, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" DOES NOT appear on this form. Applicant may still petition for a license under 37 CFR 5.12, if a license is desired before the expiration of 6 months from the filing date of the application. If 6 months has lapsed from the filing date of this application and the licensee has not received any indication of a secrecy order under 35 U.S.C. 181, the licensee may foreign file the application pursuant to 37 CFR 5.15(b).

SelectUSA

The United States represents the largest, most dynamic marketplace in the world and is an unparalleled location for business investment, innovation, and commercialization of new technologies. The U.S. offers tremendous resources and advantages for those who invest and manufacture goods here. Through SelectUSA, our nation works to promote and facilitate business investment. SelectUSA provides information assistance to the international investor community; serves as an ombudsman for existing and potential investors; advocates on behalf of U.S. cities, states, and regions competing for global investment; and counsels U.S. economic development organizations on investment attraction best practices. To learn more about why the United States is the best country in the world to develop

technology, manufacture products, deliver services, and grow your business, visit http://www.SelectUSA.gov or call +1-202-482-6800.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
17/714,965	04/06/2022	Roy Shkedi	14-TV1-CON7

CONFIRMATION NO. 7004

26362
LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, AZ 85260

FORMALITIES LETTER

OC000000133104241

Date Mailed: 04/12/2022

NOTICE TO FILE MISSING PARTS OF NONPROVISIONAL APPLICATION

FILED UNDER 37 CFR 1.53(b)

Filing Date Granted

Items Required To Avoid Abandonment:

An application number and filing date have been accorded to this application. The item(s) indicated below, however, are missing. Applicant is given **TWO MONTHS** from the date of this Notice within which to file all required items below to avoid abandonment. Extensions of time may be obtained by filing a petition accompanied by the extension fee under the provisions of 37 CFR 1.136(a).

- The statutory basic filing fee is missing.
- The application search fee must be submitted.
- The application examination fee must be submitted.
- Additional claim fees of $ **3100** as an undiscounted entity, including any required multiple dependent claim fee, are required. Applicant must submit the additional claim fees or cancel the additional claims for which fees are due.
- Surcharge as set forth in 37 CFR 1.16(f) must be submitted.
  The surcharge is due for any one of:
  - late submission of the basic filing fee, search fee, or examination fee,
  - late submission of inventor's oath or declaration,
  - filing an application that does not contain at least one claim on filing, or
  - submission of an application filed by reference to a previously filed application.

SUMMARY OF FEES DUE:

The fee(s) required within **TWO MONTHS** from the date of this Notice to avoid abandonment is/are itemized below. No entity status discount is in effect. If applicant is qualified for small entity status, a written assertion of small entity status must be submitted to establish small entity status. (See 37 CFR 1.27). If applicant is qualified for micro entity status, an acceptable Certification of Micro Entity Status must be submitted to establish micro entity status. (See 37 CFR 1.29 and forms PTO/SB/15A and 15B.)

- $ **320** basic filing fee.
- $ **160** surcharge.
- $ **700** search fee.
- $ **800** examination fee.
- $ **3100** for **31** total claims over 20.

- $(**0**) previous unapplied payment amount.
- $ **5080** TOTAL FEE BALANCE DUE.

Replies must be received in the USPTO within the set time period or must include a proper Certificate of Mailing or Transmission under 37 CFR 1.8 with a mailing or transmission date within the set time period. For more information and a suggested format, see Form PTO/SB/92 and MPEP 512.

Replies should be mailed to:

Mail Stop Missing Parts
Commissioner for Patents
P.O. Box 1450
Alexandria VA 22313-1450

Registered users of EFS-Web may alternatively submit their reply to this notice via EFS-Web, including a copy of this Notice and selecting the document description "Applicant response to Pre-Exam Formalities Notice". https://sportal.uspto.gov/authenticate/AuthenticateUserLocalEPF.html

For more information about EFS-Web please call the USPTO Electronic Business Center at 1-866-217-9197 or visit our website at http://www.uspto.gov/ebc.

If you are not using EFS-Web to submit your reply, you must include a copy of this notice.

Questions about the contents of this notice and the requirements it sets forth should be directed to the Office of Data Management, Application Assistance Unit, at **(571) 272-4000** or **(571) 272-4200** or **1-888-786-0101**.

/ptnguyen/

PATENT APPLICATION FEE DETERMINATION RECORD

Substitute for Form PTO-875

Application or Docket Number: 17/714,965

APPLICATION AS FILED - PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA	SMALL ENTITY RATE($)	SMALL ENTITY FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY FEE($)
BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A			N/A	320
SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A			N/A	700
EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A			N/A	800
TOTAL CLAIMS (37 CFR 1.16(i))	51 minus 20=	* 31			OR	x 100 =	3100
INDEPENDENT CLAIMS (37 CFR 1.16(h))	3 minus 3 =	*				x 480 =	0.00
APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).						0.00
MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))							0.00
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL			TOTAL	4920

APPLICATION AS AMENDED - PART II

AMENDMENT A

	(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	SMALL ENTITY RATE($)	SMALL ENTITY ADDITIONAL FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY ADDITIONAL FEE($)
Total (37 CFR 1.16(i))	*	Minus	**	=	x =		OR	x =	
Independent (37 CFR 1.16(h))	*	Minus	***	=	x =		OR	x =	
Application Size Fee (37 CFR 1.16(s))									
FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
					TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

AMENDMENT B

	(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE($)	ADDITIONAL FEE($)		RATE($)	ADDITIONAL FEE($)
Total (37 CFR 1.16(i))	*	Minus	**	=	x =		OR	x =	
Independent (37 CFR 1.16(h))	*	Minus	***	=	x =		OR	x =	
Application Size Fee (37 CFR 1.16(s))									
FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
					TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".
The "Highest Number Previously Paid For" (Total or Independent) is the highest found in the appropriate box in column 1.

To: Louis@valuablepatents.com,donald@valuablepatents.com,dsalavi@northwestpatent.com
From: PAIR_eOfficeAction@uspto.gov
Cc: PAIR_eOfficeAction@uspto.gov
Subject: Private PAIR Correspondence Notification for Customer Number 26362

Apr 12, 2022 05:06:19 AM

Dear PAIR Customer:

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, AZ 85260
UNITED STATES

The following USPTO patent application(s) associated with your Customer Number, 26362 , have new outgoing correspondence. This correspondence is now available for viewing in Private PAIR.

The official date of notification of the outgoing correspondence will be indicated on the form PTOL-90 accompanying the correspondence.

Disclaimer:
The list of documents shown below is provided as a courtesy and is not part of the official file wrapper. The content of the images shown in PAIR is the official record.

Application	Document	Mailroom Date	Attorney Docket No.
17714965	APP.FILE.REC	04/12/2022	14-TV1-CON7
	NTC.MISS.PRT	04/12/2022	14-TV1-CON7

To view your correspondence online or update your email addresses, please visit us anytime at https://sportal.uspto.gov/secure/myportal/privatepair.

If you have any questions, please email the Electronic Business Center (EBC) at EBC@uspto.gov with 'e-Office Action' on the subject line or call 1-866-217-9197 during the following hours:

Monday - Friday 6:00 a.m. to 12:00 a.m.

Thank you for prompt attention to this notice,

UNITED STATES PATENT AND TRADEMARK OFFICE

PATENT APPLICATION INFORMATION RETRIEVAL SYSTEM

UNITED STATES PATENT AND TRADEMARK OFFICE
COMMISSIONER FOR PATENTS
P.O.BOX 1450
ALEXANDRIA VA 22313-1451

PRESORTED
FIRST-CLASS MAIL
U.S. POSTAGE PAID
POSTEDIGITAL
NNNNN

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, AZ 85260

Courtesy Reminder for
Application Serial No: 17/714,965

Attorney Docket No: 14-TV1-CON7
Customer Number: 26362
Date of Electronic Notification: 04/12/2022

This is a courtesy reminder that new correspondence is available for this application. If you have not done so already, please review the correspondence. The official date of notification of the outgoing correspondence will be indicated on the form PTOL-90 accompanying the correspondence.

An email notification regarding the correspondence was sent to the following email address(es) associated with your customer number:

Louis@valuablepatents.com
donald@valuablepatents.com
dsalavi@northwestpatent.com

To view your correspondence online or update your email addresses, please visit us anytime at **https://ppair-my.uspto.gov/pair/PrivatePair.**
If you have any questions, please email the Electronic Business Center (EBC) at EBC@uspto.gov or call 1-866-217-9197.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor	:	Shkedi, Roy	Art Unit	:	2425 (in parent)
Serial No.	:	17/714,965	Examiner	:	Ekpo, Nnenna (in parent)
Filing Date	:	04/06/2022	Conf. No.	:	7004
Title	:	Targeted television advertisements based on online behavior			

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Filed via EFS-Web on July 29, 2022

RESPONSE TO FORMALITIES NOTICE

Dear Sir:

In response to the Notice of Missing Parts dated April 12, 2022, applicant requests entry of the following amendments and consideration of the following remarks. An information disclosure statement accompanies this paper.

Amendments to the specification begin on page **2**.

Amendments to the claims begin on page **8**.

Remarks are on page **15**.

The following fees are submitted concurrently:

Statutory basic filing fee – $320.00

Search fee – $700.00

Examination fee – $800.00

Surcharge under 37 C.F.R. § 1.16(f) – $160.00

Excess claims fees ((26 less 20) total claims x $100/claim) – $600.00

Extension of time (for the requested two months) – $640.00

TOTAL FEES = $3,220.00

Amendments

IN THE SPECIFICATION:

At page 1, before the "background" header, please add the following section:

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation of application Serial No. 16/925,655, filed July 10, 2020, now U.S. Patent 11,303,973, which is a continuation of application Serial No. 16/228,479, filed December 20, 2018, now U.S. Patent 10,715,878, which is a continuation of application Serial No. 15/803,631, filed November 3, 2017, now U.S. Patent 10,178,442, which is a continuation of application Serial No. 15/179,356, filed June 10, 2016, now U.S. Patent 9,813,778, which is a continuation of application Serial No. 14/246,525, filed April 7, 2014, now U.S. Patent 9,369,779, which is a division of application Serial No. 13/097,885, filed April 29, 2011, now U.S. Patent 8,695,032, which is a division of application Serial No. 12/860,666, filed August 20, 2010, now U.S. Patent 8,281,336, which is a continuation of application Serial No. 11/736,544, filed April 17, 2007, now U.S. Patent 7,861,260.

Please amend the following paragraphs as indicated:

[0012] Router – any device that acts as a junction between networks, to buffer and transfer data between or among them. For example, a router can be employed to connect a local area network (LAN) to the Internet, thereby enabling online user interface devices connected to the LAN to share a connection to the Internet through the router. The router receives data from devices on the LAN and transmits them to the Internet directed to their corresponding destinations, and receives data from the Internet ~~Internet~~ and directs them to the corresponding devices on the LAN.

[0028] Various systems are used currently for targeting advertisements based on user/viewer/customer behavior. Many of these rely on the collection of personally

identifiable information (PII) to enable correlation of the person exhibiting the behavior and advertisements targeted at that person. There are some examples wherein advertisements can be targeted without collecting PII, but in such examples it is typically the case that the medium of the behavior and that of the advertisement are the same. For example, many grocery stores hand out so-called "club cards" that need not be linked to PII. A shopper presents the card at checkout to receive various discounts, thereby allowing the store to link the list of purchased items to the card. As the system "learns" the shopper's purchasing habits, the system begins issuing coupons targeted at [[at]] purchases that the shopper has made previously or that the system predicts the shopper may wish to make based on past purchases. In another example, online advertisements are readily targeted based on an Internet user's online activities without using PII. The use of cookies enables an ad server to recognize an Internet site visitor who has been previously presented with ads by the ad server or who has conducted searches or accessed content at sites linked to the ad server. The ad server can target future advertising to the site visitor based on that previous activity. A user who has searched for airline tickets to southern California on an online travel site, for instance, can later receive targeted online advertisements for Disneyland, delivered perhaps while visiting some other online site, to the user's computer from an ad server that collected the user's search information from the online travel site.

[0036] ~~Figs. 7-11 and 12A-C~~ Figs. 7-10, 11A-11C, and 12 illustrate schematically a system for targeting television advertisements based on online behavior in which online access and television service are provided by different providers.

[0044] ISP/TVP 20 reports to CAS 40 that user XY123 is accessing the Internet using online access IP address aaa.bbb.ccc.ddd and has STB 36 using STB IP address eee.fff.ggg.hhh. These associations can be stored on CAS 40, or only by ISP/TVP 20, as desired. If user XY123 engages in online activity from IP address aaa.bbb.ccc.ddd (for

example, accessing an online site 50, as in Fig. 4), then online site 50 can electronically communicate (as in Fig. 5) an indication (electronic indicia) of the user's online activity to CAS 40, which in turn can cause a television advertisement to be directed to the STB IP address that is associated with online IP address aaa.bbb.ccc.ddd. In this example the online site 50 acts as a profile provider. Without relying on PII, CAS 40 knows that IP address aaa.bbb.ccc.ddd is associated with user XY123 and that user XY123 has a set-top box with IP address eee.fff.ggg.hhh, and CAS 40 causes the television advertisement to be directed to IP address eee.fff.ggg.hhh that indicates STB 36 (as in Figs. 6A-6C). In some implementations, the television advertisements are delivered to STB 36 using an ad server controlled by ISP/TVP 20, using either an IP-based protocol (an example of which is shown in the referenced figures) or a television delivery system (e.g., a dedicated cable advertisement channel). In one such implementation, that ad server report to CAS 40 the user pseudonym (XY123) and CAS 40 provides that ad server with TV ads targeting user XY123 based on profile information collected from IP addresses reported by ISP/TVP 20 as belonging to XY123. In another such implementation, that ad server reports the IP addresses used by each of a multitude of users of Internet services provided by ISP/TVP 20 along with the times and dates that those addresses were used, and CAS 40 provides that ad server with TV ads targeting each of the multitude of Internet users based on profile information received from the reported IP addresses at the corresponding reported times and dates. Personal privacy is maintained in these implementations because (1) no PII is used by CAS 40 in these implementations, and (2) although ISP/TVP 20 has the capacity to identify the users, the ad server of ISP/TVP 20 does not learn the profile information collected for and received by CAS 40.

[0045] The particular television advertisement thus directed typically is selected for delivery to the user's STB 36 based on the user's online behavior (e.g., merely accessing

the online site 50, viewing a particular web page, performing a particular search, clicking on a particular link or ad, online shopping or purchases, or other activity). Information about observed online behavior (i.e., indications or indicia of online activity) or collected demographic information can be used by the CAS or by the profile provider to direct a targeted TV ad. In the latter case, the CAS has just the IP address of the profiled user and the identity of the profile provider but not the observed online behavior or demographic information collected by the profile provider, so the CAS can direct a TV ad to the user merely by enabling the profile provider to arrange for the delivery of a targeted TV ad based on the observed online behavior or demographic information available to the profile provider. Examples of profile providers are given above.

[0046] The profile provider (such as the proprietor of online site 50) can send an electronic communication to the CAS, preferably automatically such as by having the profile provider's server pre-programmed to send such electronic communication, which electronic communication contains information that the CAS will use in causing a selected TV ad to be transferred to the STB at the proper time. The electronic communication received by the CAS contains information derived from the interaction of the user with an online site (i.e., online activity), which can be online site 50 or another online site for which the profile provider has access to data about user interactions. The electronic communication can take a variety of forms, including: (1) an indicator of the TV ad that should be shown, such as an URL or ad title given by the profile provider to retrieve the correct TV ad from the profile provider's ad server, or even the actual TV ad; (2) an indicator (indicia) of information about the user (likely not PII), such as (a) an IP address pointing to the user, preferably with time and date of visits; (b) a description of profile information about the user, or (c) a code or keyword used by the profile provider to fetch a visitor's profile from a database; or (3) an

indicator (indicia) describing user activity, such as a code or keyword indicating (a) certain categories of users, such as users likely interested in mortgages or users likely interested in travel, etc., or (b) more simply, what the visitor searched for on the online site, such as "car insurance." The electronic communication can be sent for each instance of a user interacting with the online site, or a list can be created and transmitted containing information about a number of user contacts. The electronic communication can include any combination of items 1-3 above; for example, a list of visitor IP addresses (preferably with date/time of visits) can be usefully paired with an indicator, in the form of a URL or code, identifying specific TV ads that the profile provider would like to deliver to each of the listed visitors.

[0065] For example, consider a common scenario where a laptop, used both at home and in the office, is linked to a home STB as a result of it being used at home. The day after the association was created, the laptop is taken to work, where it is used to visit a home mortgage site (using a different LAN or online access device than the one used at home). The home mortgage site redirects the laptop to the CAS along with a request to transmit a mortgage-related TV ad from one of the site's advertisers to the laptop user. The CAS, reading its own cookie that was placed on the laptop, checks its database and finds that the laptop is linked to a certain STB. In the evening, when the laptop user returns home and turns on the TV, the user will see the mortgage-related TV ad, even if ~~TV ad, even if~~ the user left the laptop at work that day.

[0076] As described previously herein, it is often the case that more than one computer or other online user interface device is connected to the Internet through a single modem or router, and all devices sharing that connection will appear to an online site to be connected through the same IP address at any given time. For example, in a home scenario, a parent and a child in the same household might each use their own computer to access the Internet, with the parent visiting sites pertaining to news or

finance and the child visiting sites pertaining to pop music or television. (The same is true in a business setting, where many computers may access websites through a common modem or router.) Each such site visited by any person on the LAN reports to the CAS the IP addresses and corresponding times and dates of each of those visits (which are the same at any given time for both computers). The STB in the same household reports to the CAS the common IP address used for online access. The CAS can now match the STB to the various sites visited. Television advertisements can be selected based not only on requests made by the visited sites (as already described) but also on content being viewed on the television when the TV advertisement is delivered. For example, if programming is being viewed on the television that is typically targeted at adults (e.g., news, financial reports, or adult drama) then television advertisements requested by the news or finance online sites are presented. Alternatively, if programming is being viewed on the television that is more typically targeted at children or teenagers, then television advertisements requested by the pop music or fashion online sites are presented. Both types of ads can be directed to the STB and stored for later viewing (during viewing of suitable television content) or can be directed to the STB and viewed as the suitable television content is viewed. The same discrimination between types of sites visited using multiple online access ~~multiple online access~~ devices can be employed regardless of whether the computers in the household are fixed or mobile devices, through the techniques discussed previously herein.

Amendments

IN THE CLAIMS:

Please amend the claims as follows:

Claims 1 – 51 (Canceled)

52. (New) A computer system comprising at least one hardware server computer, wherein the computer system is programmed and connected to perform a method of using common IP addresses to identify associations among multiple electronic devices and using an identified association to cause cross-device action, the method comprising:

(a) with the computer system, creating an electronic database containing a multitude of associations, each indicating an association between device identifiers of two or more Internet-accessing electronic devices of a multitude of Internet-accessing electronic devices, wherein:

(i) each of the associations is determined using stored data relating to Internet access by various of the multitude of devices, the stored data comprising (A) a device identifier of one of the Internet-accessing devices, (B) an IP address through which the one of the devices accessed the Internet, and (C) a timestamp identifying when the one of the devices accessed the Internet;

(ii) at least some of the associations indicate association between at least two devices of the two or more Internet-accessing electronic devices, which association is determined by using the stored data to determine that both of the at least two devices have accessed the Internet via a common IP address at respective times satisfying a predetermined time-relation criterion;

(iii) at least some of the associations indicate association between at least two devices of the two or more Internet-accessing electronic devices,

which association is determined without use of personally identifiable information; and

(iv) the computer system determining the associations accesses the Internet without using any of the common IP addresses of the associations identified in part (ii); and

(b) when the computer system and a first device identified in a first association are in communication with one another, with the computer system using a device identifier of the first device and user profile information associated with a device identifier of a second device identified in the first association to cause an automatic action to be taken with respect to the first device, wherein:

(i) the first association is one of the associations identified in part (a)(ii); and

(ii) the first association was stored in the electronic database described in part (a) prior to the computer system and the first device being in communication with one another as described in part (b).

53. (New) The computer system of claim 52 wherein the predetermined time-relation criterion comprises the criterion that one of the two or more Internet-accessing electronic devices referenced in part (a)(ii) accessed the Internet within a specified length of time after a time when another of the two or more Internet-accessing electronic devices referenced in part (a)(ii) accessed the Internet.

54. (New) The computer system of claim 52 wherein the predetermined time-relation criterion comprises the criterion that one of the two or more Internet-accessing electronic devices referenced in part (a)(ii) accessed the Internet between two times when another one of the two or more Internet-accessing electronic devices referenced in part (a)(ii) accessed the Internet.

55. (New) The computer system of claim 52 wherein the first association of part (b) is one of the associations identified in part (a)(iii).

56. (New) The computer system of claim 52 wherein the second device of part (b) comprises a mobile device that accessed the Internet both via the common IP address of the first association and via at least one different IP address, and wherein the user profile information of part (b) was collected when the mobile device accessed the Internet other than via the common IP address.

57. (New) The computer system of claim 56 wherein the user profile information of part (b) was collected after the computer system has determined the first association, and wherein the first association of part (b) is one of the associations identified in part (a)(iii).

58. (New) The computer system of claim 56 wherein the predetermined time-relation criterion comprises the criterion that one of the two or more Internet-accessing electronic devices referenced in part (a)(ii) accessed the Internet between two times when another one of the two or more Internet-accessing electronic devices referenced in part (a)(ii) accessed the Internet.

59. (New) The computer system of claim 58 wherein:

(A) the first device of part (b) is a television set-top box;

(B) the second device of part (b) is an online user interface device;

(C) the user profile information of part (b) was collected after the computer system has determined the first association; and

(D) the first association of part (b) is one of the associations identified in part (a)(iii).

60. (New) The computer system of claim 52 wherein the user profile information of part (b) was collected after the computer system has determined the first association.

61. (New) The computer system of claim 52 wherein the automatic action of part (b) comprises selecting an advertisement, from among a plurality of possible

advertisements, based at least in part on at least a portion of the user profile information associated with the device identifier of the second device, and causing the selected advertisement to be directed to the second device.

62. (New) The computer system of claim 61 wherein part (b) takes place when the first device is redirected to the computer system to receive delivery of an advertisement.

63. (New) The computer system of claim 52 wherein at least one of the device identifier of the first device of part (b) and the device identifier of the second device of part (b) is stored in a cookie readable by the computer system.

64. (New) The computer system of claim 52 wherein the device identifier of the first device of part (b) and the device identifier of the second device of part (b) include no personally identifiable information pertaining to a user of the first device or the second device.

65. (New) An article comprising a tangible data storage medium that is not a transitory propagating signal encoding computer-readable instructions that, when applied to a computer system comprising at least one hardware server computer, instruct the computer system to perform a method comprising:

(a) with the computer system, creating an electronic database containing a multitude of associations, each indicating an association between device identifiers of two or more Internet-accessing electronic devices of a multitude of Internet-accessing electronic devices, wherein:

(i) each of the associations is determined using stored data relating to Internet access by various of the multitude of devices, the stored data comprising (A) a device identifier of one of the Internet-accessing devices, (B) an IP address through which the one of the devices accessed the Internet, and (C) a timestamp identifying when the one of the devices accessed the Internet;

(ii) at least some of the associations indicate association between at least two devices of the two or more Internet-accessing electronic devices, which association is determined by using the stored data to determine that both of the at least two devices have accessed the Internet via a common IP address at respective times satisfying a predetermined time-relation criterion;

(iii) at least some of the associations indicate association between at least two devices of the two or more Internet-accessing electronic devices, which association is determined without use of personally identifiable information; and

(iv) the computer system determining the associations accesses the Internet without using any of the common IP addresses of the associations identified in part (ii); and

(b) when the computer system and a first device identified in a first association are in communication with one another, with the computer system using a device identifier of the first device and user profile information associated with a device identifier of a second device identified in the first association to cause an automatic action to be taken with respect to the first device, wherein:

(i) the first association is one of the associations identified in part (a)(ii); and

(ii) the first association was stored in the electronic database described in part (a) prior to the computer system and the first device being in communication with one another as described in part (b).

66. (New) The article of claim 65 wherein the predetermined time-relation criterion comprises the criterion that one of the two or more Internet-accessing electronic devices referenced in part (a)(ii) accessed the Internet within a specified length of time after a

time when another of the two or more Internet-accessing electronic devices referenced in part (a)(ii) accessed the Internet.

67. (New) The article of claim 65 wherein the predetermined time-relation criterion comprises the criterion that one of the two or more Internet-accessing electronic devices referenced in part (a)(ii) accessed the Internet between two times when another one of the two or more Internet-accessing electronic devices referenced in part (a)(ii) accessed the Internet.

68. (New) The article of claim 65 wherein the first association of part (b) is one of the associations identified in part (a)(iii).

69. (New) The article of claim 65 wherein the second device of part (b) comprises a mobile device that accessed the Internet both via the common IP address of the first association and via at least one different IP address, and wherein the user profile information of part (b) was collected when the mobile device accessed the Internet other than via the common IP address.

70. (New) The article of claim 69 wherein the user profile information of part (b) was collected after the computer system has determined the first association, and wherein the first association of part (b) is one of the associations identified in part (a)(iii).

71. (New) The article of claim 69 wherein the predetermined time-relation criterion comprises the criterion that one of the two or more Internet-accessing electronic devices referenced in part (a)(ii) accessed the Internet between two times when another one of the two or more Internet-accessing electronic devices referenced in part (a)(ii) accessed the Internet.

72. (New) The article of claim 71 wherein:

(A) the first device of part (b) is a television set-top box;

(B) the second device of part (b) is an online user interface device;

(C) the user profile information of part (b) was collected after the computer system has determined the first association; and

(D) the first association of part (b) is one of the associations identified in part (a)(iii).

73. (New) The article of claim 65 wherein the user profile information of part (b) was collected after the computer system has determined the first association.

74. (New) The article of claim 65 wherein the automatic action of part (b) comprises selecting an advertisement, from among a plurality of possible advertisements, based at least in part on at least a portion of the user profile information associated with the device identifier of the second device, and causing the selected advertisement to be directed to the second device.

75. (New) The article of claim 74 wherein part (b) takes place when the first device is redirected to the computer system to receive delivery of an advertisement.

76. (New) The article of claim 65 wherein at least one of the device identifier of the first device of part (b) and the device identifier of the second device of part (b) is stored in a cookie readable by the computer system.

77. (New) The article of claim 65 wherein the device identifier of the first device of part (b) and the device identifier of the second device of part (b) include no personally identifiable information pertaining to a user of the first device or the second device.

Remarks

Assignee presents new claims 52-77 for prosecution in this application. Claims 52 and 65 are independent claims, respectively, a computer system claim and a *Beauregard* claim. A continuation application is being filed to pursue method counterparts.

Amendments to the Specification

Assignee amends the specification to correct obvious typographical errors and to conform certain wording at the examiner's suggestion in parents. These amendments to the specification were entered in parent applications and approved as not containing any "new matter." Additional amendments to the specification update the "related applications" section and hence also do not contain any "new matter."

Amendments to the Claims

The claims are variants of claims of U.S. Patents 11,202,973 (parent), 10,715,878 (grandparent), and 8,677,398 (ancestor). Comments pertinent to specification support may be found in the Remarks filed 12/9/2020 in parent application (SN 16/925,655), which may be relevant as to parallel claim language, but further support exists as well.

Applicant respectfully submits that the claims should be in condition for prompt allowance. Please feel free to telephone the undersigned if it would in any way advance prosecution of this application.

Respectfully submitted,

INTENT IQ, LLC

by its attorney

Dated: July 29, 2022

/Louis J. Hoffman/
Louis J. Hoffman
Reg. No. 38,918

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane, Suite 2
Scottsdale, Arizona 85260
(480) 948-3295

Electronic Patent Application Fee Transmittal

Application Number:	17714965
Filing Date:	06-Apr-2022
Title of Invention:	Targeted television advertisements based on online behavior
First Named Inventor/Applicant Name:	Roy Shkedi
Filer:	Louis J. Hoffman
Attorney Docket Number:	14-TV1-CON7

Filed as Large Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
UTILITY APPLICATION FILING	1011	1	320	320
UTILITY SEARCH FEE	1111	1	700	700
UTILITY EXAMINATION FEE	1311	1	800	800
Pages:				
Claims:				
CLAIMS IN EXCESS OF 20	1202	6	100	600
Miscellaneous-Filing:				
LATE FILING FEE FOR OATH OR DECLARATION	1051	1	160	160

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				
Extension - 2 months with $0 paid	1252	1	640	640
Miscellaneous:				
		Total in USD ($)		**3220**

Electronic Acknowledgement Receipt	
EFS ID:	46308223
Application Number:	17714965
International Application Number:	
Confirmation Number:	7004
Title of Invention:	Targeted television advertisements based on online behavior
First Named Inventor/Applicant Name:	Roy Shkedi
Customer Number:	26362
Filer:	Louis J. Hoffman
Filer Authorized By:	
Attorney Docket Number:	14-TV1-CON7
Receipt Date:	29-JUL-2022
Filing Date:	06-APR-2022
Time Stamp:	18:49:16
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	CARD
Payment was successfully received in RAM	$3220
RAM confirmation Number	E20227SI52091181
Deposit Account	
Authorized User	
The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:	

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1		22-07-29-IIQ-TV1_CON7-RT_NtcMissingParts-ClsSpecRelApps.pdf	154667 3f4c38884516fa031904e4b04dbe433b8384c5a2	yes	15

Multipart Description/PDF files in .zip description

Document Description	Start	End
Applicant Response to Pre-Exam Formalities Notice	1	1
Specification	2	7
Claims	8	14
Applicant Arguments/Remarks Made in an Amendment	15	15

Warnings:

Information:

2	Fee Worksheet (SB06)	fee-info.pdf	51086 4ff89f161c0176894f7845b4676d8d43ae601ead	no	2

Warnings:

Information:

Total Files Size (in bytes):	205753

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PATENT APPLICATION FEE DETERMINATION RECORD

Substitute for Form PTO-875

Application or Docket Number: 17/714,965

APPLICATION AS FILED - PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA	SMALL ENTITY RATE($)	SMALL ENTITY FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY FEE($)
BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A			N/A	320
SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A			N/A	700
EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A			N/A	800
TOTAL CLAIMS (37 CFR 1.16(i))	26 minus 20=	* 6			OR	x 100 =	600
INDEPENDENT CLAIMS (37 CFR 1.16(h))	2 minus 3 =	*				x 480 =	0.00
APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).						0.00
MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))							0.00
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL			TOTAL	2420

APPLICATION AS AMENDED - PART II

		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	SMALL ENTITY RATE($)	SMALL ENTITY ADDITIONAL FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY ADDITIONAL FEE($)
AMENDMENT A	Total (37 CFR 1.16(i))	*	Minus	**	=	x =		OR	x =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	x =		OR	x =	
	Application Size Fee (37 CFR 1.16(s))									
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	SMALL ENTITY RATE($)	SMALL ENTITY ADDITIONAL FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY ADDITIONAL FEE($)
AMENDMENT B	Total (37 CFR 1.16(i))	*	Minus	**	=	x =		OR	x =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	x =		OR	x =	
	Application Size Fee (37 CFR 1.16(s))									
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".
The "Highest Number Previously Paid For" (Total or Independent) is the highest found in the appropriate box in column 1.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING or 371(c) DATE	GRP ART UNIT	FIL FEE REC'D	ATTY.DOCKET.NO	TOT CLAIMS	IND CLAIMS
17/714,965	04/06/2022		2580	14-TV1-CON7	26	2

CONFIRMATION NO. 7004

UPDATED FILING RECEIPT

26362
LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, AZ 85260

OC000000135527281

Date Mailed: 08/02/2022

Receipt is acknowledged of this non-provisional utility patent application. The application will be taken up for examination in due course. Applicant will be notified as to the results of the examination. Any correspondence concerning the application must include the following identification information: the U.S. APPLICATION NUMBER, FILING DATE, NAME OF FIRST INVENTOR, and TITLE OF INVENTION. Fees transmitted by check or draft are subject to collection.

Please verify the accuracy of the data presented on this receipt. If an error is noted on this Filing Receipt, please submit a written request for a corrected Filing Receipt, including a properly marked-up ADS showing the changes with strike-through for deletions and underlining for additions. If you received a "Notice to File Missing Parts" or other Notice requiring a response for this application, please submit any request for correction to this Filing Receipt with your reply to the Notice. When the USPTO processes the reply to the Notice, the USPTO will generate another Filing Receipt incorporating the requested corrections provided that the request is grantable.

Inventor(s)

Roy Shkedi, Forest Hills, NY;

Applicant(s)

Intent IQ, LLC, Long Island City, NY;

Power of Attorney: The patent practitioners associated with Customer Number 26362

Domestic Priority data as claimed by applicant

This application is a CON of 16/925,655 07/10/2020 PAT 11303973
which is a CON of 16/228,479 12/20/2018 PAT 10715878
which is a CON of 15/803,631 11/03/2017 PAT 10178442
which is a CON of 15/179,356 06/10/2016 PAT 9813778
which is a CON of 14/246,525 04/07/2014 PAT 9369779
which is a DIV of 13/097,885 04/29/2011 PAT 8695032
which is a DIV of 12/860,666 08/20/2010 PAT 8281336
which is a CON of 11/736,544 04/17/2007 PAT 7861260

Foreign Applications for which priority is claimed (You may be eligible to benefit from the **Patent Prosecution Highway** program at the USPTO. Please see http://www.uspto.gov for more information.) - None.
Foreign application information must be provided in an Application Data Sheet in order to constitute a claim to foreign priority. See 37 CFR 1.55 and 1.76.

Permission to Access Application via Priority Document Exchange: Yes

Permission to Access Search Results: Yes

Applicant may provide or rescind an authorization for access using Form PTO/SB/39 or Form PTO/SB/69 as appropriate.

If Required, Foreign Filing License Granted: 04/11/2022

The country code and number of your priority application, to be used for filing abroad under the Paris Convention, is **US 17/714,965**

Projected Publication Date: 11/10/2022

Non-Publication Request: No

Early Publication Request: No

Title

Targeted television advertisements based on online behavior

Preliminary Class

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications: No

PROTECTING YOUR INVENTION OUTSIDE THE UNITED STATES

Since the rights granted by a U.S. patent extend only throughout the territory of the United States and have no effect in a foreign country, an inventor who wishes patent protection in another country must apply for a patent in a specific country or in regional patent offices. Applicants may wish to consider the filing of an international application under the Patent Cooperation Treaty (PCT). An international (PCT) application generally has the same effect as a regular national patent application in each PCT-member country. The PCT process **simplifies** the filing of patent applications on the same invention in member countries, but **does not result** in a grant of "an international patent" and does not eliminate the need of applicants to file additional documents and fees in countries where patent protection is desired.

Almost every country has its own patent law, and a person desiring a patent in a particular country must make an application for patent in that country in accordance with its particular laws. Since the laws of many countries differ in various respects from the patent law of the United States, applicants are advised to seek guidance from specific foreign countries to ensure that patent rights are not lost prematurely.

Applicants also are advised that in the case of inventions made in the United States, the Director of the USPTO must issue a license before applicants can apply for a patent in a foreign country. The filing of a U.S. patent application serves as a request for a foreign filing license. The application's filing receipt contains further information and guidance as to the status of applicant's license for foreign filing.

Applicants may wish to consult the USPTO booklet, "General Information Concerning Patents" (specifically, the section entitled "Treaties and Foreign Patents") for more information on timeframes and deadlines for filing foreign patent applications. The guide is available either by contacting the USPTO Contact Center at 800-786-9199, or it can be viewed on the USPTO website at http://www.uspto.gov/web/offices/pac/doc/general/index.html.

For information on preventing theft of your intellectual property (patents, trademarks and copyrights), you may wish to consult the U.S. Government website, http://www.stopfakes.gov. Part of a Department of Commerce initiative, this website includes self-help "toolkits" giving innovators guidance on how to protect intellectual property in specific

countries such as China, Korea and Mexico. For questions regarding patent enforcement issues, applicants may call the U.S. Government hotline at 1-866-999-HALT (1-866-999-4258).

LICENSE FOR FOREIGN FILING UNDER

Title 35, United States Code, Section 184

Title 37, Code of Federal Regulations, 5.11 & 5.15

GRANTED

The applicant has been granted a license under 35 U.S.C. 184, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" followed by a date appears on this form. Such licenses are issued in all applications where the conditions for issuance of a license have been met, regardless of whether or not a license may be required as set forth in 37 CFR 5.15. The scope and limitations of this license are set forth in 37 CFR 5.15(a) unless an earlier license has been issued under 37 CFR 5.15(b). The license is subject to revocation upon written notification. The date indicated is the effective date of the license, unless an earlier license of similar scope has been granted under 37 CFR 5.13 or 5.14.

This license is to be retained by the licensee and may be used at any time on or after the effective date thereof unless it is revoked. This license is automatically transferred to any related applications(s) filed under 37 CFR 1.53(d). This license is not retroactive.

The grant of a license does not in any way lessen the responsibility of a licensee for the security of the subject matter as imposed by any Government contract or the provisions of existing laws relating to espionage and the national security or the export of technical data. Licensees should apprise themselves of current regulations especially with respect to certain countries, of other agencies, particularly the Office of Defense Trade Controls, Department of State (with respect to Arms, Munitions and Implements of War (22 CFR 121-128)); the Bureau of Industry and Security, Department of Commerce (15 CFR parts 730-774); the Office of Foreign AssetsControl, Department of Treasury (31 CFR Parts 500+) and the Department of Energy.

NOT GRANTED

No license under 35 U.S.C. 184 has been granted at this time, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" DOES NOT appear on this form. Applicant may still petition for a license under 37 CFR 5.12, if a license is desired before the expiration of 6 months from the filing date of the application. If 6 months has lapsed from the filing date of this application and the licensee has not received any indication of a secrecy order under 35 U.S.C. 181, the licensee may foreign file the application pursuant to 37 CFR 5.15(b).

SelectUSA

The United States represents the largest, most dynamic marketplace in the world and is an unparalleled location for business investment, innovation, and commercialization of new technologies. The U.S. offers tremendous resources and advantages for those who invest and manufacture goods here. Through SelectUSA, our nation works to promote and facilitate business investment. SelectUSA provides information assistance to the international investor community; serves as an ombudsman for existing and potential investors; advocates on behalf of U.S. cities, states, and regions competing for global investment; and counsels U.S. economic development organizations on investment attraction best practices. To learn more about why the United States is the best country in the world to develop

technology, manufacture products, deliver services, and grow your business, visit http://www.SelectUSA.gov or call +1-202-482-6800.

To: Louis@valuablepatents.com,donald@valuablepatents.com,dsalavi@northwestpatent.com
From: PAIR_eOfficeAction@uspto.gov
Cc: PAIR_eOfficeAction@uspto.gov
Subject: Private PAIR Correspondence Notification for Customer Number 26362

Aug 02, 2022 10:18:45 AM

Dear PAIR Customer:

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, AZ 85260
UNITED STATES

The following USPTO patent application(s) associated with your Customer Number, 26362 , have new outgoing correspondence. This correspondence is now available for viewing in Private PAIR.

The official date of notification of the outgoing correspondence will be indicated on the form PTOL-90 accompanying the correspondence.

Disclaimer:
The list of documents shown below is provided as a courtesy and is not part of the official file wrapper. The content of the images shown in PAIR is the official record.

Application	Document	Mailroom Date	Attorney Docket No.
17714965	APP.FILE.REC	08/02/2022	14-TV1-CON7

To view your correspondence online or update your email addresses, please visit us anytime at https://sportal.uspto.gov/secure/myportal/privatepair.

If you have any questions, please email the Electronic Business Center (EBC) at EBC@uspto.gov with 'e-Office Action' on the subject line or call 1-866-217-9197 during the following hours:

Monday - Friday 6:00 a.m. to 12:00 a.m.

Thank you for prompt attention to this notice,

UNITED STATES PATENT AND TRADEMARK OFFICE

PATENT APPLICATION INFORMATION RETRIEVAL SYSTEM

PTO/SB/08a (01-22)
Approved for use through 05/31/2024. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (**Not for submission under 37 CFR 1.99)**		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Kekpo, Nnenna
	Attorney Docket Number	TV1-CON7

U.S.PATENTS						
Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	7843923		2010-11-30	Baum	
	2	7899862		2011-03-11	Appelman et al.	
	3	6034689		2000-03-07	White et al.	
	4	7958224		2011-06-07	Amjadi	
	5	8131271		2012-03-06	Ramer et al.	
	6	6915271		07/05/2005	Meyer et al.	
	7	7251628		07/31/2007	Edlund et al.	
	8	5835896		11/10/1998	Fisher et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Kekpo, Nnenna
	Attorney Docket Number	TV1-CON7

	9	7356045		2008-04-08	Satapati et al.	
	10	8180674		2012-05-15	Madhavan	
	11	9495682		2016-11-15	Vallier	
	12	7908645		2011-03-15	Varghese	
	13	7299195		2007-11-20	Tawakol et al.	
	14	7805332		2010-09-28	Wilson	
	15	7894448		2011-02-22	Lillibridge et al.	
	16	8296437		2012-10-23	Pankratov	
	17	7853255		2010-12-14	Karaoguz et al.	
	18	8040819		2011-10-18	Harrington et al.	
	19	7039699		2006-05-02	Narin et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Kekpo, Nnenna
	Attorney Docket Number	TV1-CON7

	20	11301898		2022-04-12	Shkedi	
	21	11303973		2022-04-12	Shkedi et al.	
	22	11368763		2022-06-21	Shkedi et al.	

If you wish to add additional U.S. Patent citation information please click the Add button.

U.S.PATENT APPLICATION PUBLICATIONS

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20070157247		07/05/2005	Cordray et al.	
	2	20080113674		2008-05-15	Baig	
	3	20070067459		2007-03-22	Baal-Haness	
	4	20030229893		2003-12-11	Sgaraglino	
	5	20080281697		2008-11-13	Whitehead	
	6	20060069616		2006-03-30	Bau	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Kekpo, Nnenna
	Attorney Docket Number	TV1-CON7

	7	20040215515		2004-10-28	Perry	
	8	20050086109		2005-04-21	McFadden et al.	
	9	20070061331		2007-03-15	Ramer et al.	
	10	20040147265		2004-07-29	Kelley et al.	
	11	20010020242		2001-09-06	Gupta et al.	
	12	20070260514		2007-11-08	Burdick et al.	
	13	20060242294		2006-10-26	Damick et al.	
	14	20040107287		2004-06-03	Ananda et al	
	15	20030207685		2003-11-06	Rankin	
	16	20040133689		2004-07-08	Vasisht	
	17	20020116313		2002-08-22	Detering	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Kekpo, Nnenna
	Attorney Docket Number	TV1-CON7

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] ¡	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button

NON-PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	Fletcher et al., Practical Web Traffic Analysis: Standard, Privacy, Techniques, Results (2002); Chapters 1 and 3	☐
	2	Microsoft Computer Dictionary (5th ed. 2002); pp. vii-viii, 444	☐
	3	ZWICKY, Building Internet Firewalls (2d ed. 2000): table of contents, ch. 5	☐
	4	Hahn, The Internet Complete Reference (2d ed. 1996); chapters 10, 25, 26	☐
	5	P. Srisuresh et al., Request for Comments (RFC) 2663, IP Network Translator (NAT) Terminology and Considerations, August 1999 (“RFC 2663”)	☐
	6	Request for Comments (RFC) 793, Transmission Control Protocol, DARPA Internet Program Protocol Specification, September 1981 (“RFC 793”)	☐

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Kekpo, Nnenna
	Attorney Docket Number	TV1-CON7

	7	Robbin Zeff & Brad Aronson, Advertising on the Internet (2d ed. 1999) ("Zeff"); chapters 1, 2, 4, 6, 7, 9, and 10	☐
	8	RFC 950, "Internet Standard Subnetting Procedure," August 1985; pp. 1-18	☐
	9	RFC 1531, "Dynamic Host Configuration Protocol," October 1993; pp. 1-39	☐
	10	Definition of "modem," H. Newton, Newton's Telecom Dictionary, CMP Books, 2000	☐

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Kekpo, Nnenna
	Attorney Docket Number	TV1-CON7

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Louis J. Hoffman/	Date (YYYY-MM-DD)	2022-08-02
Name/Print	Louis J. Hoffman	Registration Number	38918

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

TK
5105.888
P739
2002
AUXILIARY
US

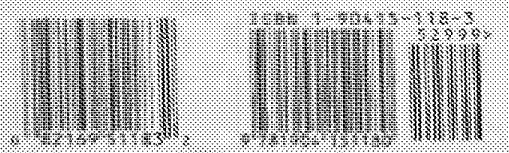
Meta's Exhibit No. 1011
Page 001

Summary of Contents

Introduction		1
Chapter 1:	Analytics Techniques	9
Chapter 2:	Database Log Analysis	39
Chapter 3:	Privacy	61
Chapter 4:	BBC News Online	87
Chapter 5:	eBay	111
Chapter 6:	ASPToday	139
Index		163



Visit www.glasshaus.com



Practical Web Traffic Analysis:

Standards, Privacy, Techniques, Results

Peter Fletcher

Alex Poon

Ben Pearce

Peter Comber

© 2002 glasshaus

Published by glasshaus Ltd,
Arden House,
1102 Warwick Road,
Acocks Green,
Birmingham,
B27 6BH, UK

Printed in the United States
ISBN 1-904151-18-3

Practical Web Traffic Analysis:

Standards, Privacy, Techniques, Results

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical articles or reviews.

The authors and publisher have made every effort in the preparation of this book to ensure the accuracy of the information. However, the information contained in this book is sold without warranty, either express or implied. Neither the authors, glasshaus nor its dealers or distributors will be held liable for any damages caused or alleged to be caused either directly or indirectly by this book.

Cover Image

James Chew is currently studying Multimedia Computing in Melbourne, Australia. Graphic Design started out just as a hobby, but quickly grew to become somewhat of an obsession in extending his knowledge and skills. Apart from design, James plays guitar and draws inspiration from music. James also enjoys designing, and redesigning his website *http://chewman.plastiqueweb.com* which is an online folio of work he has done.

James can be contacted via his website or through e-mail, *chewman@studiowhiz.com*.



labour-saving devices for web professionals

© 2002 glasshaus

Trademark Acknowledgements

glasshaus has endeavored to provide trademark information about all the companies and products mentioned in this book by the appropriate use of capitals. However, glasshaus cannot guarantee the accuracy of this information.

eBay™ screenshots used by permission of eBay™ Inc.

BBC News is a trademark of the British Broadcasting Corporation and is used under license. Screenshots and logos used by permission of the BBC.

ASPToday screenshots used by permission of ASPToday.

Credits

Authors
Peter Fletcher
Alex Poon
Ben Pearce
Peter Comber

Additional Material
SmartGirl.org
The Onion
NYTimes.com

Technical Reviewers
Dick Bennett
Jon Duckett
Richard Foan
Mark Horner
Tim Luoma
David Werthaimer

Proof Reader
Agnes Wiggers

Commissioning Editors
Peter Fletcher
Amanda Kay

Technical Editors
Mark Waterhouse
Alessandro Ansa

Managing Editor
Liz Toy

Publisher
Viv Emery

Project Manager
Sophie Edwards

Production Coordinators
Pip Wonson
Rachel Taylor

Cover
Dawn Chellingworth

Cover Image
James Chew

Indexer
Bill Johncocks

About the Authors

Peter Comber



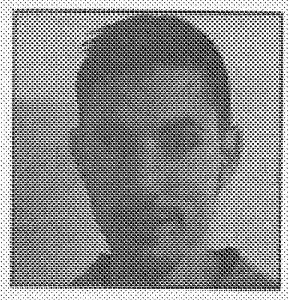
It is nearly seven years since Pete Comber started working with the Internet. Having just started working in direct marketing, like most junior execs he quickly became frustrated with the fragile, incomplete, and temperamental customer databases upon which the ambitious DM strategies tended to exist. With the Internet boom in the late 1990s, he became convinced that there was massive potential for the Internet to provide large volumes of data not just about what people purchased, but about what they looked at, what they discarded, and potentially granting some insight into how people came to make their buying decisions. Having analyzed Internet data for four years for the UK's biggest motoring web site, he joined the BBC at the end of 2001 in the hope that would be able to provide some useful insights into how people use the BBC News Online service.

Peter Fletcher



Peter Fletcher has been working in web development since 1997, via a degree in Philosophy and Theatre Studies and an MSc in Cognitive Science. Interests, professional and recreational, include using the Web as a distributed communications medium, analyzing web traffic data, and working with experimental performance group Stan's Cafe. His personal projects are documented at *www.joyfeed.com*. He is now a freelance web consultant and writer, dividing his time between Birmingham and Barcelona.

Ben Pearce



There's plenty about me later on, so I'd like to just say some thank-yous...

I would firstly like to thank Pete Fletcher for giving me the opportunity to talk about my work in this book. I hope it is useful and interesting reading for you.

I would also like to thank my wonderful wife Ange for all her support (not to mention the cakes!), and of course Jesus Christ who is my inspiration in everything.

Alex Poon



Alex Poon, originally from Baton Rouge, Louisiana, now lives with his wife Buffy and son Tyler in Northern California. He has a PhD from Stanford University. He was one of the first engineers at eBay, and started off as the de facto "UI guy" back in 1997. He started eBay's first user-interface group at eBay in 1999, then later ran the Advanced Technologies Group, during which his team implemented eBay's first web analytics system. Although Alex recently left eBay after five years, the company was happy to have him describe its analytics process in his own words.

To Buffy, who means everything to me.

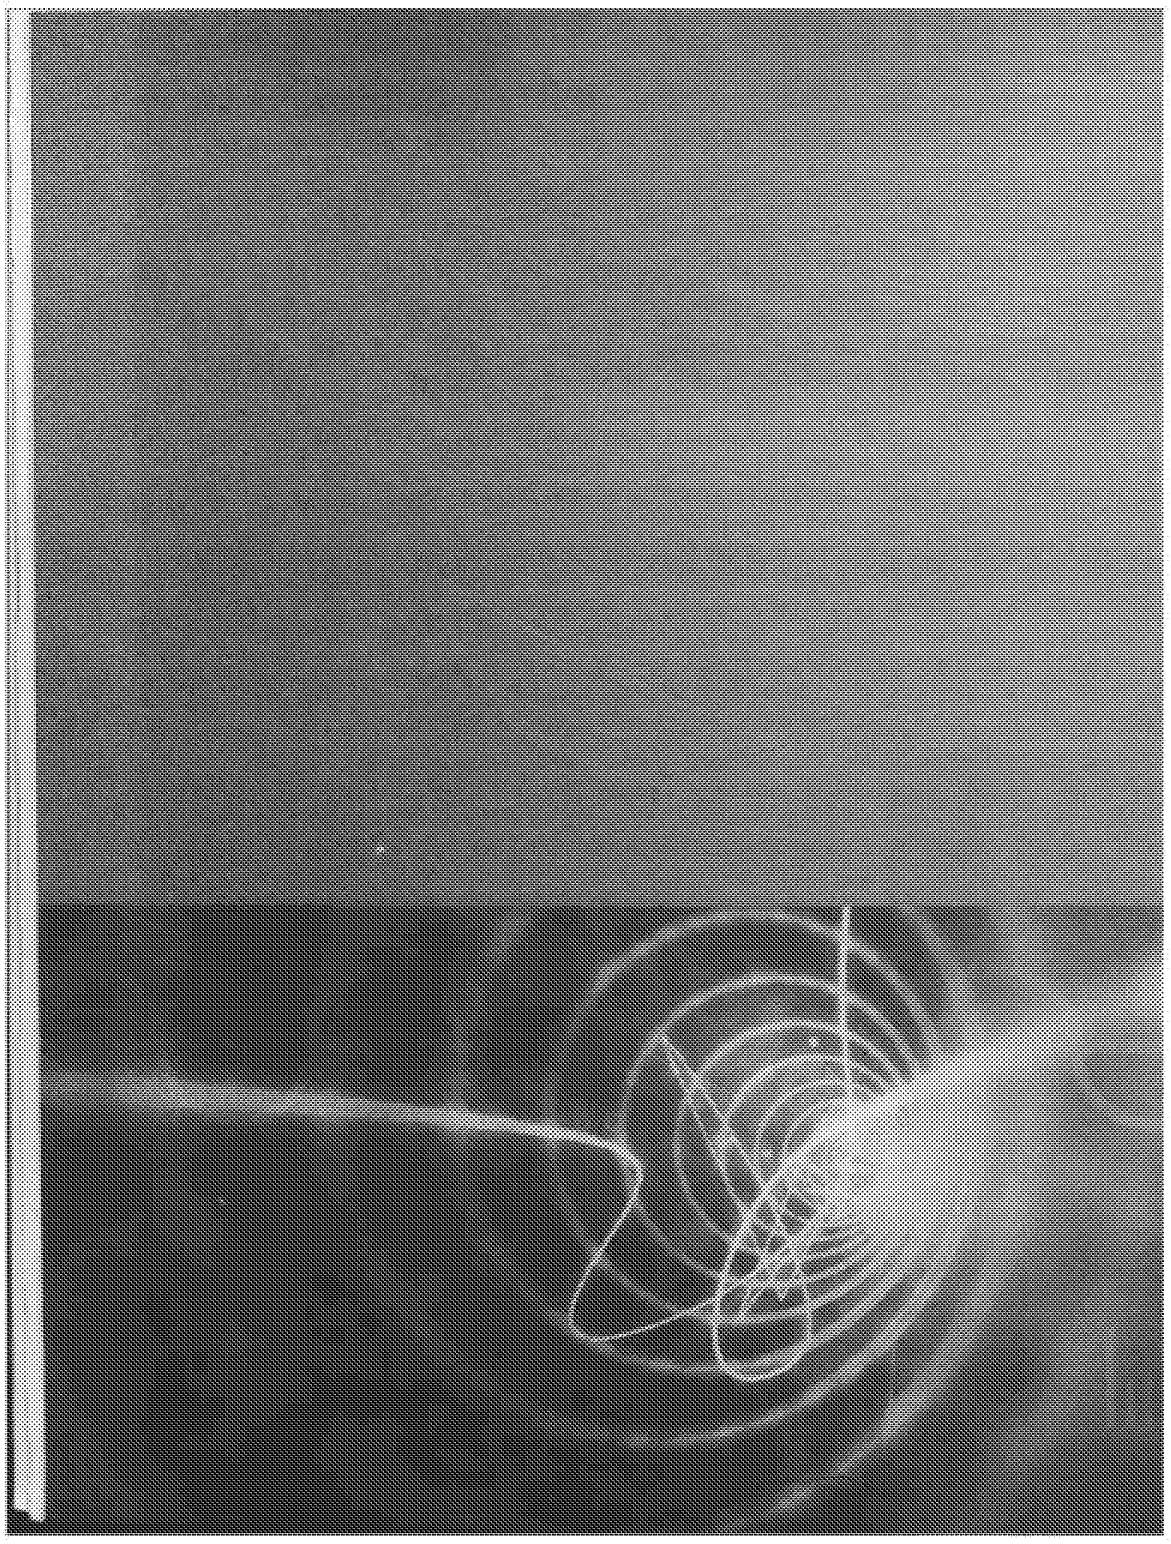


Table of Contents

Introduction 1

What's It all About? 1

Who is this Book for? 3

What's Inside? 3

Support and Feedback 6

Web Support 6

Analytics Techniques 9

1. Comparative Overview of Analysis Methods 9

Server Logs 9

Advantages 10

Disadvantages 10

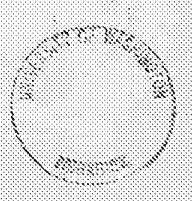
Conclusion 11

Panel Surveys 11

Advantages 11

Disadvantages 12

Conclusion 13



Browser Analysis: Tagging, or Web Bugs....13
Advantages....13
Disadvantages....14
Conclusion....15
The Skeptical Analyst....15
Log Files....15
What Does the Log Not Tell Us?....17
What Can We Say for Certain?....20
Analog....20
Cookies....29
Standards for Web Traffic Analysis....31
Metrics....32
Summary....36

2. Database Log Analysis....39
Dealing with the Token....40
Token Creation....40
Cookie Management....41
The Database....43
Admin Tables....44
Data Tables....44
The Import Logs Module....44
The Import Queries....46

Meta's Exhibit No. 1011
Page 0010

Starting the Import Process (frmRunImport) 48
Specifying the Format (spcLogFile) 48
Adding the Cookie Data (qryAppendCookies) 49
Adding the Page View Data (qryAppendPageViews) 50
Adding the Referrer Data (qryAppendReferers) 52
Adding the Click-Through Data (qryAppendLinkClickThroughs) 53
Clearing the tblLogs (qryDeleteLogsContent) 53
Finding Our Largest Referrer (qryRefererByDate) 54
Page Views by Date 55
Unique Cookie Visitors by Day 58
Total Clicks on a Particular Tracking Link 58
Summary 59

3. Privacy 61
Concerns 61
Cookies 61
Data Privacy 63
Legislation 64
The EU Directive and "Safe Harbor" 64
COPPA 66
Other Organizations 67
The Platform for Privacy Preferences Project (P3P) 68
APPEL 70

Meta's Exhibit No. 1011
Page 0011

Case Studies: Real Life Privacy 72
SmartGirl: www.SmartGirl.org 72
Personal Information 73
User Tracking and Cookies 74
Opting In, Opting Out 74
Additional Measures and Security 74
The Onion: www.theonion.com 75
Personal Information 76
User Tracking and Cookies 77
Data Analysis 77
Opting In, Opting Out 79
The New York Times: www.nytimes.com 79
Personal Information 80
User Tracking and Cookies 80
Data Analysis 81
Opting In, Opting Out 82
Seal Programs 82
Summary 83

4. BBC News Online 87
Background and Overview 87
The Traffic Information We Gather 90
The Importance of Page Impressions 91
Overview of Technical Architecture 94

Meta's Exhibit No. 1011
Page 0012

Three Requirements of Our Traffic Analysis 94
League Table Generator.....95
Audited Log Analysis.....99
Use of Cookies.....101
Detailed Visitor Analysis.....101
Drilling Down into User Behavior: The One-Hit Wonder Phenomenon.....103
Analyzing Extreme Demand: September 11, 2001.....106
News Service or Archive Service?.....107
Summary.....109

5. eBay.....111
What is eBay?.....111
How eBay Makes Money.....112
Why eBay Needs Web Traffic Analysis.....112
Measuring the Completion Rate of Key Processes.....113
Measuring the Effectiveness of Marketing Initiatives.....113
Measuring the Effectiveness of Searching and Browsing.....114
Gathering Technographic Data.....114
Gathering Anonymous Visitor Statistics.....114
eBay's Web Traffic Analysis System.....115
Requirements.....115
Page View Reports.....115
Unique Visitor Reports.....116
Page Flow Analysis.....117
Technographic Reports.....118
Page Property Reports.....119

Meta's Exhibit No. 1011
Page 0013

Client-Side JavaScript 120
How the System Works 121
JavaScript, Cookies, and Web Bugs 121
Random Sampling 122
Respecting Users' Privacy 123
"Long Distance" Tracking 123
How Page Names and Page Properties Are Set 124
Putting It to Real Use 126
Selling Flow 127
The Old Flow Versus the New Flow 127
The Analysis 129
Bidding Analysis 131
The Four Finding Methods 132
The Goal 134
Our Analysis 135
Summary 136

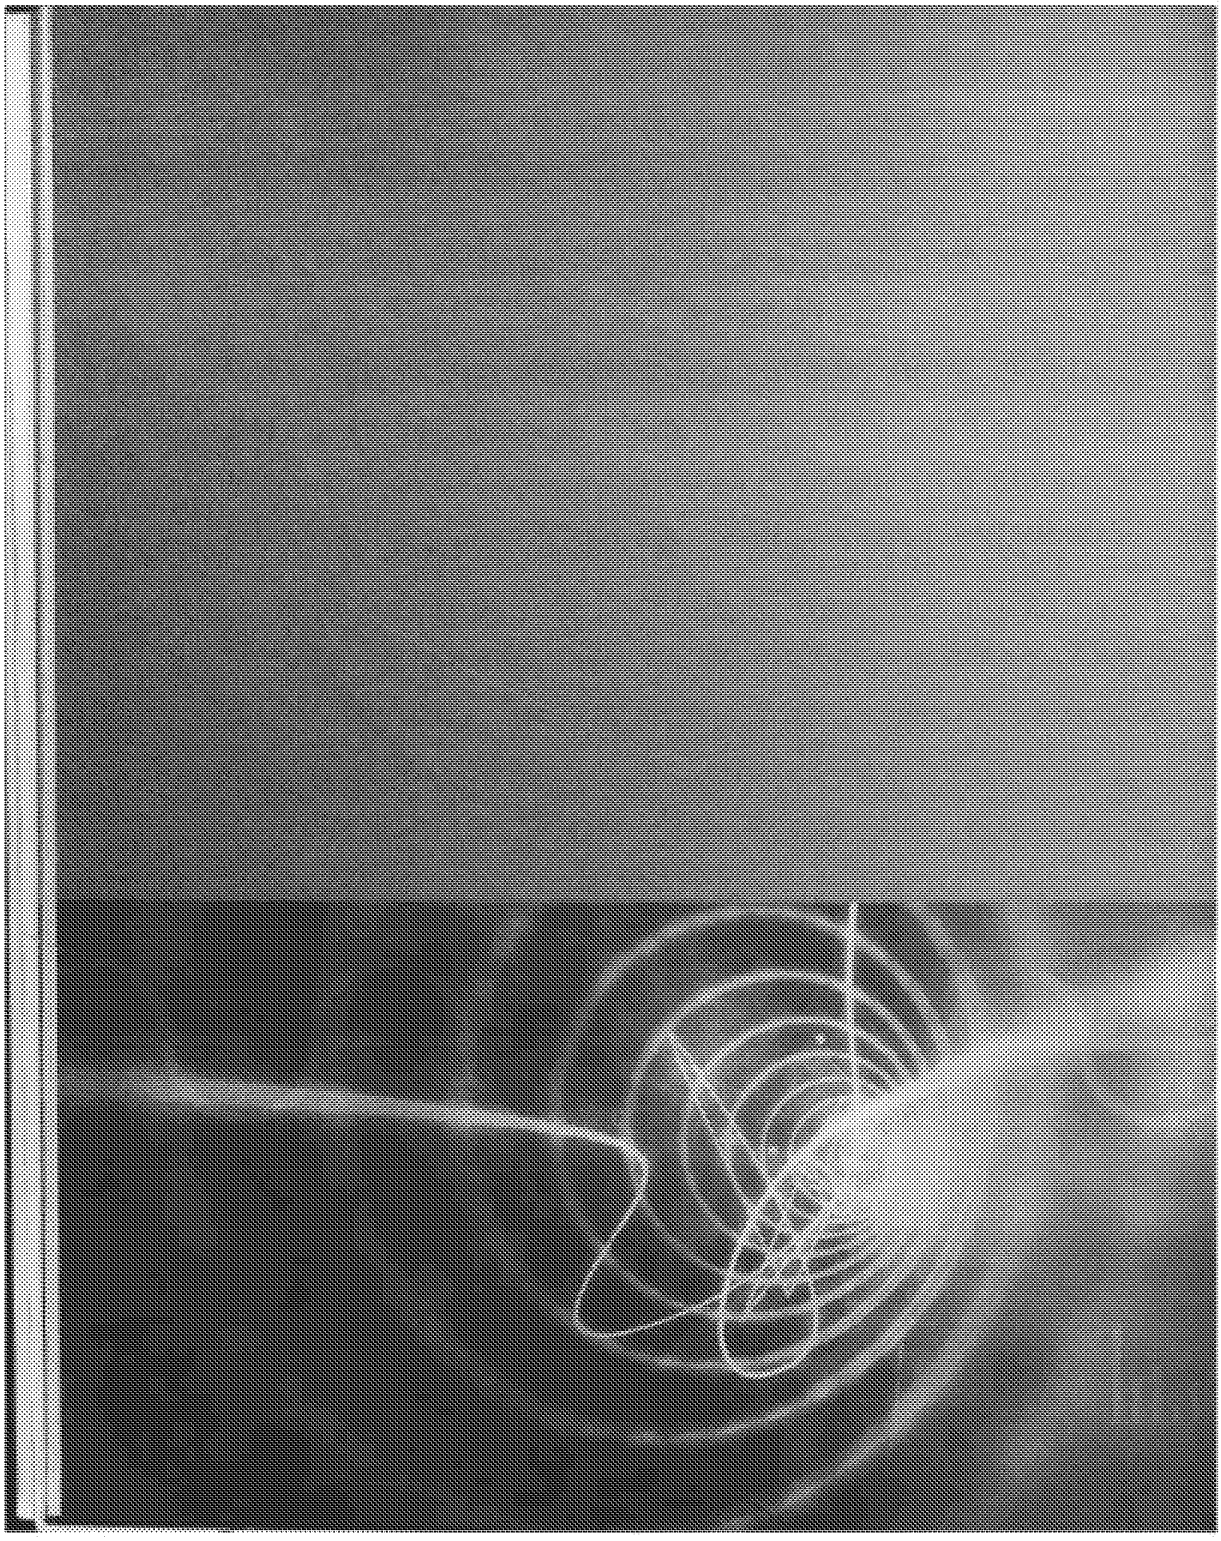
6. ASPToday 139
Content and Visitor Profile 139
History and Background 140
Role of Web Traffic Analysis 142
Problems with Early Log Analysis 143
Cookie Data Analysis 144

Meta's Exhibit No. 1011
Page 0014

Overview of Web Hosting Architecture.....145
Overview of Business Logic Architecture.....145
Database Structure.....148
Example One: Regular Visitors, Frequent Viewers.....150
Example Two: Topic Interest.....154
Example Three: Monitoring Campaigns and Schemes 156
Example Four: PDF Site Analysis.....158
Example Five: Editorial Strategy.....158
Summary.....160

Index.....163

Meta's Exhibit No. 1011
Page 0015



Web Traffic Analysis Introduction

Knowing how much traffic your site gets, and which parts get more than others, can be vitally important to understanding what visitors are getting from your site. With a bit of analysis of your server logs, you could discover that people are using your site in a way you didn't expect. Once you get a better idea of usage patterns for your site, you can begin to work on ironing out any problem areas such as that page that most shoppers never get beyond.

So, if you've ever wondered how to find out which day of the week your site gets most visits, this is the book for you. In fact if you've ever wondered how much traffic your site gets, then this is definitely the book for you. In the coming pages we look at the most important aspects of analyzing web traffic, including a comprehensive discussion of the key terms and their usage.

What's it all About?

To get the full benefit from *Chapter 1*, there are things you should know before we get there. You'll have noticed that we didn't say "If you want to know how many people have visited your site, this is the book for you," and with good reason. As we'll learn in *Chapter 1*, we can't really know this mystical figure for a number of reasons. Chiefly, because the site is accessed via HTTP which only recognizes requests for information and doesn't maintain a connection with a visitor, we can at best only get an indication of how many people have visited the site in any given time period. Instead, it has become widely accepted that a different figure be used, one which we can measure accurately, and that is the number of page requests.

In fact, this is all we can really measure with any hope of accuracy – pretty much everything else is extrapolation and educated guesswork.

Don't worry that you won't be able to tell your advertisers that you don't know how many people visit your site each day, because they shouldn't be expecting you to tell them. They should be interested in how many page impressions your site generates per day, and with a little help from this book you'll be able to tell them.

Despite having said that there is only one thing we can measure with real accuracy, there are a number of other elements to analyzing your web traffic. The first of these is knowing where to find evidence of your traffic, and how to interpret it. As we'll see, initially in *Chapter 1* but also throughout the book, this information comes to you in the form of log files generated by the web server hosting your site. Every time the server receives a request for a page, be it a simple HTML file, a picture, or a dynamic server-side scripting file, it makes a record of the request. This record includes the date and time of the request, where it came from, the file requested, and the manner in which it was requested. Importantly, if there was a cookie involved in the request, the name of the cookie is recorded as well. Counting the unique cookies in each log is as close as we get to knowing how many people visited the site – not everyone accepts cookies in their browser, for a variety of reasons. Some people may visit the site from more than one computer, and some people may even use more than one browser on the same computer.

Once we have the log file, we need to know what to do with it. Rather than have someone with a pencil and some paper read through the whole thing tallying up the page impressions and unique cookies, there's a selection of software available to do the hard work for you. Some of this software is very expensive, and some of it, like the main tool we'll be using in this book, is so cheap it's free! Specifically, we'll be looking at how a tool called Analog can be used to interpret your web server's log files. If you feel up for a challenge, you could write your own log file analysis utility, a common choice. We'll be taking a look at how one of these works, the very one in fact that is used to analyze the log files for our own site, *http://www.glasshaus.com*.

Of course, as anyone who has done any research involving statistics will know, having some numbers doesn't mean you'll know what's happening. Interpretation, whether we like it or not, has some part to play, as well as application. In the selection of case studies that make up the second half of the book, we'll see some of the benefits that can be gleaned from in-depth analysis of your site traffic. If you know where people stop trying to pay for your products, you know where you should really improve your site.

Meta's Exhibit No. 1011
Page 0018

Who is this Book for?

This book is for you if you need to learn more about web traffic analysis, whether you've been told to present some figures to wow the board or you've been told to get the figures so someone else can use them to wow the board. It's for people who want to know what to put in traffic reports, what you can know about your visitors, what you're not allowed to know, and of course how to get the information in the first place.

We're not going to assume that you know vast amounts about statistics, or are capable of hand-coding Amazon.com. What we do assume is an intermediate level of knowledge about the Web: basic scripting experience is useful for the technical parts, as is some familiarity with Access, and you'll need to know (or be prepared to learn) the basics of HTTP and how the Web works.

What do I Need to Begin?

For this book, a web site and access to its log files are a good start. Once you have your log files, you'll want to look at them with either Analog (*http://www.analog.cx/*), or the hand-built system available as part of the code download from *http://www.glasshaus.com*, which is built in Access 2000 (so you'll need a PC for those parts). We also use Excel 2000 for further analysis of the data from the hand-built system.

What's Inside?

As we've already seen, *Chapter 1* covers all the basics of web traffic analysis; from data analysis methods and theory to tools for log file analysis. If you find terms in later chapters that aren't explained there, *Chapter 1* is the place to look.

Reach is a common measure used in audience research, it measures the "market share" that any given site has claimed

The server is unable to log where the user goes after leaving the site

Applying standards and employing auditors is not a replacement for proper intelligent interrogation of the data being presented

Meta's Exhibit No. 1011
Page 0019

Chapter 2 takes a look at a hand-built web traffic analysis system, using Access 2000 and Excel 2000. We introduce the database in which log files are turned into easy-to-understand numbers, and look at how the information that goes in there is gathered, including setting cookies.

> *One of the key facets of this database solution is the ability of the analyst or administrator to specify exactly which data are required.*

> *The general principles that it exemplifies can be applied to almost any database and coding solution*

Chapter 3 tackles the thorny issue of privacy -- it's all very well being able to track a visitor across your site, but some of the ways of doing this invade their privacy, and can expose us to legal problems. As well as examining what laws there are to be aware of, we talk about some real sites and discuss their privacy policy and how it impacts their gathering of web traffic metrics.

Chapters 4, 5, and *6* are case studies, in which the people responsible for web traffic analysis for three large sites explains how they gather their data, and what they do with it.

> *There is no single "correct" attitude towards cookies, the situation changes depending on the circumstances*

> *It is especially important to engender trust in the mind of the user*

In *Chapter 4*, Peter Comber tells us how the BBC monitors the traffic on their vast content site, *News Online* (*http://news.bbc.co.uk*). As well as the differences in size, we see how useful it is to have frequent access to traffic reports, and how they indicate which stories people are most interested in.

> We have learned that most "web metrics" - including page impressions - are extremely context sensitive.

> *There was an assumption that, as the archive of articles was both so rich and so large, there must be a huge number of page impressions drawn from this part of the site*

Meta's Exhibit No. 1011
Page 0020

In *Chapter 5* Alex Poon gives us the low-down on how *eBay* (*http://www.ebay.com*) deal with their own unique problems. Since knowing when people stopped trying to buy something is vitally important to any commerce venture, eBay used detailed web traffic analysis to discover the usage patterns of their visitors.

> *For every page on the eBay site, we wanted the ability to see how users got to that page, and where users went after it.*
>
> *We found it cost-prohibitive to track every single page view from every single eBay user. Instead, we implemented a random sampling algorithm that tracks only one out of every 100 users to the eBay site*

Finally, in *Chapter 6*, Ben Pearce talks about the development of web traffic analysis for the *ASPToday* site (*http://www.asptoday.com*), from its inception in February 1999 to the site's latest version in August 2002. Here we gain an insight into the ways that web traffic analysis can help to drive the redevelopment of a site.

> *The site depends upon accurate intelligence to inform the key business decisions that have influenced its evolution from a simple static web site to a dynamic e-commerce application.*
>
> *Learning from early experiences of log analysis, the technical team and business team worked together to determine what information was required*

Meta's Exhibit No. 1011
Page 0021

Support and Feedback

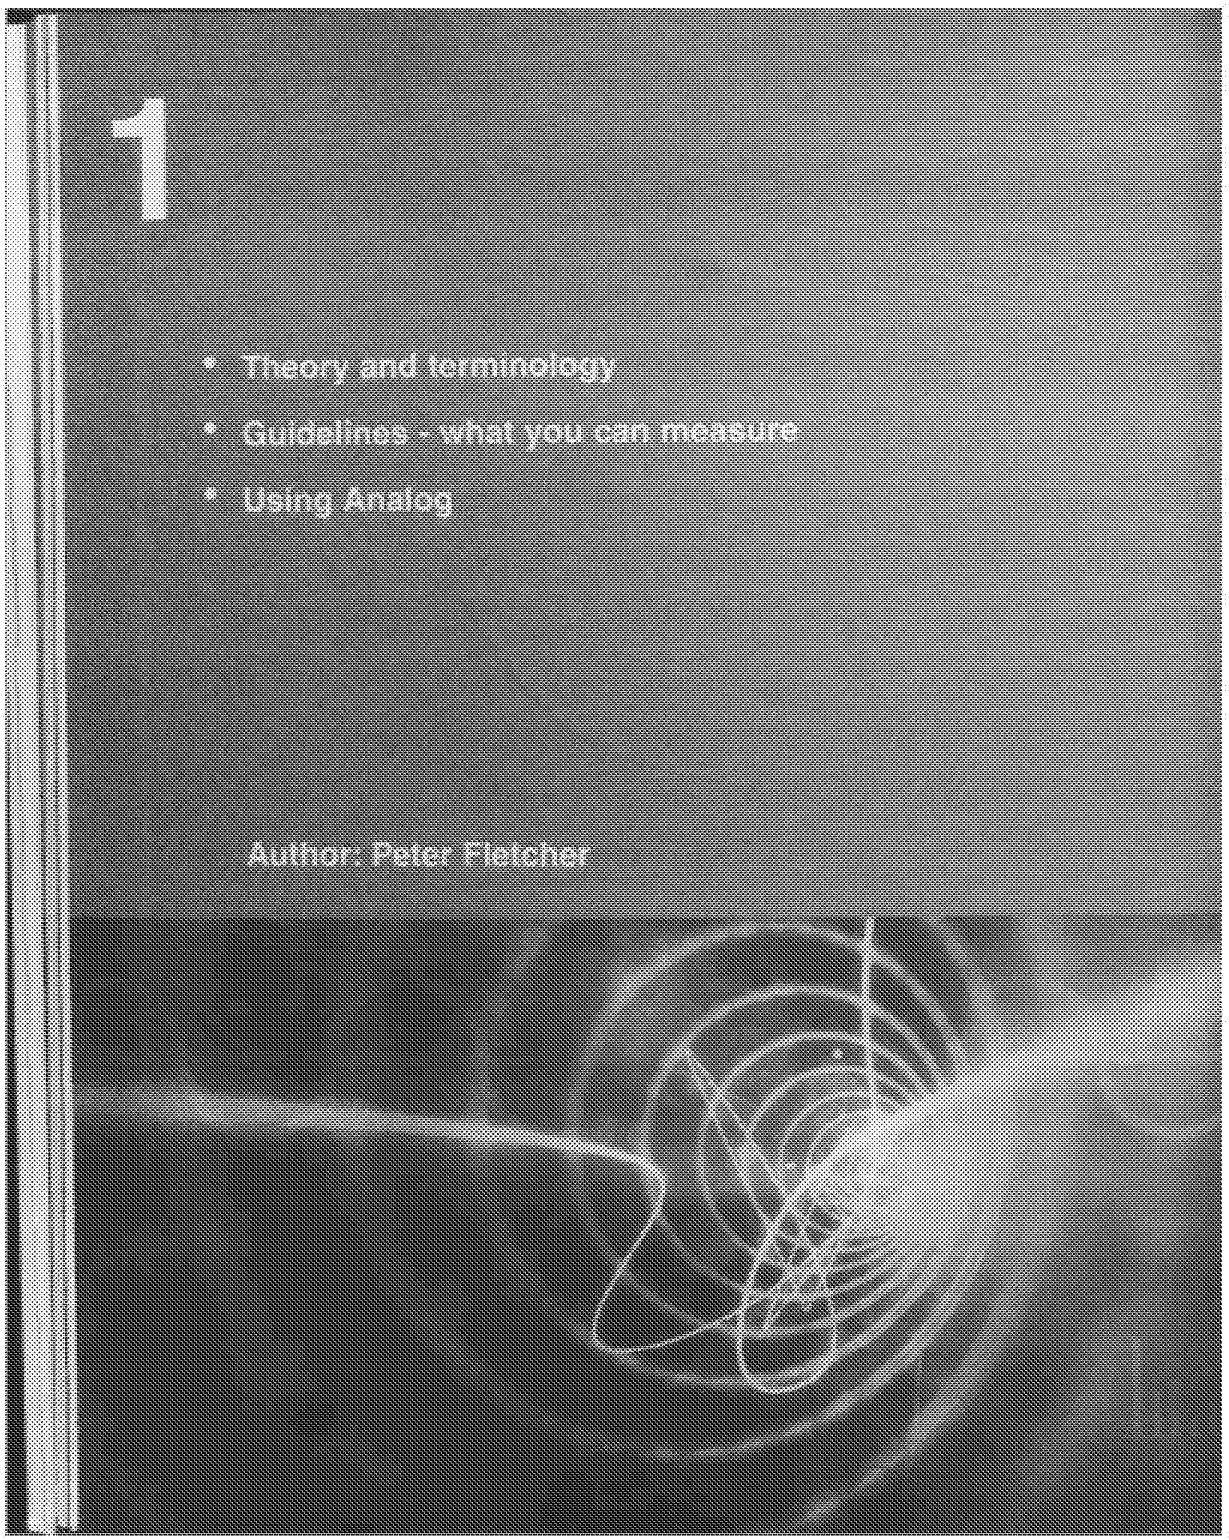
Although we aim for perfection, the sad fact of book publication is that a few errors will slip through. We would like to apologize for any errors that have reached this book despite our best efforts. If you spot an error, please let us know about it using the e-mail address *support@glasshaus.com*. If it's something that will help other readers, then we'll put it up on the errata page at *http://www.glasshaus.com*.

This e-mail address can also be used to access our support network. If you have trouble running any of the code in this book, or have a related question that you feel the book didn't answer, please mail your problem to the above address quoting the title of the book (*Practical Web Traffic Analysis*), the last 4 digits of its ISBN (1183), and the relevant chapter and page number.

Web Support

You'll want to go and visit our web site, at *http://www.glasshaus.com*. It features a freely downloadable compressed version of the code for this book, in both `.zip` and `.sit` formats. You can also find details of all our other published books, author interviews, and more.

Meta's Exhibit No. 1011
Page 0022


1
• Theory and terminology
• Guidelines - what you can measure
• Using Analog
Author: Peter Fletcher

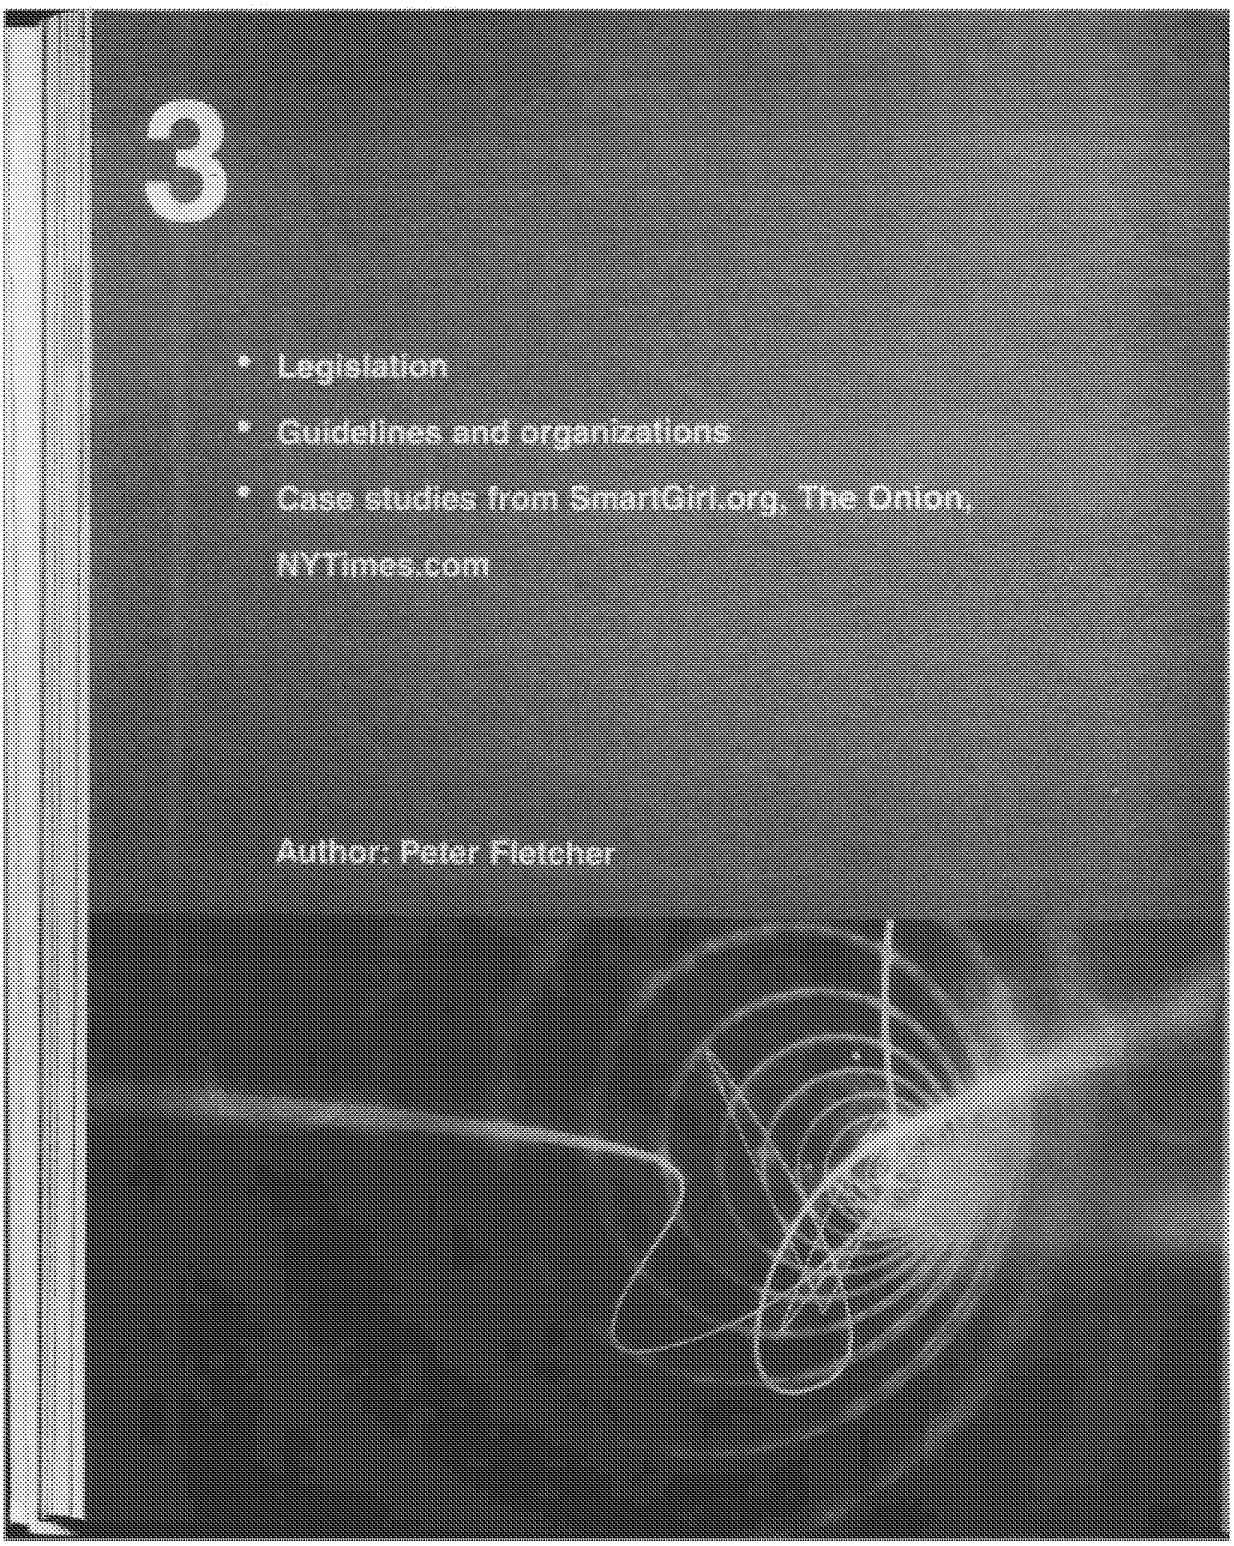
Analytics Techniques

We want to know who is coming to our web site, whether they have been here before, where they came from, what they are looking at, what route they take through the site, how long they spend on each page, and why they leave. Which, if any, of these objectives are possible?

In this chapter, we will define the landscape of Web Traffic Analysis, beginning with the various methods available to us, followed by a discussion of some of the advantages and disadvantages of these methods. Then we go back to basics, and determine exactly what it is that we can know, and what we cannot know, and the extent to which it is safe to make some assumptions. Having established these principles, we examine the emerging standards for web traffic analysis, and discuss whether there are any good "rules" that we should follow.

Comparative Overview of Analysis Methods

There are three broad types of web traffic analysis, defined by the source of the data that is being recorded and analyzed. These are data collected from activity on the web server, data collected from panel surveys, and data collected from the client or browser using "tags" or "web bugs". We can look at these in turn to determine their relative merits.


3
• Legislation
• Guidelines and organizations
• Case studies from SmartGirl.org, The Onion, NYTimes.com
Author: Peter Fletcher

Privacy

So far we have been discussing the technical aspects of traffic analysis and data collection, but privacy is becoming a very important part of building and managing a modern web site. Now we review the main elements and concerns of the privacy issue, the legislative constraints that are being developed, and the various protocols, organizations, and programs that are trying to make privacy less daunting and problematic both for web professionals and users. We then put this in context, by looking at three case studies, which demonstrate how actual web sites are meeting the challenges that privacy presents to both the business of the Web, and user confidence. Finally, we wrap up with a summary of best practices and guidelines, suggesting the issues to cover in a privacy policy.

Concerns

Concerns over privacy can be categorized into several interlinked areas, which we will look at in turn. They are, briefly: the identification of individuals' behavior (often connected to the use of cookies), the collection and abuse of personal data (such as e-mail addresses), the security of the data once collected, and the ability to verify – and delete – incorrect data being stored.

Cookies

As use of the Web has grown, users have gradually become more aware of the fact that site owners can gather and store information about their movements through a site. A fear of the implications of this has been projected onto the main technical innovation that makes it possible: the cookie. Opinion about this humble text file varies from the entirely benign view that cookies do not identify individuals and can be disabled by the user, to the more extreme belief that they are a grave threat to our personal privacy, and should be avoided at all costs. What is the correct attitude towards the privacy implications of cookies?

There is no single "correct" attitude towards cookies, the situation changes depending on the circumstances

There is no single "correct" attitude towards cookies, and, like much else that we discuss in this book, the situation changes depending on the circumstances: the type of site, the user, what is done with the cookie and the information that it generates, whether the cookie is "permanent" or deleted at the end of the browser session, and the ability of the user to opt out of any data collection. The basic fact of the cookie, however, is that it does not by default identify a known individual to the web site; it is merely an anonymous token that can be used to help the site deliver personalized pages to the user, or count the number of unique users (as defined in *Chapter 1*) that are visiting the site. This is complemented by the fact that users can disable cookies if they so desire, thus destroying any ability on the part of the site to track them, for whatever purpose.

Cookies are often used to identify an individual if the site that is using them requires some form of registration, either during a session in which the user has logged in, or from one session to the next, removing the obligation from the user to sign in each time on the same machine. For example, The New York Times on the Web *(http://www.nytimes.com/)* uses cookies to identify you between visits:



Before registering, the site doesn't know who we are

After we've registered and asked the site to remember us, it knows who we are next time we visit -- in fact it welcomes us at the right-hand side of the header.

If we disable cookies before registering, or remove the cookies after registering, *NYTimes.com* no longer recognizes us.

Meta's Exhibit No. 1011
Page 0027

Though we have established that cookies are not innately dangerous, there is a lot of uncertainty in the browsing public, and not all of it due to paranoia or misinformation, about how cookies work and what they can or cannot reveal. Concerns were raised, for example, when ad-serving company DoubleClick™ purchased Abacus Direct, a large catalog database firm. It was believed by some that this would allow DoubleClick to cross-reference its cookie data with the profiling information contained in the Abacus database. For more information on this case, we suggest that you visit the Electronic Privacy Information Center (*http://www.epic.org/*) and DoubleClick (*http://www.doubleclick.net/*).

Third-Party Cookies

These concerns came about because of third-party ad-servers' ability to place cookies on users' machines as they browse different sites. As we know from *Chapter 1* a cookie may only be retrieved by the server it was set by, which prevents any given web site from tracking the path of an individual user across the Web. Users visiting sites that use third-party ad-serving technology, however, might also receive a cookie from that server. This cookie is commonly referred to as a "Third-Party Cookie", as it is set not by the site the user is explicitly visiting, but by a completely different server. This same third-party server therefore technically has the ability to track its cookies as each individual user visits sites that use that ad-server, potentially building up a "profile" (albeit an anonymous one) of that user's surfing habits. This profile can be used to target ads at a user based on their surfing habits. Opinion is divided on this, as some feel this is a genuine benefit to the user, increasing the relevance of the information being served, while others consider it an intrusion into their private business. The key here, perhaps, is choice. DoubleClick, a provider of third-party ad-serving solutions, has introduced an "opt-out" cookie system, details of which are contained in their privacy policy: *http://www.doubleclick.com/us/corporate/privacy/privacy/default.asp*. Specifically, the policy states "**DoubleClick does not use your name, address, e-mail address, or phone number to deliver Internet ads. DoubleClick does use information about your browser and Web surfing to determine which ads to show your browser**".

Cookies have been designed to work in such a way that they do not of themselves pose a threat to privacy, but it is possible to deploy them in ways with which many users might feel uncomfortable. For this reason, it is important to be explicit in your privacy policy about exactly what cookies you set, and what you intend to do with the data collected as a result. We will cover this in more detail later in this chapter.

Data Privacy

The other broad area of concern for privacy is the collection, and integrity, of personal or sensitive information: this might include e-mail addresses, residential addresses, credit card details, purchasing history, medical information, and so on. The potential danger here is not that sites will be collecting information without the consent of the user, as it is fairly clear what is going on when a web site asks us to enter our address, but rather what happens to the data once we have clicked "submit".

The factors that we need to consider are:

- What will the web site do with my personal information?
- How can I be sure that the web site will not use this information for purposes other than that to which I have consented?
- Will the web site intentionally distribute my personal information?
- Will the web site keep my information secure, and take steps to ensure that it is not unintentionally distributed?

As before, these are all key points that should be covered in any privacy policy, but there are also third-party agencies that can help us to gain the trust of our users, and these are also the main concerns of the various legislative bodies that have intervened in the interests of citizens' privacy. We will discuss some of the organizations that exist shortly, but first we're going to examine the more formal legislative controls that have been introduced.

Legislation

This section is intended to give a helpful overview of the legal constraints that are being introduced, concentrating on the situation in the US and the EU, with some suggestions for further research. Privacy legislation, like most legislation, is legally complex and technical, and circumstances will change between businesses and jurisdictions. Clearly, when considering your site's privacy policy and its terms and conditions, you should consider taking formal legal advice.

In this overview, we will consider the implications and intentions of two key legal interventions: the EU Data Protection Directive, and the US Children's Online Privacy Protection Act (COPPA). For details of the EU legislation, visit: *http://europa.eu.int/comm/internal_market/en/dataprot/index.htm*, and for more details on COPPA visit: *http://www.ftc.gov/bcp/conline/pubs/buspubs/coppa.htm*.

The EU Directive and "Safe Harbor"

The EU Directive binds all 15 member-states of the EU to safeguard the privacy of their citizens according to a basic set of principles, the intention being to harmonize and standardize the level of data protection and privacy across all of the member states. In principle, the legislation is designed to assert the following conditions:

- It covers those who collect, hold, or transmit personal data
- It puts them under an obligation to collect data only for "specified, explicit, and legitimate purposes"

Meta's Exhibit No. 1011
Page 0029

- It obliges them to only hold data that is "relevant, accurate, and up-to-date"
- It mandates that the collection and retention of the data should be transparent, meaning it should be clear that the data collection and storage is being conducted, and that the individual has the clear right to opt out of this process

In particular, the Directive asserts the rights:

- To access the data being held
- To know where that data originated
- To rectify inaccurate data
- To legal recourse in the event of illegal use of the data
- To withhold permission to use the data for specific purposes (for example, sending spam or junk mail)

Furthermore, in the case of sensitive information, defined as "ethnic or racial origin, political or religious beliefs, trade union membership or data concerning health or sexual life", this can only be processed with the explicit consent of the individual (they have to OPT IN to this process).

There are various exceptions and provisos which refer to the rights and conditions, allowing a certain degree of flexibility in the drafting of this legislation by member states (for example, for legitimate "journalistic, artistic, or literary purposes"), and safeguards are required to be put in place to monitor this process. This flexibility is overseen by independent "data supervisory authorities" in each member state.

The EU Directive is a formal guideline of principles for member states to implement. Each state is free to draft legislation covering the Directive in its own manner, so long as it conforms to the Directive.

The impact of this legislation is that organizations (or, indeed, individuals) are not entitled to collect or process data about citizens of the EU without constraint. Data collectors must make clear what information is being collected, for what purpose, and provide the means for individuals to check the accuracy of the data held (and correct it if necessary), and opt out of certain uses of the data.

Whilst legislation within member-states specifically covers the use of data within the EU, and also the transfer of data between member-states, there is an assumption that EU citizens ought to be able to enjoy the same privacy and data protection rights regardless of the geographical destination of the data they provide. This therefore has an impact not only on web sites that operate from within the EU, but any web site that technically "transfers data" from inside of the EU zone. In effect, any web site hosted outside of this area is transferring data when it receives a request from an EU resident.

In order to avoid the potential for disputes between US companies and the EU, the US Department of Commerce and the European Commission created the Safe Harbor privacy framework, whereby US companies can verify that they will treat data collected from within the EU zone according to the standards outlined in the EU Directive. US companies can certify that they conform to the principles of the Directive, though the process is entirely voluntary, and can be undertaken either by joining a self-regulatory program (we look at one in the next section), or by developing their own program that conforms to the Safe Harbor framework.

The advantages to an organization conforming to the Safe Harbor framework are chiefly that they will accordingly be able to continue to do business with EU citizens and companies without interruption, and that any claims brought by EU citizens against US companies will, in most cases, be heard in the US. More information on the Safe Harbor framework can be found at the US Department of Commerce web site: *http://www.export.gov/safeharbor*.

COPPA

The other significant piece of US legislation that we need to cover is the Children's Online Privacy Protection Act, known as COPPA. This Act sets out the legal framework for the collection of "personal information" from children under the age of 13. It applies to all web sites or "online services" that are aimed at children under 13 and that gather personal information from their users, and also to general web sites that have "actual knowledge" that they are collecting personal information from children.

The clear aim of this legislation is to safeguard the privacy of young children who may not be aware of the potential dangers of disclosing personal information over the Internet.

The clear aim of this legislation is to safeguard the privacy of young children who may not be aware of the potential dangers of disclosing personal information over the Internet. The Act covers information that individually identifies the child, such as name, address, e-mail, telephone, or anything that would allow someone to contact the individual. It also includes information such as hobbies and, importantly, information gathered by the use of cookies or other tracking devices "when they are tied to individually identifiable information" – that is, the cross-referencing of cookie data, such as pages visited by an individual, with that individual's contact information.

The Act places various constraints on the web sites to which it applies, the principal one being that the site must publish a clear and unambiguous privacy policy, and this should have a "prominent" link on both the home page of the site and any other page where personal information is collected. The Act specifies that the policy should state:

- Who is collecting the data – including the name, address, and e-mail address of at least one contact, and the names of any other parties ("operators") involved with the collection of data on the site

Meta's Exhibit No. 1011
Page 0031

Page 180 of 1032

- What data is collected and how – included the use of cookies if appropriate
- Specifically what this information will be used for (direct marketing, sending an opt-in newsletter, etc.)
- Whether the information will be distributed to third parties, and if so for what purpose and with what safeguards. Also, parents must have the right to opt out of this third-party distribution of their child's information
- That no unnecessary information may be collected for the purposes of participation in the activities of the web site
- The parent's right to view, and if necessary remove, their child's personal information, and remove consent for further collection, as well as clear instructions as to how to carry this out

There are also a number of suitable definitions, clarifications, and exceptions to these rules, for example detailing what can be collected without the explicit consent of a parent. Nevertheless, the basic thrust of the Act is clear: when dealing with young children, the responsibilities of the web site for privacy and data protection are heightened, and explicit care must be taken to outline what is collected, why, and how the child or parent can intervene in the process.

Other Organizations

As the whole landscape of privacy, data protection, and legislation is complicated enough for web professionals, it is especially important to engender trust in the mind of the user. To this end, there are various organizations whose aim is to increase this sense of trust, and work with partner web sites to bring it about. Two such groups are BBBOnline (*http://www.bbbonline.org/*), run by the Council of Better Business Bureaus, and TRUSTe (*http://www.truste.org/*), an independent non-profit organization founded by the Electronic Frontier Foundation (EFF) and the CommerceNet Consortium. Both of these organizations have similar objectives and programs.

users have the right to know what is being collected about them

The general reason for the existence of organizations like BBBOnline and TRUSTe is to generate a greater sense of trust in Internet users, and the explicit policies that they advocate are broadly in line with the principles of the legislation that we have examined above. In short, users have the right to know what is being collected about them, what will be done with this information, including whether it will be further distributed, and the right to opt out of certain uses of the data, for example for direct marketing campaigns. The key to this is transparency: a privacy policy that states exactly what is being done, and what the user can do about it.

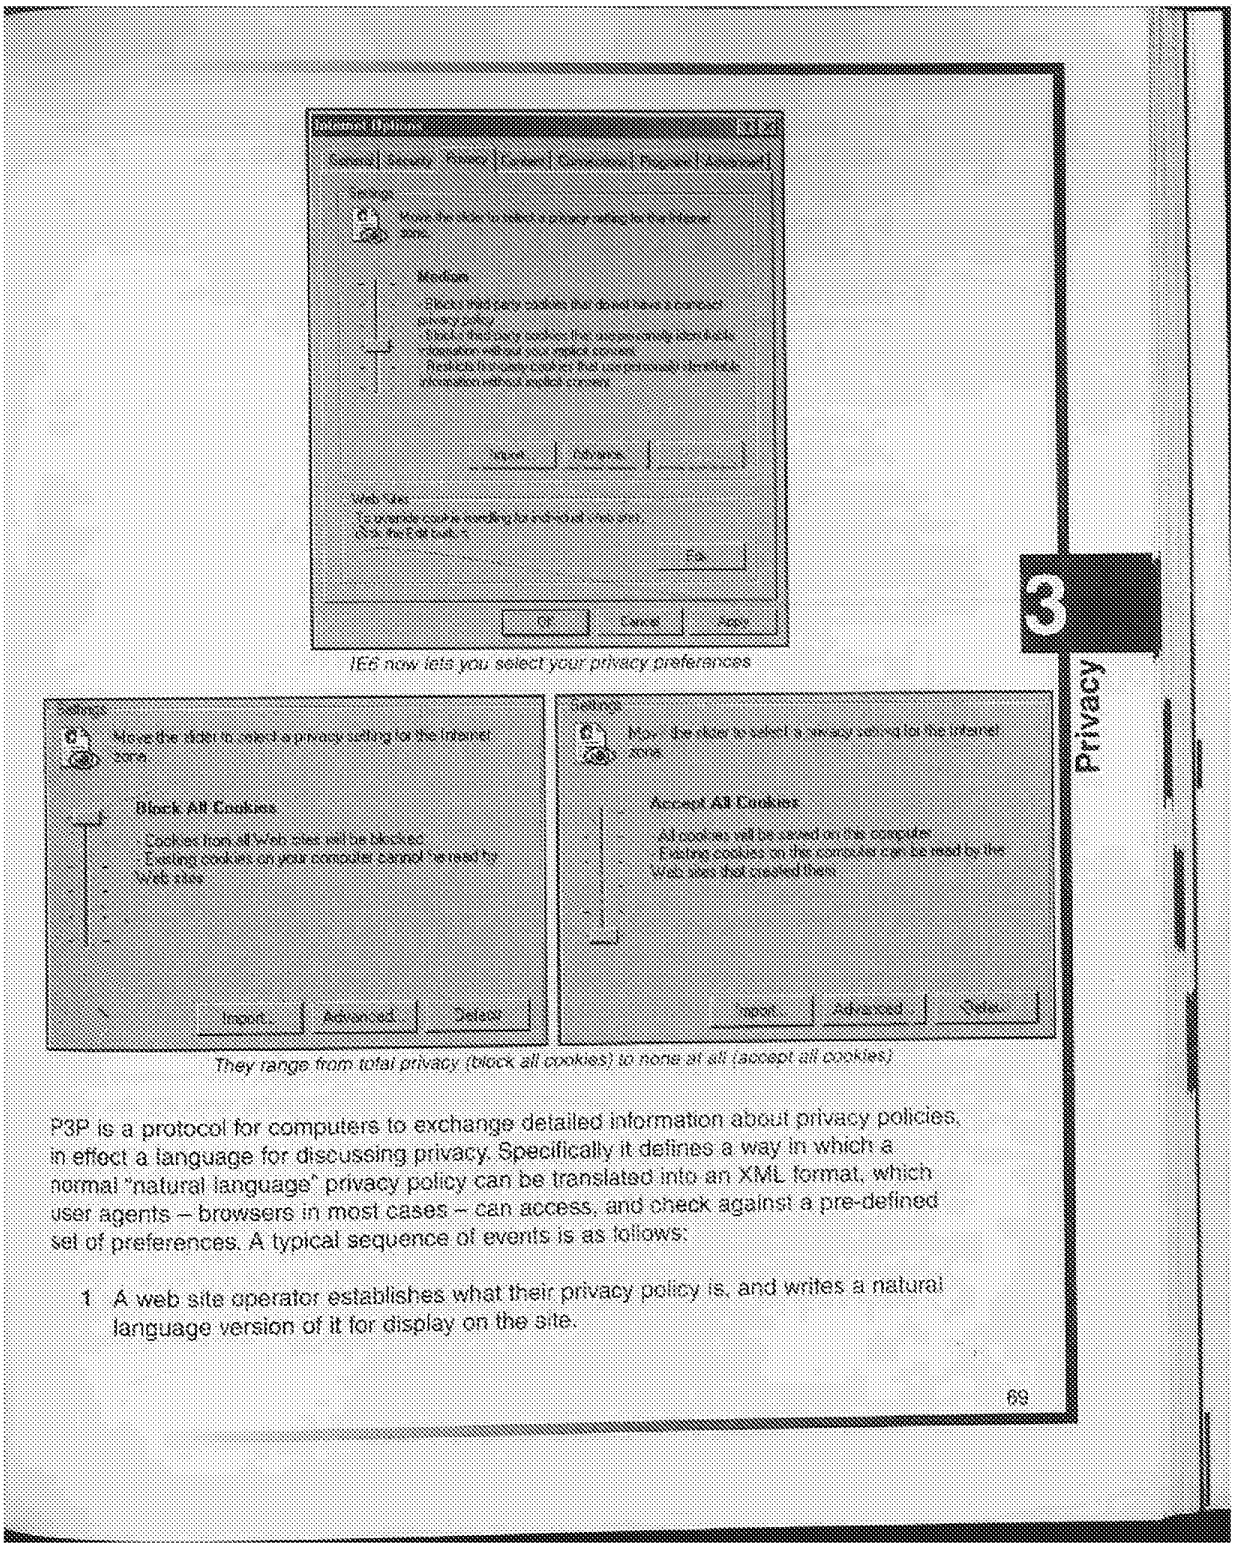
With this objective of establishing clear and authoritative privacy policies, both of these organizations offer "Seal Programs", effectively underwriting the privacy policies of affiliated sites, and offering their "Seal" to indicate this. Whilst the details of how to join these programs vary, the basic principle is that the web site submits its privacy policy and data collection practices for review by the program. If the site meets the requirements of the program, both in terms of the policy itself and its general data practices, then it can be awarded the Seal (there is a fee for this service).

These programs also undertake to arbitrate in the case of any visitor entering into a dispute with the web site in question, and the site can be suspended or expelled from the program if it fails to live up to its promises, a facility designed to increase the level of trust and confidence in the web site's users. In addition to the main Privacy Seal Program, BBBOnline and TRUSTe also offer other programs, validating such things as children's privacy, and Safe Harbor compliance.

The Platform for Privacy Preferences Project (P3P)

A further attempt to clarify and simplify privacy for both web site operators and users has been made by the W3C in the form of the Platform for Privacy Preferences Project, commonly known as P3P. In essence, this is a project to standardize the information contained within privacy policies, and to encapsulate that information in a **machine-readable form**. The idea behind this is that if privacy policies can be made available in such a way as to be understandable to the user's computer, then this removes from the user the task of reading and interpreting the written version, effectively automating the task, once the initial preferences have been set. The idea is that this would make it far more likely that web sites' policies will be reviewed each time data is transferred.

This is an emerging standard for the construction of privacy policies, and with the credible backing of the W3C, so it is worth examining in a little more detail. Let's now take a brief look at how P3P works, what the web site operators and users need to do to use it, and discuss what effect this will, and will not, have on online privacy. There is some built-in support for P3P in Internet Explorer 6, Netscape 7, and as an additional feature for Mozilla (see *http://www.mozilla.org/projects/p3p/* and *http://lxr.mozilla.org/mozilla/source/extensions/p3p/*). Internet Explorer 6 now has a *Privacy* tab on the *Internet Options* dialog, allowing a simple selection of privacy preferences.



IE6 now lets you select your privacy preferences

They range from total privacy (block all cookies) to none at all (accept all cookies)

3 Privacy

P3P is a protocol for computers to exchange detailed information about privacy policies, in effect a language for discussing privacy. Specifically it defines a way in which a normal "natural language" privacy policy can be translated into an XML format, which user agents – browsers in most cases – can access, and check against a pre-defined set of preferences. A typical sequence of events is as follows:

1 A web site operator establishes what their privacy policy is, and writes a natural language version of it for display on the site.

69

2. The operator then translates this into an XML file using the P3P standard. There are various tools being developed to make this process easier.
3. Another XML file which contains details of where to find the policy is then placed in a "well-known location", which the W3C recommends is `/w3c/p3p.xml`, where the `w3c` directory is under the root of the site.
4. A user agent with P3P preferences set visits the web site, locates the privacy description and automatically downloads it, parses the contents, and compares it with the user preferences.
5. Based on this comparison, the browser can automatically elect to exchange data (a cookie, for example), with the web site, or warn the user that the site does not meet the standards specified in the preferences file in some way.

IBM is developing a P3P Editing tool, currently in beta, to ease the process of translating a written privacy policy into a machine-readable format (XML file). At the time of writing, this tool is free to download and use, and is available at: *http://www.alphaworks.ibm.com/tech/p3peditor*. A list of other implementations is available from the W3C web site: *http://www.w3.org/P3P/implementations*.

The location details do not necessarily have to be placed on the web site as detailed above. The specification also allows a reference to the location of the privacy details to be placed in the HTTP headers from the site, or in the HTML using a `<link>` tag. The process outlined above, however, represents the W3C's expectation of how most web site operators will deploy the reference in `/w3c/p3p.xml`, as this is probably the most straightforward and transparent method.

APPEL

The W3C are also working on a language for describing user privacy preferences, called APPEL (A P3P Preferences Exchange Language). The idea of **APPEL** is that it will be used to create sets of preference rules (called rulesets), which can then be imported into the browser (say) for making automated decisions regarding the exchange of data with web sites or other remote machines. The W3C do not expect individuals to learn APPEL to set their preferences, but rather use tools or ready-made rulesets to make the process easier. Existing rulesets, most likely created by a third party, can be imported into a browser, and indeed exchanged between different browsers operated by the user.

So now it seems we have at least an emerging standard for describing our privacy policy and automating the process of checking it. What is this designed to achieve? The wider adoption and implementation of P3P should:

Meta's Exhibit No. 1011
Page 0035

- Make privacy policies more accountable. If browsers are automatically checking the details of a web site's policy, and acting accordingly on the preset user preferences, this will not only force sites to declare their policies unambiguously, but also ensure that the policies themselves are in line with their users' expectations. The principle is that privacy will become a more transparent process, meaning that users will be able to see clearly where they stand, and make informed decisions as a result.
- Enable users to have more control over their privacy and use of their personal information. It is unlikely that users currently check privacy policies in any detail (if at all) when visiting a site or exchanging data, largely due to the time needed to locate the policy and check it for potential problems. Automating the system, so long as the preference-setting part of the process is clear, should increase awareness and control.

It is important to make clear, though, what P3P cannot do:

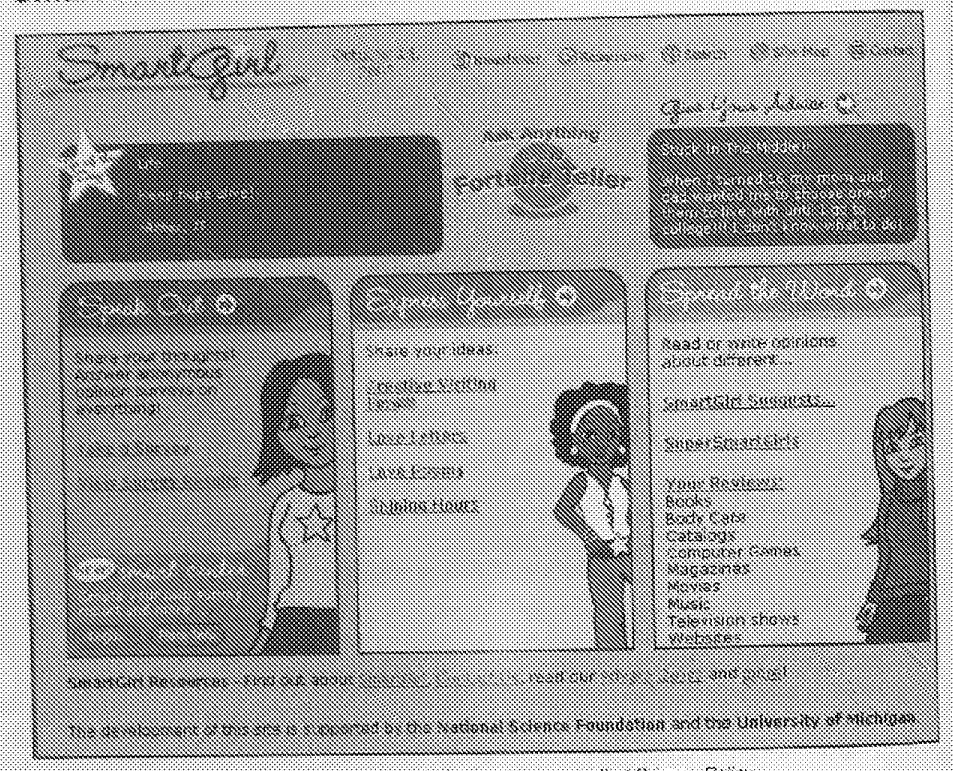
- It has no power to ensure that web sites actually abide by the privacy policy that they publish. There is nothing in the W3C specification that prevents a site from using P3P in an attempt to deceive their visitors. It could be argued that, as the system is automated, this may increase the potential for abuse, because users will begin to "trust" the browser to tell them when something is wrong. In fact, they should be trusting the web site, not P3P. In this respect, nothing has changed, as a site has always had the power to publish a misleading, or indeed false, privacy policy. This is where organizations such as TRUSTe or BBBOnline can be useful, and in fact P3P does contain a DISPUTE element in which dispute-resolution information can be described.
- P3P is not designed to supersede or replace statutory regulations, but rather complement them. The intention is that P3P will assist in the assertion of legal frameworks, and encourage a better understanding and conformity with them, but it cannot enforce them. Proper regulation and P3P go hand in hand.
- There is no provision for encryption or the security of data within the P3P specification. P3P does not replace the need for encryption and secure data exchange, but again complements it, by allowing the user to get a clear picture of what is being done, and to respond accordingly.

In summary, P3P is an important development in the privacy landscape. It has the authority of the W3C behind it, and has the support of many leading industry groups and companies. The adoption of P3P at the moment is somewhat nascent, but as newer browsers include elements of the specification in their upgrades, more users will become aware and able to take advantage of its preference checking process. In this respect, it will place the onus on web site operators – owners, managers, and developers – to create a clear privacy policy and implement it in P3P.

Case Studies: Real Life Privacy

We have so far covered something of the theory and the rules concerning online privacy, but the picture would not be complete without taking a look at how real sites are reacting to the situation. In this section we will discuss three different sites: *The Onion*, The New York Times, and *SmartGirl*. We will look at each site in turn, examining what it needs to do to maintain its business, what analysis is done, and how it explains this in its privacy policy. The information presented here is taken from interviews conducted with the help of the sites in question.

SmartGirl: www.SmartGirl.org



The front page of SmartGirl.org, with easy to find Privacy Policy

SmartGirl.org is a web site written for and by girls, many under 13 years of age, and therefore has to comply with COPPA. It is worth looking at how their exemplary privacy statement details what the site does and does not record, and the measures in place both to ensure compliance with the law, and create a safe and interesting place for girls to participate.

Meta's Exhibit No. 1011
Page 0037

The policy begins with a statement of the site's purpose:

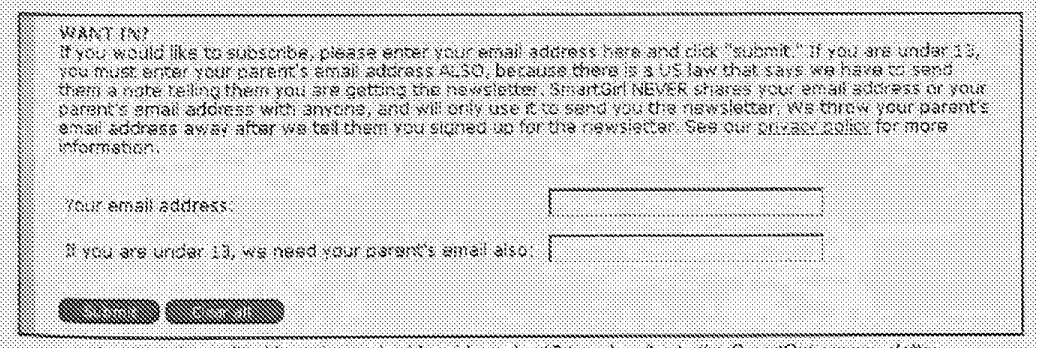
> This site is owned by the University of Michigan and is operated by the Institute for Research on Women and Gender. Our mission is to create a place on the Internet where girls can feel free and safe to express their opinions and views on everything from the latest movies and the hottest fashion trends to everyday issues like school, relationships and navigating adolescence. Girls visiting the site can fill out reviews on different products (such as books, movies and CDs), take anonymous surveys, play games, or submit love poems, love letters, and other content to be posted on the site. Girls who become members can participate in the SmartCenter.

It ends by disclosing the name and address of the site's operators, with both postal and e-mail addresses for contact in the event of any queries. In between, there is a very explicit description of SmartGirl's privacy procedures.

Personal Information

This is defined as any information that can personally identify an individual, such as an e-mail address. The policy states that no visitor to the site is required to disclose such information for the purposes of participating in the site, and any e-mail address provided is done so entirely voluntarily. Specifically, *SmartGirl.org* will never post, publish, rent, sell, or intentionally distribute personally identifiable information. All work submitted, or survey completed, is anonymous – users can include an e-mail address for the purposes of receiving notification as to whether their work will be published, but once used this e-mail address is no longer linked to the individual user.

If the user is under 13 years of age, then their parent's consent is also required for their submission of an e-mail address, for example to join the newsletter, or to receive notification if they have submitted work. This is requested at the same time, and a message is sent to the parent, which is then used to notify the parent that the child's e-mail address has been received, and give them the opportunity to remove it.



A parental e-mail address is required for girls under 13 to subscribe to the SmartGirl.org newsletter.

Parents can review their child's postings to the site, and change or delete their child's e-mail address or membership at any time. The privacy statement outlines the simple procedure for doing this.

User Tracking and Cookies

SmartGirl.org uses cookies for the following purposes:

- To track – anonymously – the traffic patterns on the site.
- To eliminate duplicate survey submissions.
- To allow members access to members-only content.

Crucially, cookies used for one purpose are not linked to cookies used for another purpose. For example, separate cookies are used for verifying survey submissions and traffic pattern analysis, and the two cookies are not identifiable as belonging to the same user. It is not possible to identify an individual user from the cookies stored on their computer. Other information, supplied by the web hosting company, is equally anonymous, and, similarly, released to third parties (such as advertisers) in aggregate form only. This information includes:

- Browser type and operating system (User-Agent)
- Sites referring users to *SmartGirl.org*
- IP address and Internet domains from which users are accessing the site
- Dates and times of distribution of traffic to the site
- Average length of stay on the site
- Pages visited

Opting In, Opting Out

As is implied above, the e-mail newsletter is strictly an opt-in list. Users voluntarily submit their e-mail address to be included on the newsletter and, in the event of the user being under 13 years of age, *SmartGirl.org* informs their parent of this, in case they wish to remove their child from the list.

Additional Measures and Security

In addition to the procedures detailed above, the site has strict policies regarding the security of the data collected.

Meta's Exhibit No. 1011
Page 0039

Our Users' privacy and any personal information we may have about them are kept as secure as we are able, although like every other business entity, we cannot guarantee that this information will not be stolen. Passwords to access data at *SmartGirl* are only given to those trained to handle them safely. All *SmartGirl* staff members have signed confidentiality agreements that legally obligate them not to release e-mail addresses or any other information about visitors.

No personal information submitted by users is used for any purpose other than the stated reason for collection: e-mail addresses collected for the newsletter are used only to send the newsletter, and not for targeted mailings, for example. Furthermore, no personally identifiable information is ever distributed outside of the Institute for Research on Women and Gender Studies.

e-mail addresses collected for the newsletter are used only to send the newsletter, and not for targeted mailings

The Onion: www.theonion.com



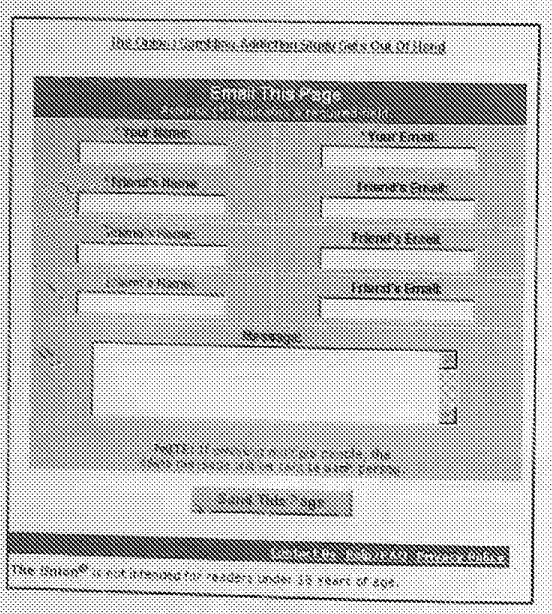
The well-known satirical news site.

3 Privacy

The Onion is in the publishing business, which means that it makes its money by selling the print edition of the paper, and the advertising space therein. This is supplemented by its online business, which includes advertising sales, personals, and *The Onion store*, selling branded merchandising and other suitable wares. The online activity is important not only as a revenue stream in its own right, but also as a highly successful branding tool for the print publication. *The Onion* is aimed at an adult audience, with a notice declaring that it is not intended for readers under 18 years of age. Even though its content is irreverent and humorous, it takes online privacy very seriously.

Personal Information

The Onion collects and records personal information from customers, but only in order to fulfill orders that are made. This information includes the customer name, address, phone, e-mail (for order confirmation), and billing address if it is different from the shipping address. Credit card numbers are not stored by *The Onion*, but transferred to the card verification agent, *authorize.com*, who processes the order and confirms the result. If the transaction succeeded then the fulfillment details are forwarded to the fulfillment company for shipping. Users are able to register in order to speed up future interaction with the store, though this is entirely optional. Registration does, however, give the user the opportunity to review and amend any personal data that is stored about them on the site.



Meta's Exhibit No. 1011
Page 0041

The only other collection of personal information on the site is the logging of e-mail addresses for the "E-mail this Story" feature. This is done as a security feature to protect *The Onion* from spam allegations. This information is not used in any other way.

Currently *The Onion* does not provide advertisers with demographic information gleaned from their campaigns on the site. Prior to a campaign, however, they are supplied with demographic data gathered from third-party sources (@plan, Nielsen) to help show potential clients the value of advertising with *The Onion*.

User Tracking and Cookies

Cookies are used on *The Onion* to maintain state during individual sessions. The information that this generates is then used to monitor visits to the site, though as the cookies are temporary – they expire at the end of the session – return visits by unique individuals are not tracked. An ad-tracking system – AdJuggler – is used to monitor general user responses to the ad campaigns on the site, in the form of ad impressions and click-throughs, but more detailed user interaction data is not recorded. The main purpose of the cookies used in advertising is to determine the frequency of the ads, and they use the cookie to control this for each unique visitor.

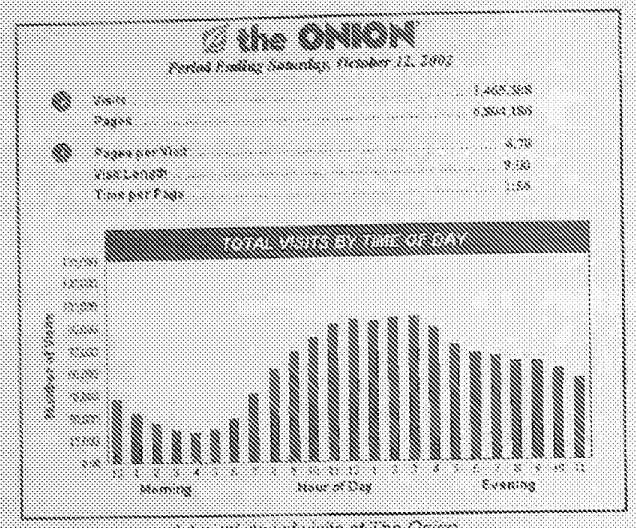
The main purpose of the cookies used in advertising is to determine the frequency of the ads...

Cookies are also used for maintaining state as a customer moves through the shopping cart process. This cookie uses a unique anonymous user ID to keep each user distinct from others; the only data stored here is the expiration date and the name of the cookie, and no personal information is stored in these cookies. Each cookie is set using the ad system, or a piece of JavaScript embedded on each page.

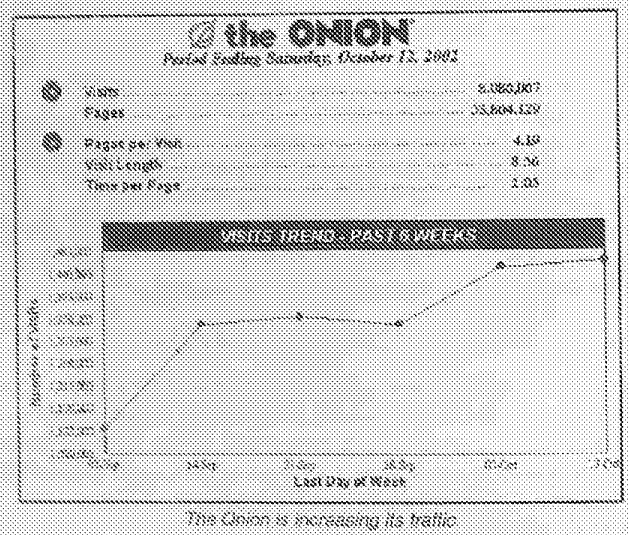
...no personal information is stored in these cookies

Data Analysis

The *Onion* outsources its log analysis to **I/Pro**, which gathers the information from the logs and sends it back in reports. I/Pro stores its extracted data in a database, but the final format can be `.xls` or `.csv` files, or a proprietary format for I/Pro Netline users. In addition to this, Analog is also used within the development team to generate more precise reports, for example, gathering browser version and OS data (from the User-Agent field). The I/Pro analysis is primarily to track page impressions, visits, and visiting organizations, resolved from the client IP address.



A typical day of visits at The Onion



The Onion is increasing its traffic

This log analysis is not formally audited by an outside organization (such as ABC Interactive), but I/Pro implements the standards as part of their service, for example by filtering out robot traffic and "false" impressions. They say: "We have a rigorous IP exclusions list that removes most known spiders and robots. This list is maintained and updated on a regular basis. With the new Netline release, we will also be excluding spiders and robots based on the User Agent in [*The Onion's*] logs."

Opting In, Opting Out

The Onion does not currently operate any lists that would require an opt-out, however if they were to operate such lists they would favor a "double opt-in" approach. On occasion, advertisers have used contest pages to entice consumers to sign up for e-mail lists. These are opt-in lists for the third party, and this fact is specifically spelled out on these pages with a link to the advertisers' privacy policy. Sometimes this information is collected by *The Onion*, and then delivered to the advertiser at the end of the campaign, but no e-mail from *The Onion* is ever sent to people who choose to sign up for these contests.

Double Opt-in

A single opt-in typically occurs when a user joins a mailing list by giving their e-mail address to the web site or list-manager. The drawback with this process is that the address could be submitted inadvertently by its owner or by a third party without the owner's consent. A double opt-in system does not add the e-mail address to the list on the first occasion, but rather triggers a message to the e-mail address itself, informing the owner that a request has been made to join the list, and soliciting a response. If no response is forthcoming, the e-mail address remains off the list. It is only added once the owner responds to the list's message, confirming that they wish to join. Members of the list have therefore had to "opt-in" twice before they are added.

3 Privacy

The New York Times: www.nytimes.com

The well-established online news site.



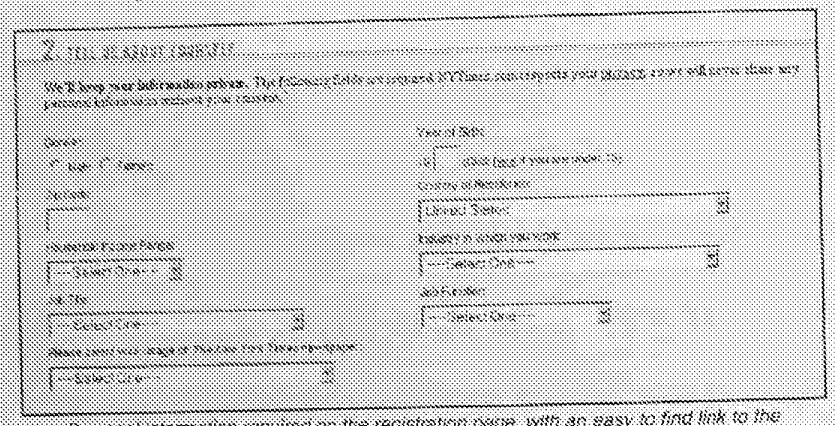
79

New York Times Digital is the name of the Internet division that includes *NYTimes.com*, *Boston.com*, and an archive distribution business.

New York Times Digital is in the "news and information business", its mission being to "enhance society by creating the premier quality network for everyone seeking the best news, information and interaction through digital media". The audience for *NYTimes.com* is global, with 85% outside of New York City, and 20% from outside the US.

Personal Information

NYTimes.com collects personal information about its registered users when they first register on the site, including details such as household income, Zip code (for US residents), and employment information. Notice that a statement referring the reader to the privacy policy is clearly visible on the following form:



Personal information required on the registration page, with an easy to find link to the NYTimes.com privacy statement

This information is only shared with third parties (such as advertisers) in an aggregated form, giving broad demographic trends for the site, except when required by law. The personal information is also available to its owner via the membership page on the site, where users can view and modify the data, and this is further password protected (users have to enter their password even if they have the "*remember me*" cookie set).

The demographic data is categorized by @Plan, and the categories used can be viewed on the *NYTimes.com*'s online media kit, at *http://www.nytimes.com/adinfo*.

User Tracking and Cookies

NYTimes.com uses cookies for authentication, tracking usage, targeting and controlling ad serving, including identifying whether a given user is allowed access to "premium" areas of the site. The terms under which this is conducted are detailed in the site's privacy statement, which includes a detailed FAQ on cookies. It is a condition of the registration process that users consent to this tracking using cookies, and again, this is clearly spelt out in the privacy policy.

Meta's Exhibit No. 1011
Page 0045

The privacy policy makes clear that their advertisers also use "clear GIFS", or "Web Beacons" – we've referred to these previously as "tags" – to monitor the ad-serving systems used on the site. The policy explains that this is happening, and outlines what is involved.

Data Analysis

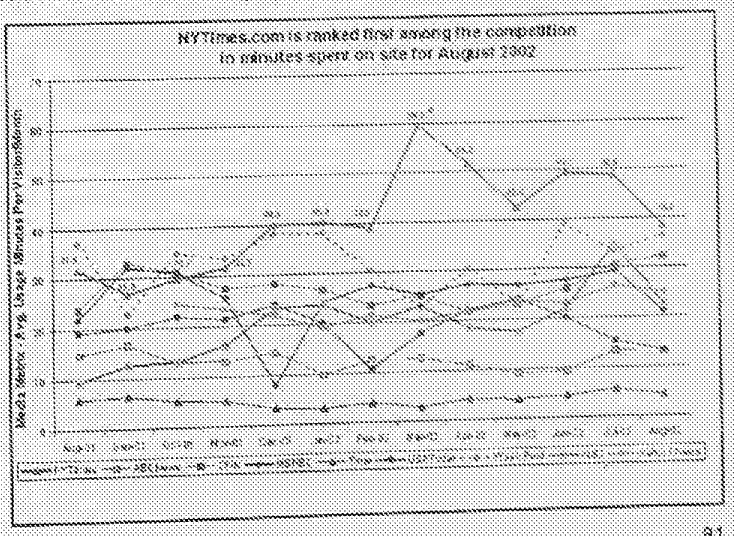
NYTimes.com uses an in-house log analysis system that parses the logs and reports the information, and also imports the data into an Oracle database. This allows them to conduct more specific queries on the data as the need arises. The output from the system can be in the form of HTML reports, which can be generated using a browser-based administration system, and Excel sheets and charts.

The system can track and report on:

- Usage patterns on the site
- Return visits for registered users
- Display of and response to advertising material
- Time spent by users on the site
- Visits referred from e-mails, including "*Today's Headlines*", which is sent to 3.3 million people every day
- International visitors (reported by comScore Networks)

Time spent on the site is considered to be a key indicator of visitor loyalty, and *NYTimes.com* currently leads the US news category in this respect.

NYTimes.com compared with other US news sites. Data from comScore Media Metrix



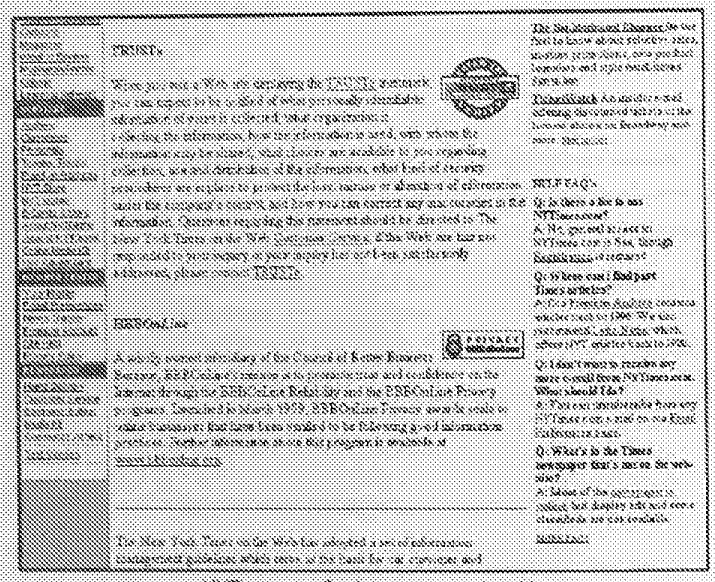
Meta's Exhibit No. 1011
Page 0046

Opting In, Opting Out

All mailing lists on *NYTimes.com* are maintained on a strictly opt-in basis. When a new user registered, they are sent an e-mail to confirm this, and another two weeks later with additional information about the site. In the event of the user not returning to the site in a year, then they are sent one further e-mail. Membership of the mailing lists are available to the registered users via their membership page, where they can select or deselect checkboxes to join or leave each list.

Seal Programs

NYTimes.com is a member of two Seal programs, having joined TRUSTe in 1997, and BBBOnline soon afterwards. The reason for wanting to join these programs was that The New York Times is a trusted brand in print media, and it was important to continue this trust online. The Seal Programs were an excellent way of maintaining this, as they are recognized representing industry standards.



NYTimes.com Seal program membership

The process of joining the programs was straightforward, although some slight modifications to the privacy statement were implemented to ensure full compliance. As privacy is an evolving subject, constant review and revision is important. Any changes that are made to the policy must involve notification to both programs, and adjustments made based on their amendments.

Meta's Exhibit No. 1011
Page 0047

Licenses from each program must be renewed annually, and approval requires the completion of a self-assessment report and questionnaire. Some time and resources are therefore required to manage the membership of the programs. *NYTimes.com* has a privacy committee with representatives from product development, marketing, PR, legal, and customer services departments, and this group meets regularly to discuss any privacy issues, industry developments, or policy changes. Online privacy is taken very seriously, and does require dedicated resources.

Summary

Privacy is increasingly important to the business of building and managing web sites as the general public (the users) become more aware of the issues and potential dangers to their privacy, and legislators begin to step in and rule on the subject. The overriding principle for the web site owner, however, is that good privacy is good business. By increasing the clarity of our privacy policies, we can increase the trust of those who use our sites, and so increase the potential for a more fruitful experience for everyone concerned, whether we are operating e-mail lists, publishing news services, or selling jam.

For this reason, it is advisable to have a privacy policy that is appropriate for our target audience. Owners of sites aimed at young children, for example, must exercise extreme care in their handling of users' information, but all sites that wish to develop a constructive relationship with their visitors (even those aimed at grown-ups) can benefit from taking privacy seriously. Many users will accept that certain information can legitimately be collected and used, to improve the site or service available, or to offer personalized content. Few web users will tolerate an incautious or dishonest attitude to their privacy rights, however, regardless of the legislation that may or may not apply.

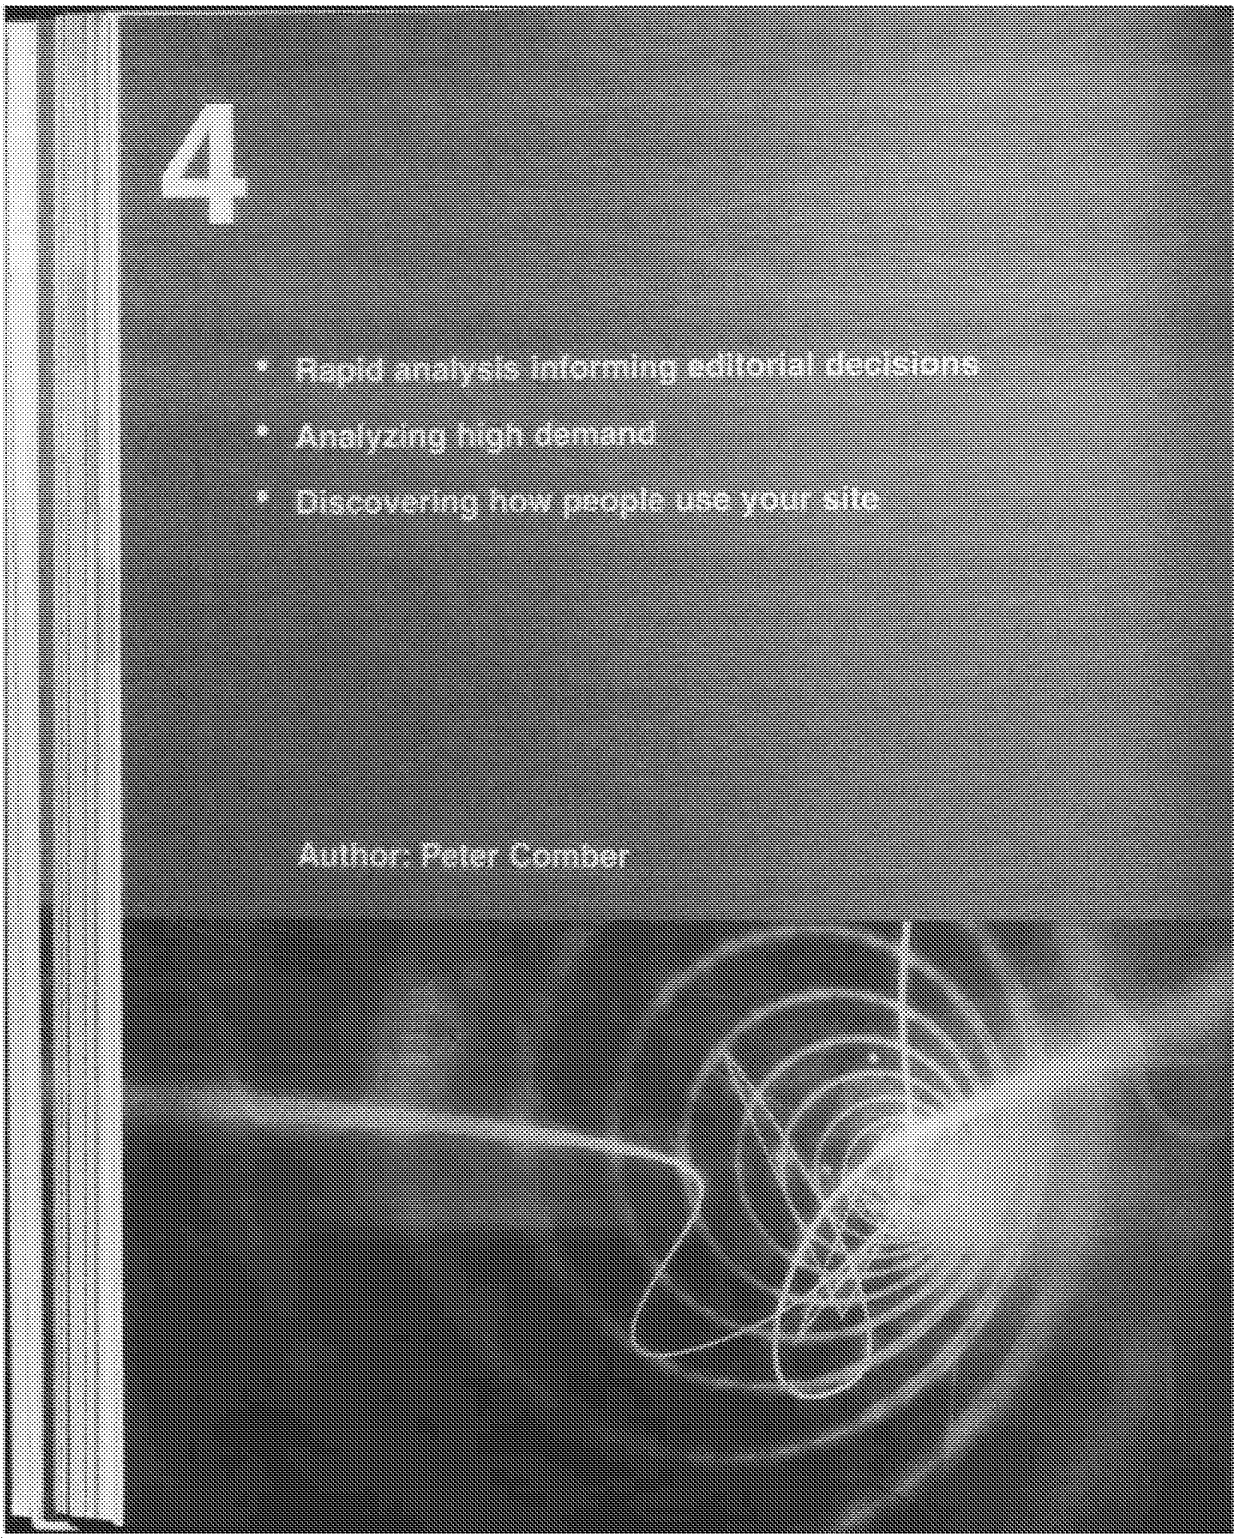
Privacy policies should in most cases have at least the following attributes:

- It must be clearly available. This usually means at least a link from the home page of the site, and also any other page where information is being collected, such as an e-mail registration page, or an online survey.
- It must say who is the operator of the site, and particularly contact details in the event of a query or dispute regarding the site's privacy practices, including any accountability mechanisms (such as TRUSTe or BBBOnline).
- It must say what information is being collected, what purpose this information will serve, whether it will be distributed to other organizations or used for any purpose other than that for which it was submitted, and the consequences for the individual if they refuse to disclose certain information. For example, they will not be able to register for a "premium" service if they refuse to submit their postal address.

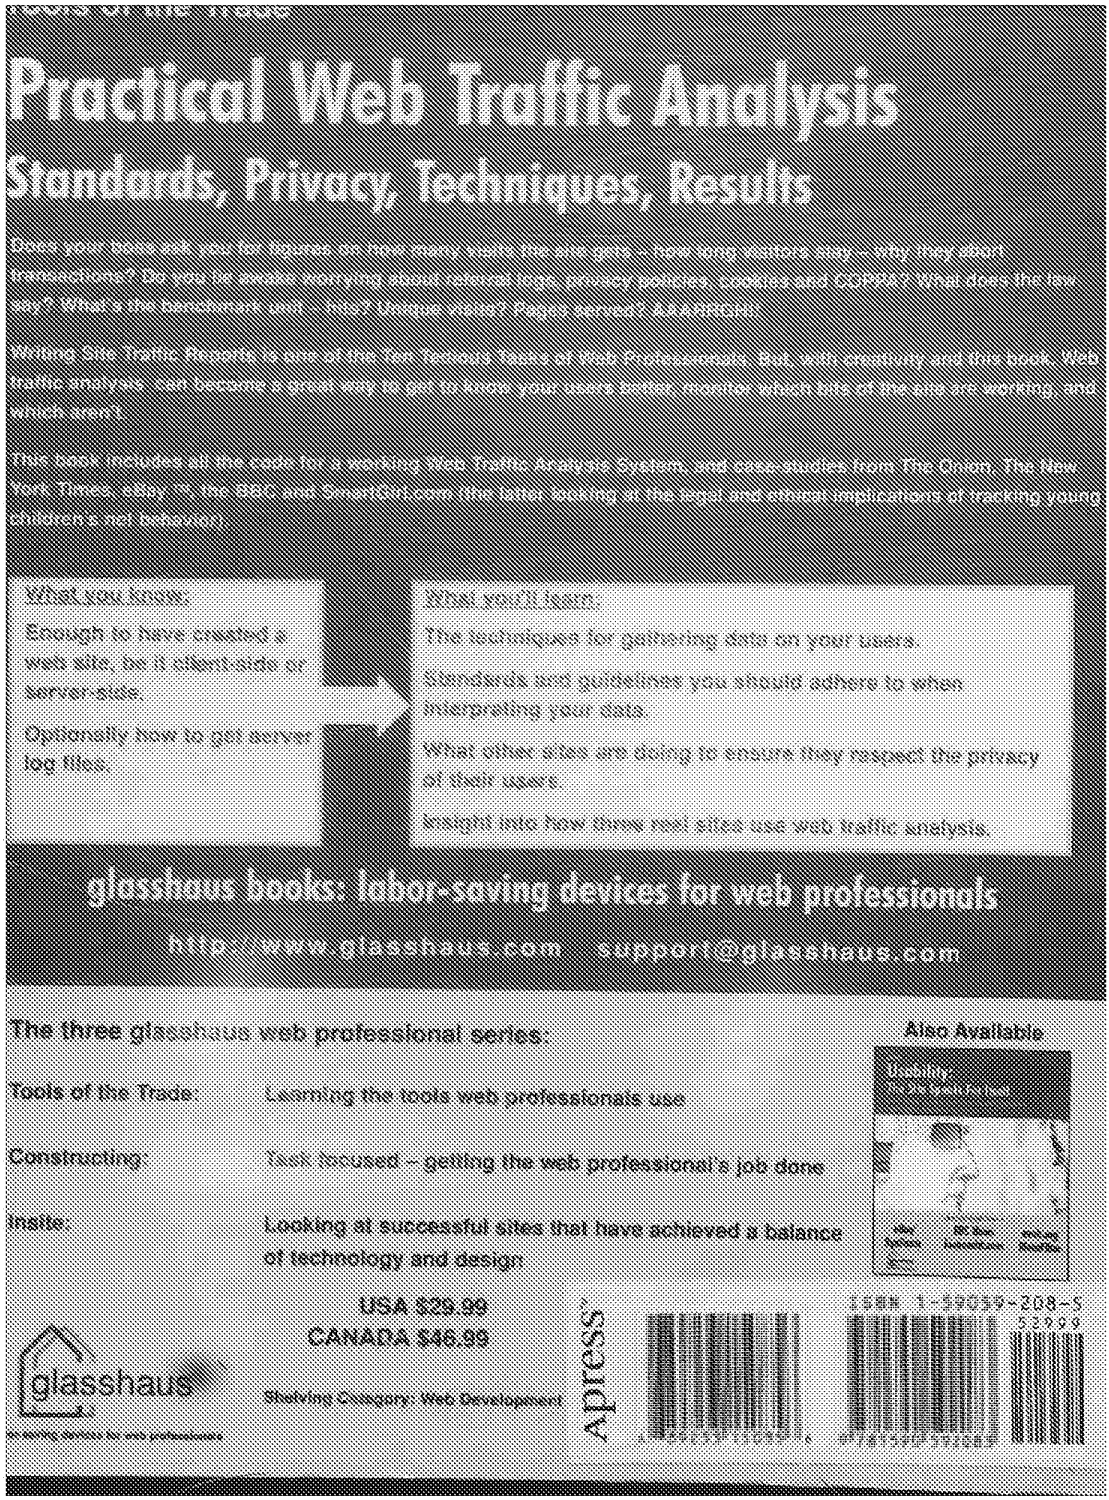
- It must say what measures the site takes to ensure the integrity and security of personal information that is stored, and the method for the individual to review, verify, and delete inaccurate or superfluous information.
- It must say which information the visitor has the right to opt out of, and which requires an opt-in (most commonly relevant to e-mail mailing lists, and the distribution of personal details to third parties or "carefully selected partners").
- It must say if cookies are used, and is so, give details. There is a lot of suspicion regarding the use of cookies, much of it unfounded, so the more explicit you can be the better. Explain what the cookie is, what it is used for, what it is NOT used for, and why it is in the user's interests to allow the cookie to be set.
- It must say where third parties are used, for example for the serving of ads on the site, or other partner programs. It must be explicit, and refer to the third parties' privacy policies if possible.

Emerging standards such as P3P are a useful tool for web site operators and users to talk to each other about privacy, but they are not a replacement for understanding the issues involved, and getting the policy right. Seal programs and similar self-regulatory schemes are another method available to us to increase trust and interaction, but again they are no replacement for understanding the needs and the concerns of our visitors, and ensuring that the requirements of the web site conform to them. Privacy is ultimately about openness.

Meta's Exhibit No. 1011
Page 0049



Meta's Exhibit No. 1011
Page 0050


Practical Web Traffic Analysis
Standards, Privacy, Techniques, Results
What you know:
Enough to have created a web site, be it client-side or server-side.
Optionally how to get server log files.
What you'll learn:
The techniques for gathering data on your users.
Standards and guidelines you should adhere to when interpreting your data.
What other sites are doing to ensure they respect the privacy of their users.
Insight into how three real sites use web traffic analysis.
glasshaus books: labor-saving devices for web professionals
http://www.glasshaus.com support@glasshaus.com
The three glasshaus web professional series:
Tools of the Trade: Learning the tools web professionals use
Constructing: Task focused – getting the web professional's job done
Inside: Looking at successful sites that have achieved a balance of technology and design
USA $29.99
CANADA $46.99
Shelving Category: Web Development
glasshaus
Also Available
Usability
Apress

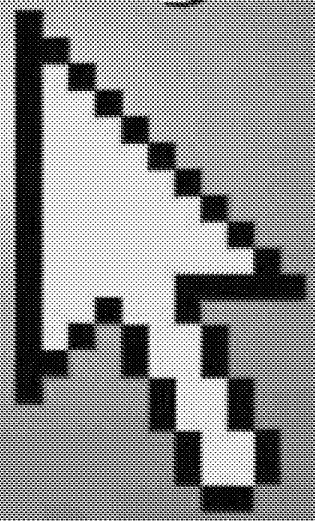
Microsoft

OVER 10,000 NEW TERMS

Microsoft Computer Dictionary

Fifth Edition

- *Fully updated with the latest technologies, terms, and acronyms*
- *Easy to read, expertly illustrated*
- *Definitive coverage of hardware, software, the Internet, and more!*



Meta's Exhibit No. 1012
Page 001

PUBLISHED BY
Microsoft Press
A Division of Microsoft Corporation
One Microsoft Way
Redmond, Washington 98052-6399

Copyright © 2002 by Microsoft Corporation

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

Library of Congress Cataloging-in-Publication Data
Microsoft Computer Dictionary.--5th ed.
p. cm.
ISBN 0-7356-1495-4
1. Computers--Dictionaries. 2. Microcomputers--Dictionaries.

AQ76.5. M52267 2002
004'.03--dc21 200219714

Printed and bound in the United States of America.

2 3 4 5 6 7 8 9 QWT 7 6 5 4 3 2

Distributed in Canada by H.B. Fenn and Company Ltd.

A CIP catalogue record for this book is available from the British Library.

Microsoft Press books are available through booksellers and distributors worldwide. For further information about international editions, contact your local Microsoft Corporation office or contact Microsoft Press International directly at fax (425) 936-7329. Visit our Web site at www.microsoft.com/mspress. Send comments to *mspinput@microsoft.com*.

Active Desktop, Active Directory, ActiveMovie, ActiveStore, ActiveSync, ActiveX, Authenticode, BackOffice, BizTalk, ClearType, Direct3D, DirectAnimation, DirectDraw, DirectInput, DirectMusic, DirectPlay, DirectShow, DirectSound, DirectX, Encourage, FoxPro, FrontPage, Hotmail, IntelliEye, IntelliMouse, IntelliSense, JScript, MapPoint, Microsoft, Microsoft Press, Mobile Explorer, MS-DOS, MSN, Music Central, NetMeeting, Outlook, PhotoDraw, PowerPoint, SharePoint, UltimateTV, Visio, Visual Basic, Visual C++, Visual FoxPro, Visual InterDev, Visual J++, Visual SourceSafe, Visual Studio, Win32, Win32s, Windows, Windows Media, Windows NT, Xbox are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Other product and company names mentioned herein may be the trademarks of their respective owners.

The example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious. No association with any real company, organization, product, domain name, e-mail address, logo, person, place, or event is intended or should be inferred.

Acquisitions Editor: Alex Blanton
Project Editor: Sandra Haynes

Body Part No. X08-41929

Introduction

The *Microsoft Computer Dictionary, Fifth Edition* is designed to be a comprehensive and authoritative source of definitions for computer-related terms and abbreviations. The dictionary includes terms drawn from a wide variety of topics relevant to computer users, including software, hardware, networking, data storage, graphics, games, information processing, the Internet and the World Wide Web, gaming, history, jargon and slang, organizations, programming, and standards.

Although this book covers nearly every aspect of computing, it does not include entries on most companies or on most makes and models of computers, nor does it contain entries on most application software products. The few exceptions to this rule of thumb are key companies and products that have a historical or universal importance within the computing industry.

This dictionary emphasizes terminology that the average computer user will encounter in documentation, online help, computer manuals, marketing and sales materials, the popular media, and the computer trade press. Because most computer users operate personal computers and desktop systems at home, work, or both, the majority of the entries in this dictionary cover the terminology used in describing and working with these systems. However, some specialized or highly technical language is included that pertains to areas of industry, academia, software and hardware development, and research. These terms have been included because they have a bearing on more common computer terminology or because they are of historical significance.

Changes in the Fifth Edition

The fifth edition of the *Microsoft Computer Dictionary* has been revised and expanded to include over 10,000 entries, reflecting the many advances in the computer field and including several areas that have come into prominence in the public eye, such as networking, Web authoring, and new technologies, such as .NET. The content from the Year 2000 appendix has been integrated into the body of the dictionary and a new appendix on emoticons and instant messaging symbols has been added.

Order of Presentation

Entries are alphabetized by letter. Spaces are ignored, as are characters such as hyphens and slashes; for example, *Baudot code* falls between *baud* and *baud rate*, and *machine-independent* falls between *machine identification* and *machine instruction*. Numbers and symbols are located at the beginning of the book and are listed in ascending ASCII order. If an entry begins with a letter or letters but contains a number, it is listed alphabetically, according to the initial letter(s), and then according to ASCII order. Thus, V20 precedes V.2x, and both precede VAB.

Entries

Entries are of two types: main entries, which contain full definitions, and synonymous cross-references, which contain *See* references to the appropriate main entries. Synonymous cross-references are generally secondary or less common ways of referring to a main entry. The definition at the main entry can be substituted as a definition for the synonymous cross-reference.

Format

Information in each main entry is presented in a consistent format: entry name in boldface, spelling variants (if any), part of speech, definition, illustration or table reference (if any), acronym (if any), alternative names (if any), and cross-references (if any).

Meta's Exhibit No. 1012
Page 003

Main Entries

Entries that are acronyms or abbreviations for one or more words or concatenations of two or more words have those words spelled out at the beginning of the definition. The letters in these words or phrases that make up the acronym, abbreviation, or concatenation are in boldface.

When a main entry is spelled exactly the same as another main entry, the two entries are differentiated by the use of a superscript numeral after each term. These entries are called homographs, and they are generally different parts of speech. For example,

e-mail[1] *(noun)*
e-mail[2] *(verb)*

Spelling Variants

When a main entry has one or more variations in the way it is spelled, each spelling variant follows the main entry, after the word *or*.

Parts of Speech

Entries are broken down into four parts of speech, in addition to prefixes, abbreviated as follows:

n. noun
vb. verb
adj. adjective
adv. adverb

Definitions

Each of the more than 10,000 entries is written in clear, standard English. Many go beyond a simple definition to provide additional detail and to put the term in context for a typical computer user. When an entry has more than one sense or definition, the definitions are presented in a numbered list, to make it easier to distinguish the particular, sometimes subtle, variations in meaning.

Illustration and Table References

Some entries have affiliated illustrations or tables that aid in defining the entry. In most cases, illustrations and tables appear on the same page as the entries to which they apply. In some instances, however, page layout requirements have forced them to a subsequent page. Entries with illustrations or tables usually have references at the end of the definition for an entry, in the following formats:

See the illustration.
See the table.

Acronyms

Some terminology in the computer field, particularly computer standards and Internet slang, can be shortened to form acronyms. Sometimes the acronym is the more common way to refer to the concept or object; in these cases, the acronym is the main entry. In other cases, the acronym is not as commonly used as the words or phrase for which it stands. In these cases, the words or phrase constitute the main entry. The acronym is given after the definition for these entries in the following format:

Acronym:

Alternative Names

Some items or concepts in the computer field can be referred to by more than one name. Generally, though, one way is preferred. The preferred terminology is the main entry. Alternative names are listed after any acronyms; otherwise they are listed after the definition in the following format:

Also called:

Cross-References

Cross-references are of three types: *See*, *See also*, and *Compare*. A *See* reference is used in an entry that is a synonymous cross-reference and simply points to another entry that contains the information sought. A *See also* reference points to one or more entries that contain additional or supplemental information about a topic and follows any acronyms or alternative names after the definition. A *Compare* reference points to an entry or entries that offer contrast and follows any *See also* references; otherwise it follows any acronyms or alternative names after the definition.

Meta's Exhibit No. 1012
Page 004

redirection *n.* The process of writing to or reading from a file or device different from the one that would normally be the target or the source. For example, the MS-DOS or OS/2 command *dir >prn* redirects a directory listing from the screen to the printer. *Compare* pipe.

redirector *n.* Software on a client computer that intercepts requests for information and, when appropriate, directs them to the network. Redirectors can either be built into the client operating system or be part of an added networking package.

redlining *n.* A feature of a word processing application that marks changes, additions, or deletions made to a document by a coauthor or editor. The purpose of redlining is to produce a record of the changes made to a document during the course of its development.

redraw *n. See* refresh (definition 1).

reduce *vb.* In a graphical user interface, to decrease the size of a window. A user can reduce a window either by clicking the appropriate button in the title bar or by clicking the mouse on the border of the window and dragging the border toward the middle of the window. *See also* maximize, minimize.

Reduced Instruction Set Computing *n. See* RISC.

redundancy *n.* Using one or more servers on a Web site to perform identical tasks. If one of the servers crashes, another server assumes its tasks. Redundancy ensures that the Web site will continue to function if one of the servers stops working.

redundancy check *n. See* CRC, LRC.

redundant code *n.* Code that duplicates a function performed elsewhere—for example, code to sort a list that has already been sorted.

reengineer *vb.* To rethink and redefine processes and procedures. In the context of computer systems, to reengineer means to change the way work is done in order to maximize the benefits of new technology.

reengineering *vb.* **1.** With regard to software, changing existing software to strengthen desirable characteristics and remove weaknesses. **2.** With regard to corporate management, using information technology principles to address the challenges posed by a global economy and to consolidate management of a rapidly expanding work force.

reentrant code *n.* Code written so that it can be shared by several programs at the same time. When a program is executing reentrant code, another program can interrupt the execution and can then start or continue execution of that same code. Many operating-system routines are written to be reentrant so that only one copy needs to reside in memory to serve all executing applications. *See also* relocatable code.

refactoring *n.* An optimization process in object-oriented programming intended to improve the design or structure of a program without changing its functionality. The goal of refactoring is to make the program clearer and easier to work with—in part by removing duplication, abstracting common behaviors, and refining class hierarchies—and to improve the extensibility and reusability of existing code.

reference[1] *n.* A data type in the C++ programming language. A reference must be initialized with a variable name. The reference then becomes an alias for that variable but actually stores the address of the variable.

reference[2] *vb.* To access a variable, such as an element in an array or a field in a record.

reference parameter *n.* A parameter in which the address of a variable, rather than the explicit value, is passed to the called routine. *See also* parameter.

reference type *n.* A data type that is represented by a reference (similar to a pointer) to the type's actual value. If a reference type is assigned to a variable, that variable references (or "points to") the original value. No copy is made. Reference types comprise classes, interfaces, delegates, and boxed value types. *See also* data type, value type.

reflecting software *n. See* reflector.

reflective liquid-crystal display *n.* A liquid crystal display that is not equipped with edge light or backlight to enhance readability but rather depends on reflecting ambient light, making it difficult to read in brightly lit environments such as the outdoors. *Also called:* reflective LCD.

reflective routing *n.* In wide area networks, the process of using a reflector to distribute data, thereby reducing the load of the network server. *See also* reflector.

reflector *n.* A program that sends messages to a number of users upon receipt of a signal from a single user. A common type of reflector is an e-mail reflector, which forwards any e-mail sent to it to the multiple recipients currently on its list. *See also* multiple recipients. *Compare* mail reflector.

reformat *vb.* **1.** In applications, to change the look of a document by altering stylistic details, such as font, layout, indention, and alignment. **2.** In data storage, to prepare for

Meta's Exhibit No. 1012
Page 005

Electronic Acknowledgement Receipt

EFS ID:	46328649
Application Number:	17714965
International Application Number:	
Confirmation Number:	7004
Title of Invention:	Targeted television advertisements based on online behavior
First Named Inventor/Applicant Name:	Roy Shkedi
Customer Number:	26362
Filer:	Louis J. Hoffman/Michelle Lisa
Filer Authorized By:	Louis J. Hoffman
Attorney Docket Number:	14-TV1-CON7
Receipt Date:	02-AUG-2022
Filing Date:	06-APR-2022
Time Stamp:	23:59:47
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Transmittal Letter	22-08-02-TV1-CON7-IDS_Cover.pdf	494115 fe4b185cd5231d1c3f7f26567eb484b16f18ac51	no	2
Warnings:					

Information:					
2	Information Disclosure Statement (IDS) Form (SB08)	22-08-02-TV1-CON7-Supp-IDS.pdf	477532 8aab3b9d875a701ece35a18cf443b33de721ce73	no	8
Warnings:					
Information:					
This is not an USPTO supplied IDS fillable form					
3	Non Patent Literature	Fletcher-P1-FILE.pdf	22366783 214865ff22bf0f9fa9bf8e517d536c1e14fe0b74	no	30
Warnings:					
Information:					
4	Non Patent Literature	Fletcher-P2-FILE.pdf	20766228 8d4fecbd657108ca1b08efac3216ac59201ba33a	no	21
Warnings:					
Information:					
5	Non Patent Literature	NPL-Excerpts-Microsoft-Computer-Dictionary.pdf	3056697 8fb4eef578841f97835ebc1677494b63992de7db	no	5
Warnings:					
Information:					
		Total Files Size (in bytes):	47161355		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor	:	Shkedi, Roy	Art Unit	:	3693 (in parent)
Serial No.	:	17/714,965	Examiner	:	Ekpo, Nnenna (in parent)
Filing Date	:	04/06/2022	Conf. No.	:	7004
Title	:	Targeted television advertisements based on online behavior			

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Filed via EFS-Web on August 2, 2022

INFORMATION DISCLOSURE STATEMENT

Dear Sir:

Assignee discloses the references listed on attached form PTO/SB/08. Assignee submits herewith copies of all NPL references that were not previously cited in this application.

Also, assignee wishes to direct the examiner to the co-pending applications and proceedings in the table below, as well as the references cited in those proceedings. All listed applications and the listed patent share the same first inventor:

Serial No.	Filing Date	Earliest Parent Filing Date
17/714,936	04/06/2022	06/14/2007
17/877,870	07/29/2022	04/17/2007
17/843,311	06/17/2022	08/03/2011 (prov.)
IPR2022-00953	05/24/2022	N/A - re: USP 8,566,164 (same family as SN 17/401,078, previously cited)
IPR2022-00959	05/13/2022	N/A - re: USP 8,595,069 (same family as SN 17/401,078, previously cited)
IPR2022-00960	05/13/2022	N/A - re: USP 10,321,198 (same family as SN 17/401,078, previously cited)
IPR2022-00773	03/28/2022	Re: parent USP 8,677,398

IPR2022-01236	07/05/2022	Re: parent USP 8,677,398
IPR2022-01315	07/27/2022	Re: parent USP 10,715,878

If the Office has any questions, please feel free to contact assignee's undersigned attorney of record.

Respectfully submitted,

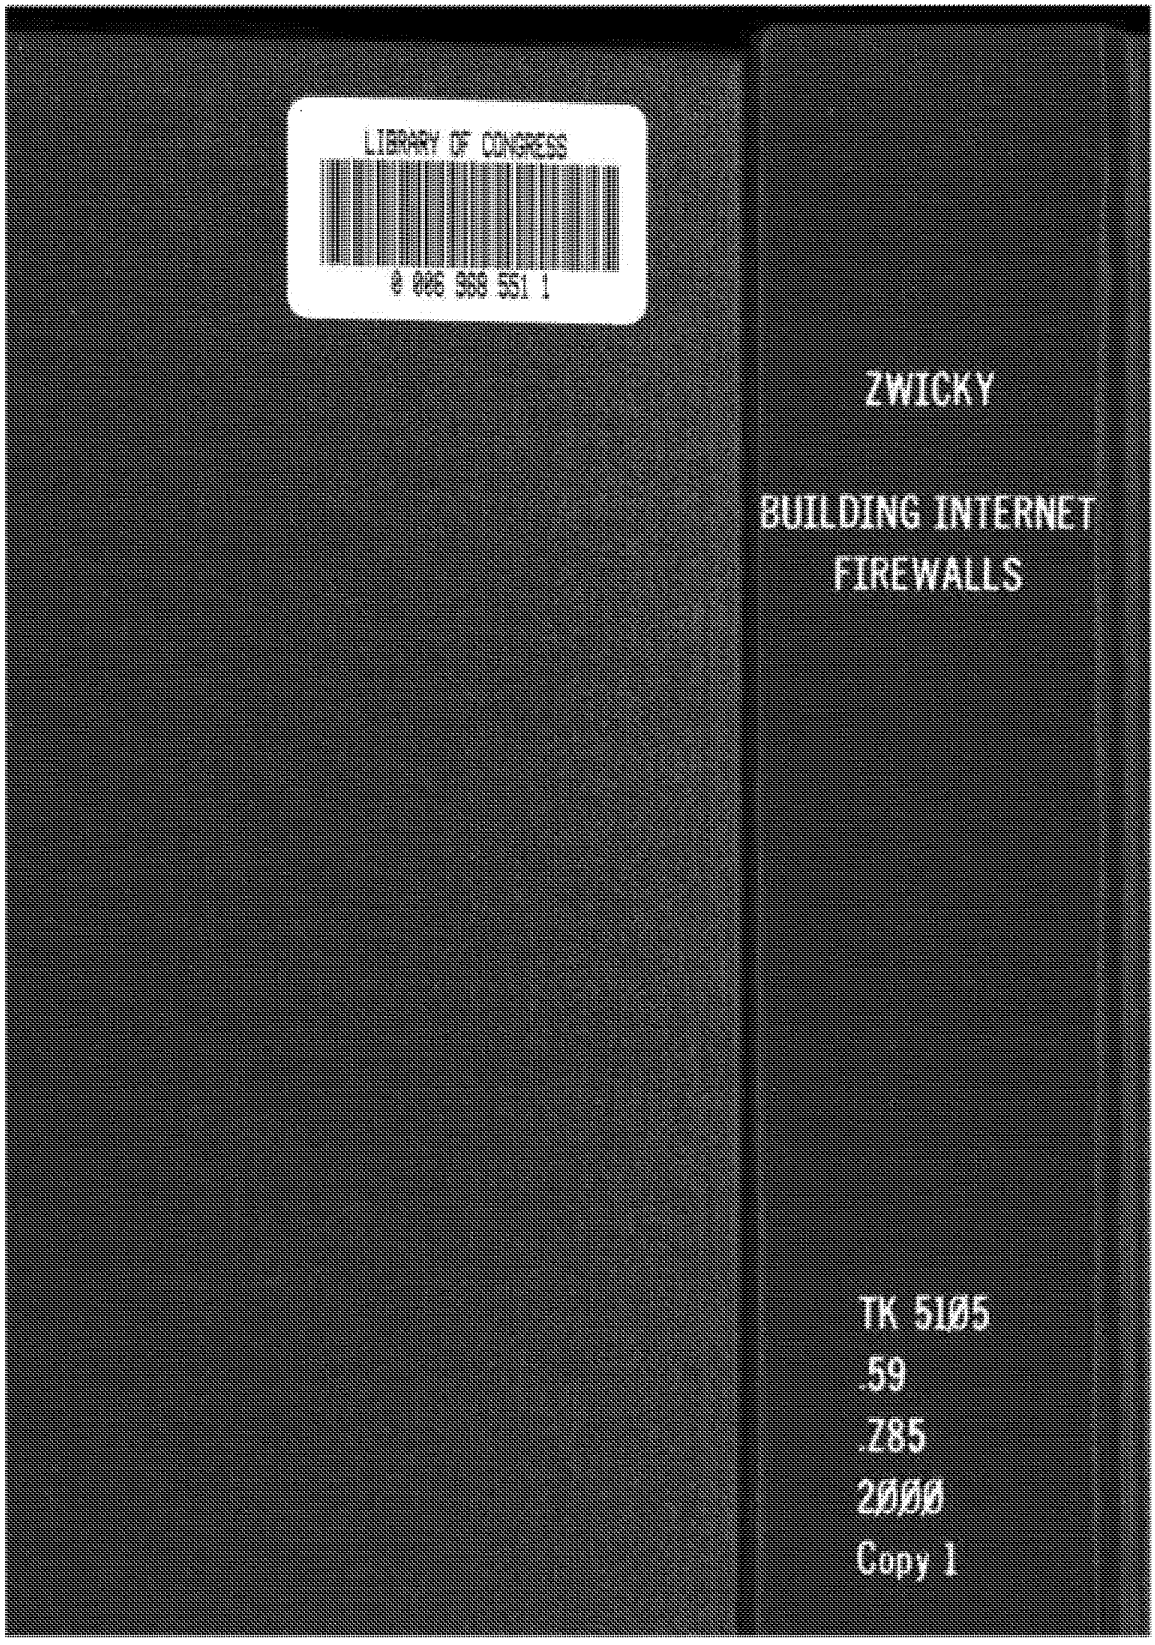
INTENT IQ, LLC.

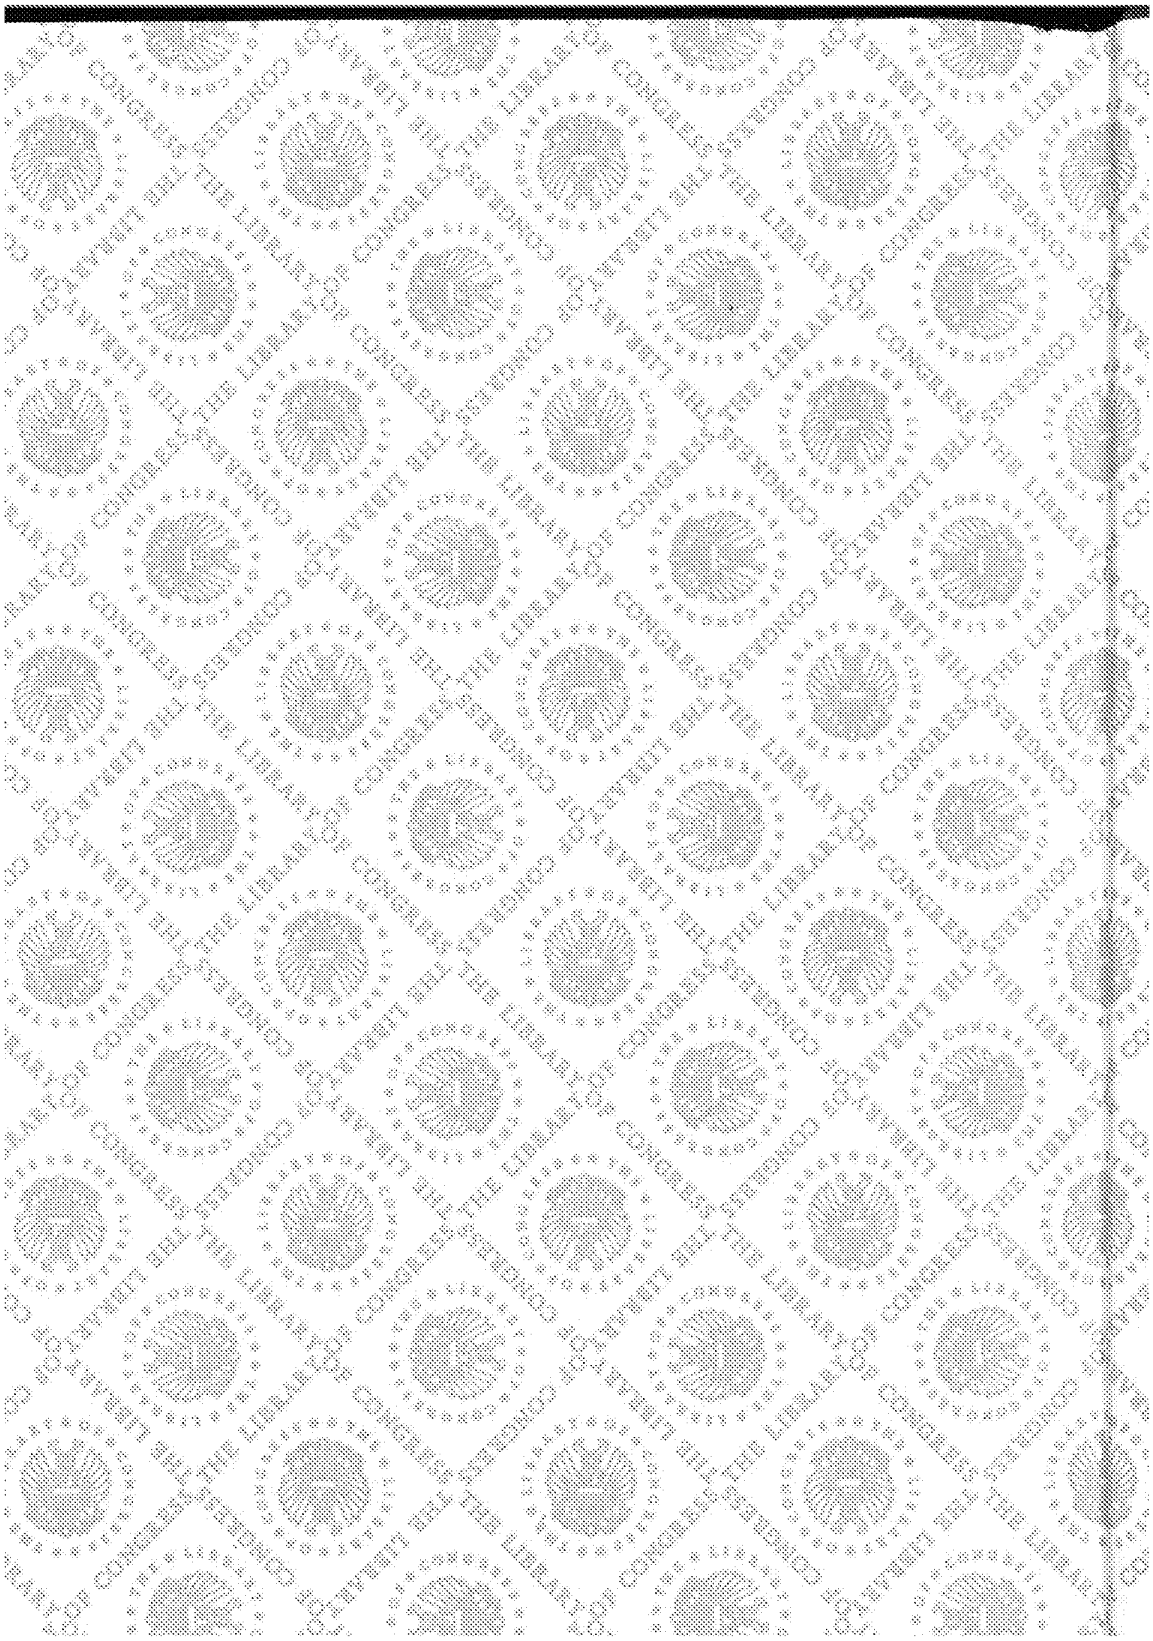
by its attorney

Dated: August 2, 2022

/Louis J. Hoffman/
Louis J. Hoffman
Reg. No. 38,918

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, Arizona 85260
(480) 948-3295


LIBRARY OF CONGRESS
0 006 959 551 1
ZWICKY
BUILDING INTERNET FIREWALLS
TK 5105
.59
.Z85
2000
Copy 1

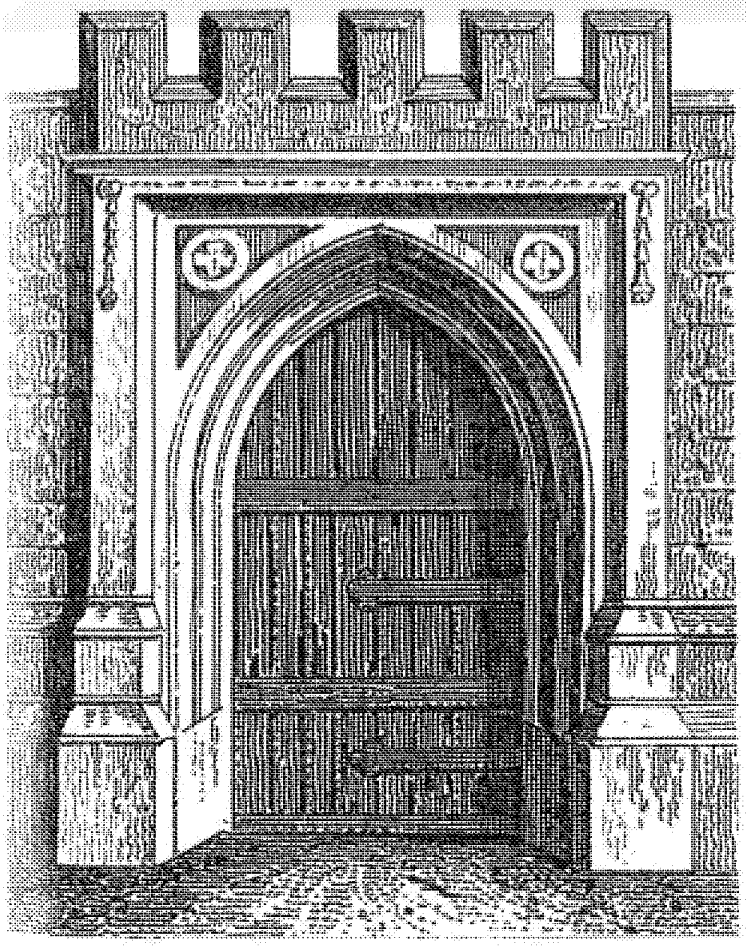


Internet and Web Security

2nd Edition

Covers Unix, Windows NT, and Linux

Building Internet Firewalls



Elizabeth D. Zwicky, Simon Cooper
& D. Brent Chapman

O'REILLY®

Meta's Exhibit No. 1006
Page 003

Building Internet Firewalls

Building Internet Firewalls

Second Edition

Elizabeth D. Zwicky, Simon Cooper &
D. Brent Chapman

O'REILLY®

Beijing · Cambridge · Farnham · Köln · Paris · Sebastopol · Taipei · Tokyo

Meta's Exhibit No. 1006
Page 005

Building Internet Firewalls, Second Edition
by Elizabeth D. Zwicky, Simon Cooper, and D. Brent Chapman

Copyright © 2000 O'Reilly & Associates, Inc. All rights reserved.
Printed in the United States of America.

Published by O'Reilly & Associates, Inc., 101 Morris Street, Sebastopol, CA 95472.

Editor: Deborah Russell

Production Editor: Nancy Crumpton

Production Coordinator: Madeleine Newell

Cover Designer: Edie Freedman

TK5105
.59
.Z85
2000
copy 1

Printing History:

April 1995:	First Edition.
November 1995:	Minor corrections.
June 2000:	Second Edition.

Nutshell Handbook, the Nutshell Handbook logo, and the O'Reilly logo are registered trademarks of O'Reilly & Associates, Inc. Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States and other countries. O'Reilly & Associates, Inc. is independent of Sun Microsystems.

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and O'Reilly & Associates, Inc. was aware of a trademark claim, the designations have been printed in caps or initial caps. The association between the image of a Gothic doorway and the topic of Internet firewalls is a trademark of O'Reilly & Associates, Inc.

While every precaution has been taken in the preparation of this book, the publisher assumes no responsibility for errors or omissions, or for damages resulting from the use of the information contained herein.

Library of Congress Cataloging-in-Publication Data

Zwicky, Elizabeth D.
Building Internet firewalls / Elizabeth D. Zwicky, Simon Cooper & D. Brent Chapman.--
2nd ed. p. cm.
ISBN 1-56592-871-7
1. Computer networks--Security measures. 2. Internet (Computer network)--Security measures. I. Cooper, Simon. II. Chapman, D. Brent. III. Title.

TK5105.59.Z85 2000
005.8--dc21 00-039954

ISBN: 1-56592-871-7
[M]

LIBRARY OF CONGRESS
DEC 2 1 2000
COPYRIGHT OFFICE

Table of Contents

Preface ..... *xiii*

I. Network Security *1*

1. Why Internet Firewalls? *3*
What Are You Trying to Protect? *4*
What Are You Trying to Protect Against? *7*
Who Do You Trust? *16*
How Can You Protect Your Site? *17*
What Is an Internet Firewall? *21*
Religious Arguments *28*

2. Internet Services *33*
Secure Services and Safe Services *35*
The World Wide Web *35*
Electronic Mail and News *40*
File Transfer, File Sharing, and Printing *43*
Remote Access *48*
Real-Time Conferencing Services *51*
Naming and Directory Services *52*
Authentication and Auditing Services *54*
Administrative Services *55*
Databases *58*
Games *58*

Meta's Exhibit No. 1006
Page 007

3. *Security Strategies* 59
Least Privilege 59
Defense in Depth 61
Choke Point 62
Weakest Link 63
Fail-Safe Stance 64
Universal Participation 67
Diversity of Defense 68
Simplicity 70
Security Through Obscurity 71

II. Building Firewalls 73

4. *Packets and Protocols* 75
What Does a Packet Look Like? 75
IP 79
Protocols Above IP 85
Protocols Below IP 93
Application Layer Protocols 94
IP Version 6 94
Non-IP Protocols 96
Attacks Based on Low-Level Protocol Details 97

5. *Firewall Technologies* 102
Some Firewall Definitions 102
Packet Filtering 104
Proxy Services 110
Network Address Translation 114
Virtual Private Networks 119

6. *Firewall Architectures* 122
Single-Box Architectures 122
Screened Host Architectures 126
Screened Subnet Architectures 128
Architectures with Multiple Screened Subnets 133
Variations on Firewall Architectures 137
Terminal Servers and Modem Pools 148
Internal Firewalls 149

Meta's Exhibit No. 1006
Page 008

7. *Firewall Design* 157
Define Your Needs 157
Evaluate the Available Products 159
Put Everything Together 162

8. *Packet Filtering* 165
What Can You Do with Packet Filtering? 166
Configuring a Packet Filtering Router 171
What Does the Router Do with Packets? 173
Packet Filtering Tips and Tricks 178
Conventions for Packet Filtering Rules 180
Filtering by Address 183
Filtering by Service 185
Choosing a Packet Filtering Router 190
Packet Filtering Implementations for General-Purpose Computers 203
Where to Do Packet Filtering 214
What Rules Should You Use? 216
Putting It All Together 216

9. *Proxy Systems* 224
Why Proxying? 225
How Proxying Works 226
Proxy Server Terminology 231
Proxying Without a Proxy Server 232
Using SOCKS for Proxying 233
Using the TIS Internet Firewall Toolkit for Proxying 237
Using Microsoft Proxy Server 238
What If You Can't Proxy? 239

10. *Bastion Hosts* 241
General Principles 242
Special Kinds of Bastion Hosts 243
Choosing a Machine 244
Choosing a Physical Location 248
Locating Bastion Hosts on the Network 249
Selecting Services Provided by a Bastion Host 250
Disabling User Accounts on Bastion Hosts 253
Building a Bastion Host 255
Securing the Machine 256
Disabling Nonrequired Services 259

Meta's Exhibit No. 1006
Page 009

Operating the Bastion Host 269
Protecting the Machine and Backups 270

11. Unix and Linux Bastion Hosts 273
Which Version of Unix? 273
Securing Unix 275
Disabling Nonrequired Services 278
Installing and Modifying Services 288
Reconfiguring for Production 291
Running a Security Audit 295

12. Windows NT and Windows 2000 Bastion Hosts 297
Approaches to Building Windows NT Bastion Hosts 297
Which Version of Windows NT? 298
Securing Windows NT 299
Disabling Nonrequired Services 301
Installing and Modifying Services 313

III. Internet Services 315

13. Internet Services and Firewalls 317
Attacks Against Internet Services 319
Evaluating the Risks of a Service 327
Analyzing Other Protocols 334
What Makes a Good Firewalled Service? 336
Choosing Security-Critical Programs 339
Controlling Unsafe Configurations 347

14. Intermediary Protocols 349
Remote Procedure Call (RPC) 349
Distributed Component Object Model (DCOM) 358
NetBIOS over TCP/IP (NetBT) 359
Common Internet File System (CIFS) and Server Message Block (SMB) 361
Common Object Request Broker Architecture (CORBA) and Internet Inter-Orb Protocol (IIOP) 365
ToolTalk 367
Transport Layer Security (TLS) and Secure Socket Layer (SSL) 368
The Generic Security Services API (GSSAPI) 373
IPsec 373

Meta's Exhibit No. 1006
Page 0010

Remote Access Service (RAS) 377
Point-to-Point Tunneling Protocol (PPTP) 378
Layer 2 Transport Protocol (L2TP) 381

15. *The World Wide Web* 384
HTTP Server Security 385
HTTP Client Security 390
HTTP 397
Mobile Code and Web-Related Languages 406
Cache Communication Protocols 412
Push Technologies 415
RealAudio and RealVideo 417
Gopher and WAIS 419

16. *Electronic Mail and News* 423
Electronic Mail 423
Simple Mail Transfer Protocol (SMTP) 430
Other Mail Transfer Protocols 441
Microsoft Exchange 442
Lotus Notes and Domino 443
Post Office Protocol (POP) 445
Internet Message Access Protocol (IMAP) 448
Microsoft Messaging API (MAPI) 450
Network News Transfer Protocol (NNTP) 450

17. *File Transfer, File Sharing, and Printing* 454
File Transfer Protocol (FTP) 455
Trivial File Transfer Protocol (TFTP) 468
Network File System (NFS) 470
File Sharing for Microsoft Networks 479
Summary of Recommendations for File Sharing 482
Printing Protocols 483
Related Protocols 487

18. *Remote Access to Hosts* 488
Terminal Access (Telnet) 488
Remote Command Execution 491
Remote Graphical Interfaces 507

19. *Real-Time Conferencing Services* 520
Internet Relay Chat (IRC) 520
ICQ 523
talk 525
Multimedia Protocols 528
NetMeeting 533
Multicast and the Multicast Backbone (MBONE) 535

20. *Naming and Directory Services* 539
Domain Name System (DNS) 539
Network Information Service (NIS) 563
NetBIOS for TCP/IP Name Service and Windows Internet Name Service 565
The Windows Browser 576
Lightweight Directory Access Protocol (LDAP) 583
Active Directory 585
Information Lookup Services 586

21. *Authentication and Auditing Services* 591
What Is Authentication? 592
Passwords 596
Authentication Mechanisms 600
Modular Authentication for Unix 604
Kerberos 609
NTLM Domains 615
Remote Authentication Dial-in User Service (RADIUS) 622
TACACS and Friends 625
Auth and identd 627

22. *Administrative Services* 630
System Management Protocols 630
Routing Protocols 637
Protocols for Booting and Boot-Time Configuration 644
ICMP and Network Diagnostics 647
Network Time Protocol (NTP) 654
File Synchronization 658
Mostly Harmless Protocols 661

23. Databases and Games 664
Databases 664
Games 678

24. Two Sample Firewalls 681
Screened Subnet Architecture 681
Merged Routers and Bastion Host Using General-Purpose Hardware 704

IV. Keeping Your Site Secure 721

25. Security Policies 723
Your Security Policy 724
Putting Together a Security Policy 731
Getting Strategic and Policy Decisions Made 734
What If You Can't Get a Security Policy? 741

26. Maintaining Firewalls 742
Housekeeping 742
Monitoring Your System 746
Keeping up to Date 758
How Long Does It Take? 762
When Should You Start Over? 762

27. Responding to Security Incidents 764
Responding to an Incident 764
What to Do After an Incident 774
Pursuing and Capturing the Intruder 775
Planning Your Response 778
Being Prepared 787

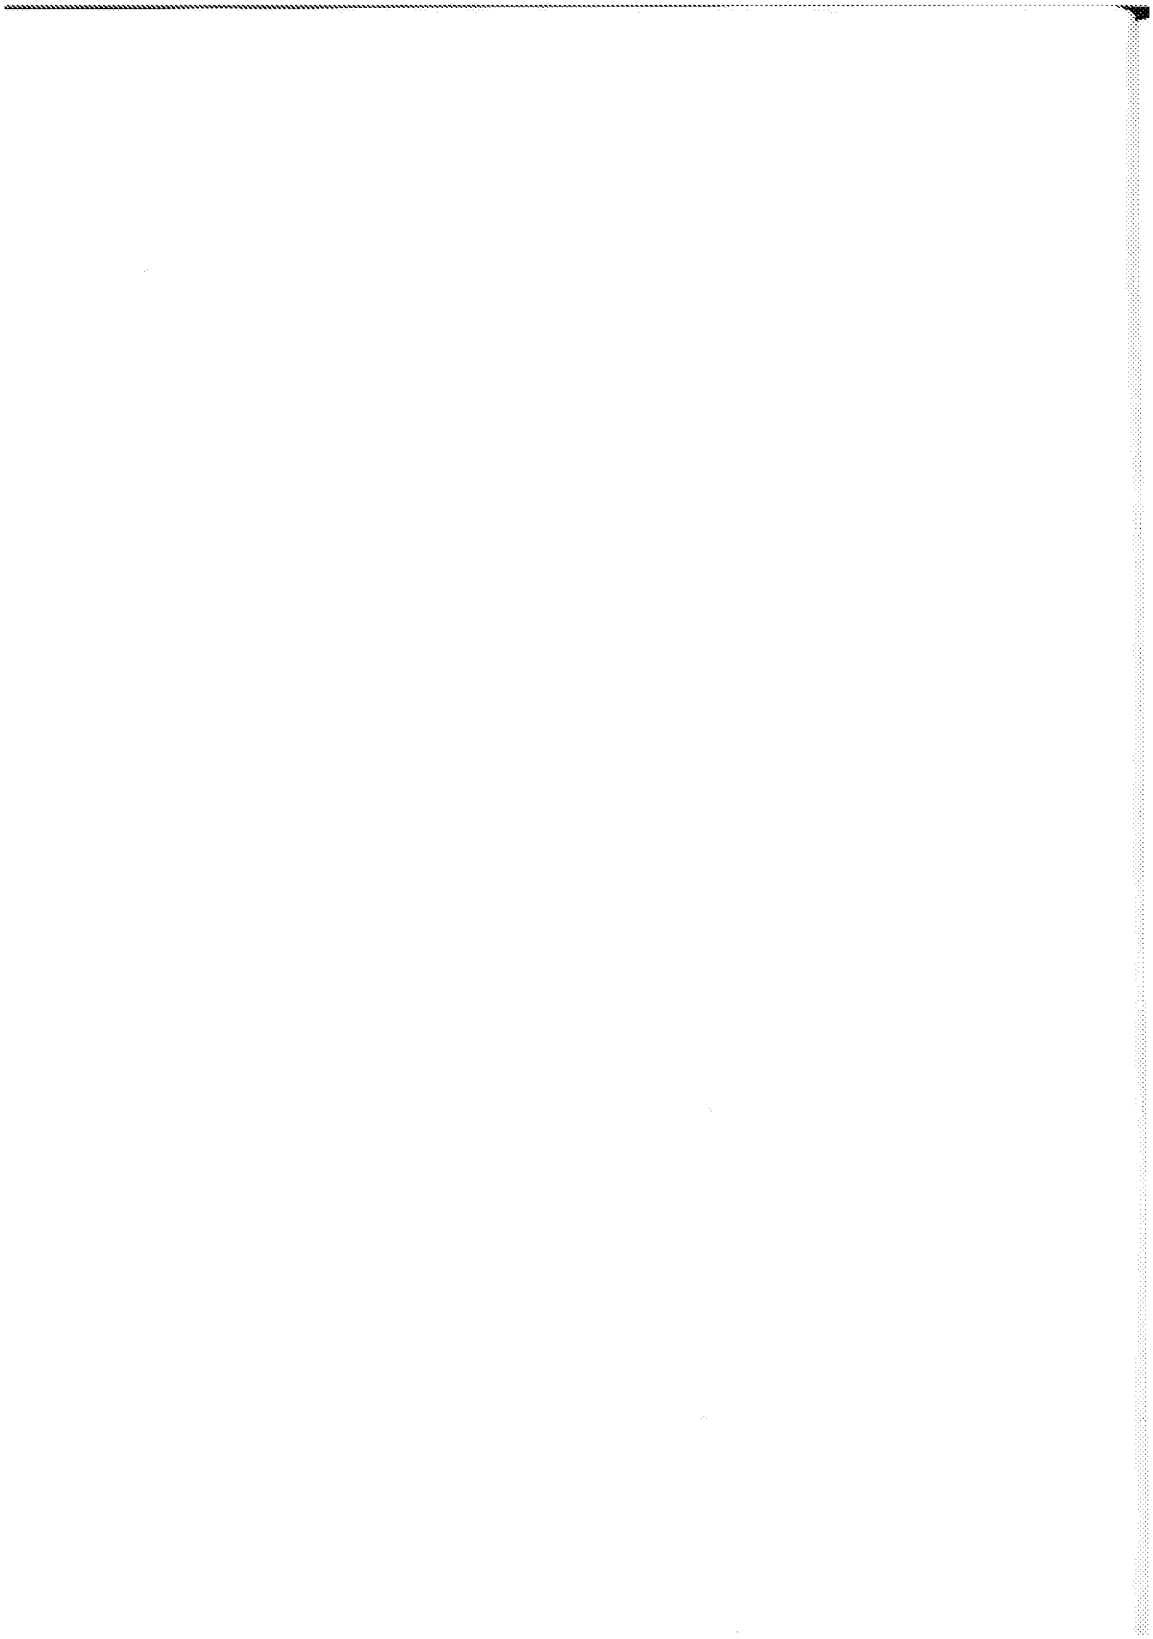
V. Appendixes 795

A. Resources 797

B. Tools 813

C. Cryptography 823

Index 849



Preface

This book is a practical guide to building your own firewall. It provides step-by-step explanations of how to design and install a firewall at your site and how to configure Internet services such as electronic mail, FTP, the World Wide Web, and others to work with a firewall. Firewalls are complex, though, and we can't boil everything down to simple rules. Too much depends on exactly what hardware, operating system, and networking you are using at your site, and what you want your users to be able to do and not do. We've tried to give you enough rules, examples, and resources here so you'll be able to do the rest on your own.

What is a firewall, and what does it do for you? A firewall is a way to restrict access between the Internet and your internal network. You typically install a firewall at the point of maximum leverage, the point where your network connects to the Internet. The existence of a firewall at your site can greatly reduce the odds that outside attackers will penetrate your internal systems and networks. The firewall can also keep your own users from compromising your systems by sending dangerous information—unencrypted passwords and sensitive data—to the outside world.

The attacks on Internet-connected systems we are seeing today are more serious and more technically complex than those in the past. To keep these attacks from compromising our systems, we need all the help we can get. Firewalls are a highly effective way of protecting sites from these attacks. For that reason, we strongly recommend you include a firewall in your site's overall Internet security plan. However, a firewall should be only one component in that plan. It's also vital that you establish a security policy, that you implement strong host security, and that you consider the use of authentication and encryption devices that work with the firewalls you install. This book will touch on each of these topics while maintaining its focus on firewalls.

Meta's Exhibit No. 1006
Page 0015

Scope of This Book

This book is divided into five parts.

Part I, *Network Security*, explores the problem of Internet security and focuses on firewalls as part of an effective strategy to address that problem.

- Chapter 1, *Why Internet Firewalls?*, introduces the major risks associated with using the Internet today; discusses what to protect, and what to protect against; discusses various security models; and introduces firewalls in the context of what they can and can't do for your site's security.
- Chapter 2, *Internet Services*, outlines the services users want and need from the Internet, and summarizes the security problems posed by those services.
- Chapter 3, *Security Strategies*, outlines the basic security principles an organization needs to understand before it adopts a security policy and invests in specific security mechanisms.

Part II, *Building Firewalls*, describes how to build firewalls.

- Chapter 4, *Packets and Protocols*, describes the basic network concepts firewalls work with.
- Chapter 5, *Firewall Technologies*, explains the terms and technologies used in building firewalls.
- Chapter 6, *Firewall Architectures*, describes the major architectures used in constructing firewalls, and the situations they are best suited to.
- Chapter 7, *Firewall Design*, presents the process of designing a firewall.
- Chapter 8, *Packet Filtering*, describes how packet filtering systems work, and discusses what you can and can't accomplish with them in building a firewall.
- Chapter 9, *Proxy Systems*, describes how proxy clients and servers work, and how to use these systems in building a firewall.
- Chapter 10, *Bastion Hosts*, presents a general overview of the process of designing and building the bastion hosts used in many firewall configurations.
- Chapter 11, *Unix and Linux Bastion Hosts*, presents the details of designing and building a Unix or Linux bastion host.
- Chapter 12, *Windows NT and Windows 2000 Bastion Hosts*, presents the details of designing and building a Windows NT bastion host.

Part III, *Internet Services*, describes how to configure services in the firewall environment.

- Chapter 13, *Internet Services and Firewalls*, describes the general issues involved in selecting and configuring services in the firewall environment.

- Chapter 14, *Intermediary Protocols*, discusses basic protocols that are used by multiple services.
- Chapter 15, *The World Wide Web*, discusses the Web and related services.
- Chapter 16, *Electronic Mail and News*, discusses services used for transferring electronic mail and Usenet news.
- Chapter 17, *File Transfer, File Sharing, and Printing*, discusses the services used for moving files from one place to another.
- Chapter 18, *Remote Access to Hosts*, discusses services that allow you to use one computer from another computer.
- Chapter 19, *Real-Time Conferencing Services*, discusses services that allow people to interact with each other online.
- Chapter 20, *Naming and Directory Services*, discusses the services used to distribute information about hosts and users.
- Chapter 21, *Authentication and Auditing Services*, discusses services used to identify users before they get access to resources, to keep track of what sort of access they should have, and to keep records of who accessed what and when.
- Chapter 22, *Administrative Services*, discusses other services used to administer machines and networks.
- Chapter 23, *Databases and Games*, discusses the remaining two major classes of popular Internet services, databases and games.
- Chapter 24, *Two Sample Firewalls*, presents two sample configurations for basic firewalls.

Part IV, *Keeping Your Site Secure*, describes how to establish a security policy for your site, maintain your firewall, and handle the security problems that may occur with even the most effective firewalls.

- Chapter 25, *Security Policies*, discusses the importance of having a clear and well-understood security policy for your site, and what that policy should and should not contain. It also discusses ways of getting management and users to accept the policy.
- Chapter 26, *Maintaining Firewalls*, describes how to maintain security at your firewall over time and how to keep yourself aware of new Internet security threats and technologies.
- Chapter 27, *Responding to Security Incidents*, describes what to do when a break-in occurs, or when you suspect that your security is being breached.

Meta's Exhibit No. 1006
Page 0017

Part V, *Appendixes*, consists of the following summary appendixes:

- Appendix A, *Resources*, contains a list of places you can go for further information and help with Internet security: World Wide Web pages, FTP sites, mailing lists, newsgroups, response teams, books, papers, and conferences.
- Appendix B, *Tools*, summarizes the best freely available firewall tools and how to get them.
- Appendix C, *Cryptography*, contains background information on cryptography that is useful to anyone trying to decrypt the marketing materials for security products.

Audience

Who should read this book? Although the book is aimed primarily at those who need to build firewalls, large parts of it are appropriate for everyone who is concerned about Internet security. This list tells you what sections are particularly applicable to you:

System administrators
: You should read the entire book.

Senior managers
: You should read at least Part I of the book. The chapters in Part I will introduce you to the various types of Internet threats, services, and security approaches and strategies. These chapters will also introduce you to firewalls and describe what firewalls can and cannot do to enforce Internet security. You should also read Chapter 5, which provides an overview of firewall technologies. In addition, Appendix A will tell you where to go for more information and resources.

Information technology managers and users
: You should read all of the chapters we've cited for the managers in the previous category. In addition, you should read Part III, which explains the kinds of issues that may arise at your site over time—for example, how to develop a security policy, keep up to date, and react if someone attacks your site.

Although this book provides general concepts of firewalls appropriate to any site, it focuses on "average" sites: small to large commercial or educational sites. If you are setting up a personal firewall, you may wish to read just Part I, Chapter 5, and the service chapters appropriate to the services you wish to run. If you are setting up a firewall for an extremely large site, all of the chapters will be useful to you, but you may find that you need to use additional techniques.

Meta's Exhibit No. 1006
Page 0018

Platforms

To a large extent, this book is platform-independent. Because most of the information provided here consists of general principles, most of it should be applicable to you, regardless of what equipment, software, and networking you are using. The most platform-specific issue is what type of system to use as a bastion host. People have successfully built bastion hosts (which we describe in Chapter 10) using all kinds of computers, including Unix systems, Windows NT machines, Macintoshes, VMS VAXes, and others.

Having said this, we must acknowledge that this book is strongly oriented towards Unix (including Linux), with Windows NT as a major secondary theme. There are several reasons for this orientation. First, these operating systems are the dominant operating systems in the Internet world. Unix is still the predominant operating system for Internet servers, although Windows NT is a strong presence. Another reason is, of course, that our own experience is primarily in the Unix world; we have entered the world of Windows NT only recently, as it started to intersect with the world of the Internet. Although we do speak Windows NT, we do so with a strong Unix accent.

Linux, while it is not strictly speaking Unix, is a close relative of the Unix we have spent our careers working with. In many cases, it is truer to the Unix tradition than commercial operating systems entitled to use the Unix trademark. While we do mention Linux by name in some places, you should bear in mind that all of our statements about Unix are meant to include Linux except when we explicitly state otherwise.

Similarly, when we mention "Windows NT", unless we explicitly mention versions, we mean both Windows NT 4 and Windows 2000. Windows 2000 is a direct descendant of Windows NT 4 and behaves like it in most important respects. We call out differences where appropriate (although you should bear in mind that Windows 2000 was being released as this book went to press; both the operating system and the world's experience with it are bound to have changed by the time you read this).

Products

It's impossible to give a complete list of commercial and publicly available products in this book because new products are constantly being introduced and capabilities are constantly being added to existing products. Instead, we concentrate on discussing generic features and capabilities, and the consequences of having—or not having—particular capabilities, so that you can make your own evaluation of the products currently available to you. We do periodically mention individual

products, some commercial and some publicly available, particularly when there are striking features of well-known products. This is not intended to be an endorsement of the products we mention, or a slight to products that we omit.

Examples

Writing a book of this nature requires a large number of examples with hostnames and addresses in them. In order to avoid offending or inconveniencing people, we have attempted to use only names and addresses that are not in use. In most cases, we have used names and addresses that are reserved and cannot be publicly registered. In particular, this is why most of the example hosts in this book are in the ".example" domain (reserved for this use in RFC 2606). In a few cases where we needed large numbers of hostnames and felt that using the reserved example namespace would be confusing, we have used names that can be registered; we have attempted to use names that are not currently registered and do not seem likely to be registered. We apologize to anybody who inadvertently uses one of these names and is inconvenienced.

We also apologize to those readers who have memorized the entire reserved IP address space, and find it upsetting that many of our illustrations show reserved IP addresses in use over the Internet. This is, of course, impossible in practice, and we show it only to avoid attracting undesirable attention to addresses that can be accessed over the Internet.

Conventions Used in This Book

The following conventions are used in this book:

Italic
: Used for file and directory names and URLs, and for the first mention of new terms under discussion.

`Constant width`
: Used for code examples.

`Constant width italic`
: In some code examples, indicates an element (e.g., a filename) that you supply.

The following icon is used in this book:



Indicates a tip, suggestion, or general note.

Comments and Questions

We have tested and verified the information in this book to the best of our ability, but you may find that features have changed (or even that we have made mistakes!). Please let us know about any errors you find, as well as your suggestions for future editions, by writing to:

O'Reilly & Associates
101 Morris Street
Sebastopol, CA 95472
800-998-9938 (in the U.S. or Canada)
707-829-0515 (international or local)
707-829-0104 (FAX)

You can also send us messages electronically. To be put on the mailing list or request a catalog, send email to:

info@oreilly.com

To ask technical questions or comment on the book, send email to:

bookquestions@oreilly.com

We have a web site for this book, where we'll list examples, errata, and any plans for future editions. You can access this page at:

http://www.oreilly.com/catalog/fire2

For more information about this book and others, see the O'Reilly web site:

http://www.oreilly.com

Acknowledgments for the Second Edition

As unlikely as it may seem, we still had no idea how much time and effort the second edition would take when we started working on it; what we expected to be a relatively simple effort has turned into a marathon. Even the smallest revision requires many hands, and a fully new edition requires what seems like a cast of thousands.

Thanks to those who reviewed the second edition and made helpful comments: Steve Beaty, David LeBlanc, Phil Cox, Eric Pearce, Chuck Phillips, Greg Rose, and Wietse Venema—and to Bruce Schneier and Diana Smetters who read Appendix C on a four-hour turnaround! Thanks to the entire editorial and production team at O'Reilly, especially project manager Madeleine Newell and production editor Nancy Crumpton.

Elizabeth says: My thanks to my friends, family, and colleagues for their patience and aid; my monomaniacal interest in network protocols coupled with emotional instability and intermittent overwork have required more than a reasonable and customary amount of tolerance. I am particularly indebted to Arnold Zwicky, Diana Smetters, Jeanne Dusseault, and Brent Chapman. Special thanks are due to my second father, Jacques Transue, who required me to take slow and calm breaks from writing. Thanks to Debby Russell and Sue Miller at O'Reilly for their deft, patient, and calm job of editing; and to Simon, who expected a simple writing project, got his life disrupted for more than a year and a half, and kept working anyway, even though we insisted on spelling everything in American instead of proper English. And thanks to the many O'Reilly people who helped to produce this book.

Simon says: I would like to thank my colleagues, my friends, and my family for their understanding and support during this project. Particular thanks go to Beryl Cooper, Mel Pleasant, Landon Curt Noll, Greg Bossert, James R. Martin II, Alesia Bischoff, and Cherry Mill for their encouragement and patience. A special mention goes to my ice hockey teammates—thanks for such an active alternative to writing. Enormous thanks to Elizabeth for asking me to coauthor and for coaching me through the process. Finally, thanks to Debby, Sue, and the staff of O'Reilly for putting this book into the hands of our readers.

Acknowledgments for the First Edition

Note: We've preserved these acknowledgments for the first edition because we continue to be grateful to the people who helped us with that edition. Note, however, that several parts of the first edition (e.g., the foreword and the TCP/IP appendix) are no longer included in the book.

When we set out to write this book, we had no idea that it would consume so much time and energy. We would never have succeeded without the help of many people.

Special thanks to Ed DeHart and Craig Hunt. Ed worked with Brent in the early stages of this book and wrote the foreword to it; we appreciate all that he has done to help. TCP/IP is essential for understanding the basics of firewall construction, and Craig Hunt, author of *TCP/IP Network Administration* (O'Reilly & Associates) has kindly let us excerpt much of that book's Chapters 1 and 2 in this book's Appendix C so readers who do not already have a TCP/IP background can get a jump start.

Thanks to all those who reviewed drafts of the book before publication and made helpful suggestions: Fred Avolio, Steve Bellovin, Niels Bjergstrom, Rik Farrow, Simson Garfinkel, Eliot Lear, Evi Nemeth, Steve Simmons, Steve Romig, Gene Spaf-

ford, Phil Trubey, and Mark Verber. Thanks as well to Eric Allman for answering many Sendmail questions and Paul Traina for answering many Cisco questions.

Thanks to all the people at O'Reilly & Associates who turned this manuscript into a finished book: to Mary Anne Weeks Mayo, the wonderful and patient project manager/copyeditor for the book; Len Muellner, Ellen Siever, and Norm Walsh, who converted the book from Word to SGML and contributed their tool-tweaking prowess; Chris Reilley, who created the many excellent diagrams; Edie Freedman, who designed the cover, and Nancy Priest, who designed the interior layout; John Files and Juliette Muellner, who assisted with production; Seth Maislin, who prepared the index; and Sheryl Avruch and Kismet McDonough-Chan, who did the final quality control on the book.

Brent says: I would like to extend personal thanks to my friends and family, for keeping me going for a year and a half while I worked on the book; to my staff at Great Circle Associates, for keeping my business going; to the many hundreds of folks who've attended my Internet Security Firewalls Tutorial, for providing the impetus for this whole endeavor (and for keeping my bills paid!); and to the many thousands of subscribers to the Firewalls mailing list on the Internet, for providing a stimulating environment to develop many of the ideas found in this book. I also owe a lot of thanks to Debby Russell, our editor at O'Reilly & Associates, for all her help and guidance, and to our technical reviewers, for all their wonderful comments and suggestions. Most of all, though, I'd like to thank my very good friend and coauthor, Elizabeth Zwicky, without whose collaboration and encouragement this book probably never would have been finished, and certainly wouldn't have been as good.

Elizabeth says: My thanks go to my friends, my family, and my colleagues at Silicon Graphics, for an almost infinite patience with my tendency to alternate between obsessing about the book and refusing to discuss anything even tangentially related to it. I'd like to particularly thank Arnold Zwicky, Diana Smetters, Greg Rose, Eliot Lear, and Jeanne Dusseault for their expert moral support (often during similar crises of their own). But the most thanks for this effort have to go to Debby and Brent, for giving me a chance to be part of an unexpected but extremely rewarding project.

5

Firewall Technologies

In Part I, we introduced Internet firewalls and summarized what they can and cannot do to improve network security. In this chapter, we present major firewalls concepts. What are the terms you will hear in discussions of Internet firewalls? What are the components that can be put together to build these common firewall architectures? How do you evaluate a firewall design? In the remaining chapters of this book, we'll describe these components and architectures in detail.

Some Firewall Definitions

You may be familiar with some of the following firewall terms, and some may be new to you. Some may seem familiar, but they may be used in a way that is slightly different from what you're accustomed to (though we try to use terms that are as standard as possible). Unfortunately, there is no completely consistent terminology for firewall architectures and components. Different people use terms in different—or, worse still, conflicting—ways. Also, these same terms sometimes have other meanings in other networking fields; the following definitions are for a firewalls context.

Here are some very basic definitions; we describe these terms in greater detail elsewhere:

Firewall
: A component or set of components that restricts access between a protected network and the Internet, or between other sets of networks.

Host
: A computer system attached to a network.

Meta's Exhibit No. 1006
Page 0024

Page 233 of 1032

Bastion host

A computer system that must be highly secured because it is vulnerable to attack, usually because it is exposed to the Internet and is a main point of contact for users of internal networks. It gets its name from the highly fortified projections on the outer walls of medieval castles.*

Dual-homed host

A general-purpose computer system that has at least two network interfaces (or homes).

Network address translation (NAT)

A procedure by which a router changes data in packets to modify the network addresses. This allows a router to conceal the addresses of network hosts on one side of it. This technique can enable a large number of hosts to connect to the Internet using a small number of allocated addresses or can allow a network that's configured with illegal or unroutable addresses to connect to the Internet using valid addresses. It is not actually a security technique, although it can provide a small amount of additional security. However, it generally runs on the same routers that make up part of the firewall.

Packet

The fundamental unit of communication on the Internet.

Packet filtering

The action a device takes to selectively control the flow of data to and from a network. Packet filters allow or block packets, usually while routing them from one network to another (most often from the Internet to an internal network, and vice versa). To accomplish packet filtering, you set up a set of rules that specify what types of packets (e.g., those to or from a particular IP address or port) are to be allowed and what types are to be blocked. Packet filtering may occur in a router, in a bridge, or on an individual host. It is sometimes known as *screening*.†

Perimeter network

A network added between a protected network and an external network, in order to provide an additional layer of security. A perimeter network is sometimes called a *DMZ*, which stands for *De-Militarized Zone* (named after the zone separating North and South Korea).

---

* Marcus Ranum, who is generally held responsible for the popularity of this term in the firewalls professional community, says, "Bastions . . . overlook critical areas of defense, usually having stronger walls, room for extra troops, and the occasional useful tub of boiling hot oil for discouraging attackers".

† Some networking literature (in particular, the BSD Unix release from Berkeley) uses the term "packet filtering" to refer to something else entirely (selecting certain packets off a network for analysis, as is done by the *etherfind* or *tcpdump* programs).

Proxy
: A program that deals with external servers on behalf of internal clients. Proxy clients talk to proxy servers, which relay approved client requests on to real servers, and relay answers back to clients.

Virtual private network (VPN)
: A network where packets that are internal to a private network pass across a public network, without this being obvious to hosts on the private network. In general, VPNs use encryption to protect the packets as they pass across the public network. VPN solutions are popular because it is often cheaper to connect two networks via public networks (for instance, getting them both Internet connections) than via private networks (like traditional leased-line connections between the sites).

The next few sections briefly describe the major technologies associated with firewalls: packet filtering, proxy services, network address translation, and virtual private networks.

There are legitimate questions about how to distinguish between packet filtering and proxying, particularly when dealing with complex packet filtering systems and simple proxies. Many people believe that systems that pay attention to individual protocols and/or modify packets should not be considered packet filters, and may even refer to these systems as *transparent proxies*. In fact, these systems don't behave much like older, simpler packet filtering systems, and it's a good idea not to apply generalizations about packet filtering to them blindly. On the other hand, they don't behave much like proxying systems, either.

Similarly, a number of proxying systems provide *generic proxies*, which essentially function like packet filters, accepting all traffic to a given port without analyzing it. It's advisable to pay close attention to the individual technology a product uses, without making assumptions based on whether it claims to be a packet filter or a proxy. However, many systems still are clearly packet filters or clearly proxies, so it is worth understanding what these technologies are and how they work.

Packet Filtering

Packet filtering systems route packets between internal and external hosts, but they do it selectively. They allow or block certain types of packets in a way that reflects a site's own security policy, as shown in Figure 5-1. The type of router used in a packet filtering firewall is known as a *screening router*.

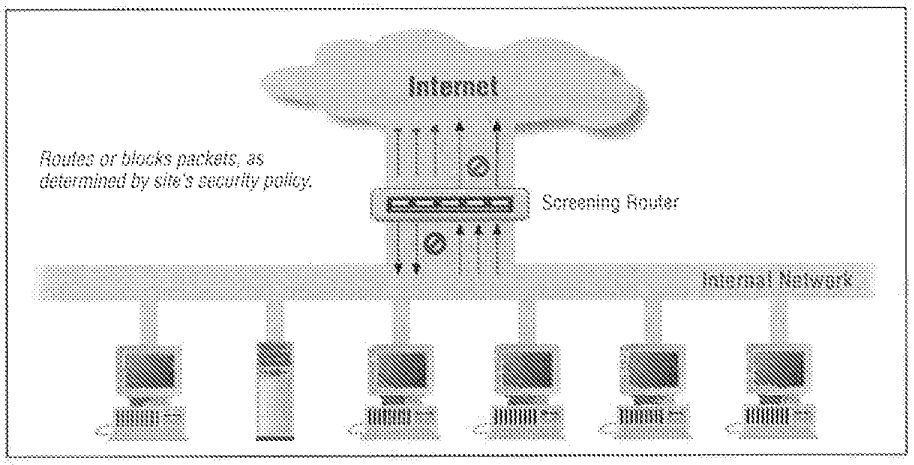


Figure 5-1. Using a screening router to do packet filtering

As we discuss in Chapter 8, *Packet Filtering*, every packet has a set of headers containing certain information. The main information is:

- IP source address
- IP destination address
- Protocol (whether the packet is a TCP, UDP, or ICMP packet)
- TCP or UDP source port
- TCP or UDP destination port
- ICMP message type
- Packet size

The router can also look past the packet headers at data further on in the packet; this allows it, for instance, to filter packets based on more detailed information (like the name of the web page that somebody is requesting) and to verify that packets appear to be formatted as expected for their destination port. The router can also make sure that the packet is valid (it actually is the size that it claims to be and is a legal size, for instance), which helps catch a number of denial of service attacks based on malformed packets.

In addition, the router knows things about the packet that aren't reflected in the packet itself, such as:

- The interface the packet arrives on
- The interface the packet will go out on

Meta's Exhibit No. 1006
Page 0027

Finally, a router that keeps track of packets it has seen knows some useful historical facts, such as:

- Whether this packet appears to be a response to another packet (that is, its source was the destination of a recent packet and its destination is the source of that other packet)
- How many other packets have recently been seen to or from the same host
- Whether this packet is identical to a recently seen packet
- If this packet is part of a larger packet that has been broken into parts (fragmented)

To understand how packet filtering works, let's look at the difference between an ordinary router and a screening router.

An ordinary router simply looks at the destination address of each packet and picks the best way it knows to send that packet towards that destination. The decision about how to handle the packet is based solely on its destination. There are two possibilities: the router knows how to send the packet towards its destination, and it does so; or the router does not know how to send the packet towards its destination, and it forgets about the packet and returns an ICMP "destination unreachable" message, to the packet's source.

A screening router, on the other hand, looks at packets more closely. In addition to determining whether or not it *can* route a packet towards its destination, a screening router also determines whether or not it *should*. "Should" or "should not" are determined by the site's security policy, which the screening router has been configured to enforce.

Packet filtering may also be performed by devices that pay attention only to "should" and "should not" and have no ability to route. Devices that do this are *packet filtering bridges*. They are rarer than packet filtering routers, mostly because they are dedicated security devices that don't provide all the other functions that routers do. Most sites would rather add features to routers that they need anyway, instead of adding a dedicated device. However, being a dedicated device provides advantages for packet filtering bridges; in particular, they are harder to detect and attack than packet filtering routers. They provide the same general features that we discuss for packet filtering routers.

Once it has looked at all the information, a straightforward packet filtering router can do any of the following things:

- Send the packet on to the destination it was bound for.
- Drop the packet—just forget it, without notifying the sender.
- Reject the packet—refuse to forward it, and return an error to the sender.

- Log information about the packet.
- Set off an alarm to notify somebody about the packet immediately.

More sophisticated routers might also be able to do one or more of these things:

- Modify the packet (for instance, to do network address translation).
- Send the packet on to a destination other than the one that it was bound for (for instance, to force transactions through a proxy server or perform load balancing).
- Modify the filtering rules (for instance, to accept replies to a UDP packet or to deny all traffic from a site that has sent hostile packets).

The fact that servers for particular Internet services reside at certain port numbers lets the router block or allow certain types of connections simply by specifying the appropriate port number (e.g., TCP port 23 for Telnet connections) in the set of rules specified for packet filtering. (Chapter 8, *Packet Filtering*, describes in detail how you construct these rules.)

Here are some examples of ways in which you might program a screening router to selectively route packets to or from your site:

- Block all incoming connections from systems outside the internal network, except for incoming SMTP connections (so that you can receive electronic mail).
- Block all connections to or from certain systems you distrust.
- Allow electronic mail and FTP services, but block dangerous services like TFTP, the X Window System, RPC, and the "r" services (*rlogin*, *rsh*, *rcp*, etc.). (See Chapter 13, *Internet Services and Firewalls*, for more information.)

Packet filtering devices that keep track of packets that they see are frequently called *stateful packet filters* (because they keep information about the state of transactions). They may also be called *dynamic packet filters* because they change their handling of packets dynamically depending on the traffic they see. Devices that look at the content of packets, rather than at just their headers, are frequently called *intelligent packet filters*. In practice, almost all stateful packet filters also are capable of looking at the contents of packets, and many are also capable of modifying the contents of packets, so you may see all these capabilities lumped together under the heading "stateful packet filtering". However, something can legitimately be called a "stateful packet filter" without having the ability to do advanced content filtering or modification.

A packet filtering system is also a logical place to provide virtual private network or network address translation services. Since the packet filter is already looking at all of the packets, it can easily identify packets that are intended for a destination

that is part of the virtual private network, encrypt those packets, and encapsulate them in another packet bound for the appropriate destination.

Advantages of Packet Filtering

Packet filtering has a number of advantages.

One screening router can help protect an entire network

One of the key advantages of packet filtering is that a single, strategically placed packet filtering router can help protect an entire network. If only one router connects your site to the Internet, you gain tremendous leverage on network security, regardless of the size of your site, by doing packet filtering on that router.

Simple packet filtering is extremely efficient

Because simple packet filtering requires paying attention only to a few packet headers, it can be done with very low overhead. Proxying is a fairly time-consuming operation, and adding proxying means directing connections through another program, usually on a machine that otherwise wouldn't be necessary to the routing process. Packet filtering takes place on a machine that was already in the critical path, and introduces a much smaller delay.

However, there is no free lunch; the more work your packet filters do, the slower they will be. If your packet filters behave like proxies, doing complicated data-driven operations that require keeping track of multiple packets, they will tend to perform like proxies as well.

Packet filtering is widely available

Packet filtering capabilities are available in many hardware and software routing products, both commercial and freely available over the Internet. Most sites already have packet filtering capabilities available in the routers they use.

Most commercial router products include packet filtering capabilities. Packet filtering capabilities are also available for a number of general-purpose computers. These are discussed further in Chapter 8, *Packet Filtering*.

Disadvantages of Packet Filtering

Although packet filtering provides many advantages, there are some disadvantages to using packet filtering as well.

Current filtering tools are not perfect

Despite the widespread availability of packet filtering in various hardware and software packages, packet filtering is still not a perfect tool. The packet filtering capabilities of many of these products share, to a greater or lesser degree, common limitations:

- The packet filtering rules tend to be hard to configure. Although there is a range of difficulty, it mostly runs from slightly mind-twisting to brain-numbingly impossible.
- Once configured, the packet filtering rules tend to be hard to test.
- The packet filtering capabilities of many of the products are incomplete, making implementation of certain types of highly desirable filters difficult or impossible.
- Like anything else, packet filtering packages may have bugs in them; these bugs are more likely than proxying bugs to result in security problems. Usually, a proxy that fails simply stops passing data, while a failed packet filtering implementation may allow packets it should have denied.

Packet filtering reduces router performance

Doing packet filtering places a significant extra load on a router. As we discussed previously, more complex filters place more load on the router, but in some cases, simply turning on packet filtering on a given interface can also cost you a lot of performance on some routers, because the filtering is incompatible with certain caching strategies commonly used for performance enhancement. Cisco's "fastpath" functionality is an example of this; normally, fastpath can perform basic routing functions completely on the interface card, without involving the main CPU, but using some forms of filtering requires involving the main CPU for each packet, which is much slower. What enables/disables fastpath depends on the hardware and software version.

Some policies can't readily be enforced by normal packet filtering routers

The information that a packet filtering router has available to it doesn't allow you to specify some rules you might like to have. For example, packets say what host they come from but generally not what user. Therefore, you can't enforce restrictions on particular users. Similarly, packets say what port they're going to but not what application; when you enforce restrictions on higher-level protocols, you do it by port number, hoping that nothing else is running on the port assigned to that protocol. Malicious insiders can easily subvert this kind of control.

This problem is eased by using more intelligent packet filters; however, in each case, you have to give up some of the advantages of normal packet filtering. For

instance, a packet filter can insist that users authenticate themselves before sending packets, and then it can filter packets by username. However, this removes the transparency advantage of normal packet filtering. A packet filter can also do protocol validity checking, but this is less than perfect and also increases filtering overhead.

Proxy Services

In general, a *proxy* is something or someone who does something on somebody else's behalf. For instance, you may give somebody the ability to vote for you by proxy in an election.

Proxy services are specialized application or server programs that take users' requests for Internet services (such as FTP and Telnet) and forward them to the actual services. The proxies provide replacement connections and act as gateways to the services. For this reason, proxies are sometimes known as *application-level gateways*.* In this book, when we are talking about proxy services, we are specifically talking about proxies run for security purposes, which are run on a firewall host: either a dual-homed host with an interface on the internal network and one on the external network, or some other bastion host that has access to the Internet and is accessible from the internal machines.

You will also run into proxies that are primarily designed for network efficiency instead of for security; these are *caching proxies*, which keep copies of the information for each request that they proxy. The advantage of a caching proxy is that if multiple internal hosts request the same data, the data can be provided directly by the proxy. Caching proxies can significantly reduce the load on network connections. There are proxy servers that provide both security and caching; in general, they are better at one purpose than the other.

Proxy services sit, more or less transparently, between a user on the inside (on the internal network) and a service on the outside (on the Internet). Instead of talking to each other directly, each talks to a proxy. Proxies handle all the communication between users and Internet services behind the scenes.

Transparency is the major benefit of proxy services. It's essentially smoke and mirrors. To the user, a proxy server presents the illusion that the user is dealing directly with the real server. To the real server, the proxy server presents the illu-

* Firewall terminologies differ. Whereas we use the term *proxy service* to encompass the entire proxy approach, other authors refer to *application-level gateways* and *circuit-level gateways*. Although there are small differences between the meanings of these various terms, which we'll explore in Chapter 9, *Proxy Systems*, in general our discussion of proxies refers to the same type of technology other authors mean when they refer to these gateway systems.

sion that the real server is dealing directly with a user on the proxy host (as opposed to the user's real host).

How do proxy services work? Let's look at the simplest case, where we add proxy services to a dual-homed host. (We'll describe these hosts in some detail in the "Dual-Homed Host" section in Chapter 6, *Firewall Architectures*.)



Proxy services are effective only when they're used in conjunction with a mechanism that restricts direct communications between the internal and external hosts. Dual-homed hosts and packet filtering are two such mechanisms. If internal hosts are able to communicate directly with external hosts, there's no need for users to use proxy services, and so (in general) they won't. Such a bypass probably isn't in accordance with your security policy.

As Figure 5-2 shows, a proxy service requires two components: a proxy server and a proxy client. In this illustration, the *proxy server* runs on the dual-homed host (as we discuss in Chapter 9, *Proxy Systems*, there are other ways to set up a proxy server). A *proxy client* is a special version of a normal client program (e.g., a Telnet or FTP client) that talks to the proxy server rather than to the "real" server out on the Internet; in some configurations, normal client programs can be used as proxy clients. The proxy server evaluates requests from the proxy client and decides which to approve and which to deny. If a request is approved, the proxy server contacts the real server on behalf of the client (thus the term *proxy*) and proceeds to relay requests from the proxy client to the real server, and responses from the real server to the proxy client.

In some proxy systems, instead of installing custom client proxy software, you'll use standard software but set up custom user procedures for using it. (We'll describe how this works in Chapter 9.)

There are also systems that provide a hybrid between packet filtering and proxying where a network device intercepts the connection and acts as a proxy or redirects the connection to a proxy; this allows proxying without making changes to the clients or the user procedures.

The proxy server doesn't always just forward users' requests on to the real Internet services. The proxy server can control what users do because it can make decisions about the requests it processes. Depending on your site's security policy, requests might be allowed or refused. For example, the FTP proxy might refuse to let users export files, or it might allow users to import files only from certain sites. More sophisticated proxy services might allow different capabilities to different hosts, rather than enforcing the same restrictions on all hosts.

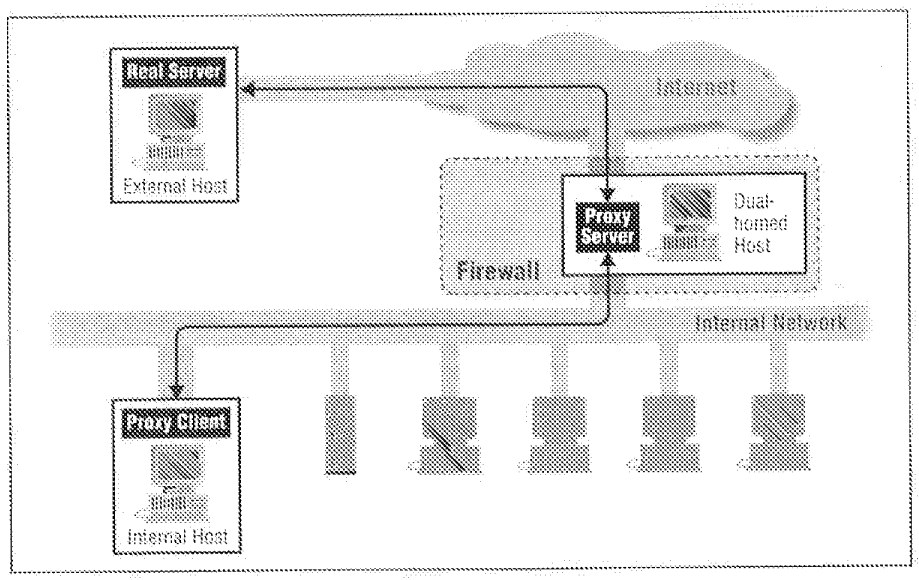


Figure 5-2. Using proxy services with a dual-homed host

Some proxy servers do in fact just forward requests on, no matter what they are. These may be called *generic proxies* or *port forwarders*. Programs that do this are providing basically the same protections that you would get if you had a packet filter in place that was allowing traffic on that port. You do not get any significant increase in security by replacing packet filters with proxies that do exactly the same thing (you gain some protection against malformed packets, but you lose by adding an attackable proxying program).

Some excellent software is available for proxying. SOCKS is a proxy construction toolkit, designed to make it easy to convert existing client/server applications into proxy versions of those same applications. The Trusted Information Systems Internet Firewall Toolkit (TIS FWTK) includes proxy servers for a number of common Internet protocols, including Telnet, FTP, HTTP, *rlogin*, X11, and others; these proxy servers are designed to be used in conjunction with custom user procedures. See the discussion of these packages in Chapter 9, *Proxy Systems*.

Many standard client and server programs, both commercial and freely available, now come equipped with their own proxying capabilities or with support for generic proxy systems like SOCKS. These capabilities can be enabled at runtime or compile time.

Most proxy systems are used to control and optimize outbound connections; they are controlled by the site where the clients are. It is also possible to use proxy systems to control and optimize inbound connections to servers (for instance, to

balance connections among multiple servers or to apply extra security). This is sometimes called *reverse proxying*.

Advantages of Proxying

There are a number of advantages to using proxy services.

Proxy services can be good at logging

Because proxy servers can understand the application protocol, they can allow logging to be performed in a particularly effective way. For example, instead of logging all of the data transferred, an FTP proxy server can log only the commands issued and the server responses received; this results in a much smaller and more useful log.

Proxy services can provide caching

Since all requests are passing through the proxy service anyway, the proxy can provide caching, keeping local copies of the requested data. If the number of repeat requests is significant, caching can significantly increase performance and reduce the load on network links.

Proxy services can do intelligent filtering

Since a proxy service is looking at specific connections, it is frequently able to do filtering more intelligently than a packet filter. For instance, proxy services are much more capable of filtering HTTP by content type (for instance, to remove Java or JavaScript) and better at virus detection than packet filtering systems.

Proxy systems can perform user-level authentication

Because a proxy system is actively involved in the connection, it is easy for it to do user authentication and to take actions that depend on the user involved. Although this is possible with packet filtering systems, it is much more difficult.

Proxy systems automatically provide protection for weak or faulty IP implementations

As a proxy system sits between a client and the Internet, it generates completely new IP packets for the client. It can therefore protect clients from deliberately malformed IP packets. (You just need a proxy system that isn't vulnerable to the bad packets!)

Disadvantages of Proxying

There are also some disadvantages to using proxy services.

Proxy services lag behind nonproxied services

Although proxy software is widely available for the older and simpler services like FTP and Telnet, proven software for newer or less widely used services is harder to find. There's usually a distinct lag between the introduction of a service and the availability of proxying servers for it; the length of the lag depends primarily on how well the service is designed for proxying. This makes it difficult for a site to offer new services immediately as they become available. Until suitable proxy software is available, a system that needs new services may have to be placed outside the firewall, opening up potential security holes. (Some services can be run through generic proxies, which will give at least minimal protection.)

Proxy services may require different servers for each service

You may need a different proxy server for each protocol, because the proxy server may need to understand the protocol in order to determine what to allow and disallow, and in order to masquerade as a client to the real server and as the real server to the proxy client. Collecting, installing, and configuring all these various servers can be a lot of work. Again, you may be able to use a generic proxy, but generic proxies provide only the same sorts of protection and functionality that you could get from packet filters.

Products and packages differ greatly in the ease with which they can be configured, but making things easier in one place can make it harder in others. For example, servers that are particularly easy to configure can be limited in flexibility; they're easy to configure because they make certain assumptions about how they're going to be used, which may or may not be correct or appropriate for your site.

Proxy services usually require modifications to clients, applications, or procedures

Except for services designed for proxying, you will need to use modified clients, applications, and/or procedures. These modifications can have drawbacks; people can't always use the readily available tools with their normal instructions.

Because of these modifications, proxied applications don't always work as well as nonproxied applications. They tend to bend protocol specifications, and some clients and servers are less flexible than others.

Network Address Translation

Network address translation (NAT) allows a network to use one set of network addresses internally and a different set when dealing with external networks. Network address translation does not, by itself, provide any security, but it helps to

conceal the internal network layout and to force connections to go through a choke point (because connections to untranslated addresses will not work, and the choke point does the translation).

Like packet filtering, network address translation works by having a router do extra work. In this case, not only does the router send packets on, but it also modifies them. When an internal machine sends a packet to the outside, the network address translation system modifies the source address of the packet to make the packet look as if it is coming from a valid address. When an external machine sends a packet to the inside, the network address translation system modifies the destination address to turn the externally visible address into the correct internal address. The network address translation system can also modify the source and destination port numbers (this is sometimes called Port and Address Translation or PAT). Figure 5-3 shows a network address translation system modifying only addresses, while Figure 5-4 shows port and address translation.

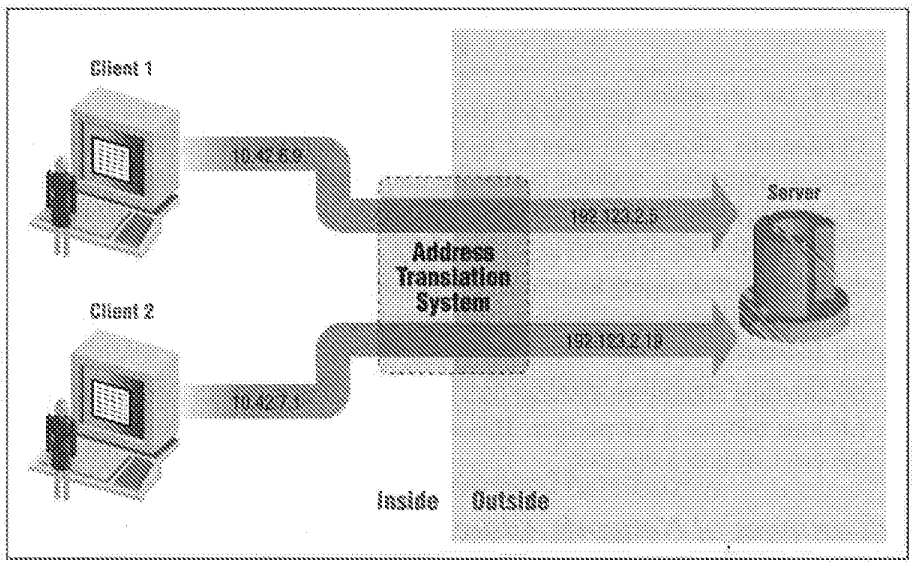


Figure 5-3. Network address translation

Network address translation systems can use different schemes for translating between internal and external addresses:

- Allocate one external host address for each internal address and always apply the same translation. This provides no savings in address space, and it slows down connections, it is normally a temporary measure used by sites that have been using illegal address spaces but are in the process of moving to using valid addresses.

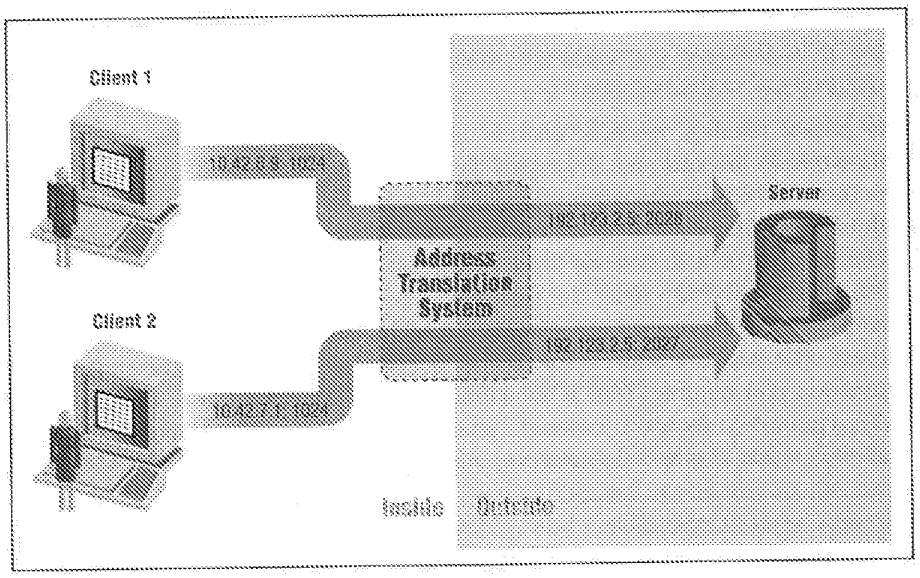


Figure 5-4. Port and address translation

- Dynamically allocate an external host address each time an internal host initiates a connection, without modifying port numbers. This limits the number of internal hosts that can simultaneously access the Internet to the number of available external addresses.
- Create a fixed mapping from internal addresses to externally visible addresses, but use port mapping so that multiple internal machines use the same external addresses.
- Dynamically allocate an external host address and port pair each time an internal host initiates a connection. This makes the most efficient possible use of the external host addresses.

Advantages of Network Address Translation

The main purpose of network address translation is to economize on address space, but it can also have some security advantages.

Network address translation helps to enforce the firewall's control over outbound connections

Since individual hosts have addresses that won't work on the external network, they require the assistance of the network address translation system to connect. If a host finds a way to connect to the Internet without going through the address translation, the connection won't work.

Network address translation can help restrict incoming traffic

Depending on how you configure a network address translation system, it can provide stronger restrictions on incoming traffic than packet filtering. A network address translation system that's doing dynamic translation will allow only packets that are part of a current interaction initiated from the inside. This is similar to the protection that a dynamic packet filter offers, but the changing IP addresses put stronger time constraints on attackers. Not only can they attack only certain ports, but if they wait too long, the address translation will have gone away, and the entire address will have disappeared or been given to another host.

Many people assume that all network address translation systems provide this sort of protection, but this is not true. If you configure a network address translation system to do static translations, it may provide no restrictions at all on incoming traffic. Even doing dynamic translations, the simplest implementations allocate an entire externally visible address to the internal host and translate all traffic sent to that address. This does limit the time that an attacker has, but otherwise provides no protection at all.

Network address translation helps to conceal the internal network's configuration

The less an attacker knows about you, the better off you are. A network address translation system makes it much more difficult for an attacker to determine how many computers you have, what kind of machines they are, and how they're arranged on the network. Note, however, that many protocols leak useful information (for instance, they may include the client's IP address or hostname in places where the network address translation system doesn't need to change it). When we discuss the network address translation properties of protocols, we attempt to mention leaks of this sort.

Disadvantages of Network Address Translation

While network address translation is a very useful way of conserving network address space, it presents some problems.

Dynamic allocation requires state information that is not always available

It is very easy for a network address translation system to tell whether or not a host has stopped using a TCP connection, but there's no way to know at the packet header level whether a UDP packet is part of an ongoing conversation or is an isolated event. This means that a network address translation system has to guess how long it should keep a particular translation. If it guesses incorrectly, responses may be lost or delivered to unexpected hosts.

Embedded IP addresses are a problem for network address translation

Network address translation systems normally translate the addresses in the headers of packets (see Chapter 4, *Packets and Protocols*, for more information about packet layout). Some protocols also hide addresses in other places, and in order to find those addresses, the network address translator has to understand the protocol enough to find and modify the address, while preserving the validity of the packet. Most network address translation systems are capable of doing this for at least some protocols (for instance, FTP) but not for all protocols.

Network address translation interferes with some encryption and authentication systems

Systems for encrypting data often attempt to ensure the integrity of the data, so that the systems that are communicating know that packets have not been tampered with in transit. Network address translation is a form of tampering with the data in transit. If the protocol that's being translated does not protect the data that the network address translation system modifies, it will work. Otherwise, the integrity checking will be violated, and connections will fail. In most cases, protocols that do not have embedded IP addresses are compatible (the packet headers are not part of the protocol's protected data). The major exception to this rule is IPsec, which protects the entire packet, including headers. Network address translation is almost guaranteed to fail for protocols that combine embedded IP addresses with data integrity protection.

Dynamic allocation of addresses interferes with logging

If you are logging information after the network address translation happens, the logs will show the translated addresses, and you will have to correlate the logs with information from the network address translation system to figure out what internal system is actually involved. For instance, if you have a screened subnet architecture (discussed in Chapter 6, *Firewall Architectures*), and you are doing network address translation on the interior router, the translated addresses will be in logs from the exterior router or from a caching web proxy server on the screened subnet. Although log correlation is theoretically possible, it may be difficult, and clock synchronization will be critical.

Dynamic allocation of ports may interfere with packet filtering

Packet filtering systems pay attention to source and destination port numbers in order to try to figure out what protocol a packet should be using. Changing the source port may change the packet's acceptability. In most cases, this is not a problem because address translation systems are translating for clients, which are usually allowed to use any port above 1023. However, if ports above 1023 are translated to ports below 1023, traffic may be dropped.

Virtual Private Networks

A *virtual private network* (VPN) is a way of employing encryption and integrity protection so that you can use a public network (for instance, the Internet) as if it were a private network (a piece of cabling that you control). Making a private, high-speed, long-distance connection between two sites is much more expensive than connecting the same two sites to a public high-speed network, but it's also much more secure. A virtual private network is an attempt to combine the advantages of a public network (it's cheap and widely available) with some of the advantages of a private network (it's secure).

Fundamentally, all virtual private networks that run over the Internet employ the same principle: traffic is encrypted, integrity protected, and encapsulated into new packets, which are sent across the Internet to something that undoes the encapsulation, checks the integrity, and decrypts the traffic.

Virtual private networks are not exactly a firewall technology, but we discuss them here for several reasons:

- If you're using virtual private networking, you need to be careful about how it interacts with the firewall. In many cases, the firewall can't control traffic that comes in over the virtual network, which makes it a way to avoid the firewall controls and open new insecurities.
- A firewall is a convenient place to add virtual private networking features.
- We will frequently mention virtual private networking as a way to provide remote services that cannot be provided securely using other firewall techniques.

Where Do You Encrypt?

Virtual private networks depend on encryption. That encryption can be done as a transport method, where a host decides to encrypt traffic when it is generated, or as a tunnel, where traffic is encrypted and decrypted somewhere in between the source and the destination. The question of where you do the encryption and decryption relative to your packet filtering is an important one. If you do the encryption and decryption inside the packet filtering perimeter (i.e., on your internal net), then the filters just have to allow the encrypted packets in and out. This is especially easy if you're doing tunneling, because all the tunneled packets will be addressed to the same remote address and port number at the other end of the tunnel (the decryption unit). On the other hand, doing the encryption and decryption inside your filtering perimeter means that packets arriving encrypted are not subject to the scrutiny of the packet filters. This leaves you vulnerable to attack from the other site if that site has been compromised.

If you do the encryption and decryption outside the packet filtering perimeter (i.e., on your perimeter net or in your exterior router), then the packets coming in from the other site can be subjected to the full scrutiny of your packet filtering system. On the other hand, they can also be subjected to the full scrutiny of anyone who can read traffic on your perimeter net, including intruders.

Key Distribution and Certificates

As with any encryption and integrity protection system, key distribution can be a very sticky problem. A number of choices are available, including sharing keys or using a public key system; see Appendix C, *Cryptography*, for descriptions of these systems and the advantages and disadvantages of each.

Advantages of Virtual Private Networks

Most of the advantages of virtual private networks are economic; it's cheaper to use shared public networks than it is to set up dedicated connections, whether those are leased lines between sites or modem pools that allow individual machines to connect to a central site. On the other hand, virtual private networks also provide some security advantages.

Virtual private networks provide overall encryption

A virtual private network conceals all the traffic that goes over it. Not only does it guarantee that all the information is encrypted, but it also keeps people from knowing which internal machines are being used and with what protocols. You can protect information from snooping by using individual encrypted protocols, but attackers will still have some idea what machines are talking and what kind of information they're exchanging (for instance, if you use an encrypted mail protocol, they will know that things are being mailed). A virtual private network conceals more information.

Virtual private networks allow you to remotely use protocols that are difficult to secure any other way

Some protocols are extremely difficult to provide securely through a firewall. For instance, a number of protocols used on Microsoft systems are based on SMB, which provides a wide variety of services with different security implications over the same ports and connections. Packet filtering and proxying both have trouble adding security to SMB. Virtual private networking provides a way to give remote access for these protocols without letting people attack them from the Internet at large.

Disadvantages of Virtual Private Networks

Although virtual private networks are an important security tool, they also present problems in a firewall environment.

Virtual private networks involve dangerous network connections

A virtual private network runs over an actual network, which is presumably not a private network. The hosts on the virtual private network must be connected to that actual network, and if you're not careful, they will be vulnerable to attack from that network. For instance, if you use a virtual private network to provide connectivity to your internal network for mobile users who connect to the Internet, their machines may be attacked from the Internet.

Ideally, a virtual private network system will disable all other uses of the network interface. It's important to choose a system that will allow you to force this on the remote system. It's not good enough to have a system where the remote system is able to turn off other uses because the user on the remote system may turn networking back on. It's very tempting as a way to get rapid access to Internet resources.

Virtual private networks extend the network you must protect

When you attach something via a virtual private network, you are making it part of your internal network. If a machine on the virtual private network is broken into, the attacker will then be able to use the virtual private network to attack the rest of your site, from something that's treated as if it were inside of your local network. Virtual private networking is commonly used to give access to machines that are much more vulnerable than those that are physically on the network—for instance, laptops that are carried around in public, home machines that curious children have physical access to, and machines owned by other sites with interests and policies that are not identical to yours.

Even if the virtual private network disables other uses of the network interface it is running over, the machine may have other network interfaces. This can make it into a gateway between your network and others, inside your network's security perimeter.

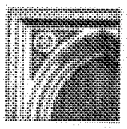
Because of this, you want to be careful how you attach the virtual private network to your real private network, and how you secure the remote end. It may not be appropriate to make the virtual private network a seamless part of your internal network. Consider putting in a subsidiary firewall or at least special intrusion detection to watch for problems.

Networking/Computer Security

O'REILLY®

LIBRARY OF CONGRESS
0 002 174 454 1
LIBRARY OF CONGRESS
0 006 968 551 1

Building Internet Firewalls



In the years since the publication of the first edition... onet use has exploded and e-commerce has become a daily part of business and personal life. As Internet usage has grown, so have the secu... fers to IP forgeries, to defacing web sites, to attacks tha... asing distributed denial of service.

Firewalls are a very effective way to protect your system from n... and they're a critical component of today's computer networks. ...ig *Internet Firewalls*, this second edition is a practical and detailed guide to designing and building firewalls and to configuring Internet services to work with firewalls. This much expanded edition covers Linux and Windows NT, as well as Unix. It describes a variety of firewall technologies (packet filtering, proxying, network address translation, virtual private networks) and architectures, and more than one hundred Internet services from email and file transfer, to web services and scripting languages, to naming, authentication, and database protocols.

The book also contains chapters on security policies, cryptography, maintaining firewalls, and responding to security incidents, as well as summaries of helpful firewall tools and other resources.

"In 1991, I described everything you (then) needed to know about firewalls in 11 pages of Practical Unix Security. *Times have definitely changed . . . and there is a lot more to know about firewalls. This book tells it all, in encyclopedic detail, and explains step-by-step where and how to apply that information. If you're computing on a network, you can't afford to be without this book."*

—**Professor Gene Spafford**
Purdue University CERIAS (Center for Education and Research in Information Assurance and Security)

"It is always easier to design a secure network properly than it is to correct deficiencies once it has been implemented. Building Internet Firewalls *lays out the fundamental principles and the detailed information you need to design a network correctly. This book is an essential reference for anyone involved in operational security."*

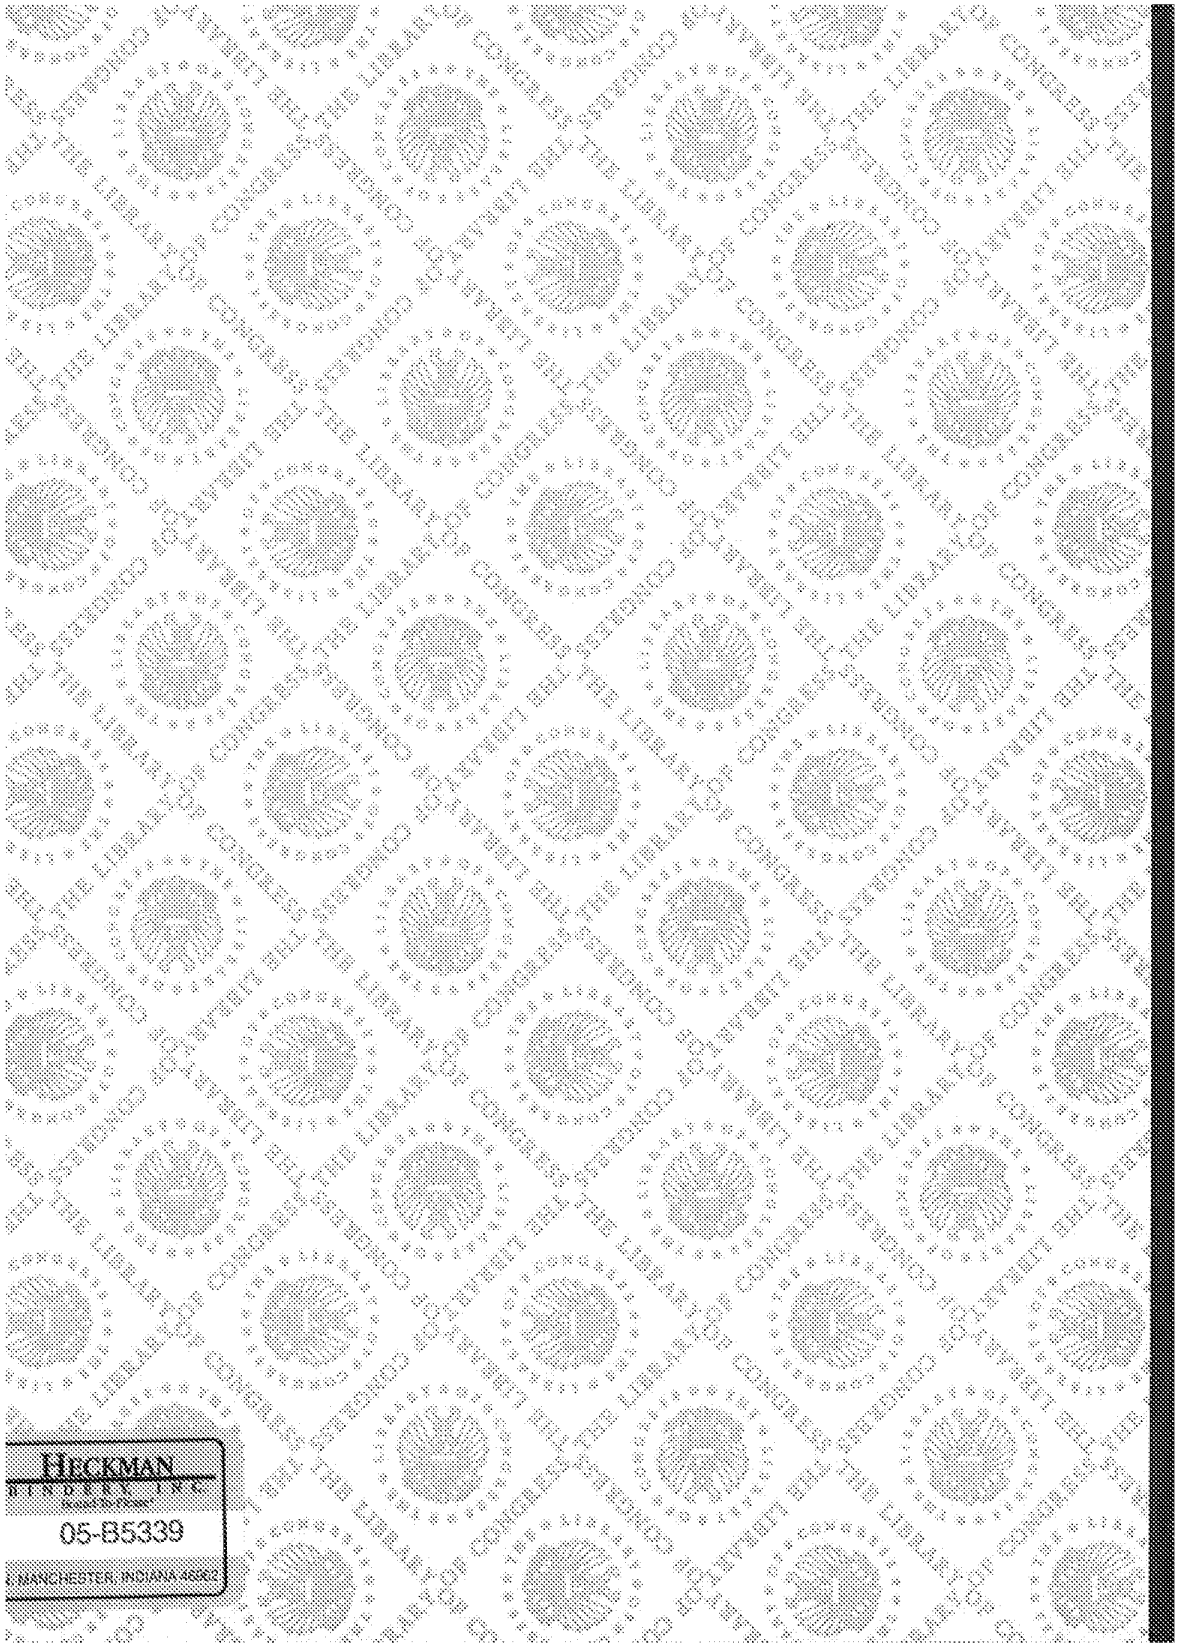
—David LeBlanc
Senior Technologist, ITG Information Security, Microsoft Corporation

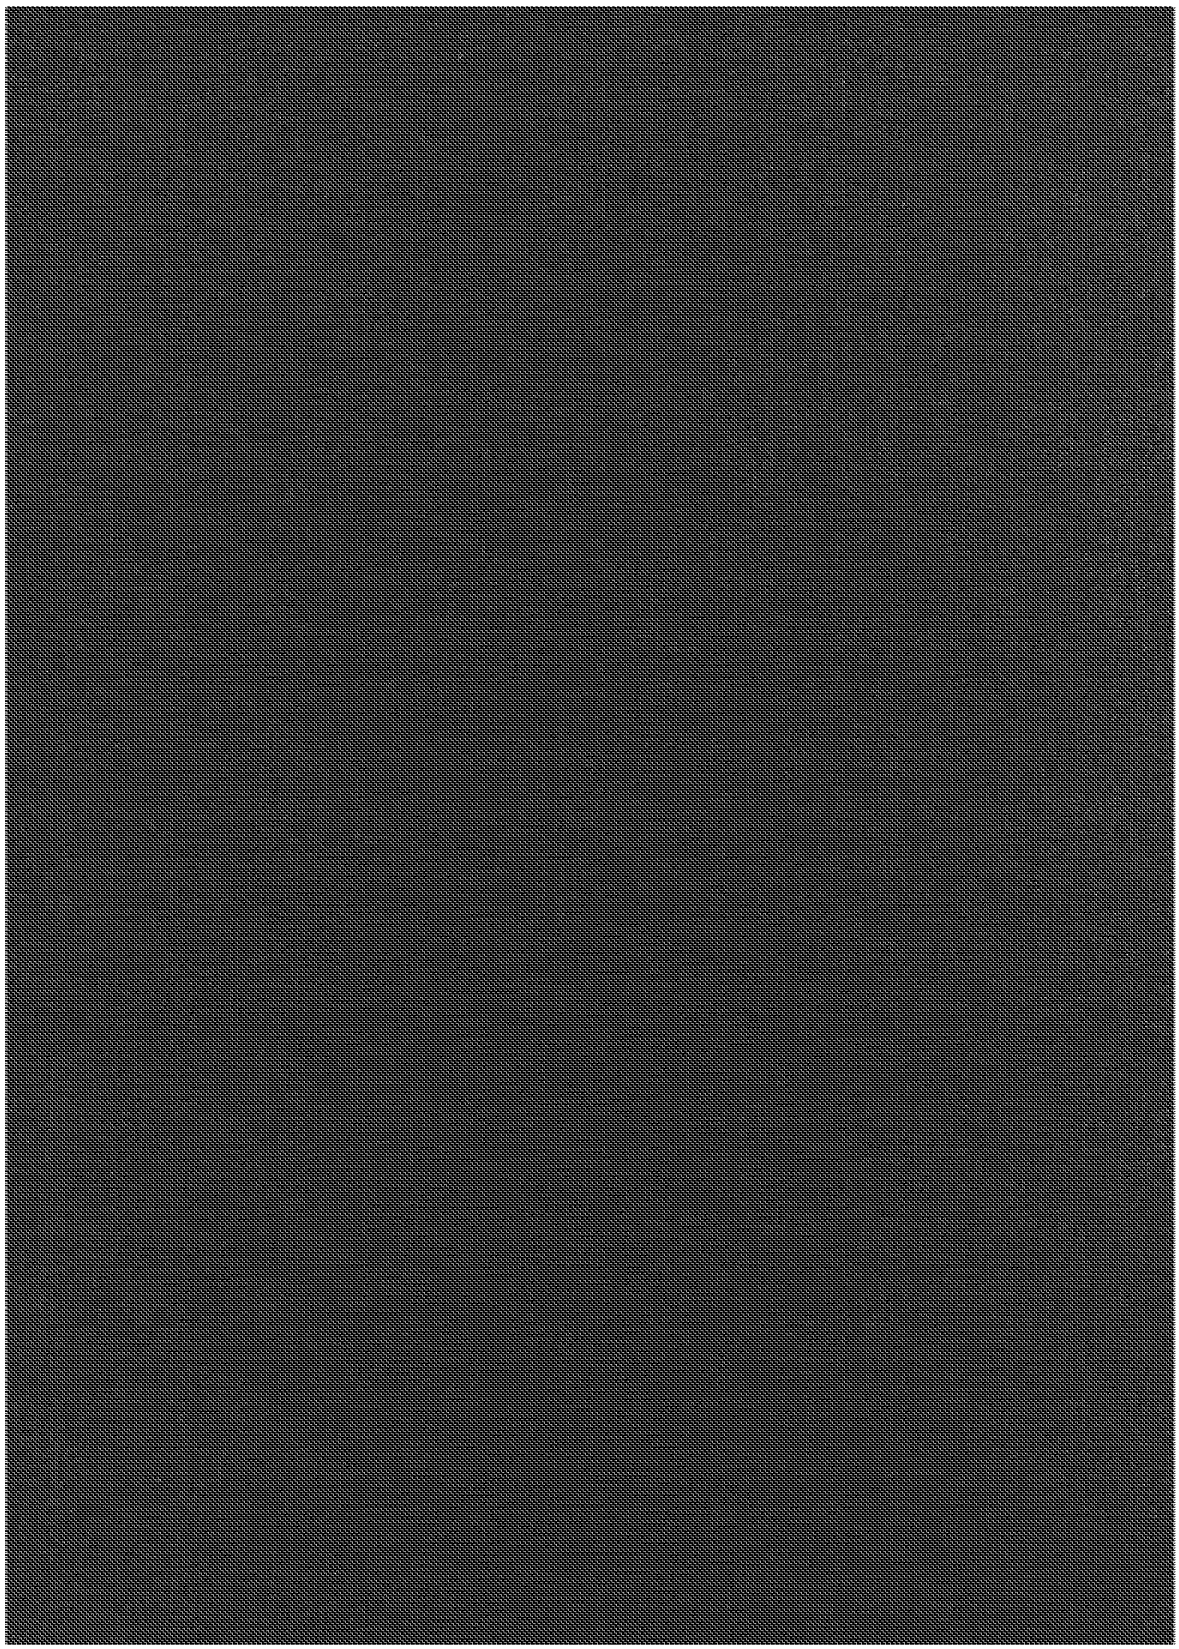
"So much more than a book about firewalls! There's great information about management, threats, security in general, design for defense, policy, detection, reaction, and more—as well as all the detail needed to effectively implement and maintain a best-practice firewall."

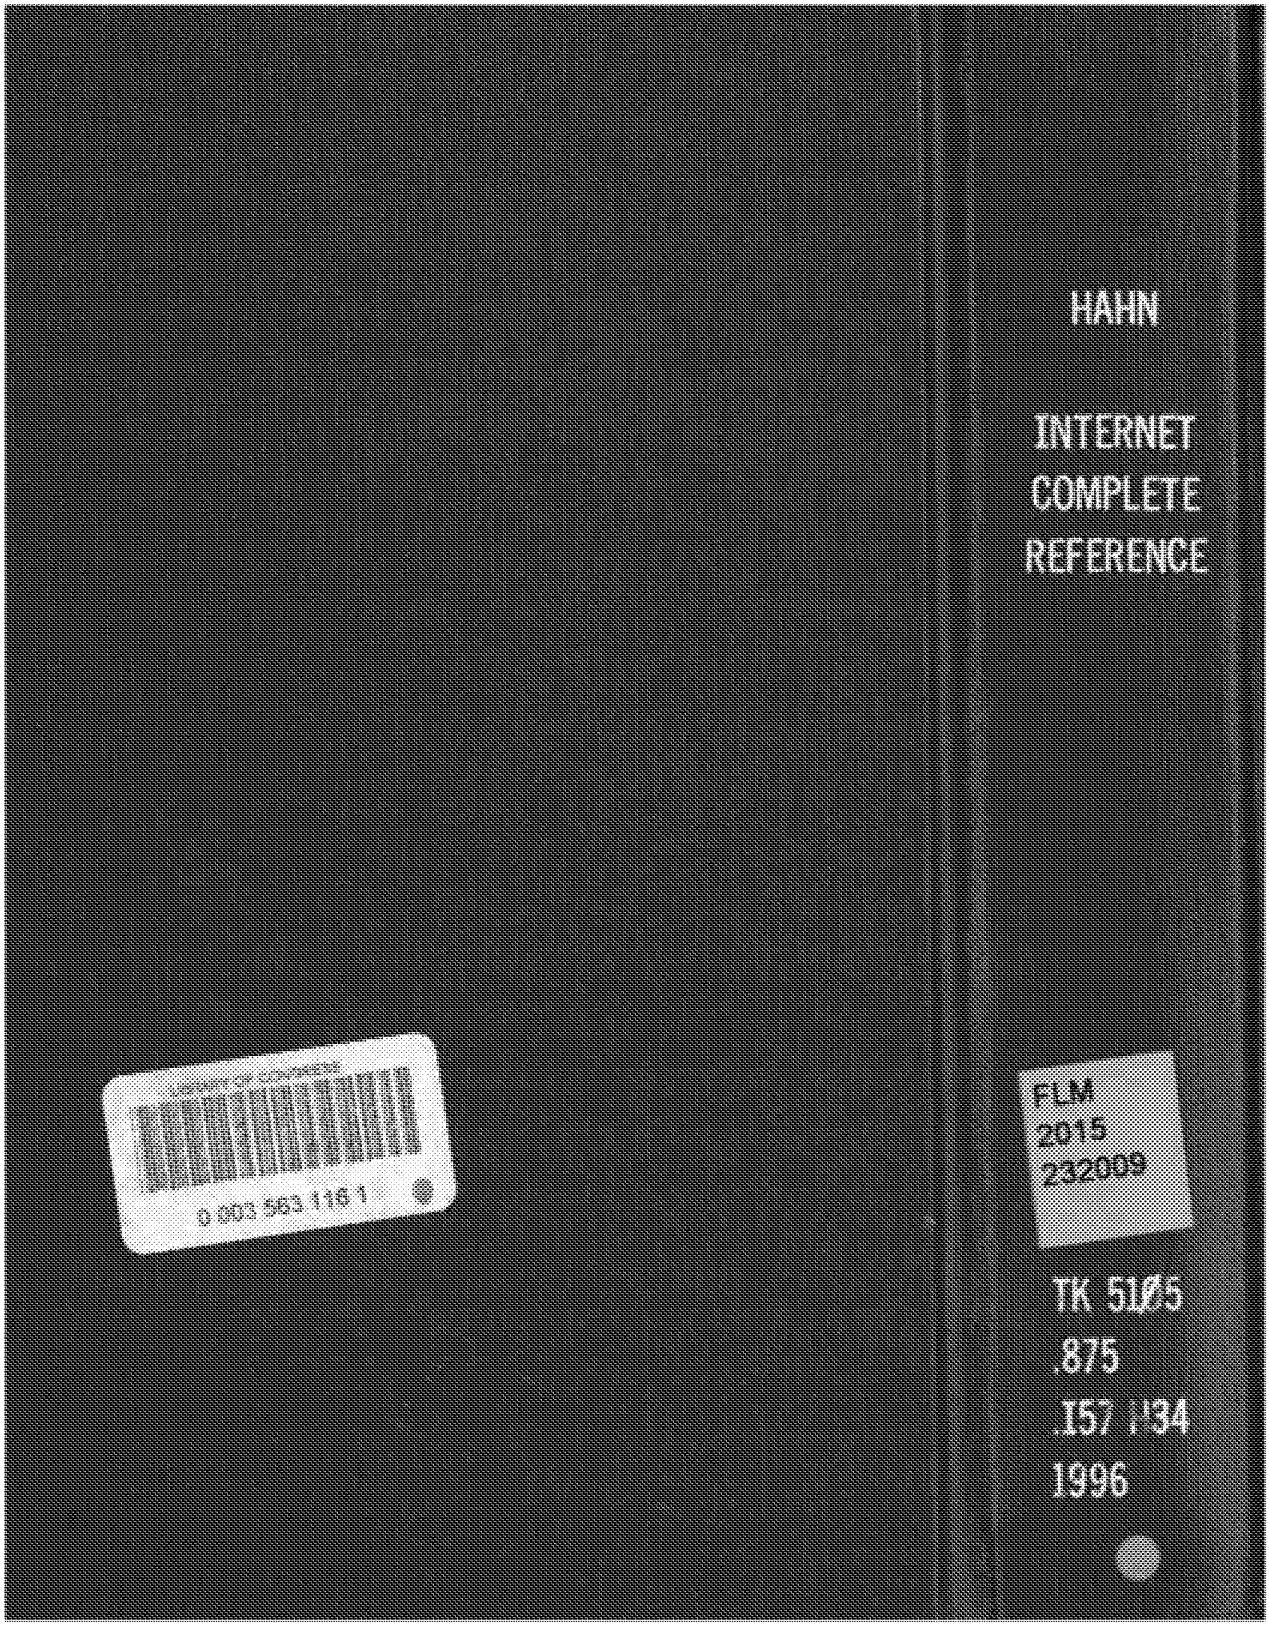
—Greg Rose
Principal Engineer, QUALCOMM Incorporated

ISBN 1-56592-871-7
US $44.95
CAN $65.95
90000
9 781565 928718
6 36920 92871 3

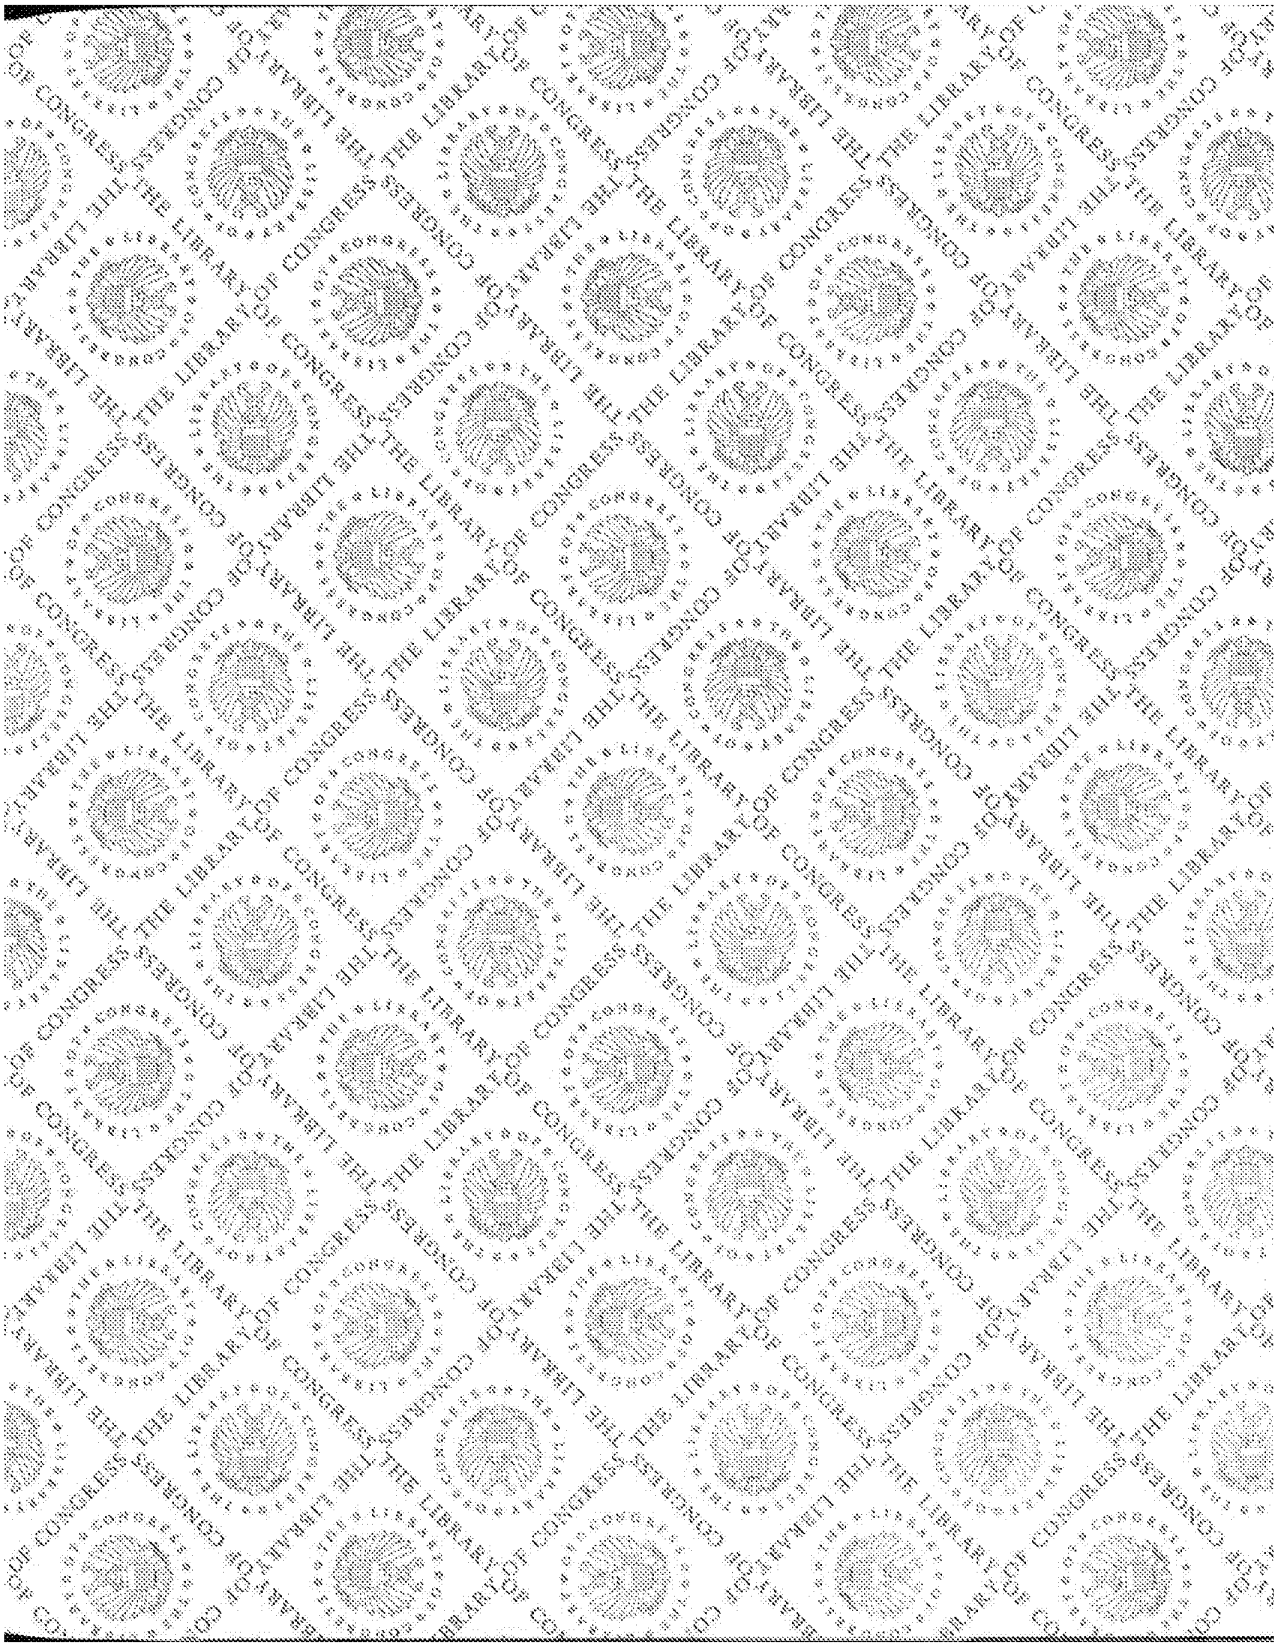
Visit O'Reilly on the Web at *www.oreilly.com*

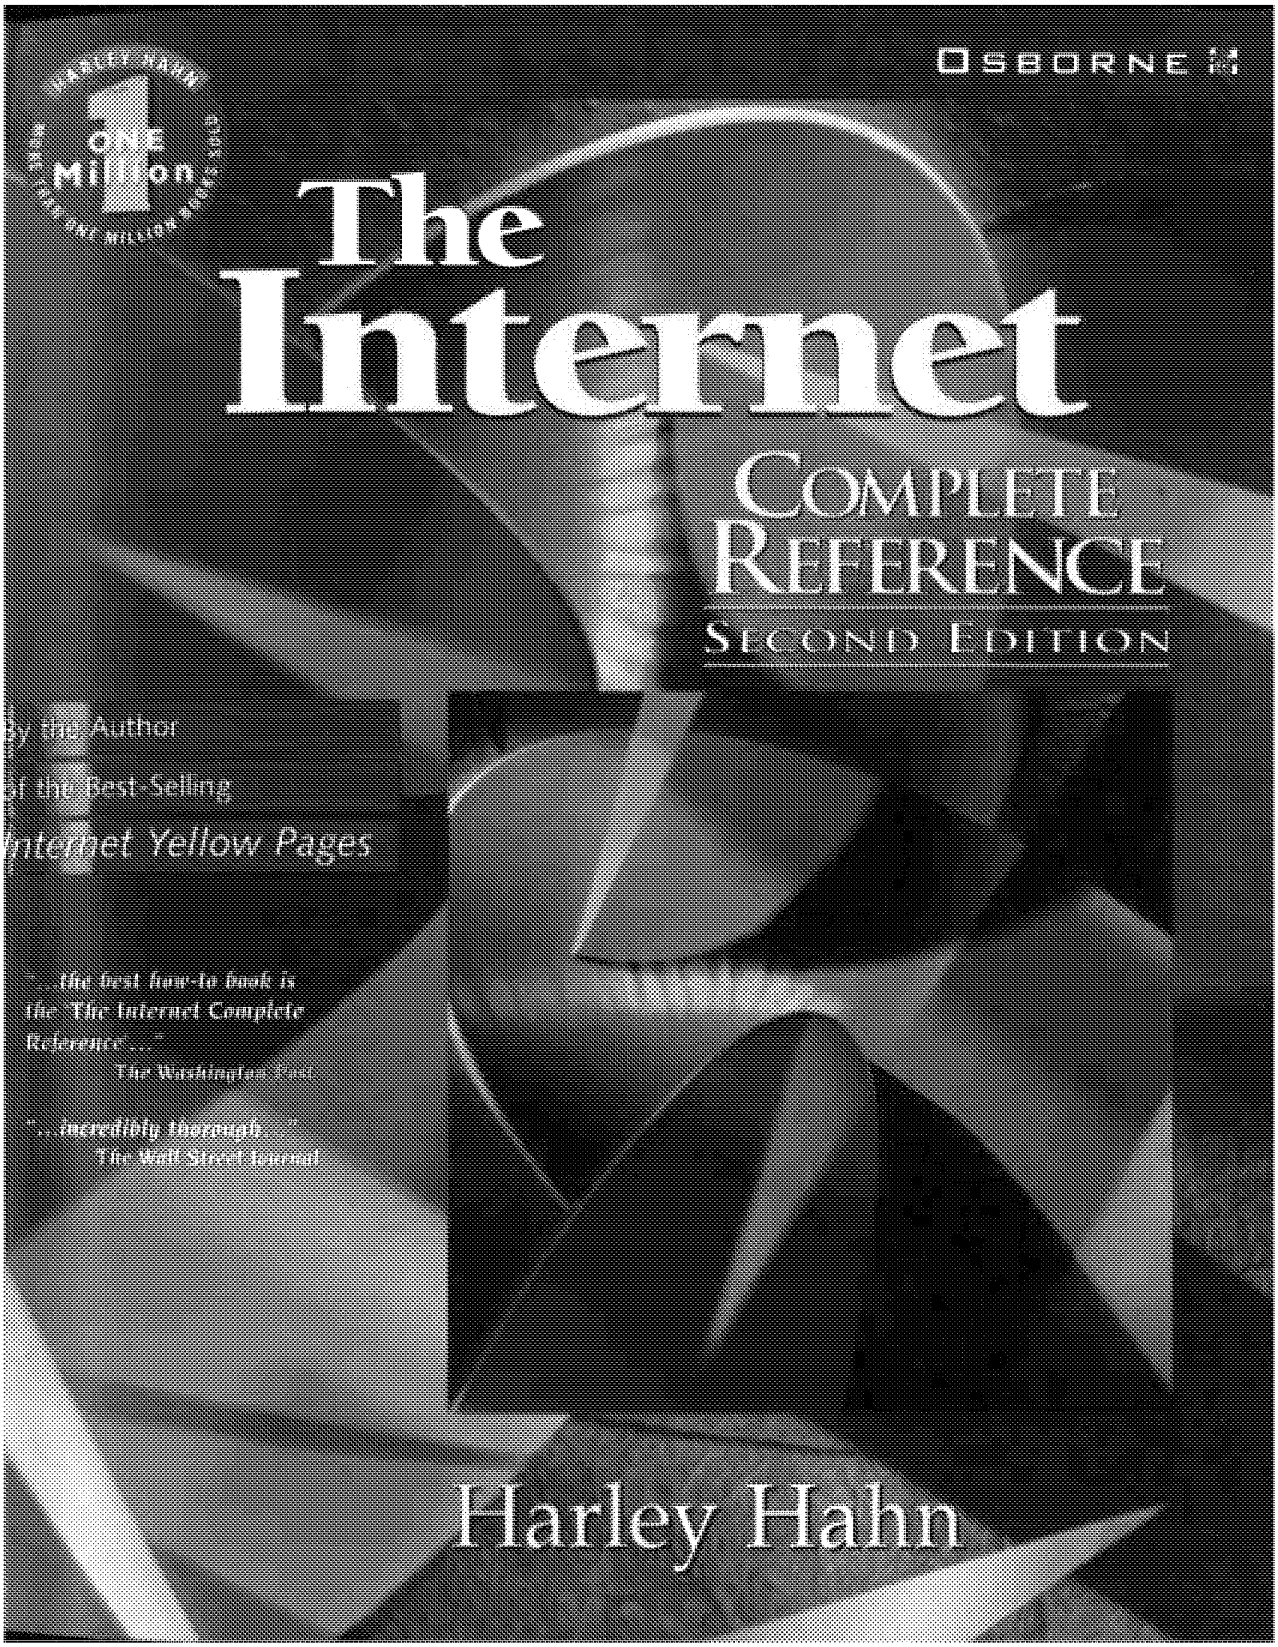

HECKMAN
BINDERY INC.
05-B5339
MANCHESTER, INDIANA 46962





Meta's Exhibit No. 1008
Page 001


THE LIBRARY OF CONGRESS



Meta's Exhibit No. 1008
Page 003

A PERSONAL NOTE FROM HARLEY HAHN

The Internet is, by far, the greatest and most significant achievement in the history of mankind. What? Am I saying that the Internet is more impressive than the pyramids? More beautiful than Michelangelo's David? More important to mankind than the wondrous inventions of the industrial revolution?

Yes, yes and yes.

Do I expect you to believe this? Of course not—not right now anyway. However, for years now, people have been connecting computers into networks, and now we have the Internet: a worldwide network connecting millions of computers and millions of people. What is amazing is that, within a few short years, the Internet has changed our civilization permanently and has introduced us to two completely unexpected ideas.

First, hundreds of thousands of people have been laboring to build the Internet. They have worked alone, in small groups, and within organizations, like so many ants in a global anthill. Many of these people are only doing their jobs but, just like the ants, they serve the common good while having no conception of the order and of the compelling forces that drive their work.

Deep inside us, there is an irresistible genetic program to follow. There is a voice that we hear only subconsciously and only as a species. A voice that commands us to take our computers, connect them into networks and... communicate.

This leads to the second great idea: When we connect computers we invariably create something that is well beyond the sum of its parts. Although there are millions of computers on the Internet, the computers are important only in that we use them to follow our instructions and to store our data. What is more important is something we cannot yet understand: the answer to the question, what happens when millions of people gather in a safe place to talk and to share? That is what the Internet is about, and that is what we are just beginning to find out.

I sense we are near the beginning of a great and important change in human affairs. There is a reason we built the Internet and, as human beings, we have an obligation to learn how to use it and to participate.

Now, about this book.

Before we start—you and I—I want to make sure you understand an important point: the Internet is easy to use, but is difficult to use well. I have filled this book with technical details, important details you must understand and master in order to use and enjoy the Internet. Still, don't for a moment think that you need to be a computer expert. The Internet is not for nerds, but just as surely it is not for dummies: it is for those people who are willing to think and to learn and to contribute.

If you ever have one of those days where everything seems to go wrong, and the computer just won't cooperate, and you can't understand what is happening, take a moment and try to see things in a larger sense. As a human being, you share the birthright of intelligence, curiosity and, above all, the ability to learn. However, these gifts are not free. By your very nature, you can not only learn, you are compelled to learn if you are to remain happy and fulfilled. Still, as the saying goes, there is no royal road to knowledge.

To use the Internet, you will have to expend some time and some effort (actually, a great deal of time and effort). But if you do, here is what I promise you:

- I promise I will stay with you.
- I promise I will guide you.
- I promise I will teach you the new words and the new ideas and the new skills you need to use the Internet well.

I have taken a great deal of time to organize this book so as to lead you from one topic to another in a way that will make sense to you (at least, in retrospect). However, you do not need to read everything from start to finish (although I think you should). Actually, this book comes with only two simple instructions:

1. Start by reading chapter 1.
2. When in doubt, do something enjoyable that requires effort.

I know what it is like to be faced with a large, inconceivable mass of unknowable information. What you are holding in your hand is much more than a computer book. It is a connection—a real connection—between you and me. Your job is to put in the effort and the time. My job is to open the doors.

I wrote this book to be your companion and, I promise you, it will be a good one. Take a few moments and skim through the pages. This is not your ordinary computer book, not by a long shot. There is a long journey ahead of you, and I will be with you every step of the way.

Buy this book. I am on your side.

Harley Hahn

Harley Hahn's

The Internet Complete Reference, Second Edition

Harley Hahn

Osborne McGraw-Hill
Berkeley New York St. Louis San Francisco
Auckland Bogotá Hamburg London Madrid
Mexico City Milan Montreal New Delhi Panama City
Paris São Paulo Singapore Sydney
Tokyo Toronto

Meta's Exhibit No. 1008
Page 006

Osborne McGraw-Hill
2600 Tenth Street
Berkeley, California 94710
U.S.A.

For information on translations or book distributors outside the U.S.A., or to arrange bulk purchase discounts for sales promotions, premiums, or fundraisers, please contact Osborne **McGraw-Hill** at the above address.

The Internet Complete Reference, Second Edition

Copyright © 1996 by Harley Hahn. All rights reserved. Printed in the United States of America. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher, with the exception that the program listings may be entered, stored, and executed in a computer system, but they may not be reproduced for publication.

1234567890 DOC 99876

ISBN 0-07-882138-X

TK5105
.875
I57H34
1996

Executive Editor
Scott Rogers

Editorial Assistant
Daniela Dell'Orco

Project Editor
Janet Walden

Copy Editor
Lunaea Hougland

Proofreader
Pat Mannion

Indexer
Wendy Murdock

Computer Designer
Richard Whitaker

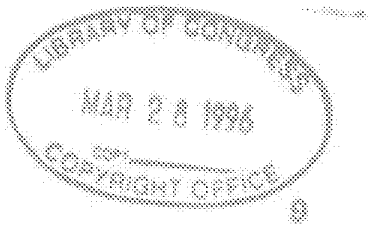
Quality Control Specialist
Joe Scuderi

Illustrator
Marla J. Shelasky

Cover Design
Compass Marketing

Information has been obtained by Osborne **McGraw-Hill** from sources believed to be reliable. However, because of the possibility of human or mechanical error by our sources, Osborne McGraw-Hill, or others, Osborne **McGraw-Hill** does not guarantee the accuracy, adequacy, or completeness of any information and is not responsible for any errors or omissions or the results obtained from use of such information.

96-14849



About the Author

Harley Hahn is an internationally recognized writer, analyst and consultant. He is the author of fourteen books, including the best-selling *The Internet Complete Reference* and *The Internet Yellow Pages*, and the highly regarded *Assembler Inside & Out*. His Unix books include *The Unix Companion*, *A Student's Guide to Unix*, *Peter Norton's Guide to Unix* and *Mastering Xenix on the IBM PC AT*.

Hahn has a degree in Mathematics and Computer Science from the University of Waterloo, Canada; and a graduate degree in Computer Science from the University of California at San Diego. Before becoming a professional writer, he studied medicine at the University of Toronto Medical School.

Hahn enjoys writing computer books, because "I get to sleep in, and I like telling people what to do."

Hahn does not live in a converted farmhouse in Connecticut with his wife, three children, and a Labrador Retriever named Rolf. Nor does he commute frequently to New York.

His favorite pajamas are green.

To Wendy

Contents at a Glance

1	Introduction	1
2	Understanding the Internet	9
3	A Tour of the Internet	23
4	Hardware Requirements to Connect to the Internet	33
5	Software Requirements and Internet Service Providers	49
6	Internet Addressing	75
7	Mail	93
8	Using Mail from a Shell Account: `pine` and `pico`	125
9	Understanding the Web	163
10	Using the Web	173
11	The Web: Advanced Topics	201
12	Using the Web from a Shell Account: `lynx`	229
13	Introduction to Usenet	249
14	Reading and Posting Usenet Articles	271

15 Using Usenet From a Shell Account: `tin` 297
16 Gopher, Veronica and Jughead 315
17 Using Gopher from a Shell Account: `gopher` 333
18 Anonymous Ftp 349
19 Using Anonymous Ftp from a Shell Account: `ftp` . . . 367
20 Archie 387
21 File Types Used on the Internet; Downloading Software . 419
22 Mailing Lists 451
23 Telnet: Logging In to a Remote Computer 469
24 Using Telnet from a Shell Account: `telnet` 477
25 Talk Facilities 489
26 Using Talk from a Shell Account: `talk`, `ytalk` 505
27 IRC (Internet Relay Chat) 521
28 Muds 549
A List of Top-Level Internet Domains 573
Glossary 579
Harley Hahn's Catalog of Internet Resources 603
Index 787

Contents

List of Figures xxiii
Acknowledgments xxvii

1 Introduction 1
What is the Internet? 2
Using the Internet 5
How to Use This Book 5
Do You Need to Know Unix? 7
Should You Learn Unix? 7

2 Understanding the Internet 9
Our Friend the Network 10
The Secret of the Internet: Clients and Servers 12
X Window and X Clients 14
Hosts and Terminals 18
What Is TCP/IP? 20

3 A Tour of the Internet 23
Mail 24

The Web 24
Web Search Engines 25
Usenet 25
Gopher 26
Veronica and Jughead 27
Anonymous ftp 27
Archie 28
Mailing Lists 29
Telnet (Remote Connection) 29
Talk Facilities 30
Internet Relay Chat 30
Muds and Other Imaginary Places 31

4 Hardware Requirements to Connect to the Internet 33
What Do You Need to Use the Internet? 34
Accessing the Internet via Your Local Network 34
Accessing the Internet via the Telephone System 35
Choosing a Computer 37
Choosing a Modem 38
ISDN 41
How ISDN Is Organized 42
What Do You Need in Order to Use ISDN? 44

5 Software Requirements and Internet Service Providers 49
Hosts and Terminals 50
Unix Hosts, VT-100 Terminals, and Telnet 51
Shell Accounts and PPP Accounts 52
The Nature of a PPP Account 54
The Nature of a Shell Account 55
Which Type of Account Is Right for You? 57
The True Story of the Three Brothers 57
PPP and SLIP 58
Internet Service Providers 60
Online Information Services 61
Choosing an Internet Service Provider 63
What Software Do You Need? 66
TCP/IP and PPP Software 68
Installing and Configuring TCP and PPP 69
Where to Get Internet Client Programs 72

6 Internet Addressing 75
Standard Internet Addresses 76
Understanding a Domain Name: Sub-Domains . . . 77
Variations on the Standard Internet Address Format 78
Top-Level Domains 81
Pseudo-Internet Addresses 83
IP Addresses and DNS 84
UUCP 86
UUCP Addresses and Bang Paths 87
Simplified UUCP Addressing 88
Sending Mail to Other Networks 89

7 Mail 93
SMTP: The Basis for the Internet Mail System 94
How You Access the Mail System 95
Which Mail Program Should You Use? 96
An Example of Moving Mail Around the Net 98
Understanding Mail Headers 99
Signatures 103
Mail Addresses 104
Sending Mail 106
Sending Copies of a Message 109
Reading Mail 111
Hints for Managing Your Mail 112
Replying to a Message 113
Forwarding and Bouncing Mail 114
How Mail Is Stored (Folders) 116
The Difference Between Text and Binary Data . . . 117
Using Mime to Mail Binary Data as an Attachment 119
Hints for Practicing Safe Mail 120
Smileys 122

8 Using Mail from a Shell Account: `pine` and `pico` 125
A Quick History of `pine` and `pico` 126
Using `pico` by Itself 128
Orientation to `pine` 128
Starting `pine` to Send a Message 130
Composing a Message: The Screen 131
Composing a Message: General Commands 133

Composing a Message: Moving the Cursor 135
Composing a Message: Editing Commands 136
Composing a Message: Advanced Commands 140
Sending Your Message 142
Editing the Header 143
Starting `pine` to Read Your Mail 145
The Main Menu 147
The Index 150
Selecting a Message to Read 151
Reading a Message 152
Deleting a Message 155
Replying to a Message 155
Forwarding and Bouncing a Message 156
Saving, Printing, and Piping a Message 157
Using the Address Book 158
Customizing `pine` 159

9 Understanding the Web 163
Orientation to the Web 164
How Did the Web Start? 165
How the World Wide Web Turned into "The Web" 166
What the Web Really Is 167
The Web Is the Fourth Attempt 167
Why the Web Grew 169

10 Using the Web 173
What Is It Like to Use the Web? 174
Links, URLs and Hyperspace 177
Words to Avoid 180
Web Pages and Home Pages 181
Telling Your Browser Which URL to Use 183
Reading a Web Page 184
Links, Forms and Image Maps 185
The History List and the Bookmark List 188
Images 189
Sounds and Video 190
Updating Information Automatically 192
Web Directories and Search Engines 193
Customizing Your Working Environment 196
Distributed Applications: Java 198

11 The Web: Advanced Topics 201
URLs and Schemes 202
URLs and Host Names 203
URLs and Port Numbers 204
Pathnames 205
Using Your Browser to Access Anonymous Ftp . . . 208
Using Your Browser to Access a Gopher 209
Using Your Browser to Read Usenet Articles 210
Using Your Browser to Send Mail Messages 211
Using Your Browser to Initiate a Telnet Session . . . 212
Hint for Using URLs 213
Why Hypertext Is Important 216
Hypertext and HTML 217
An Example of HTML 219
The Nature of HTML 222
Making Your Own Home Page 223
Putting Your Home Page Onto the Net 226
Default File and Directory Names 227
Unix File Permissions 228

12 Using the Web from a Shell Account: `lynx` 229
Why `lynx` Is So Cool 230
A Quick History of `lynx` 231
Starting 231
Stopping 232
Displaying Help 233
Generally Useful Keys 233
How `lynx` Displays a Page 234
Working with URLs 235
Navigating 237
Using Your Numeric Keypad to Navigate 239
Special Keys for `vi` and `emacs` Users 241
The History List 242
The Bookmark List 242
Using Forms 243
Saving, Mailing and Printing Files 243
Accessing Anonymous Ftp and Gopher Sites 245
Reading Usenet Newsgroups 246
Sending Mail 247
Customizing `lynx`: Options 247

13 Introduction to Usenet 249
What Is Usenet? 250
Basic Terminology 251
Newsreaders 251
News Clients and Servers 252
What Is the Difference Between Usenet and the Internet? 254
Who Runs Usenet? 255
How Is the News Transported? 256
How Big Is Usenet? 257
The Mainstream and Alternative Hierarchies 258
The Cultural, Organizational and Regional Hierarchies 261
How Many Different Newsgroups Are There? 262
Newsgroup Naming Conventions 262
Moderated Newsgroups 265
What Are the Most Popular Newsgroups? 266
How Are New Newsgroups Started? 267
Frequently Asked Question Lists (FAQs) 270

14 Reading and Posting Usenet Articles 271
The Format of a News Article 272
The Header 272
The Signature 279
Followup Articles 281
Commonly Used Acronyms 284
Usenet Slang 284
Foo, Bar and Foobar 287
Rtfm 287
Smileys 289
How Does Your News Server Keep Track of All the Articles? 291
The Format of Your `newsrc` File 292
Posting Your Own Articles 294
Flames and Good Behavior 294
The Final Word 295

15 Using Usenet From a Shell Account: `tin` 297
A Quick History of `tin` 298
Overview of the `tin` Newsreader 298

Customizing Your Working Environment 300
Starting and Stopping 300
The Newsgroup Selection List 303
Selecting a Newsgroup 304
The Thread Selection List 305
Selecting a Thread to Read 306
Reading Articles 306
Posting an Article 308
Commands That Are Always Available 310
Commands to Use While Selecting a Newsgroup 311
Commands to Use While Selecting a Thread to Read 312
Commands to Use While Reading an Article 313

16 Gopher, Veronica and Jughead 315
What Is the Gopher? 316
A Quick History of the Gopher 317
Which Gopher Client Should You Use? 318
Traveling Through Gopherspace 320
Gopher Addresses 321
History Lists and Bookmark Lists 322
Different Types of Resources 323
Saving, Printing, Mailing and Downloading 325
What Is Veronica? 325
Using Veronica 326
Specifying Words for a Veronica Search 328
Using Jughead to Search Gopherspace 330

17 Using Gopher from a Shell Account: **gopher** 333
Starting the **gopher** Program 334
Stopping, Displaying Help, Shell Commands, Redrawing the Screen 335
Understanding a **gopher** Menu 335
Using Menus 340
Advanced Menu Commands 342
Bookmarks 344
Saving, Mailing and Downloading 345
Summary of **gopher** Commands 346

18 Anonymous Ftp 349
What Is Ftp? 350
Anonymous Ftp 351
The Importance of Anonymous Ftp 352
Using a Browser for Anonymous Ftp 354
Using an Ftp Client for Anonymous Ftp 355
Files 358
Directories 360
Finding Your Way Around a New Computer 362
Requesting Anonymous Ftp Services by Mail 362

19 Using Anonymous Ftp from a Shell Account: `ftp` 367
Starting the `ftp` Program 368
A Second Way to Start the `ftp` Program 372
An Overview of the `ftp` Commands 374
The Basic `ftp` Commands 375
The Connection `ftp` Commands 377
The Directory `ftp` Commands 377
The File Transfer `ftp` Commands 381
Setting `ftp` Options 383
A Typical Anonymous Ftp Session 385

20 Archie 387
What Is Archie? 388
How Does Archie Work? 389
Using an Archie Client 390
Specifying a Search Pattern 392
Using an Archie Client From a Shell Account: `archie` 394
Using an Archie Server Interactively 398
Setting Variables With an Archie Server 399
Setting the `search` Variable 404
Performing a Search 405
Mailing the Results of an Archie Search 405
Commands to Use While Reading Archie Server Output 406
Using the `whatis` Database 407
Displaying Help Information 408
A Summary of Archie Server Commands 410
A Typical Session With Archie 412
Using Archie by Mail 416

21 File Types Used on the Internet; Downloading Software 419
The Five Important Ideas Regarding File Types . . 420
File Types 420
File Extensions for Common File Types 421
Using a Sound, Picture or Video File 421
Using a Document or Postscript File 424
Text Files and Binary Files 425
Compressed Files and Archives 426
Using Mime to Encode Binary Files for Mailing . . 427
Using Uuencode to Encode Binary Files for Usenet 428
Downloading Software for Windows and DOS . . . 430
Important File Types for Windows and DOS Users 434
Background Information About Macintosh Files 434
Downloading Software for a Macintosh 436
Important File Types for Macintosh Users 439
Working With Unix Compressed Files (`.Z` and `.gz`) 440
Working With Unix Archives (Tarfiles) 442
Compressed Tarfiles 446
Uploading and Downloading Files with a Shell Account 447

22 Mailing Lists 451
What Is a Mailing List? 452
How Are Mailing Lists Different From Usenet Newsgroups? 452
Moderated and Unmoderated Mailing Lists 454
How Mailing Lists Are Administered 454
Bitnet 455
Bitnet Mailing Lists and Usenet 456
Sending Mail to Bitnet 457
Finding Mailing Lists 458
Listserv, Listproc and Majordomo 460
The Two Types of Mailing List Addresses 460
Listserv: Subscribing and Unsubscribing 462
Listproc: Subscribing and Unsubscribing 464
Majordomo: Subscribing and Unsubscribing 466
Manually Administered Mailing Lists 466

23 Telnet: Logging In to a Remote Computer 469
What Is Telnet? 470
Basic Concepts 472
Using Your Telnet Client 472
Connecting to a Specific Port Number 474
More About Port Numbers 475

24 Using Telnet from a Shell Account: `telnet` 477
Starting the `telnet` Program 478
A Second Way to Start `telnet` 480
The `telnet` Escape Character: CTRL-] 482
Summary of Starting and Stopping `telnet` 482
Using `telnet` Commands 483
Job Control 484

25 Talk Facilities 489
Overview of Talk Facilities 490
Chat Rooms 492
Public Events 492
Privacy 492
White Pages Directory 493
A Personal Address Book 493
Embedded HTML 493
Sounds 493
Pictures 493
Video 494
File Transfer 494
Answering Machine 494
Remote Control 494
How Does It All Work? 494
What Do You Need to Use a Talk Facility? 495
Types of Internet Talk Facilities 496
The Unix `talk` Program 496
Internet Relay Chat (IRC) 496
Muds 496
Web-Based Chat Rooms 497
Bulletin Board Systems (BBSs) 497
Talkers 498
Direct Talk Connections 498
Telephone-Like Systems 499
Video Conferencing 499

Web-Based Talk Facilities 500
Guidelines for Text-based Talking 502

26 Using Talk from a Shell Account: `talk`, `ytalk` 505
Introducing the `talk` Program 506
The `talk` Screen 508
Problems with Addresses and What to Do 510
The "Checking for Invitation" Problem and What to Do 512
Refusing to Talk 513
The `ytalk` Program 515

27 IRC (Internet Relay Chat) 521
What Is IRC? 522
IRC Clients 524
IRC Clients for a Shell Account 525
Channels 527
Nicknames 528
Watching a Conversation 529
IRC Commands 531
Basic Commands 533
Help Information 537
Learning More About IRC 538
Using a Nickname 538
Commands for Talking to Specific People 539
Displaying Information About People 542
Using Direct Connections with DCC 543
Bots 545
Summary of Important IRC Commands 546
Harley Hahn's Rules for Safe IRC 546

28 Muds 549
What Is a Mud? 550
Mud Clients 554
Mud Clients for Shell Accounts 555
Social Muds and Adventure Muds 555
LPMuds, DikuMuds and TinyMuds 557
How to Choose a Mud 558
Getting Started on a Mud 560
Talking 565
Illusion and Reality 568
Getting Along with People 569

Moving Around 569
Player Killing 569
Mud Sex 570

A List of Top-Level Internet Domains 573

Glossary 579

Harley Hahn's Catalog of Internet Resources 603

Index 787

List of Figures

2-1: A local area network 11
2-2: A wide area network 12
2-3: An X Window display 17
2-4: A time-sharing system 19
3-1: Summary of internet resources 32
6-1: Organizational top-level domains 81
6-2: Examples of geographical top-level domains 82
7-1: An example of sending mail over the Internet 98
7-2: A typical mail message, showing the full header 100
7-3: Summary of time zone names 101
7-4: Composing a message with Eudora 107
7-5: A summary of incoming mail with Eudora 111
7-6: Replying to a message with Eudora 114
8-1: Composing a message with **pine** 132
8-2: **pine**: Composing a message, general commands 134
8-3: **pine**: Composing a message, moving the cursor 137
8-4: **pine**: Composing a message, editing commands 140
8-5: **pine**: Composing a message, advanced editing commands 142
8-6: **pine**: Commands to send a message 142

Meta's Exhibit No. 1008
Page 0024

8-7: `pine`: Commands to edit the header . . . 145
8-8: The `pine` main menu . . . 147
8-9: A typical `pine` index . . . 150
8-10: `pine`: Commands for using the index . . . 153
8-11: `pine`: Commands for reading a message . . . 154
8-12: `pine`: Commands for processing a message . . . 158
8-13: `pine`: commands for using your address book . . . 159
9-1: A typical gopher menu . . . 170
9-2: A typical web screen . . . 170
10-1: A typical web page . . . 175
10-2: A personal home page . . . 182
10-3: A web page containing a form . . . 186
10-4: A web page containing an image map . . . 187
10-5: Common file formats for images, sound and video . . . 191
11-1: List of schemes used within URLs . . . 203
11-2: A simple HTML file as interpreted by a browser . . . 220
11-3: A hypertext file being edited with Kenn Nesbitt's WebEdit . . . 224
12-1: `lynx`: A typical screen . . . 234
12-2: `lynx`: Navigating keys . . . 238
12-3: `lynx`: A screen with link numbers turned on . . . 240
12-4: `lynx`: An example using forms . . . 244
12-5: `lynx`: The options menu . . . 248
13-1: A graphical newsreader client . . . 252
13-2: The most important Usenet newsgroup hierarchies . . . 259
13-3: The mainstream Usenet newsgroup hierarchies . . . 260
13-4: The alternative Usenet newsgroup hierarchies . . . 261
13-5: Newsgroups devoted to jokes . . . 263
13-6: A selection of newsgroups related to sex . . . 264
13-7: The 25 most popular Usenet newsgroups . . . 268
14-1: A typical Usenet article . . . 273
14-2: Header lines for Usenet articles . . . 274
14-3: A typical Usenet followup article . . . 282
14-4: Common acronyms used on the Net . . . 285
14-5: A selection of smileys . . . 290
14-6: Usenet newsgroups for complaints . . . 295
15-1: `tin`: A typical group selection list . . . 303
15-2: `tin`: A typical thread selection list . . . 305
15-3: `tin`: A typical article . . . 307
15-4: `tin`: Commands that are always available . . . 310
15-5: `tin`: Commands to use while selecting a newsgroup . . . 311
15-6: `tin`: Commands to use while selecting a thread to read . . . 312
15-7: `tin`: Commands to use while reading an article . . . 313
16-1: Using a browser to access a gopher menu . . . 319
16-2: Using a gopher client to access a gopher menu . . . 319

17-1: **gopher**: A typical gopher menu 336
17-2: **gopher**: Menu item symbols 341
17-3: **gopher**: Summary of commands 346-348
18-1: An ftp client program 357
18-2: Summary of important ftpmail commands 364
18-3: Ftpmail servers 366
19-1: A list of all the **ftp** commands 375
19-2: Summary of the most useful **ftp** commands 376
19-3: A sample anonymous ftp session 386
20-1: A graphical archie client 390
20-2: Important symbols used in archie regular expressions 393
20-3: Summary of important options for the Unix **archie** client 396
20-4: Public archie servers 399
20-5: Archie servers: Typical output from a **show** command 400
20-6: Commands to use when viewing output with **less** 406
20-7: Archie servers: Summary of important commands 410-411
20-8: A typical session with archie and anonymous ftp 414-415
21-1: Master list of common file types 422
21-2: Common file extensions, organized by category 423
21-3: The most important file types for Windows and DOS users 434
21-4: The most important file types for Macintosh users 440
21-5: File extensions used with Unix compressed files and archives 448
22-1: Summary of important Listserv commands 464
22-2: Summary of important Listproc commands 465
22-3: Summary of important Majordomo commands 467
23-1: A telnet client program 474
24-1: Using job control with **telnet** 486
25-1: A text-oriented talk program 491
25-2: A talk program with a multiple-person conversation 491
25-3: Live video over the net, using CU-SeeMe 500
25-4: A typical web-based chat room 501
25-5: Common communication abbreviations 503
26-1: A typical **talk** screen 510
26-2: A **ytalk** screen with a two-way conversation 516
26-3: A **ytalk** screen with a five-way conversation 517
27-1: An IRC client (mIRC) 524
27-2: An IRC client from a shell account (**irc**) 525
27-3: A typical conversation with IRC 530
27-4: IRC: Basic commands 546
27-5: IRC: Commands for working with channels 547
27-6: IRC: Commands for sending messages 547
27-7: IRC: Commands for displaying information about people 548
27-8: IRC: DCC commands 548
28-1: A typical message sent from one mud person to another 551

28-2: A mud client program 554
28-3: Selecting attributes for a new character 563
28-4: Talk commands for DikuMuds and LPMuds 565
28-5: Talk commands for TinyMuds 566

Acknowledgments

My favorite part of writing a book comes at the end, when everything else is finished, and all I have to do is write the final section: the acknowledgments. However, it occurs to me that it may not be as interesting for you to read these acknowledgments as it is for me to write them. After all, I know all these people and you don't.

"Why," I can hear you say, "should I care about all the people who helped you? What is it to me that Joe Such-and-such from the Acme Computer Company gave you a Personal Widget System to use with your research?"

Good point. So, to make it more interesting and rewarding for you to read this section, I have decided to hold a contest. Here is how it works.

From now on, wherever you go—to the supermarket, to school, to a business conference, for a workout at the health club—carry this book with you. Whenever you meet a new person, check to see if his or her name is mentioned in this section. If so, have the person sign the page on which the name appears. Once you have collected the signatures of everyone in this section, send your book to me. The first person to send me a full list of all the signatures will win a nice prize. (Of course, the validity of the signatures will have to verified.)

So, now that it is worth your while to read these acknowledgments, let's push on. To start, I would like to thank my chief researcher, Wendy Murdock. Wendy knows

Meta's Exhibit No. 1008
Page 0028

Chapter Ten

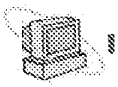
Using the Web

Meta's Exhibit No. 1008
Page 0029

In Chapter 9, I introduced the Web and explained why it is so important and useful. In this chapter, I will show you how the Web works, along with the basic techniques you need to have fun and be productive. I will cover everything you need to understand to use the Web, as well as explain what is involved in having your own web site.

What Is It Like to Use the Web?

In Chapter 9, I explained that the Web is based on hypertext, information that contains links to other information or to services. Underneath, the Web is a huge client/server system (see Chapter 2), with many web servers all over the Internet. To access those servers, you use a client program called a browser that fetches information from a server on your behalf, and then displays the information on your screen. Your browser acts as your interface into the Web, doing whatever is necessary to carry out your requests.



INTERNET RESOURCE *Look in the catalog under Internet Resources for Web: Browsers*

If you have never used the Web before, the idea of hypertext and using a browser will seem a bit strange, so let's start with a few examples. First, here is an example from regular life. You are reading an encyclopedia, looking at an article with the title "Trees". At the end of the article, you see a reference that says "For related information, see Plants". This last line is a link from the "Trees" article to the "Plants" article. To follow this link, you must fetch the appropriate volume of the encyclopedia and turn to the "Plants" article.

Now, if you were using a browser to read the "Trees" article on the Web, the reference to the "Plants" article would be highlighted in some way. For instance, the actual word "Plants" might be displayed in a different color and underlined. To follow the link, you would use your mouse to click on this word. Your web client program would then fetch the "Plants" article—no matter where it was on the Internet—and display it for you on your screen.

Here is another, more complex example. You have a friend who has set up his own web site. He gives you the address, and you decide to check it out. To do so, you start your browser and tell it to contact the address of your friend's web site. (Don't worry about the details, I will explain later.)

In order to carry out your request, your browser contacts the web server at the address you specified. This server honors the request by sending information to your browser, which displays it on your screen. You see messages and pictures designed by your friend. Within the text, most of the words are black on a gray background. However, there are certain words that are displayed in red and underlined. These words represent links that you can click on at any time.

This web site is shown in Figure 10-1. Here is the first paragraph you read:

```
I am a librarian at Stanford University.  Last year,
I went on a bicycling trip in Holland with my girlfriend
Yolanda.  While we were there, we visited my friends
Peter and Karin and their bunny Pixie.  My hobbies are
photography and collecting Elvis memorabilia.
```

This example contains several links (which I have underlined). The first is **Stanford University**. If you click on this, you will jump to the web site for Stanford University. The next link is **bicycling trip**. Clicking on this will display a long description, written by your friend, of his trip and all the places he visited. The following link is **Yolanda**. Click on this and you will find more information about your friend's girlfriend. First, you will see a picture of Yolanda, followed by the words "For more information about Yolanda, click **here**." If you do click on that link, you will jump to Yolanda's own web site.

Now that you see how this works, you might be able to guess what type of information the rest of the links point to. If you select **Peter and Karin**, you will see a picture of Peter, Karin, Yolanda and your friend, standing beside their bicycles outside a restaurant in Amsterdam. Clicking on **Pixie** shows you a picture of Pixie

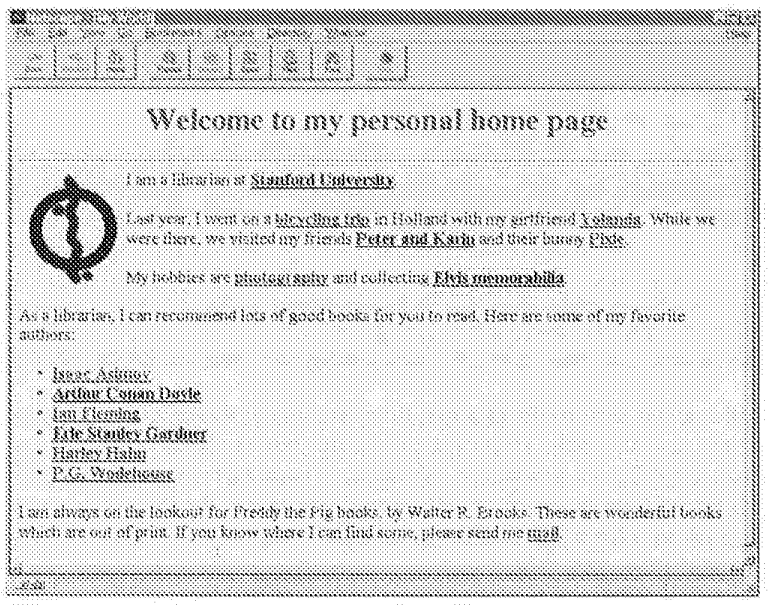


Figure 10-1. *A typical web page*

the Bunny. Clicking on **photography** shows you a small gallery of some of your friend's photos along with a few links pointing to various photographic resources on the Net. Finally, selecting **Elvis memorabilia** takes you to a list of Elvis information sites around the Net, one of which is a guided tour of Graceland (Elvis' home).

Now, the only way to really get a feel for the Web is to use it, but the foregoing example will at least give you an idea of how it works. People create text, and within that text, put links to other text, pictures and various resources. Sometimes links point to something else created by the same person (such as the link to the description of the bicycle trip, or the pictures of Peter, Karin and Pixie). Other links point to resources elsewhere on the Net (such as the links to Yolanda's web site or the Elvis information sites).

This example describes a web site created by one person for his own amusement, and there are many such sites on the Net. However, there are also many web sites created by organizations: companies, schools, government agencies and so on. When you visit such sites, you will find "official" information. For example, the AT&T web site will describe the services and products offered by AT&T. As you might expect, such web sites are better organized, but more formal than the personal ones. For example, the AT&T web site will have slick graphics and all kinds of corporate information. However, what you won't see are descriptions by an AT&T executive of her last trip to Alabama, along with pictures of her family and the pet possum she used to keep as a child.



HINT

The Web is actually a mixture of two very different types of web sites: those created by individuals and those containing information about an organization. This gives the Web a countenance that is unmatched anywhere else in our society. Thus, as you use the Web, you will encounter official, no-nonsense corporate, academic and governmental web sites coexisting peacefully with the highly personal and idiosyncratic creations of people who are free to write and exhibit anything they want.

As I will explain later in the chapter, there are lots of things you can do with your browser as you use the Web. Most of the time, though, you will be doing one of three things: reading some text, looking at a picture, or waiting for something to be copied from a web server to your computer. To be realistic, using the Web means spending a lot of time waiting. This is why, in Chapter 4, I advise you to get yourself as fast an Internet connection as possible.

As you use the Web, the waiting periods are short but frequent. To jump from one web item to another, you use your mouse, click on the next item you want to see, and then wait for it to be loaded. Thus, you will find that you develop a definite rhythm:

click, wait, read... click, wait, read... click, wait, read... The overall timing is affected by what you choose to look at—for example, pictures take much longer to copy than plain text—but just the same, there is a rhythm. If you use the Web during the day when many people are sharing the communication lines, you will wait longer. In the middle of night when there are fewer people, the response times decrease and the pulse of the Web speeds up.

Compare this rhythm (...click, wait, read...) to the rhythm of other popular menial activities. Watching television, for example, is ...click, watch, watch, watch, watch... click, watch, watch, watch, watch. Reading a book is ...turn a page, read, read, read, read... turn a page, read, read, read, read.



HINT

We are shaped by our tools more than we like to admit.

Links, URLs and Hyperspace

Before I discuss the details of using your browser, I would like to explain a few of the important terms and concepts you will encounter as you use the Web.

I have already discussed the idea of hypertext: information that contains links to other information or services. I would like to take a closer look at this idea and answer the question: what exactly is a link?

In the same way that every Internet user has a unique email address, every hypertext item on the Net has an address of its own. This only makes sense. When you tell your browser to fetch and display a particular item, the browser must know where on the Internet that item resides.

The format of such addresses is different from that used for mail addresses. One of the reasons is that the people who first developed the Web wanted it to provide access to all types of information and not just hypertext. Thus, they used an addressing system that is capable of referring to just about any type of Internet information resource. Within this system, we use a special type of address called a UNIFORM RESOURCE LOCATOR, usually referred to as a URL. (The name is pronounced as three separate letters: "U R L".)

Within this system, there is a unique URL for any hypertext item on the Net. Moreover, there are also unique URLs for non-hypertext items from other services, such as gophers, anonymous ftp sites, Usenet newsgroups and wais databases. I discuss these resources in detail in other parts of the book. The important thing to understand here is that, because browsers are designed to access any item that has a URL, you can use your browser to access information from a variety of Internet resources and not just hypertext.

Here are two examples of URLs. (For now, don't worry about the details, I will explain them later in the chapter.)

```
http://www.alan.com/
http://www.bu.edu/Games/games.html
```

Notice that both of these URLs start with **http:**. This means that the URL refers to a hypertext item. (The name **http** stands for HYPERTEXT TRANSFER PROTOCOL, the protocol used to send hypertext data over the Internet.) If a URL refers to a different type of item, you will see a different prefix. For example, a URL that starts with **gopher:** refers to a gopher item, a URL that starts with **ftp:** refers to an ftp resource, and so on. Again, don't worry about the details, I will cover them later.

Each URL refers to one specific Internet resource. To express this idea, we say that the URL POINTS to that resource. For example, the first URL in the above example points to the web site for Alan Colmes, a well-known American radio talk show host. The second URL points to a collection of interactive web games at Boston University.

As you already know, your browser retrieves hypertext data and displays it on your screen. During this process, your browser will recognize the links within the data, and display them so as to make them easy to recognize. For example, your browser may display all the links in a special color, or it may underline each link, or both. (In fact, most browsers give you some choice as to how links are to be displayed.)

Take a look at the following example of hypertext in which the links are underlined.

```
If you want to find out about a real cool guy, check
out Alan Colmes, the famous radio star.  Or, if
you have some extra time, try playing some of
the interactive web games at Boston University.
```

Suppose these two links represent the URLs in our last example. That is, suppose the first link, **Alan Colmes**, corresponds to the URL **http://www.alan.com/**, while the second link, **interactive web games**, corresponds to the following URL: **http://www.bu.edu/Games/games.html**.

If you want to see the information at either of these web sites, all you have to do is select one of the links. (With a mouse, you would click on the actual link.) Your browser will then use the corresponding URL to contact the site and ask for the data to be sent to your computer. When the data arrives, your browser will display it for you on your screen.



HINT

A nice way to think about links is to imagine that underneath each link lies a URL. Thus, each time you select a link, it is as if your browser looks under the link, finds the URL, and does whatever is necessary to fetch the information to which that URL points. With most browsers, you can set an option to have the browser display URLs. Once this option is set, whenever you move your mouse pointer over a link, your browser will show you the corresponding URL. Thus, you can move the pointer around the screen and, each time you pass over a link, your browser will show you the URL that lies "underneath".

When you access a web site by selecting a link, we say that you are FOLLOWING the link. For example, you might follow the `Alan Colmes` link to Alan's web site. Although, technically, it is the browser (a program) that is following the link, we usually talk as if we are performing the action ourselves.

For example, say that you see a guy wearing a T-shirt with a picture of Alan Colmes. You ask, "Where did you get such a radical shirt?" and he says, "I made it from a picture I got from the Web. After I found the picture, I had my browser save it to a file on a floppy disk, which I then took to a place where they printed it on a T-shirt for me." "But how did you ever get the picture in the first place?" you ask. "Oh, that was simple," he responds, "I was looking at a place that had information about talk radio, and I just followed a link to the Alan Colmes web site."

Notice the assumed metaphor. When we use our mouse to select a link, we talk as if *we* are following the link; as if using the Web is like venturing into a foreign land in which we jump from one place to another by following links.

Such observations are more than mere fancy. As human beings we tend to interpret new ideas in light of things we already understand. Thus, it is natural for a human being who is using the Web to imagine himself jumping from one Internet location to another. For this reason, we sometimes talk about NAVIGATING the Web.

The reality of it, though, is that you do not really visit various Internet locations. Each time you follow a link, your browser asks for the information to be sent to *your* computer. Thus, you don't really go anywhere: the various places you "visit" are actually brought to you. Still, if you can suspend belief long enough to think of your screen as a window into the Net, it does seem as if the view changes each time you click on a link.

The word we use to describe all of this is HYPERSPACE. Technically speaking, hyperspace consists of all the possible sites you can access using the Web. Of course, this is a metaphor: hyperspace is not a real place that exists in time and space. Still, it is intriguing to imagine that, in the future, you might connect yourself to the Net in such a way that you actually feel as if you are in a different dimension. And, as you follow the links, you might actually feel as if you are traveling in a strange, alien world. In that sense, hyperspace may one day become a real place in which we may spend all our spare hours. In the meantime, don't forget to stay grounded on Planet Earth and pay your telephone bill.

Words to Avoid

In 1982, William Gibson published a science fiction novel called *Neuromancer* in which he described a futuristic computer-mediated world called "cyberspace". At the time, Gibson was fairly ignorant about computers, so his speculation about computers, humans and the future was both annoying and fascinating.

In the early 1990s, as the Internet and the commercial online services began to grow, people started to refer to these new systems as CYBERSPACE. Later, new expressions for the same idea—the INFORMATION SUPERHIGHWAY, the INFOBAHN, and so on—became popular, especially after they were popularized by the American Senator and Vice-President Al Gore.

Then, in 1992, Jean Armour Polly wrote an article introducing the Internet for a librarians' journal. At the time, Brendon Kehoe had just published a landmark article entitled *Zen and the Art of the Internet* and Polly wanted a title that captured the same tone. She decided to choose something nautical (because the "Net" reminded her of a fishing net). After discarding a number of titles, she came up with one that invoked the idea of "fun, changing conditions, and danger". That title was *Surfing the Internet*.

To Polly (who was a qualified librarian), the word "surfing" captured the essence of using the Internet: "The fun was in trying to find something in the same way that a librarian does research. The danger was that the information might be faulty, old or incorrect. And because conditions on the Internet were always changing, you might have trouble finding your way back to a particular source. Even if you could find your way back, the original information might not be there any longer."

After publishing *Surfing the Internet*, Polly revised it and made it available for free on the Internet. The article quickly became the most popular Internet document of its time. What happened next surprised even Polly. The term SURFING became part of the popular culture, but not as Polly originally intended. Instead, the term came to refer to the activity of browsing the Internet.

Today, Jean Armor Polly is a highly respected Internet expert, sometimes referred to as "Mother Internet". She is, perhaps, the most famous woman on the Internet. (If I weren't so modest, I would tell you who the most famous man is.) However, along with William Gibson and Al Gore, she has the dubious distinction of introducing into the language a word that has become trite and overused.

So here is my advice. I encourage you to tell people that you know how to "navigate the Web" and that you spend time on "the Net". And, if you really want, you can even talk about "hyperspace" (as long as you don't use the word more than once in a single conversation). However, please don't ever let me hear you refer to "cruising the information superhighway", or "surfing the Net" or "hanging out in cyberspace". Such words are to be avoided as they will instantly mark you as someone who has no idea what he is talking about.



HINT

Stay away from meaningless words and expressions that ignorant people use to describe the Net. Remember, as one of my readers, you are a highly intelligent, knowledgeable and motivated individual—not a clueless goober.

Web Pages and Home Pages

Hypertext information is stored in files, and a URL that points to hypertext is actually the address of a specific file on a particular computer. Each time you select a link that points to a hypertext file, your browser will contact the web server at the appropriate site and request a copy of the file specified in the URL.

The contents of a single file of hypertext is called a PAGE or, more formally, a WEB PAGE. Don't get confused, however, by thinking that a web page is like a page in a book. A web page can be any size at all. Some people will design a web site to consist of a single long page. Other people will break the same information into a number of smaller pages (each in its own file), with links between the pages. Thus, a page may be only a few lines, or it may be so large as to contain many screenfuls of information.

Since pages are the logical units of hypertext, you will often hear people referring to an entire web site as a "page". For example, you might overhear a young child ask his mother, "Is it true that Santa Claus has his own web page?" This is a common way to speak, even though many web sites (including Santa Claus') actually have a number of pages.

When your browser requests and waits for a copy of the page from a web server, we say that the browser LOADS the page. (The term comes from the word "download", which means to copy information from another computer to your computer.) Thus, you might hear someone say, "I'll be with you in a moment, just let me get this web page loading."

Another term you will encounter frequently is HOME PAGE. This term has two distinct meanings. First, a home page is the main page for a particular web site. For example, many people have their own web sites which they maintain for fun. These are called home pages. (Figure 10-2 contains an example of such a home page.) With

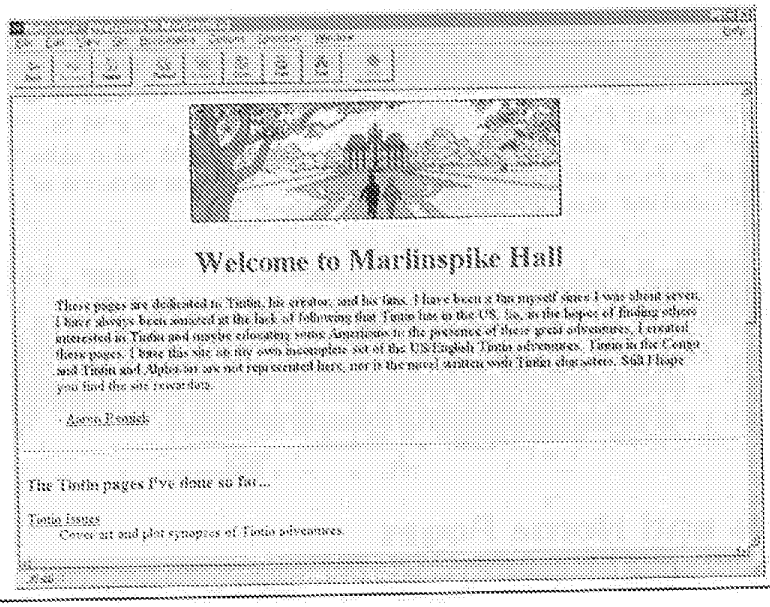


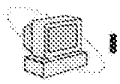
Figure 10-2. *A personal home page*

an organization (such as a company, university or government department), the home page will be less personal and will act as a main menu from which you can navigate to other pages.

Thus, if you were to take a look at IBM's home page, you would expect to find a main menu pointing to various other pages. If you were to look at Norman Sweetcream's home page, you would expect to find a more personal, idiosyncratic offering, perhaps only a single page.

The second meaning of the term "home page" has to do with your browser. Some people like to start each web session by connecting to the same web page. For example, each time you start your browser, you may want to connect to your favorite humor page to see if there are any new jokes. Most browsers allow you to specify the name of a page that should be loaded automatically each time you start the browser. Within your browser, this page is called your home page.

So don't be confused. To your browser, your "home page" is the page you want to be loaded automatically each time you start the browser. To the Internet community, your "home page" (if you have one) is a specific hypertext page, created by you, as your own personal web site. (I will discuss how to create your own home page later in the chapter.)



INTERNET RESOURCE *Look in the catalog under Internet Resources for Personal Home Pages*

Telling Your Browser Which URL to Use

Each browser works a little differently, but there are certain characteristics that all browsers have in common. As you read the next few sections, follow along with your own browser, so you can see how the details work on your particular system. In the examples, I will mention details pertaining to the Netscape browser, but things should work pretty much the same no matter which browser you are using.

(Note: In the following sections I concentrate on graphical browsers, such as the ones you would use with a PPP account running under Windows, Macintosh, OS/2 or X Window. If you have a Unix shell account, you will probably use the Lynx browser, which I describe separately in Chapter 12.)

Some browsers will automatically load a specific web page for you each time you start the program. This is a feature you can turn off or on and, by default, it may be turned on when you install the program. The web page that is loaded automatically is called your home page. (However, as I explained in the previous section, be sure not to confuse this type of "home page" with a personal home page.) If your program has this feature, and you decide to leave it turned on, you can set your home page to any URL you want. For example, if you have a favorite humor site on the Web, you can have your browser connect to it automatically, so you can check to see if there are any new jokes.

The idea of having a home page from which to start is one that seemed to make sense when browsers first became popular. At the time, the gopher system (see Chapter 16) was widely used and, with a gopher, it makes sense to have a home base. The Web, however, is a lot different and you really don't need to start from the same place each time. For this reason, most people find the auto-loading of a home page annoying and choose to turn off this feature. Still, it is there for you if you want it.

In general, it is up to you to tell your browser where you want to go. You can do this in one of three main ways. First, you can type a specific URL. To do this, you can pull down a menu item that allows you to "open" a new URL. With Netscape, for example, you pull down the "File" menu and choose the "Open Location" item. A dialog box will pop up and wait for you to enter a URL. Alternatively, your browser will have a box, probably near the top of the screen, into which you can enter a URL directly. Simply move your mouse pointer to the box and click the mouse. You can now type a new URL and press ENTER.

Another way to choose a URL is to use one of the built-in features of your browser. Most browsers come with pre-installed URLs to help you connect to a few useful resources. For instance, the Netscape browser has several buttons near the top of the screen. These buttons are assigned to specific resources and, to connect to one of these resources, all you have to do is click on a button. For example, one of the buttons says,

"What's Cool!" When you select it, you will be connected to a list of "cool" web sites maintained by the Netscape company.

The third way in which you can tell your browser what URL to follow is to choose an item from your "bookmark list". A bookmark list is a collection of URLs that you yourself have compiled. As you navigate around the Web, you will encounter sites that are especially interesting to you. Your browser will have an easy way for you to save the URLs of such sites in a special list. At any time, you can recall the list and select any item you want. I will talk about your bookmark list in more detail later in the chapter. For now, I just want you to know you can create such a list and that it is an easy way to jump immediately to a particular web site.

Reading a Web Page

Once you specify a URL, your browser will begin to load the web page at that address. You will find that many web pages contain pictures as well as text and, compared to text, pictures can take a long time to load. For this reason, most browsers will load the text first and then work on the pictures. This allows you to start reading as soon as possible.

If you change your mind about wanting the entire page, you can stop the loading whenever you want. Most browsers have some type of "Stop" button which you can click with your mouse. Alternatively, pressing the ESCAPE key will also work with most programs. If you stop loading a page in this way, your browser will display the parts of the page that were already received, so if you see what you want, there is no need to wait for the entire page to be loaded.



HINT

With most browsers, you do not have to wait until the entire page is loaded to start to select a new link. Even though the current page may still be loading, you can move your mouse pointer to a link and click on it. The browser will stop loading the current page and jump to the new one.

Once a web page is loaded, you can read it at your own speed. If it is too large to fit on your screen, you can move around by using your mouse (click on the scroll bars) or your keyboard (with the PAGEUP and PAGEDOWN keys).

Aside from reading, there are several other things you can do with a page once it is loaded. You can save it to a disk file, print it or mail a copy to someone. You can also save the URL for the page in a special collection of URLs called your "bookmark list" (explained later in the chapter). Saving a URL on your bookmark list allows you to find it again whenever you want. The details of performing these operations vary from one browser to the next. Usually, though, all you will need to do is select an item

from a pull-down menu, and specify the appropriate information (file name, mail address and so on).

Links, Forms and Image Maps

Once you have loaded a web page, you can navigate by selecting one link after another. Just move your mouse pointer to a link, click on it, and wait.

HINT

There is an easy way to tell if a word or small picture is a link. Turn on the browser option that tells it to display URLs. Now whenever you move your mouse pointer over a link, your browser will display the corresponding URL. If you are not sure whether or not something is a link, all you have to do is point to it and see if a URL appears.

Most links are simple. They will either be highlighted words or a small graphical element such as a button or a picture. There are, however, two other types of links which are more complex that I want to discuss.

The first type of special link is a FORM. A form is a facility that allows you to enter information to be sent back to the remote web site for processing. For example, say you are visiting the home page of someone who has set up a "guest book". The guest book is set up so anyone who visits the page can leave a personal message, and subsequent visitors can read what other people have written.

When you look at the home page, you will see a box into which you can type a comment. This is the form. Beside the box will be a word or a small picture you can click to send the comment once it is typed. For example, you may see a button that says "Submit". All you have to do is type your comment and then click on the "Submit" button. Figure 10-3 shows a web page that has such a form.

Originally, the idea of a form was developed so that users could submit information which would be sent to the remote web site to be processed, just like a regular form on paper. Although this is still one of the biggest uses for forms, people have found more imaginative ways to use this facility. For example, there are interactive web games in which you can make a move or specify a strategy by filling out a "form" and submitting it.

You might wonder, what happens to the information when it reaches the other end? How is it processed? The answer is, the web server that accepts the information sends it to a special program called a CGI SCRIPT. (The name stands for COMMON GATEWAY INTERFACE.) Thus, whatever you type into the form is sent to a CGI script at the other end that receives the information and processes it in an appropriate manner. This means the person who designed the form must also have set up a CGI script to handle the incoming data. As a user, you do not need to worry about the

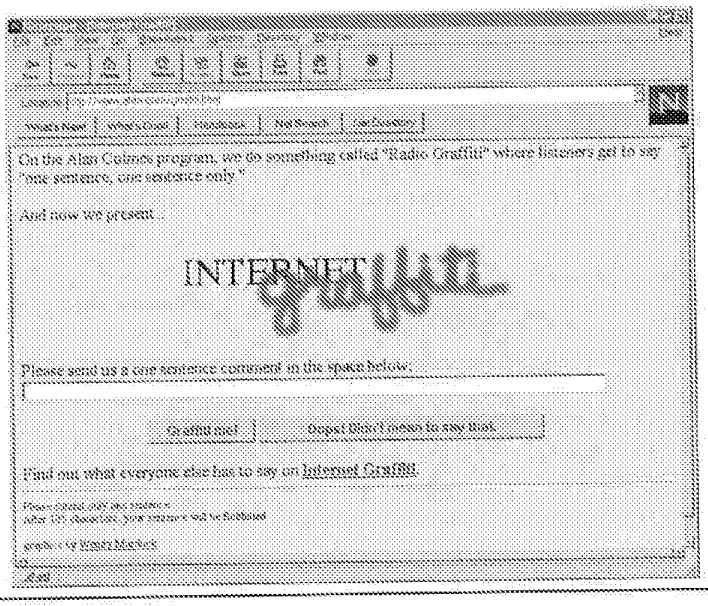


Figure 10-3. *A web page containing a form*

details. However, if you are developing a home page of your own and you want to use a form, you will have to supply a CGI script to handle the incoming data.

INTERNET RESOURCE *Look in the catalog under Internet Resources for Web: CGI (Common Gateway Interface)*

The second type of special link I want to mention is an IMAGE MAP. This is a picture of some type (typically a photo or a drawing) in which the various parts of the image act as separate links. For example, there is a weather report service that will send you a page containing a map of the United States. To request a weather report for any area of the country, all you have to do is click on that part of the map. (Figure 10-4 shows an example.) One common design is to have an image map consisting of several small pictures, each of which is actually a link to a different web page.

How does an image map work? As you move your mouse pointer around the image, your browser keeps track of the coordinates (similar to longitude and latitude on a map). When you click on a particular location, your browser will send the

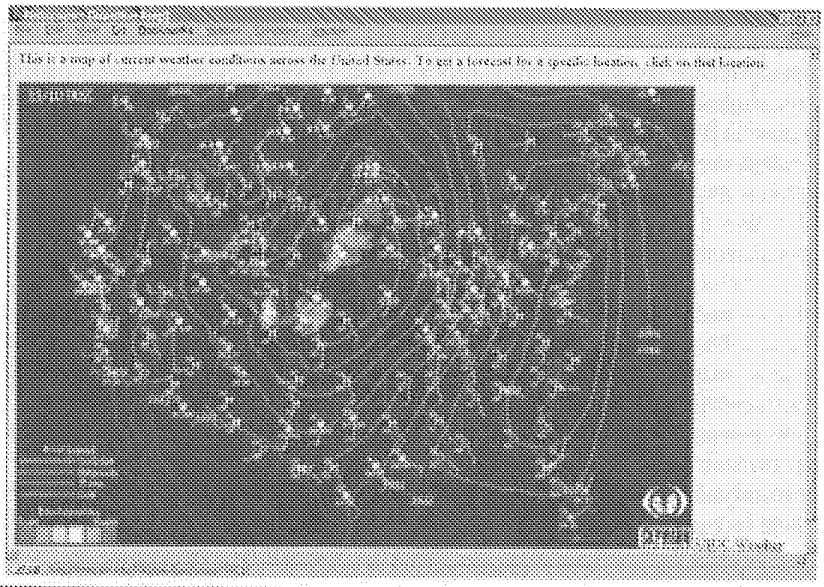


Figure 10-4. *A web page containing an image map*

coordinates of that location to the remote web server. The web server will then look at the value of the coordinates and send back the appropriate information.



HINT

One problem with image maps is there is no sure way of knowing exactly where the links are. Some browsers will show you the actual coordinates as you move your mouse pointer around the image, but there may be no other clues. Indeed, it can sometimes be hard to recognize a particular picture as being an image map.

For this reason, a good web page designer will put in some text to show you are looking at an image map. For example, you may see something like, "To visit one of our electronic stores, just click on the part of the picture that shows what you want to buy."



INTERNET RESOURCE *Look in the catalog under Internet Resources for Web: Image Map Tutorial*

The History List and the Bookmark List

Hyperspace can be a confusing place, and it is easy to jump from one link to another without remembering where you have been. Still, there will be times when you will want to revisit a link you left behind. To help you, your browser keeps a record—called the HISTORY LIST—of the links you have visited. Some browsers keep every link you have visited in the current session, regardless of how much you have branched. Other browsers keep only those links that trace your path from the first page to the current page, with no branching.

There are two ways to use the history list. First, most browsers have an easy way to move backward or forward one link at a time. Typically, there will be two buttons you can press. They will be marked "Back" and "Forward", or have left and right pointing arrows, or something similar. When you click on the "Back" button, your browser will return to the last page you loaded. If you want to move back even further, you can press the button as many times as you want. Once you have pressed the "Back" button, you can press "Forward" to display a subsequent page.

The second way to use the history list is to choose an item from it directly. Your browser will have a way to display the whole history list on your screen. You can then select any item from the list, click on it, and jump to it immediately.

The history list is useful for moving back and forth within a single web session, but there will be many times when you will want to save a web address permanently. For this reason, all browsers have a facility that will save the URL of the current page to a BOOKMARK LIST. Whenever you want, you can display the bookmark list, select an item, and jump to it directly. Unlike a history list, your bookmark list only contains URLs put there by you. Moreover, a bookmark list is permanent. Your browser will automatically save it to a file, so as to be available from one session to the next.

Most browsers have a way for you to edit your bookmark list. At the very least, you should be able to modify the description of any item on the list, as well as delete any unwanted items. Some browsers even allow you to organize the items into categories and subcategories, in the same way that you might organize files into folders or directories.



HINT

As your bookmark list grows, it can become unwieldy and filled with unwanted items. I suggest that you take a few moments every now and then to organize your treasures. You never know what the future holds. If you become famous, a major university may want your bookmark list (along with your papers) after you die. An overgrown, disorganized list will reflect badly on you, serving only to create everlasting embarrassment for your family and friends.

Images

One of the nice things about the Web is there are so many visual images. Broadly speaking, there are two types of pictures: inline images and external images. An INLINE IMAGE is a picture that is a part of a web page. For example, you will often see small pictures or icons loaded along with the text of a page. These are inline images. An EXTERNAL IMAGE is a picture that is a web page in its own right and, when loaded, is displayed in its own window.

Here is an example. Imagine an Internet art gallery which is set up to show you pictures of famous paintings. However, a single picture contains so much data that it would take a long time to load. If the gallery were set up like a real gallery—where you walk in and see all the pictures at once—it might take minutes and minutes to load the gallery home page. The solution is to design the gallery so you can preview the pictures by looking at small images (which load quickly). Then, when you decide which picture you want to see, you can tell your browser to load the full-sized version.

In order to implement this system, the web page designer will create miniature versions of each picture in the gallery. These small pictures will be included on the home page as inline images. Thus, whenever you load the home page, you will see a catalog of what is available in the gallery. Each inline image acts as a link to the file that contains the full-sized picture. So to display a full-sized picture, all you need to do is click on the corresponding small, inline image. When you do, your browser will load the picture and display it in its own window. If you want to look at another picture, you can go back to the home page and click on another inline image.



HINT

If you ever want to use the Web, and you don't care about looking at pictures, your browser will have an option to turn off the automatic loading of images. Setting this option is handy when you want to find something fast, because you don't have to wait for the pictures each time you load a new page.

If you have a shell account (see Chapter 5), another alternative is to use Lynx, a text-only browser. (See Chapter 12 for details on using Lynx.) Personally, when I am doing research and I don't care about images, I often use Lynx because it is much faster than a graphical browser.

Images are stored in various formats, the two most common being GIF and JPEG. For convenience, GIF files are given names that end in `.gif`. Similarly, JPEG files are given names that end in `.jpg`. For example, you might encounter image files named `harley.gif` or `wendy.jpg`. The GIF format was developed by Compuserve. The name GIF stands for "Graphics Interchange Format". JPEG is named after the organization that developed the format, the Joint Photographic Experts Group.

In order to display an image, your browser must be able to understand the format of the data. For example, to display the file `harley.gif`, your browser would have to be able to understand what to do with a GIF file. In web language, a program that can display a particular type of file format is called a VIEWER. Thus, in order to display GIF files, your browser must have access to a GIF viewer. To display JPEG files, your browser needs a JPEG viewer.

It is likely that your browser will have built-in viewers for GIF and JPEG. If it doesn't—or if you want to be able to display other image formats—you can use an external viewer. Simply find an appropriate program (they are readily available on the Net), install it on your computer, and then tell your browser which file formats the new viewer will handle. To do this, use the menu item in your browser that lets you change the configuration or preferences information. This tells your browser which viewer to use when it encounters a particular file format.

Sounds and Video

In addition to images, you will also find many sound and video resources on the Web. Although there are a number of different file formats, only a few are in common use.

The sound formats you are most likely to encounter are AU and WAV. AU was developed by Sun Microsystems; the name stands for "audio file". WAV was developed by Microsoft for Windows; the name stands for "waveform data". There are also other formats (which I will discuss in a minute) that are used for "real-time audio".

For video, there are three common file formats. MPEG was developed by the Motion Picture Experts Group, Quicktime was developed by Apple, and AVI was developed by Microsoft for Windows. The name AVI stands for "Audio/Visual Interleaved data". ("Interleaved" is a technical term indicating that the data is processed in alternating sequences. Don't worry about it.)

As with GIF and JPEG, each of these types of data is stored in files with special names. AU files are given names that end with `.au`; WAV file names end with `.wav`; MPEG file names end with `.mpg`; Quicktime file names end with `.qt` or `.mov` (movie); and AVI file names end with `.avi`. This means that you can tell what format a file uses just by looking at its name. For reference, the most common file formats for images, sound and video are summarized in Figure 10-5.

To listen to audio files or to watch video files, you need programs that can work with the appropriate data formats. With GIF and JPEG files, you use a viewer to display the file. Similarly, we call programs that interpret sound and video files "viewers" (even though you don't really view sound). You may see other names—for example, Netscape calls such programs "Helper Applications"—but the name "viewer" is a generally-accepted generic term.

Although some of the formats in Figure 10-5 were developed by one company or another, they are all in common use, and viewers are available for all types of computers. For example, although the AU format was developed by Sun, you can

Format	Typical File Name	Origin of Name
GIF	`filename.gif`	Graphics Interchange Format (Compuserve)
JPEG	`filename.jpg`	Joint Photographic Experts Group
AU	`filename.au`	Audio file (Sun)
WAV	`filename.wav`	Waveform data (Microsoft)
Quicktime	`filename.qt`	Quicktime (Apple)
Quicktime	`filename.mov`	Quicktime [Movie] (Apple)
AVI	`filename.avi`	Audio/Visual Interleaved data (Microsoft)
MPEG	`filename.mpg`	Motion Picture Experts Group

Figure 10-3. *Common file formats for images, sound and video*

listen to AU files on a PC or Macintosh, as long as you have a program that can interpret that file format.

With respect to sounds, there are two types: regular sounds and REAL-TIME SOUNDS. When you listen to a regular sound, the entire file is downloaded to your computer before the sound is played. With a real-time sound, the sound is played as it is being downloaded.

The advantage of a real-time sound is you do not have to wait for the entire downloading process to complete before you can start listening. However, the sound itself is not saved. It is played as soon as it is received, and the data is not retained.

Regular sounds, on the other hand, are stored in a file. Thus, you can keep a copy of a regular sound and replay it whenever you want. Moreover, regular sounds tend to be of higher quality, because in order to make a real-time sound load fast enough, some of the data is discarded.

Real-time sounds are good for transmitting large amounts of data that no one really wants to store permanently. For example, you might listen to a broadcast of a sporting event or a radio program via real-time sound. Regular sounds are used for storing more permanent data. For example, you will find many regular sounds available in web sites and anonymous ftp sites around the Internet.

If your browser does not come with a built-in viewer for any of these formats, external viewers are readily available on the Net. When you install a viewer on your computer, all you need to do is tell your browser where the program can be found, so it knows where to ask for help when it encounters that particular file format. (You set up such information using the configuration or preferences menu.)

Regular sounds use the file formats I mentioned above. Real-time sounds use a variety of different formats. The important thing is to make sure you have the correct viewer—that is, playback program—for the type of sound to which you want to listen.



HINT

Sound and video data takes up a lot of room, even more so than pictures. For this reason, such files tend to be large. Indeed, it is not uncommon to spend minutes and minutes loading a (regular) sound or a video clip, only to have it play back in just a few seconds.

In this way, listening to sounds or watching video on the Web is a lot like Life. People complain that it takes a long time to get what they want and that the quality is poor, and yet the whole thing is over much too soon.

Updating Information Automatically

Once you tell your browser to load a web page, you must wait for the information to be sent from the web server to your computer. As the information arrives, your browser (the web client) displays it for you. However, once the page is displayed, it does not change.

Nevertheless, there are times when it is desirable to display information which can be updated automatically. There are two ways in which this can be done.

First it is possible for a web page to be set up so that it tells your browser to reload itself (or a different web page) automatically. For example, a web page might tell the browser to reload the page every 30 seconds. This facility is called CLIENT PULL, because the client program (your browser) is told to "pull" information from the server without your intervention.

An alternate system allows the server to send new data on its own, without your browser having to make a request. This facility is called SERVER PUSH.

With client pull, the connection between the server and your computer is dropped as soon as the web page is finished loading (which is what happens normally). When your browser is ready to ask for more data, it reconnects to the server and issues the request.

With server push, the connection between the server and your computer is kept open indefinitely. This allows the server to send more data whenever it wants.

Client pull and server push are useful in situations where you are looking at data that changes. For example, a financial service could use server pull to update a stock market report every five minutes. An example of client pull is found with certain web-based talk facilities. (A talk facility lets you use the Web to send messages back

and forth to other people.) By using client pull, the talk system can update your window at regular intervals in order to show you any new messages that may have arrived. (For more details about talk facilities, see Chapter 25.)

Client pull and server push will not work with any web page or web server. A web page must specifically request client pull, and a server must be designed to offer server push. In addition, your browser must be aware of these features and must know how to support them. (If your browser does not support these services, the web page will still load, but it will not be updated automatically.)

Web Directories and Search Engines

There are so many items available on the Web, I guarantee you will never run out things to do—as long as you know how to find what you want. However, as I explain in Chapter 1, there is no central authority in charge of the Internet, which means there is no central directory of web sites. Anyone can create (or remove) a web site at any time.

A central directory would be wonderful, as it would allow you to search for anything you want. For example, you might want to look for a certain type of gardening information. Or you might want to get a copy of a free program which is available somewhere on the Net. How do you know where to look?

A good number of people have attempted to solve this problem by compiling a Web directory, and designing a tool to allow users to search the directory for a particular resource. The reason I say "attempted to solve the problem" is that the Web is so huge and so volatile, it is impossible to have a fully comprehensive, up-to-date directory.

Still, web directories do make excellent resources, and if you know how to use them, you can almost always find what you want. This is because web page designers tend to put in links to similar resources. Thus, even if you find only a few resources related to a particular topic, these will often lead you to all the information you want. For example, if you are looking for some specific gardening information, starting with just a few gardening pages will likely lead you to all the links you need.

As you can imagine, it would be a difficult and time-consuming job to compile a comprehensive web directory by hand and to keep it up to date. For this reason, the people who maintain such directories use special programs to do most of the work. These programs search the Web relentlessly, looking for new and updated items. (It is common to see such programs referred to by exotic names, such as "spiders", "worms", "crawlers", "robots" and so on. However, don't worry about the names; they are not important.)

There are a number of free, well-maintained directories available on the Web and, in general, there are two ways to use them. First, some directories have home pages organized into categories. To use such a directory, all you need to do is load the home page and choose the category you want. Within that category, there may be subcategories. Thus, you can select categories and subcategories, until you narrow down your search to find exactly what you want.

Other directories allow you to look for something by using a SEARCH ENGINE. To use a search engine, you use your browser to fill in a form, in which you describe what you are looking for. Once you submit the form, it is sent to the search engine, which quickly looks through the directory, finds what you want (if it exists), and sends back the results.

Regardless of how you access a directory—either by selecting categories or by submitting a request to a search engine—the results are similar. You will end up with a number of links. All you need to do is follow the links to see if, indeed, they contain the information you want.

To help you decide which links are the most useful, the search engine will show you some descriptive information. For example, you will usually see not only a set of links pointing to web pages, but the titles of those web pages as well. Some search engines may also show you part of the actual page. Still others will do their best to rank the results from most to least relevant, based on the search criteria you specified.

Using a search engine for a simple search is easy. Usually, all you have to do is type one or more words into the form, click on the "Submit" button, and wait for the results. However, if you specify a more complex search, you can often get better results.

Most search engines allow you to use various types of search criteria. For example, say you want to find all the web sites that contain a reference to "Harley Hahn". The simplest way would be to search for all the web pages in the directory which contain the two words **`Harley Hahn`**. However, if this doesn't find what you want, you might want to expand your search by looking for pages that contain either **`Harley`** or **`Hahn`**. Or you might want to search for the word **`Harley`**, followed immediately by the word **`Hahn`**. Or, perhaps, either the word **`Harley`** followed by **`Hahn`**, or the word **`Hahn`** all by itself. Various search engines allow you to formulate a search in different ways. However, there will be help information, and learning the details for a particular search engine is usually easy.

One set of rules you will often see are those allowing you to combine words using the "and" and "or" operations (like the examples in the previous paragraph). Such searches are often called BOOLEAN SEARCHES.



What's in a Name?

Boolean Search

A boolean search is one that combines search items using "and" and "or" operations. The name originates with George Boole, an English mathematician (1815-1864) who flourished in the mid-nineteenth century. Boole, along with another mathematician, Augustus de Morgan, was one of the chief developers of symbolic logic.

Now, symbolic logic (keep following because I am leading you somewhere) applies mathematical symbols to logic. Using such symbols removes the possibility of verbal ambiguities and simplifies the manipulation of logical relationships. In the history of mathematics, this is considered a big deal. For example, not long after the turn of the century, Bertrand Russell and Alfred Whitehead used and extended Boole's symbolic logic to create the first logical foundation of mathematics. Although this work has since been superseded by more modern analysis, it is considered to be one of the stunning intellectual achievements of the twentieth century.

This same use of logic and symbols was later adapted by computer scientists to create techniques used to design and analyze computing machines. In honor of George Boole, this field of study was called boolean algebra. One important facet of boolean algebra is being able to analyze expressions that involve the "and" and "or" operations. Thus, it has become the custom to use the word "boolean" to describe anything that involves such operations. Hence, it is common for searches defined by "and"s and "or"s to be called boolean searches.

Although many people talk about such searches, few people know what the name means, and even fewer know its origin. (But now you do.)

Aside from simple boolean searches, most search engines allow you to use other, more complex types of search patterns. I suggest you become familiar with three or four different search engines, and take the time to read the help information. With a little practice, you will be able to find just about anything you want.

HINT

Take some time to learn how to use several search engines well. The Web is vast beyond the understanding of any single human being, and being able to search with skill and confidence will open doors which are closed to most everyone else.

INTERNET RESOURCE *Look in the catalog under Internet Resources for Web: Search Engines*

Customizing Your Working Environment

In general, all browsers have the same purpose in life: to act as a client program—that is, your interface—as you use the Web and other Internet services. There are a great many different browsers, and each one has its own characteristics. Nevertheless, all browsers allow you to customize your working environment to some extent.

In order to use your browser well, you should take some time to explore the program itself and its options; to see what you can change and what choices you have. In most cases, you will access the customization features by using a pull-down menu. Look for a menu or submenu named "Preferences", "Options", "Settings" or something similar.

My advice is to take some time and find out how your particular browser lets you modify your working environment. In general, it is safe to change any of the settings, except those which control your actual network connection.

HINT

As you learn how to use your browser, don't be afraid to experiment. What's the worst that could happen? If you really screw things up, you might need someone else to fix the settings for you, or perhaps even reinstall the software. So what? After all, if Humphrey Bogart had been afraid to experiment with his browser, where would he be today?

To get you started, I will discuss the most common ways to customize your working environment. Your browser will have other settings, but the ones I describe here are generic.

- TEXT AND COLORS: You can specify how text should be displayed (for example, size and typestyle), as long as your system allows such controls. You can also set the background and foreground colors used by your browser.
- LINKS: You can specify how links are to be displayed on your screen. You can change their color, underline them, and so on.
- HOME PAGE: Each time your browser starts, it can automatically load a particular page, called a home page. You can turn this feature off and on, as well as specify a URL for the page to be loaded. (Note: Don't confuse the use of the term "home page" in this context with a personal home page. See the discussion earlier in the chapter.)
- WINDOWS: You can change the size and shape of the various windows used to display information. To experiment, load a page with both text and images, and see what happens to each as you change the size of the window. In general, text will wrap automatically to fit in a newly sized window, while images will not change.
- MANIPULATING WINDOWS: Learn how your particular program and GUI (graphical user interface) let you open, close, iconize (shrink) and move windows. In addition, some browsers will let you open a new window whenever you want, so you can load more than one web page at the same time, each in its own window.
- IMAGE LOADING: You can control whether or not images are loaded automatically. If you turn off this feature, things will move a lot faster, as images can take a long time to load. However, you won't see the pictures unless you ask for them specifically.
- VIEWERS: Your browser already knows how to display certain types of data. However, when you encounter a different data type (image, sound, video), your browser will have to call upon another program, called a viewer, to interpret the data. You can find various types of viewers on the Net and install them on your computer. After you do, you must tell your browser about the new viewer and what data types it handles. (Note: Sometimes viewers are referred to as HELPER APPLICATIONS.)



INTERNET RESOURCE *Look in the catalog under Internet Resources for Web: Helper Applications*

- THE TOOLBAR: Some browsers have toolbars (collections of small icons) which allow you to access frequently used operations with a single mouse click. To perform an operation, you can click on one of the small icons, rather than selecting an item from a pull-down menu. If your browser uses toolbars, you will be able to control which toolbars or icons are displayed. Some people like to show all the toolbars all the time. Other people never use them at all.

- BOOKMARK LIST: Your bookmark list is a collection of your favorite links, which you can recall and use whenever you want. Your browser allows you to customize your bookmark list by adding, deleting and editing links. Some browsers even let you organize your bookmark links into folders and subfolders. (For more information about using your bookmark list, see the discussion earlier in this chapter.)
- THE CACHE: Whenever your browser loads a new page, it keeps a copy in a storage area called the CACHE. Using the cache makes it quick to return to a page you have already visited, as your browser can display the data immediately without having to reload it from a distant web server. The cache can be either in memory (which means it disappears when you quit the program) or stored on disk (which means the data is maintained from one work session to the next). If your browser can use both types of cache, you should choose which one you want to use. (A memory cache is faster; a disk cache is more permanent.) You can also control how long items are kept in the cache before they are deleted automatically. (When your browser deletes expired items, we say that it FLUSHES the cache.)



HINT

If your browser uses a disk cache, you may end up accumulating a lot of unused cache files. (Each web page and each external image is kept in a separate file.) For example, as I was writing this chapter, I checked my disk looking for old cache files. One of the browsers I use had created 534 files taking up 4,854,820 bytes of disk space. Another of my browsers had stored 266 files using 1,125,128 bytes. Needless to say, it was the work of a moment for me to delete these superfluous files and reclaim the disk space.

Some browsers allow you to select a menu item to flush the cache immediately. If you are using a disk cache, you should probably do this once in a while, just to get rid of old abandoned files. Alternatively, you can look in the disk directory in which your browser keeps the cache files, and delete all the old files manually. Of course, you should not do this unless you are comfortable with files and directories and know what you are doing.

Distributed Applications: Java

A web page contains information in the form of hypertext, and a web browser displays that information. Once the page is displayed, it doesn't really do anything, it just sits there. However, using special technology, it is possible for web page designers to construct programs that can be sent over the Web along with the regular information. When your browser receives such a program, it can run it while it

displays the hypertext. Being able to distribute web programs in this way adds a lot of functionality to the basic system.

For example, a web program might display graphics, or show animation, or perform calculations. In fact, just about anything that can be done with a program can be attached to a regular web page and distributed over the Net. The important thing about these programs is the flexibility they bring. The programs are not installed on your computer permanently. They are loaded and executed dynamically, whenever you select the web page to which they are linked.

Such programs are sometimes referred to as DISTRIBUTED APPLICATIONS ("application" being a synonym for "program"). Distributed applications are written in a computer language that is specifically designed for creating distributed programs. The language design must ensure that it is impossible to write a program that can cause damage on someone's computer. After all, when a program is sent over the Web and executed automatically, you don't have the same level of control as when you download a program and install it yourself.

To distinguish such programs from regular software, we sometimes refer to them as APPLETS. One of the important computer languages for writing web-based applets is named Java. If you are a programmer, you can think of Java as being a lot like C++ without the pointers.

To use a Java applet, all you need is a browser that is Java compatible. If you have such a browser, you do not have to do anything special. When you select a web page that contains a Java applet, your browser will load and execute the applet automatically along with the rest of the page. If you try to load the same web page with a browser that is not Java compatible, it will display the information in the regular manner and ignore the applet.

Chapter Twenty-Five

Talk Facilities

Meta's Exhibit No. 1008
Page 0056

Imagine a group of children in a classroom, working quietly under the supervision of a gimlet-eyed schoolteacher. What do they do when the teacher turns her back? They whisper furtively to one another and pass notes, hoping the teacher won't see.

Skip ahead twenty years and the same kids are now sitting in a conference room, listening to the Vice President of Marketing go on and on about the sales projections for the winter quarter. Look closely, and you will see one person quietly slip a piece of paper to his neighbor. She opens the paper to read: "He sure is boring, isn't he? He reminds me of my Uncle Arthur."

Like the child of twenty years ago, the man in the meeting is perfectly willing to do whatever is necessary to make personal contact with another person. When he shares his feelings and ideas he feels better. Moreover, he may be interested in establishing a more personal relationship with the woman sitting next to him.

Put human beings together anywhere and, in a short time, they will have figured out a way to talk to one another. So, as you might expect, as an Internet user, there are many ways for you to talk to someone else.

In this chapter, I discuss the ideas involved in talking over the Internet and describe some of the more common services. Two of the services are so important that I have given them chapters of their own: Internet Relay Chat (Chapter 27) and Muds (Chapter 28). In addition, in Chapter 26, I explain how to use `talk`, a Unix program for shell account users. However, before you read any of these chapters, do start with this one, so I can make sure you understand the basic concepts.

Overview of Talk Facilities

When we use the word TALK on the Net, we mean any type of communication in which you can send messages back and forth with someone in real time (that is, while you wait). The "talking" may involve text messages or actual voice transmission. For example, you might tell someone, "Last night I was talking to a friend in Australia," even if you and your friend were typing back and forth with no voice. On the Net, anytime you make contact with another person, you are talking.

(You will sometimes see people use the word CHAT. Usually this refers to talking to more than one person at a time, although there is no strict definition.)

The Internet has many different talk facilities available for free. The most basic service lets you type messages back and forth with another person. Whatever you type shows up on the other person's display, and whatever he or she types shows up on your display.

A typical setup will use a window divided into two areas. Your typing is displayed in one part; the other person's typing is displayed in the other part. As you talk, you can see both sides of the conversation at once. Since there are two display areas (one for each of you), you can both type at the same time without getting mixed up. An example of this is shown in Figure 25-1.

Many such systems allow you to talk to more than one person at the same time. When you do, your program will divide the window into as many talk areas as necessary. You can see this in Figure 25-2 which shows Bill Gates talking to four people at the same time.

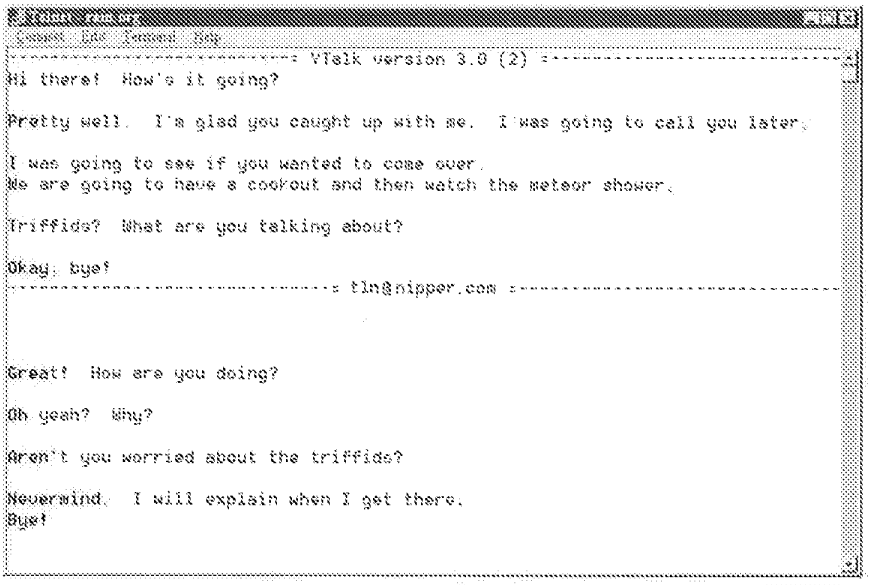


Figure 25-1. *A text-oriented talk program*

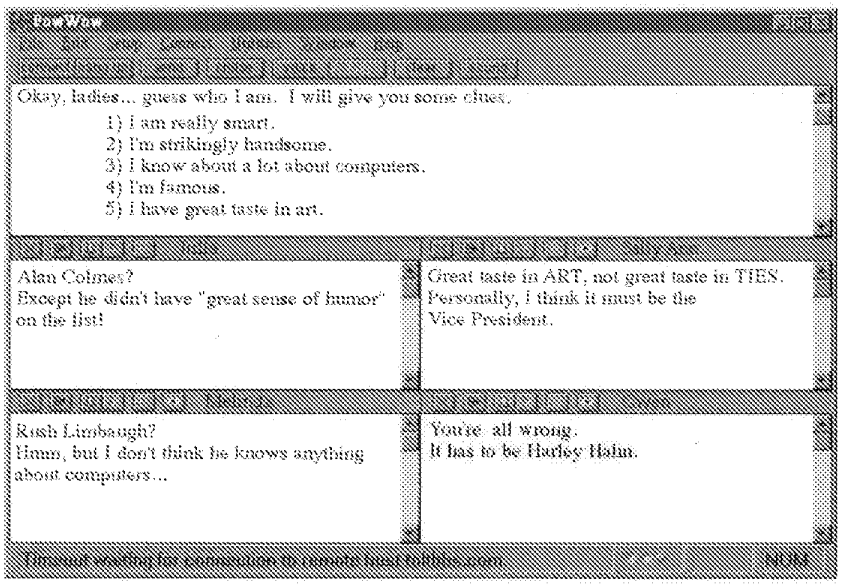


Figure 25-2. *A talk program with a multiple-person conversation*

A more complex talk facility will let you send and receive actual voice sounds, like a telephone. When you use a voice-oriented program over the Net, you can talk with people all over the world for free without paying a long distance charge. However, the quality may not be as good as with a regular telephone. Still, as a beautiful woman in California recently remarked to me, "I'm not going to pay a huge telephone bill just to talk to some stud in Germany for an hour."



HINT

To use a voice-oriented talk facility, your computer must have the appropriate hardware. It must be able to play sounds, and it must have a microphone.

These services—text and voice—are the two most important ones you will find. Some talk facilities have one or the other, some have both. In addition, there are a number of other features you may find.

Chat Rooms

A CHAT ROOM is a facility that allows more than one person to talk at the same time. Chat rooms simulate an environment in which people gather together to talk. Although chat rooms are imaginary, it is common to talk about them as if they really exist. When someone joins the conversation, we might say he or she "enters" the room. Similarly, a talk facility may have several chat rooms, each with different names. So you might, for example, enter the "Zen room". Once you do, you can see all the conversations going on in that room.

Public Events

Some talk facilities have special events in which anyone can gather in a particular chat room and listen to a guest speaker. (That room might be called the "Auditorium".) For example, there might be a scheduled time at which Bill Gates joins a particular chat room to give a speech on "Windows and Personal Enlightenment". During the speech (which Bill would type on his computer), everybody in the room could read what he was saying. Afterwards, there might be questions and answers. As people asked questions and Bill answered them, everyone could see what was happening.

Privacy

You can control whether or not a conversation is private or whether other people can join. If you are participating in a group conversation, you can initiate a temporary, private connection with one of the people (and tell secrets about everyone else). Generally speaking, if you are in a chat room, your conversation is not private.

White Pages Directory

Some talk systems keep a public directory in which you can look up someone's name (like a telephone book without the phone numbers). Once you find a name, you can send a request to that person to see if he or she wants to talk to you. Aside from the name, you may also see an email address, a link to the person's home page and other information.

A Personal Address Book

Most talk programs allow you to create your own personal address book. You can save the names and addresses of your friends on the Net, making changes whenever you want. To initiate a talk connection, all you have to do is display your address book and select the person to whom you want to talk.

Embedded HTML

HTML (Hypertext Markup Language) is the descriptive language used to build web pages (see Chapter 11). Some web-based talk facilities let you type HTML instructions. When these instructions are "read" by someone's browser, they will be interpreted as valid HTML. For example, as you are talking, you can send a link to your home page (or any web page), display buttons, make words blink, and so on. Of course, to do this, you have to know a little HTML. (It won't be long, though, until HTML is taught in schools as a second language.)

Sounds

You can send a sound effect to the person to whom you are talking. For example, you might send a laugh, a sigh, applause, and so on. Aside from pre-installed sounds, you can use sounds of your own. The other day, I was talking to someone in Sweden and he played a sound of a California surfer dude saying, "Okay, so I'll see ya later, huh?" I liked what I heard, so my Swedish friend used the file transfer facility to send me the actual sound file. I then registered the sound with my own talk program, so I could use it myself.

Pictures

You can use a picture to identify yourself. When people talk to you, they will see the picture. Some people use a photo of themselves, although this is not necessary: the picture can be any GIF or JPEG file. Some talk facilities have a library of pictures for you to choose from, if you don't have one of your own.

Video

There is a system called CU-SeeMe that allows you to send video images to another person on the Net. The video shows motion, although it is not of high quality like a television.

File Transfer

You can send a file to someone while you are talking without having to stop talking or without using a separate file transfer or mail program.

Answering Machine

If you are busy and do not wish to be disturbed, you can have your talk program send a specific message to people when they try to talk to you. When you want to start talking again, you can turn off the answering machine and take the calls yourself.

Remote Control

While you are talking to someone, you can control a program that is running on his or her computer. This will work for more than one person. For example, you could take control of someone's web browser and show them specific web pages. Or one person might take a group of people on a guided tour of the Web.

How Does It All Work?

Like all Internet services, talking on the Internet is based on a client/server system. There are two ways in which such systems are organized.

One design uses a central talk server, somewhere on the Net. Each person who uses the service runs a talk client on his own computer. The individual clients communicate by sending messages to and from the server.

For example, say you are using a centralized system to talk to a friend named Norman Sweetcream. Each time you type a line of text, your client program sends it to the server. The server relays the text to Norman's client program, which displays your message on his screen. When Norman types a line of text, his client sends it to the central server, which then sends the text to your client to display. In other words, everything you and Norman type goes through the server.

The other basic design connects people directly without using a central server. When you talk with Norman, everything you type goes right to his computer. When he responds, his messages go directly to your computer. With such systems, the program you run on your computer must be more than a talk client: It must act as both a client *and* a server.

There are advantages and disadvantages with both systems. The most obvious advantage of a decentralized system is there are no middlemen to slow down the talking. When everything has to go through a central server, things bog down and there is often a noticeable delay in the conversation.

In addition, with a decentralized system, no one has to maintain a central server. After all, why should someone (or some organization) have to pay for the hardware, software and administration time to support a talk server, just so people on the Internet can talk for free? Moreover, when someone *does* offer such a service to the Net, you must play by his or her rules or you are out.

Finally, a decentralized system is more dependable. When a central server goes down, everything stops. Such servers are attractive targets for hackers: talented programmers, who love to see how much mischief they can cause. From time to time, you will see a notice that the administrator has removed the "such and such feature" from a central server, because hackers were using it to cause trouble.

Centralized systems do have two big advantages, however. First, because the talk server is centralized, it can be more elaborate and offer more services than a server program you would run on your own computer. For example, a centralized server might create an imaginary environment (such as a mud) which the users can explore as they talk.

Second, a centralized system can provide a place for people to gather and to find one another. When you use IRC (Internet Relay Chat), for example, you will find many groups of people talking about different topics. Similarly, many people like to use muds because it allows them to meet new people and to hang out with their friends. When you use a decentralized system, there is no easy way to join a group of people or to just hang out.

What Do You Need to Use a Talk Facility?

To use an Internet talk facility, you need several things. First, your Internet connection should be either PPP or a direct network connection. If you use a shell account you will not be able to access some of the talk facilities. You will be able to use the Unix `talk` program (see Chapter 26), IRC (Chapter 27) and Muds (Chapter 28), but not much else. (For more information about Internet connections, see Chapter 5.)

Second, your Internet connection should be as fast as possible, especially if you want to use a voice facility. Thus, do not buy a modem that is slower than 28,800 bits per second. ISDN is even better. (See Chapter 4.)

Third, your computer should be capable of playing sounds. With a Macintosh, this capability is built in. With a PC, you must have a sound card or the equivalent. In addition, if you want to talk using voice, your computer must have a microphone.

Finally, you need the appropriate software. For all the talk systems I describe in this book, the software is available for free on the Net.

Types of Internet Talk Facilities

In this section, I will take you on a tour of some of the talk facilities you will encounter on the Internet. All of these are fun and useful. When you get a chance, experiment and see which talk facilities you like best.

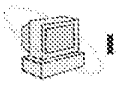
The Unix `talk` Program

The oldest of the Internet talk facilities is a program named **`talk`** that runs under Unix. The **`talk`** program provides a basic, text-based talk facility. You can type messages back and forth to another person, and all the conversations are private.

The **`talk`** program allows you to talk with anyone else on the Internet, as long as they have a **`talk`** program. Virtually every Unix computer on the Internet has a **`talk`** program, so if you have a shell account, you can use **`talk`**. If you do not have a shell account (that is, if you use PPP or a direct network connection), you need to find a client for your own computer that is compatible with Unix **`talk`**.

When you use **`talk`** with a shell account, both the client and server run on the Unix computer. Thus, when you use a compatible program on your PC or Macintosh, the program must act as both a client and a server. Such programs are available for free on the Net.

The **`talk`** program is an important Internet service, and I discuss it in detail in Chapter 26.

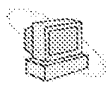


INTERNET RESOURCE *Look in the catalog under Internet Resources for Talking: **`talk`** Clients*

Internet Relay Chat (IRC)

IRC is a huge Internet-wide system which supports a large number of group conversations. It is text-based and uses a number of central talk servers around the Net to coordinate the flow of data. IRC is a lot of fun and is a great place to meet people. To use IRC you need an IRC client for your computer. There are clients for all types of computers, including Unix, so you can use IRC with a shell account.

I discuss IRC in detail in Chapter 27.



INTERNET RESOURCE *Look in the catalog under Internet Resources for Talking: IRC Clients*

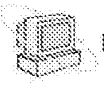
Muds

A mud is an imaginary environment in which you can talk to people and meet new friends. Moreover, some muds are set up as fantasy worlds, where you travel around and have adventures while assuming the role of an imaginary character. Many such muds are complex and require a long time to master.

You can access a mud by using a plain telnet client (see Chapter 23). However, there are specialized mud clients available as well. If you get into serious mudding, get yourself a serious mud client.

Muds can be a lot of fun and I encourage you to try a few. Many people tend to think of the Internet as being mostly the Web, but that is not true at all. If you haven't experienced a mud, you are missing out on one of life's most interesting experiences.

I discuss muds in detail in Chapter 28.

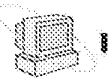


INTERNET RESOURCE *Look in the catalog under Internet Resources for*
Mud Clients
Telnet Clients

Web-Based Chat Rooms

The Web is huge, and there are a lot of talk facilities that make use of its size and popularity. To use these facilities you do not need a special talk client. All you need to do is connect to a centralized talk server, and your web browser will act as your talk client.

There are a lot of chat rooms on the Web, and they are worth checking out. However, there are some special considerations you should understand, which I discuss in detail in the next section.



INTERNET RESOURCE *Look in the catalog under Internet Resources for*
Talking: Web-based Talk

Bulletin Board Systems (BBSs)

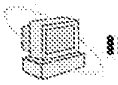
Before the Internet became popular, there were a large number of Bulletin Board Systems (BBSs) that offered dial-up online services. Today, some of these BBSs, along with many new ones, have moved to the Internet. Many BBSs are run by an individual, although some are large enough to require a group of administrators. In general, the person who runs a BBS is called the SYSOP, and the personality of an individual BBS very much depends on the sysop.

Most BBSs allow you to send messages to people who are logged in. This is a good way to talk with people and meet new friends. There are usually ongoing discussion groups to which you can post messages as well as read the messages left by others. In addition, all BBSs have archives from which you can download files or to which you can contribute files of your own.

There are general-purpose BBSs, as well as BBSs devoted to just about any topic you can imagine: meeting people, religion, hobbies, corporate and government information, X-rated talk and resources, and so on.

To access a BBS over the Internet, you use telnet (see Chapter 23). Before you can use most BBSs, you will have to register and ask for an account. Many BBSs are free,

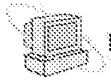
although some of them do charge a membership fee (especially the ones with the X-rated resources).



INTERNET RESOURCE *Look in the catalog under Internet Resources for Talking: Bulletin Board Systems (BBSs)*

Talkers

A TALKER is a multiuser talk facility that is easy to use and is devoted primarily to conversation. You connect to a talker by using telnet (see Chapter 23). Once you do, you can talk to anyone else who happens to be connected at the same time. The word "talker" is a descriptive term. For example, some talkers are muds, while others are BBSs; there is no strict definition. If it's easy to use, and you connect in order to talk to other people, it is a talker.



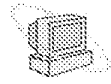
INTERNET RESOURCE *Look in the catalog under Internet Resources for Talking: Talkers*

Direct Talk Connections

Aside from the system based on the Unix `talk` program, there are other facilities that connect two people with a direct connection. Since such systems do not depend on a central server, the communication is fast and private. Some programs allow you to make connections between more than one person at a time to have a group conversation. The other people can be in the next room or on the other side of the world.

One program I particularly like is named Powwow. It not only supports text but, if you have the right hardware, voice as well. I use Powwow to send messages to my chief researcher, even though she often works in a room just down the hall. Such systems are great when you are participating in a telephone conference call. You can type secret messages back and forth to one of the people on the call without everyone else seeing what you are saying.

The Powwow system has a nice feature that is common to other talk facilities. There is a white pages server to which you can connect to find lists of people who use the same program. You can then choose someone on the list and see if they want to talk with you. In addition, while you are talking, one person can take control of the other people's web browsers (with permission, of course) and guide the other people on a tour of the Web.



INTERNET RESOURCE *Look in the catalog under Internet Resources for Talking: Direct Connections*

Telephone-Like Systems

There are a number of voice-only talk facilities that are designed to work like a telephone. In order to use the system, you must have a PPP or a direct network connection (not a shell account). You run a special program on your computer, which acts as both a client and server. Using this program, you can talk to anyone else on the Net who is running the same program. (Of course, your computer needs a microphone and sound capabilities.)

How do you find someone to talk to? There are two ways. First, you can specify the exact address of someone's computer. If the computer name in the mail address doesn't work, you can always use a numeric IP address (see Chapter 6). However, you must know this address in advance, and you must be sure the person to whom you want to talk has his telephone program running, waiting for a connection.

Second, some systems have a number of special servers on the Net (maintained by the company that sells the telephone program). You can connect to a server and tell it you are ready to talk. At the same time, you can see all the other people who are ready to talk. To talk with someone, you simply select their name from the list on the server. The server will send a signal to that person. If he agrees to talk with you, a direct connection will be made between your computer and his. From that point on, you are not using the server (and everything you say is private).

Of course, since all of this happens on the Net, you do not have to pay long distance phone charges. So why doesn't everyone use the Net instead of a telephone? There are two reasons. First, the voice quality is not as good as a real telephone. Second, you can't call someone unless he has an Internet connection, a computer with a microphone and sound capability, and the same telephone program.

However, if you like talking to people for fun, try one of these systems. In most cases, you can use the software for free in a limited way. (For example, your conversations may be limited to one minute.) If you have someone to whom you would like to talk regularly for a long time, you can pay a small fee to the company and upgrade your program, so it allows you to talk as long as you want.



INTERNET RESOURCE *Look in the catalog under Internet Resources for Talking: Telephone-like Systems*

Video Conferencing

A video conferencing system allows you to send live video and sound over the Internet. The most important such system is called CU-SeeMe. (The name is a pun. The "CU" stands for Cornell University where the original program development was done. However, when you say the name out loud, it sounds like "see you, see me".)

CU-SeeMe will run on either a Macintosh or a PC using Windows as long as you have a PPP or direct network connection (not a shell account). To watch video, all you need is the software, which is available for free on the Net. To broadcast, you need a video camera, a digitizing board and the appropriate video software.

When you watch, you will see actual video. (See Figure 25-3.) However, the quality is not all that good. (It looks like the video broadcasts from the space shuttle.) Still, it is exciting to watch people from all over the world. Although you need special equipment to send video, it is not that expensive, and there are a surprising number of people broadcasting over the Net.

By itself, the CU-SeeMe software will enable you to establish a video link between any two computers. However, by connecting to a "reflector" you can watch or broadcast to more than one site at a time.

One word of warning: A lot of schools use CU-SeeMe to let the kids talk to one another. Please remember, though, there is no censorship on the Internet. While I was working on this chapter, I saw more than one site broadcasting naked people performing quasi-sexual acts. Depending on your outlook, this could be either erotic or immoral. What I want you to appreciate is that the Internet is an open forum, and the only effective way to censor something is to control what comes into your own computer.

Web-Based Talk Facilities

The Web was not designed to allow individual users to communicate, so the talk facilities had to be fitted into the existing web client/server system. Because of this, you will often find web-based talking to be less smooth than other systems. When you



Figure 25-3. *Live video over the net, using CU-SeeMe*

talk with someone, you would like to see the words as the person is typing. With a web-based talk facility, this is not possible. Here is why.

Within the Web, you use a talk facility by connecting to a specific web page that represents a chat room. Whenever you load that page, you will see whatever messages people have typed in the last few minutes. An example of this system is shown in Figure 25-4.

How do you send a message of your own? On the Web, there is only one way to type information and send it to a server. The server must send you a web page that contains a special input area called a "form" (see Chapter 10). You type information into the form and then submit it, either by pressing the ENTER key or by clicking on a button. (Notice the input form in Figure 25-4.)

This means the information you type is not sent to the server until you submit the form. At that time, everything you typed into the form is sent all at once. The talk server receives your input in chunks.

Similarly, there is no way for a web server to send you information one character or one word at a time. All the server can do is wait for you to reload a web page and then send you updated information.

When you put this all together, it explains why talking on the Web is not a smooth process. Say you are talking to a friend. You must type your message into a form, and nothing will be sent to the server until you submit the form. At the other end, your

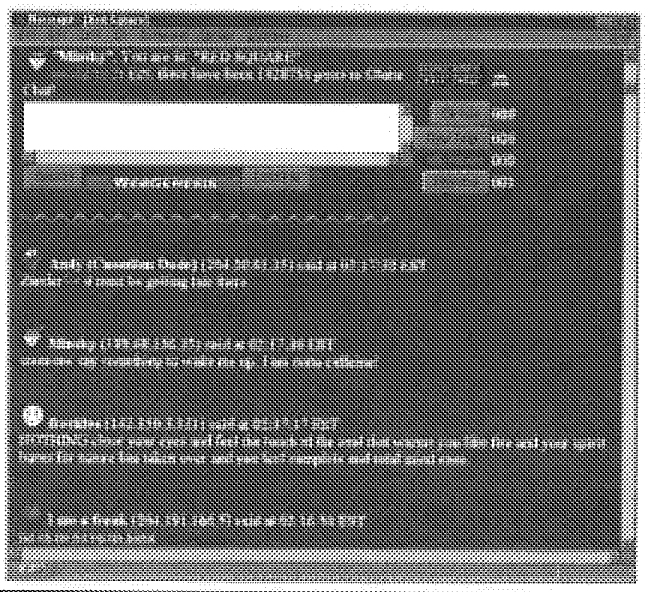


Figure 25-4. *A typical web-based chat room*

friend will not see anything until he or she reloads the web page. It is only at then that the server will be able to send any new information.

(Some talk systems mitigate this problem by using server push [see Chapter 10]. This allows the web server to update your web page automatically at regular intervals. With such systems, you can sometimes specify the interval, so you can control the time between updates.)

Thus, web-based talking is not smooth for several reasons. First, all messages must pass through a central server. Second, a message is sent to the server as one large chunk, not one character at a time. Third, you will not see any new messages until you reload the chat room web page. Finally, everyone in the room can talk at the same time. Thus, each time you reload your web page, you will see *all* the messages that have been sent by everyone in the room, since the last time you reloaded the page. As you might imagine, this can be a tad confusing.

Still, after a while, you will get used to web-based talking. It is only a matter of waiting for certain cells in your brain to change appropriately. Moreover, many of the servers are quite elaborate, offering an interesting variety of talk-related resources such as specialized chat rooms, public events, and so on. So, in spite of all the limitations, web-based talk facilities are worth investigating.



HINT

Most web-based chat systems place the newer messages at the top of the web page. Thus, to look at the messages in chronological order, you must read from the bottom up.

Guidelines for Text-based Talking

Having a text-based conversation on the Internet combines the immediacy of a face-to-face encounter with the verbosity of written language. For this reason, you will notice two significant differences between talking by computer and talking by mouth.

First, talking by computer is a lot slower, because almost everybody can speak faster than they can type (not that this is always an advantage). Second, because you are not looking at the other person, the customary nuances and inflections that comprise so much of the non-verbal aspects of communication are missing.

We make up for these limitations in two ways. First, there are a number of abbreviations that are commonly used to speed things up. Some of these are shown in Figure 25-5. These abbreviations are the spiritual descendants of the Morse-code conventions used by ham radio operators since the 1920s.

The second important convention recognizes that it is easy to offend someone (or at least have your words misconstrued) when you are not talking in person. Over the computer, we are unable to use subtle body language to indicate that what we just said was debonair and ironic, and not really as dumb and insensitive as it looks.

Abbreviation	Meaning
BBL	be back later
BCNU	be seeing you
BRB	be right back
BTW	by the way
BYE	goodbye (I am ready to stop talking)
BYE?	goodbye? (are you ready to stop talking?)
CU	see you (goodbye)
CUL	see you later
CUL8R	see you later
F2F	face to face
FWIW	for what it's worth
IMHO	in my humble opinion
JAM	just a minute
L8R	later
LOL	laugh out loud
OBTW	oh, by the way
R U THERE?	are you there?
ROTF	rolling on the floor
ROTFL	rolling on the floor laughing
SEC...	wait a second
TTFN	ta-ta for now
TTYL	talk to you later
WRT	with respect to

Figure 23-3. *Common communication abbreviations*

Thus, when there is any doubt as to whether or not the other person might be offended, it is a good idea to use a smiley to show that you are just kidding. A smiley is a little face that we draw with punctuation characters. Here is the basic smiley:

```
:-)
```

(To see the face, tilt your head to the left.) You can use a smiley whenever you say something that might bother the other person, for example:

```
I'd feel the same way as you, if I had a
head that looked like a ripe eggplant :-)
```

If you want to see a few more smileys, look at Figure 14-5 in Chapter 14.

It is also important to become aware of what is private and what is not. With some talk systems, it is possible for people to sit quietly and watch what other people are saying. Such people are called LURKERS. Just because you see only one person talking does not mean that no one else is watching.



HINT

When in doubt, assume your conversation is not private.

Some talk programs have a way for you to talk with a group, but break off temporarily to have a private conversation with just one person. If you have never used this facility, be sure to test it. It is a good idea to be completely sure that you understand how the privacy option works before you say anything you will regret.

A further consideration has to do with how the words you type are sent to the other person. With some systems, what you type is not sent until you press the ENTER key or click on a button. In such cases, you can backspace to make corrections before you send the text. On other systems, the message is sent, one character at a time, as you type it. (This is the case with the Unix `talk` program.) Thus, if you accidentally say something insulting, by the time you backspace and change it, the other person will already have seen it.

If the network or the talk server slows down, you may notice replies are delayed and become unsynchronized, leading to various amusing and incoherent situations. In such cases, we refer to the delay as LAG. Thus, you might see someone say, "There was so much lag, I thought you had forgotten all about me."



HINT

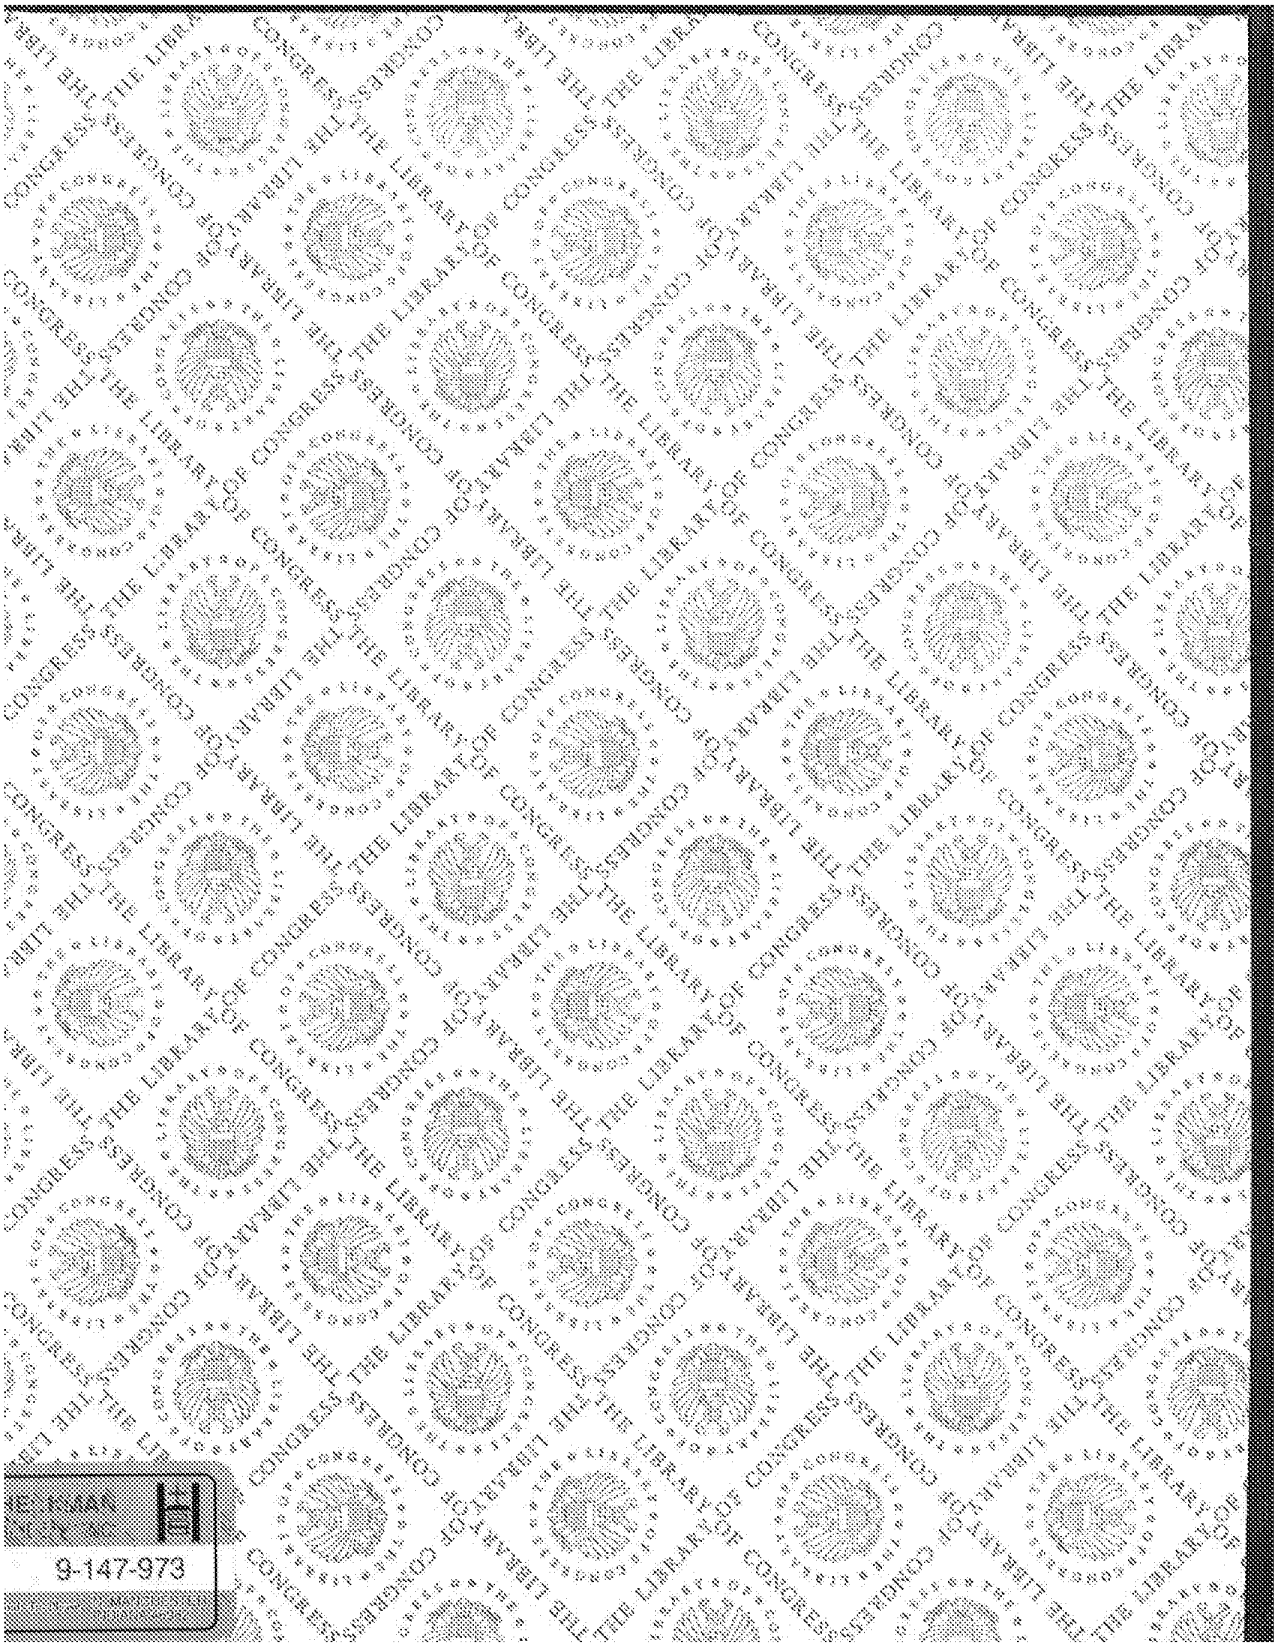
For possibly the first time in your life, spelling does not count.

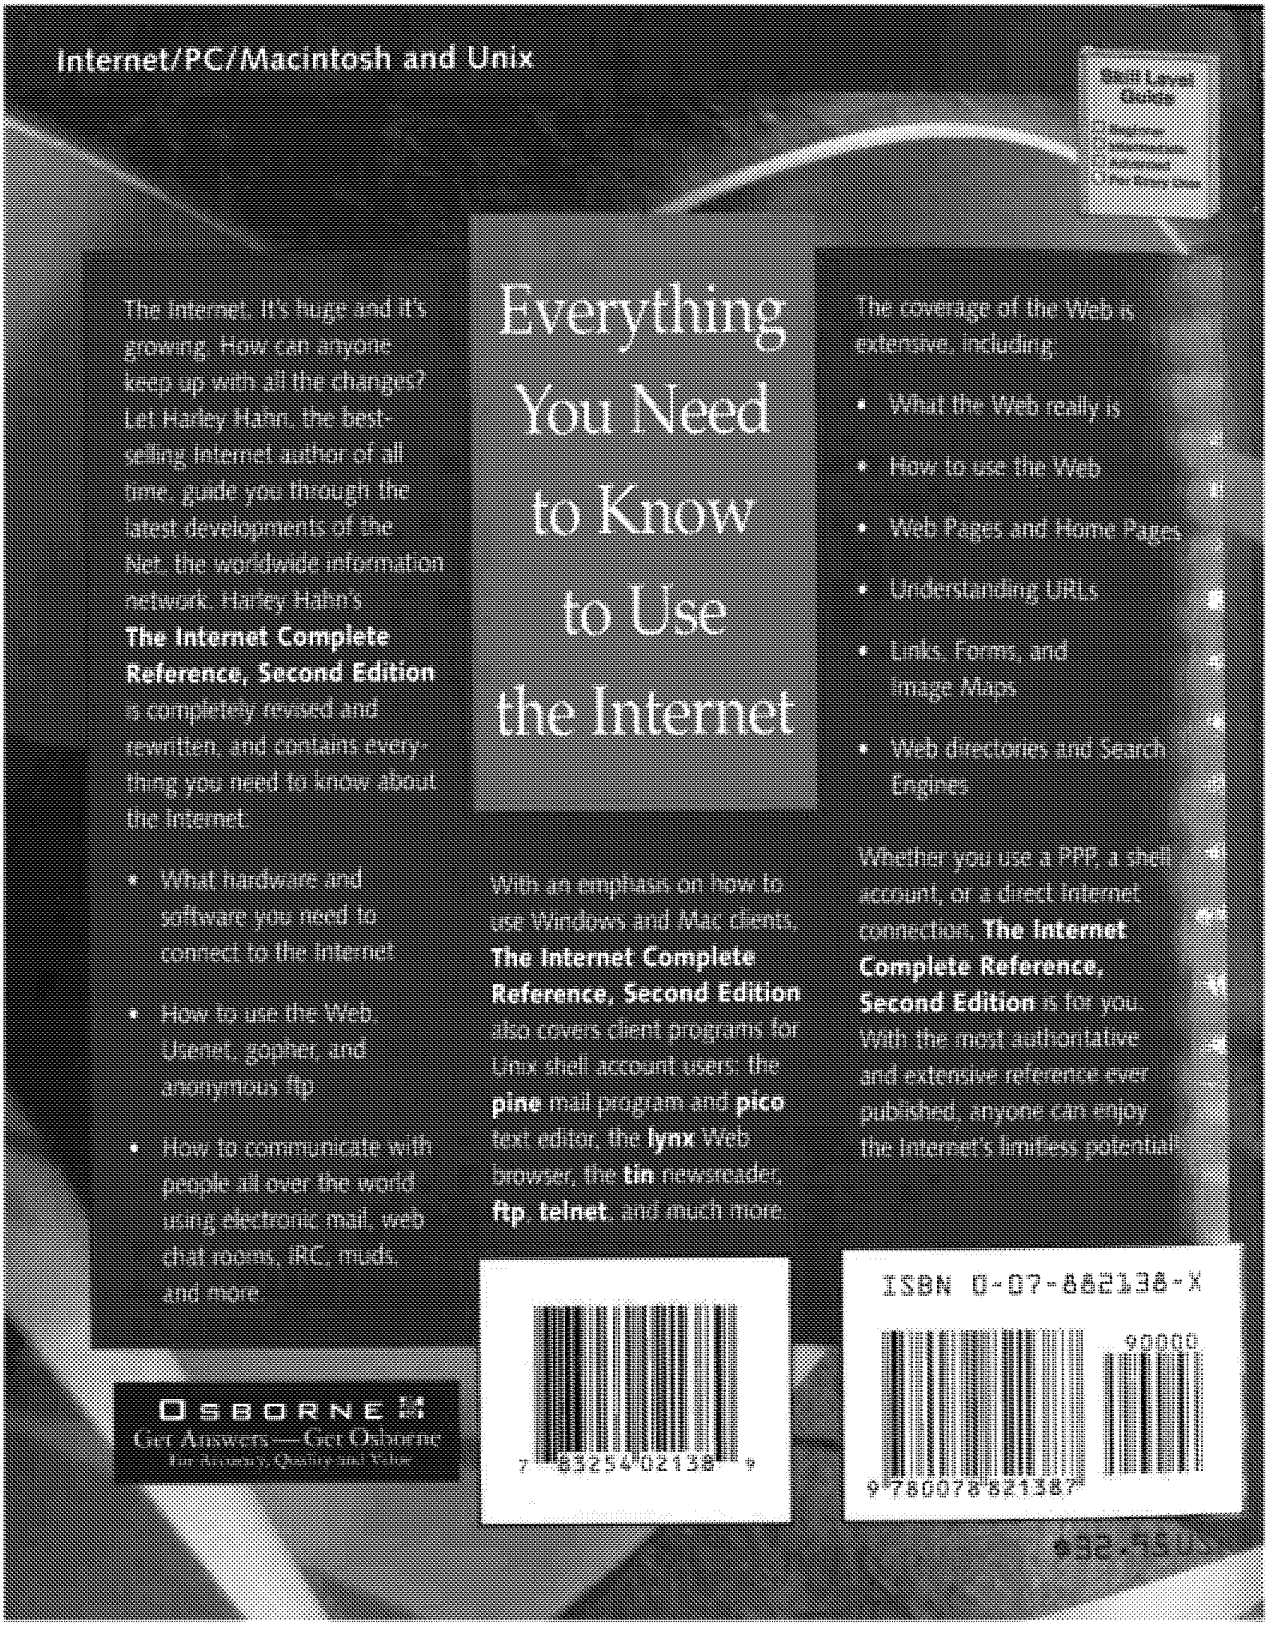
Because typing is so much slower than talking, it is perfectly acceptable to type as fast as you can and not worry about spelling mistakes. Nothing is worse than watching someone laboriously backspacing and correcting each word. Just go ahead and type and don't worry. Almost evrWything is underftandabll with kust a few wurng leitrs and, no doubt,m the other pweosn will get the gist of theta you are svaing.

Chapter Twenty-Six

Using Talk from a Shell Account: talk, ytalk




9-147-973


Internet/PC/Macintosh and Unix
Everything You Need to Know to Use the Internet
The Internet. It's huge and it's growing. How can anyone keep up with all the changes? Let Harley Hahn, the best-selling Internet author of all time, guide you through the latest developments of the Net, the worldwide information network. Harley Hahn's **The Internet Complete Reference, Second Edition** is completely revised and rewritten, and contains everything you need to know about the Internet.
• What hardware and software you need to connect to the Internet
• How to use the Web, Usenet, gopher, and anonymous ftp
• How to communicate with people all over the world using electronic mail, web chat rooms, IRC, muds, and more
With an emphasis on how to use Windows and Mac clients, **The Internet Complete Reference, Second Edition** also covers client programs for Unix shell account users: the **pine** mail program and **pico** text editor, the **lynx** Web browser, the **tin** newsreader, **ftp**, **telnet**, and much more.
The coverage of the Web is extensive, including:
• What the Web really is
• How to use the Web
• Web Pages and Home Pages
• Understanding URLs
• Links, Forms, and Image Maps
• Web directories and Search Engines
Whether you use a PPP, a shell account, or a direct Internet connection, **The Internet Complete Reference, Second Edition** is for you. With the most authoritative and extensive reference ever published, anyone can enjoy the Internet's limitless potential.
OSBORNE
Get Answers—Get Osborne
ISBN 0-07-882138-X
90000
$32.95 U.S.

Network Working Group
Request for Comments: 2663
Category: Informational

P. Srisuresh
M. Holdrege
Lucent Technologies
August 1999

IP Network Address Translator (NAT) Terminology and Considerations

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.

Preface

The motivation behind this document is to provide clarity to the terms used in conjunction with Network Address Translators. The term "Network Address Translator" means different things in different contexts. The intent of this document is to define the various flavors of NAT and standardize the meaning of terms used.

The authors listed are editors for this document and owe the content to contributions from members of the working group. Large chunks of the document titled, "IP Network Address Translator (NAT)" were extracted almost as is, to form the initial basis for this document. The editors would like to thank the authors Pyda Srisuresh and Kjeld Egevang for the same. The editors would like to thank Praveen Akkiraju for his contributions in describing NAT deployment scenarios. The editors would also like to thank the IESG members Scott Bradner, Vern Paxson and Thomas Narten for their detailed review of the document and adding clarity to the text.

Abstract

Network Address Translation is a method by which IP addresses are mapped from one realm to another, in an attempt to provide transparent routing to hosts. Traditionally, NAT devices are used to connect an isolated address realm with private unregistered addresses to an external realm with globally unique registered addresses. This document attempts to describe the operation of NAT devices and the associated considerations in general, and to define the terminology used to identify various flavors of NAT.

Meta's Exhibit No. 1013
Page 001

Page 330 of 1032

1. Introduction and Overview

The need for IP Address translation arises when a network's internal IP addresses cannot be used outside the network either because they are invalid for use outside, or because the internal addressing must be kept private from the external network.

Address translation allows (in many cases, except as noted in sections 8 and 9) hosts in a private network to transparently communicate with destinations on an external network and vice versa. There are a variety of flavors of NAT and terms to match them. This document attempts to define the terminology used and to identify various flavors of NAT. The document also attempts to describe other considerations applicable to NAT devices in general.

Note, however, this document is not intended to describe the operations of individual NAT variations or the applicability of NAT devices.

NAT devices attempt to provide a transparent routing solution to end hosts trying to communicate from disparate address realms. This is achieved by modifying end node addresses en-route and maintaining state for these updates so that datagrams pertaining to a session are routed to the right end-node in either realm. This solution only works when the applications do not use the IP addresses as part of the protocol itself. For example, identifying endpoints using DNS names rather than addresses makes applications less dependent of the actual addresses that NAT chooses and avoids the need to also translate payload contents when NAT changes an IP address.

The NAT function cannot by itself support all applications transparently and often must co-exist with application level gateways (ALGs) for this reason. People looking to deploy NAT based solutions need to determine their application requirements first and assess the NAT extensions (i.e., ALGs) necessary to provide application transparency for their environment.

IPsec techniques which are intended to preserve the Endpoint addresses of an IP packet will not work with NAT enroute for most applications in practice. Techniques such as AH and ESP protect the contents of the IP headers (including the source and destination addresses) from modification. Yet, NAT's fundamental role is to alter the addresses in the IP header of a packet.

2. Terminology and concepts used

Terms most frequently used in the context of NAT are defined here for reference.

Meta's Exhibit No. 1013
Page 002

2.1. Address realm or realm

An address realm is a network domain in which the network addresses are uniquely assigned to entities such that datagrams can be routed to them. Routing protocols used within the network domain are responsible for finding routes to entities given their network addresses. Note that this document is limited to describing NAT in IPv4 environment and does not address the use of NAT in other types of environment. (e.g. IPv6 environments)

2.2. Transparent routing

The term "transparent routing" is used throughout the document to identify the routing functionality that a NAT device provides. This is different from the routing functionality provided by a traditional router device in that a traditional router routes packets within a single address realm.

Transparent routing refers to routing a datagram between disparate address realms, by modifying address contents in the IP header to be valid in the address realm into which the datagram is routed. Section 3.2 has a detailed description of transparent routing.

2.3. Session flow vs. Packet flow

Connection or session flows are different from packet flows. A session flow indicates the direction in which the session was initiated with reference to a network interface. Packet flow is the direction in which the packet has traveled with reference to a network interface. Take for example, an outbound telnet session. The telnet session consists of packet flows in both inbound and outbound directions. Outbound telnet packets carry terminal keystrokes and inbound telnet packets carry screen displays from the telnet server.

For purposes of discussion in this document, a session is defined as the set of traffic that is managed as a unit for translation. TCP/UDP sessions are uniquely identified by the tuple of (source IP address, source TCP/UDP port, target IP address, target TCP/UDP port). ICMP query sessions are identified by the tuple of (source IP address, ICMP query ID, target IP address). All other sessions are characterized by the tuple of (source IP address, target IP address, IP protocol).

Address translations performed by NAT are session based and would include translation of incoming as well as outgoing packets belonging to that session. Session direction is identified by the direction of the first packet of that session (see sec 2.5).

Meta's Exhibit No. 1013
Page 003

Note, there is no guarantee that the idea of a session, determined as above by NAT, will coincide with the application's idea of a session. An application might view a bundle of sessions (as viewed by NAT) as a single session and might not even view its communication with its peers as a session. Not all applications are guaranteed to work across realms, even with an ALG (defined below in section 2.9) enroute.

2.4. TU ports, Server ports, Client ports

For the reminder of this document, we will refer TCP/UDP ports associated with an IP address simply as "TU ports".

For most TCP/IP hosts, TU port range 0-1023 is used by servers listening for incoming connections. Clients trying to initiate a connection typically select a source TU port in the range of 1024-65535. However, this convention is not universal and not always followed. Some client stations initiate connections using a source TU port number in the range of 0-1023, and there are servers listening on TU port numbers in the range of 1024-65535.

A list of assigned TU port services may be found in RFC 1700 [Ref 2].

2.5. Start of session for TCP, UDP and others

The first packet of every TCP session tries to establish a session and contains connection startup information. The first packet of a TCP session may be recognized by the presence of SYN bit and absence of ACK bit in the TCP flags. All TCP packets, with the exception of the first packet, must have the ACK bit set.

However, there is no deterministic way of recognizing the start of a UDP based session or any non-TCP session. A heuristic approach would be to assume the first packet with hitherto non-existent session parameters (as defined in section 2.3) as constituting the start of new session.

2.6. End of session for TCP, UDP and others

The end of a TCP session is detected when FIN is acknowledged by both halves of the session or when either half receives a segment with the RST bit in TCP flags field. However, because it is impossible for a NAT device to know whether the packets it sees will actually be delivered to the destination (they may be dropped between the NAT device and the destination), the NAT device cannot safely assume that the segments containing FINs or SYNs will be the last packets of the session (i.e., there could be retransmissions). Consequently, a session can be assumed to have been terminated only after a period of

Meta's Exhibit No. 1013
Page 004

4 minutes subsequent to this detection. The need for this extended wait period is described in RFC 793 [Ref 7], which suggests a TIME-WAIT duration of 2 * MSL (Maximum Segment Lifetime) or 4 minutes.

Note that it is also possible for a TCP connection to terminate without the NAT device becoming aware of the event (e.g., in the case where one or both peers reboot). Consequently, garbage collection is necessary on NAT devices to clean up unused state about TCP sessions that no longer exist. However, it is not possible in the general case to distinguish between connections that have been idle for an extended period of time from those that no longer exist. In the case of UDP-based sessions, there is no single way to determine when a session ends, since UDP-based protocols are application specific.

Many heuristic approaches are used to terminate sessions. You can make the assumption that TCP sessions that have not been used for say, 24 hours, and non-TCP sessions that have not been used for a couple of minutes, are terminated. Often this assumption works, but sometimes it doesn't. These idle period session timeouts vary a great deal both from application to application and for different sessions of the same application. Consequently, session timeouts must be configurable. Even so, there is no guarantee that a satisfactory value can be found. Further, as stated in section 2.3, there is no guarantee that NAT's view of session termination will coincide with that of the application.

Another way to handle session terminations is to timestamp entries and keep them as long as possible and retire the longest idle session when it becomes necessary.

2.7. Public/Global/External network

A Global or Public Network is an address realm with unique network addresses assigned by Internet Assigned Numbers Authority (IANA) or an equivalent address registry. This network is also referred as External network during NAT discussions.

2.8. Private/Local network

A private network is an address realm independent of external network addresses. Private network may also be referred alternately as Local Network. Transparent routing between hosts in private realm and external realm is facilitated by a NAT router.

RFC 1918 [Ref 1] has recommendations on address space allocation for private networks. Internet Assigned Numbers Authority (IANA) has three blocks of IP address space, namely 10/8, 172.16/12, and 192.168/16 set aside for private internets. In pre-CIDR notation, the

Meta's Exhibit No. 1013
Page 005

first block is nothing but a single class A network number, while the second block is a set of 16 contiguous class B networks, and the third block is a set of 256 contiguous class C networks.

An organization that decides to use IP addresses in the address space defined above can do so without coordination with IANA or any other Internet registry such as APNIC, RIPE and ARIN. The address space can thus be used privately by many independent organizations at the same time. However, if those independent organizations later decide they wish to communicate with each other or the public Internet, they will either have to renumber their networks or enable NAT on their border routers.

2.9. Application Level gateway (ALG)

Not all applications lend themselves easily to translation by NAT devices; especially those that include IP addresses and TCP/UDP ports in the payload. Application Level Gateways (ALGs) are application specific translation agents that allow an application on a host in one address realm to connect to its counterpart running on a host in different realm transparently. An ALG may interact with NAT to set up state, use NAT state information, modify application specific payload and perform whatever else is necessary to get the application running across disparate address realms.

ALGs may not always utilize NAT state information. They may glean application payload and simply notify NAT to add additional state information in some cases. ALGs are similar to Proxies, in that, both ALGs and proxies facilitate Application specific communication between clients and servers. Proxies use a special protocol to communicate with proxy clients and relay client data to servers and vice versa. Unlike Proxies, ALGs do not use a special protocol to communicate with application clients and do not require changes to application clients.

3. What is NAT?

Network Address Translation is a method by which IP addresses are mapped from one address realm to another, providing transparent routing to end hosts. There are many variations of address translation that lend themselves to different applications. However, all flavors of NAT devices should share the following characteristics.

Meta's Exhibit No. 1013
Page 006

a) Transparent Address assignment.
b) Transparent routing through address translation.
(routing here refers to forwarding packets, and not exchanging routing information)
c) ICMP error packet payload translation.

Below is a diagram illustrating a scenario in which NAT is enabled on a stub domain border router, connected to the Internet through a regional router made available by a service provider.

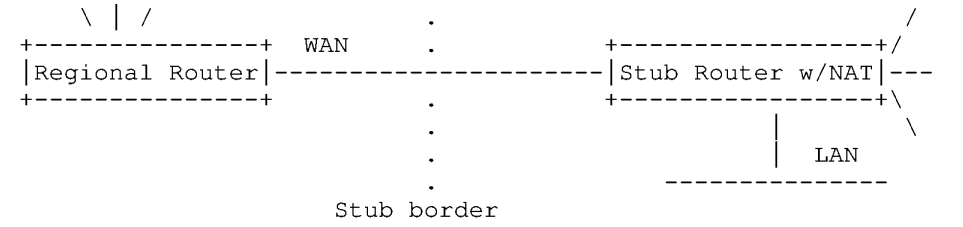


Figure 1: A typical NAT operation scenario

3.1. Transparent Address Assignment

NAT binds addresses in private network with addresses in global network and vice versa to provide transparent routing for the datagrams traversing between address realms. The binding in some cases may extend to transport level identifiers (such as TCP/UDP ports). Address binding is done at the start of a session. The following sub-sections describe two types of address assignments.

3.1.1. Static Address assignment

In the case of static address assignment, there is one-to-one address mapping for hosts between a private network address and an external network address for the lifetime of NAT operation. Static address assignment ensures that NAT does not have to administer address management with session flows.

3.1.2. Dynamic Address assignment

In this case, external addresses are assigned to private network hosts or vice versa, dynamically based on usage requirements and session flow determined heuristically by NAT. When the last session using an address binding is terminated, NAT would free the binding so that the global address could be recycled for later use. The exact nature of address assignment is specific to individual NAT implementations.

Meta's Exhibit No. 1013
Page 007

3.2. Transparent routing

A NAT router sits at the border between two address realms and translates addresses in IP headers so that when the packet leaves one realm and enters another, it can be routed properly. Because NAT devices have connections to multiple address realms, they must be careful to not improperly propagate information (e.g., via routing protocols) about networks from one address realm into another, where such an advertisement would be deemed unacceptable.

There are three phases to Address translation, as follows. Together these phases result in creation, maintenance and termination of state for sessions passing through NAT devices.

3.2.1. Address binding

Address binding is the phase in which a local node IP address is associated with an external address or vice versa, for purposes of translation. Address binding is fixed with static address assignments and is dynamic at session startup time with dynamic address assignments. Once the binding between two addresses is in place, all subsequent sessions originating from or to this host will use the same binding for session based packet translation.

New address bindings are made at the start of a new session, if such an address binding didn't already exist. Once a local address is bound to an external address, all subsequent sessions originating from the same local address or directed to the same local address will use the same binding.

The start of each new session will result in the creation of a state to facilitate translation of datagrams pertaining to the session. There can be many simultaneous sessions originating from the same host, based on a single address binding.

3.2.2. Address lookup and translation

Once a state is established for a session, all packets belonging to the session will be subject to address lookup (and transport identifier lookup, in some cases) and translation.

Address or transport identifier translation for a datagram will result in the datagram forwarding from the origin address realm to the destination address realm with network addresses appropriately updated.

Meta's Exhibit No. 1013
Page 008

3.2.3. Address unbinding

Address unbinding is the phase in which a private address is no longer associated with a global address for purposes of translation. NAT will perform address unbinding when it believes that the last session using an address binding has terminated. Refer section 2.6 for some heuristic ways to handle session terminations.

3.3. ICMP error packet translation

All ICMP error messages (with the exception of Redirect message type) will need to be modified, when passed through NAT. The ICMP error message types needing NAT modification would include Destination-Unreachable, Source-Quench, Time-Exceeded and Parameter-Problem. NAT should not attempt to modify a Redirect message type.

Changes to ICMP error message will include changes to the original IP packet (or portions thereof) embedded in the payload of the ICMP error message. In order for NAT to be completely transparent to end hosts, the IP address of the IP header embedded in the payload of the ICMP packet must be modified, the checksum field of the same IP header must correspondingly be modified, and the accompanying transport header. The ICMP header checksum must also be modified to reflect changes made to the IP and transport headers in the payload. Furthermore, the normal IP header must also be modified.

4.0. Various flavors of NAT

There are many variations of address translation that lend themselves to different applications. NAT flavors listed in the following sub-sections are by no means exhaustive, but they do capture the significant differences that abound.

The following diagram will be used as a base model to illustrate NAT flavors. Host-A, with address Addr-A is located in a private realm, represented by the network N-Pri. N-Pri is isolated from external network through a NAT router. Host-X, with address Addr-X is located in an external realm, represented by the network N-Ext. NAT router with two interfaces, each attached to one of the realms provides transparent routing between the two realms. The interface to the external realm is assigned an address of Addr-Nx and the interface to private realm is assigned an address of Addr-Np. Further, it may be understood that addresses Addr-A and Addr-Np correspond to N-Pri network and the addresses Addr-X and Addr-Nx correspond to N-Ext network.

Meta's Exhibit No. 1013
Page 009

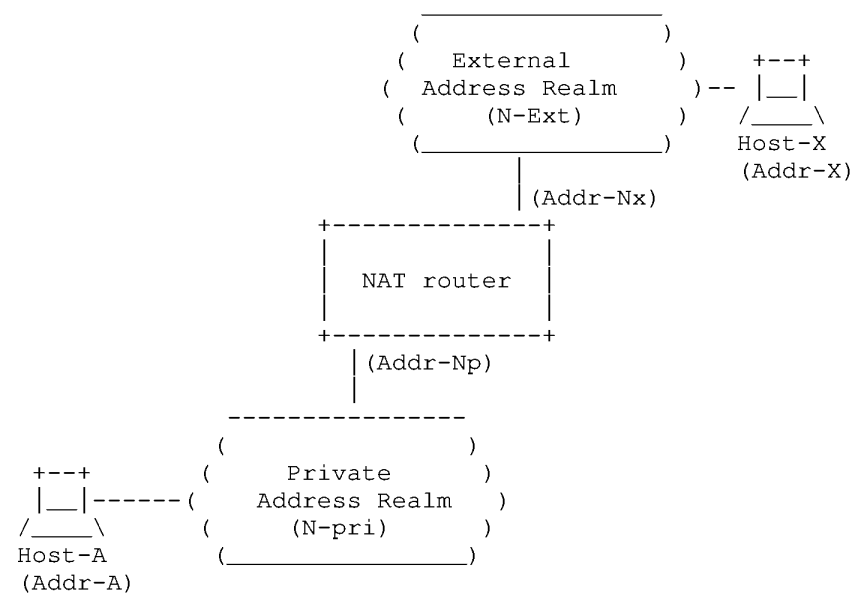


Figure 2: A base model to illustrate NAT terms.

4.1. Traditional NAT (or) Outbound NAT

Traditional NAT would allow hosts within a private network to transparently access hosts in the external network, in most cases. In a traditional NAT, sessions are uni-directional, outbound from the private network. This is in contrast with Bi-directional NAT, which permits sessions in both inbound and outbound directions. A detailed description of Bi-directional NAT may be found in section 4.2.

The following is a description of the properties of realms supported by traditional NAT. IP addresses of hosts in external network are unique and valid in external as well as private networks. However, the addresses of hosts in private network are unique only within the private network and may not be valid in the external network. In other words, NAT would not advertise private networks to the external realm. But, networks from the external realm may be advertised within the private network. The addresses used within private network must not overlap with the external addresses. Any given address must either be a private address or an external address; not both.

Meta's Exhibit No. 1013
Page 0010

A traditional NAT router in figure 2 would allow Host-A to initiate sessions to Host-X, but not the other way around. Also, N-Ext is routable from within N-Pri, whereas N-Pri may not be routable from N-Ext.

Traditional NAT is primarily used by sites using private addresses that wish to allow outbound sessions from their site.

There are two variations to traditional NAT, namely Basic NAT and NAPT (Network Address Port Translation). These are discussed in the following sub-sections.

4.1.1. Basic NAT

With Basic NAT, a block of external addresses are set aside for translating addresses of hosts in a private domain as they originate sessions to the external domain. For packets outbound from the private network, the source IP address and related fields such as IP, TCP, UDP and ICMP header checksums are translated. For inbound packets, the destination IP address and the checksums as listed above are translated.

A Basic NAT router in figure 2 may be configured to translate N-Pri into a block of external addresses, say Addr-i through Addr-n, selected from the external network N-Ext.

4.1.2. Network Address Port Translation (NAPT)

NAPT extends the notion of translation one step further by also translating transport identifier (e.g., TCP and UDP port numbers, ICMP query identifiers). This allows the transport identifiers of a number of private hosts to be multiplexed into the transport identifiers of a single external address. NAPT allows a set of hosts to share a single external address. Note that NAPT can be combined with Basic NAT so that a pool of external addresses are used in conjunction with port translation.

For packets outbound from the private network, NAPT would translate the source IP address, source transport identifier and related fields such as IP, TCP, UDP and ICMP header checksums. Transport identifier can be one of TCP/UDP port or ICMP query ID. For inbound packets, the destination IP address, destination transport identifier and the IP and transport header checksums are translated.

Meta's Exhibit No. 1013
Page 0011

A NAPT router in figure 2 may be configured to translate sessions originated from N-Pri into a single external address, say Addr-i.

Very often, the external interface address Addr-Nx of NAPT router is used as the address to map N-Pri to.

4.2. Bi-directional NAT (or) Two-Way NAT

With a Bi-directional NAT, sessions can be initiated from hosts in the public network as well as the private network. Private network addresses are bound to globally unique addresses, statically or dynamically as connections are established in either direction. The name space (i.e., their Fully Qualified Domain Names) between hosts in private and external networks is assumed to be end-to-end unique. Hosts in external realm access private realm hosts by using DNS for address resolution. A DNS-ALG must be employed in conjunction with Bi-Directional NAT to facilitate name to address mapping. Specifically, the DNS-ALG must be capable of translating private realm addresses in DNS Queries and responses into their external realm address bindings, and vice versa, as DNS packets traverse between private and external realms.

The address space requirements outlined for traditional NAT routers are applicable here as well.

A Bi-directional NAT router in figure 2 would allow Host-A to initiate sessions to Host-X, and Host-X to initiate sessions to Host-A. Just as with traditional NAT, N-Ext is routable from within N-Pri, but N-Pri may not be routable from N-Ext.

4.3. Twice NAT

Twice NAT is a variation of NAT in that both the source and destination addresses are modified by NAT as a datagram crosses address realms. This is in contrast to Traditional-NAT and Bi-Directional NAT, where only one of the addresses (either source or destination) is translated. Note, there is no such term as 'Once-NAT'.

Twice NAT is necessary when private and external realms have address collisions. The most common case where this would happen is when a site had (improperly) numbered its internal nodes using public addresses that have been assigned to another organization. Alternatively, a site may have changed from one provider to another, but chosen to keep (internally) the addresses it had been assigned by the first provider. That provider might then later reassign those addresses to someone else. The key issue in such cases is that the address of the host in the external realm may have been assigned the

Meta's Exhibit No. 1013
Page 0012

same address as a host within the local site. If that address were to appear in a packet, it would be forwarded to the internal node rather than through the NAT device to the external realm. Twice-NAT attempts to bridge these realms by translating both source and destination address of an IP packet, as the packet transitions realms.

Twice-NAT works as follows. When Host-A wishes to initiate a session to Host-X, it issues a DNS query for Host-X. A DNS-ALG intercepts the DNS query, and in the response returned to Host-A the DNS-ALG replaces the address for Host-X with one that is properly routable in the local site (say Host-XPRIME). Host A then initiates communication with Host-XPRIME. When the packets traverse the NAT device, the source IP address is translated (as in the case of traditional NAT) and the destination address is translated to Host-X. A similar translation is performed on return packets coming from Host-X.

The following is a description of the properties of realms supported by Twice-NAT. Network address of hosts in external network are unique in external networks, but not within private network. Likewise, the network address of hosts in private network are unique only within the private network. In other words, the address space used in private network to locate hosts in private and public networks is unrelated to the address space used in public network to locate hosts in private and public networks. Twice NAT would not be allowed to advertise local networks to the external network or vice versa.

A Twice NAT router in figure 2 would allow Host-A to initiate sessions to Host-X, and Host-X to initiate sessions to Host-A. However, N-Ext (or a subset of N-Ext) is not routable from within N-Pri, and N-Pri is not routable from N-Ext.

Twice NAT is typically used when address space used in a Private network overlaps with addresses used in the Public space. For example, say a private site uses the 200.200.200.0/24 address space which is officially assigned to another site in the public internet. Host_A (200.200.200.1) in Private space seeks to connect to Host_X (200.200.200.100) in Public space. In order to make this connection work, Host_X's address is mapped to a different address for Host_A and vice versa. The twice NAT located at the Private site border may be configured as follows:

Meta's Exhibit No. 1013
Page 0013

```
Private to Public : 200.200.200.0/24 -> 138.76.28.0/24
Public to Private : 200.200.200.0/24 -> 172.16.1.0/24

Datagram flow  : Host_A(Private) ->  Host_X(Public)

a) Within private network

   DA: 172.16.1.100      SA: 200.200.200.1

b) After twice-NAT translation

  DA: 200.200.200.100    SA: 138.76.28.1

Datagram flow Host_X (Public) -> Host_A (Private)

a) Within Public network

   DA: 138.76.28.1       SA: 200.200.200.100

b) After twice-NAT translation, in private network

   SA: 200.200.200.1     DA: 172.16.1.100
```

4.4. Multihomed NAT

There are limitations to using NAT. For example, requests and responses pertaining to a session must be routed via the same NAT router, as a NAT router maintains state information for sessions established through it. For this reason, it is often suggested that NAT routers be operated on a border router unique to a stub domain, where all IP packets are either originated from the domain or destined to the domain. However, such a configuration would turn a NAT router into a single point of failure.

In order for a private network to ensure that connectivity with external networks is retained even as one of the NAT links fail, it is often desirable to multihome the private network to same or multiple service providers with multiple connections from the private domain, be it from same or different NAT boxes.

For example, a private network could have links to two different providers and the sessions from private hosts could flow through the NAT router with the best metric for a destination. When one of NAT routers fail, the other could route traffic for all connections.

Meta's Exhibit No. 1013
Page 0014

Multiple NAT boxes or multiple links on the same NAT box, sharing the same NAT configuration can provide fail-safe backup for each other. In such a case, it is necessary for backup NAT device to exchange state information so that a backup NAT can take on session load transparently when the primary NAT fails. NAT backup becomes simpler, when configuration is based on static maps.

5.0. Realm Specific IP (RSIP)

"Realm Specific IP" (RSIP) is used to characterize the functionality of a realm-aware host in a private realm, which assumes realm-specific IP address to communicate with hosts in private or external realm.

A "Realm Specific IP Client" (RSIP client) is a host in a private network that adopts an address in an external realm when connecting to hosts in that realm to pursue end-to-end communication. Packets generated by hosts on either end in such a setup would be based on addresses that are end-to-end unique in the external realm and do not require translation by an intermediary process.

A "Realm Specific IP Server" (RSIP server) is a node resident on both private and external realms, that can facilitate routing of external realm packets within a private realm. These packets may either have been originated by an RSIP client or directed to an RSIP-client. RSIP-Server may also be the same node that assigns external realm addresses to RSIP-Clients.

There are two variations to RSIP, namely Realm-specific Address IP (RSA-IP) and Realm-Specific Address and Port IP (RSAP-IP). These variations are discussed in the following sub-sections.

5.1. Realm Specific Address IP (RSA-IP)

A Realm Specific Address IP (RSA-IP) client adopts an IP address from the external address space when connecting to a host in external realm. Once an RSA-IP client assumes an external address, no other host in private or external domain can assume the same address, until that address is released by the RSA-IP client.

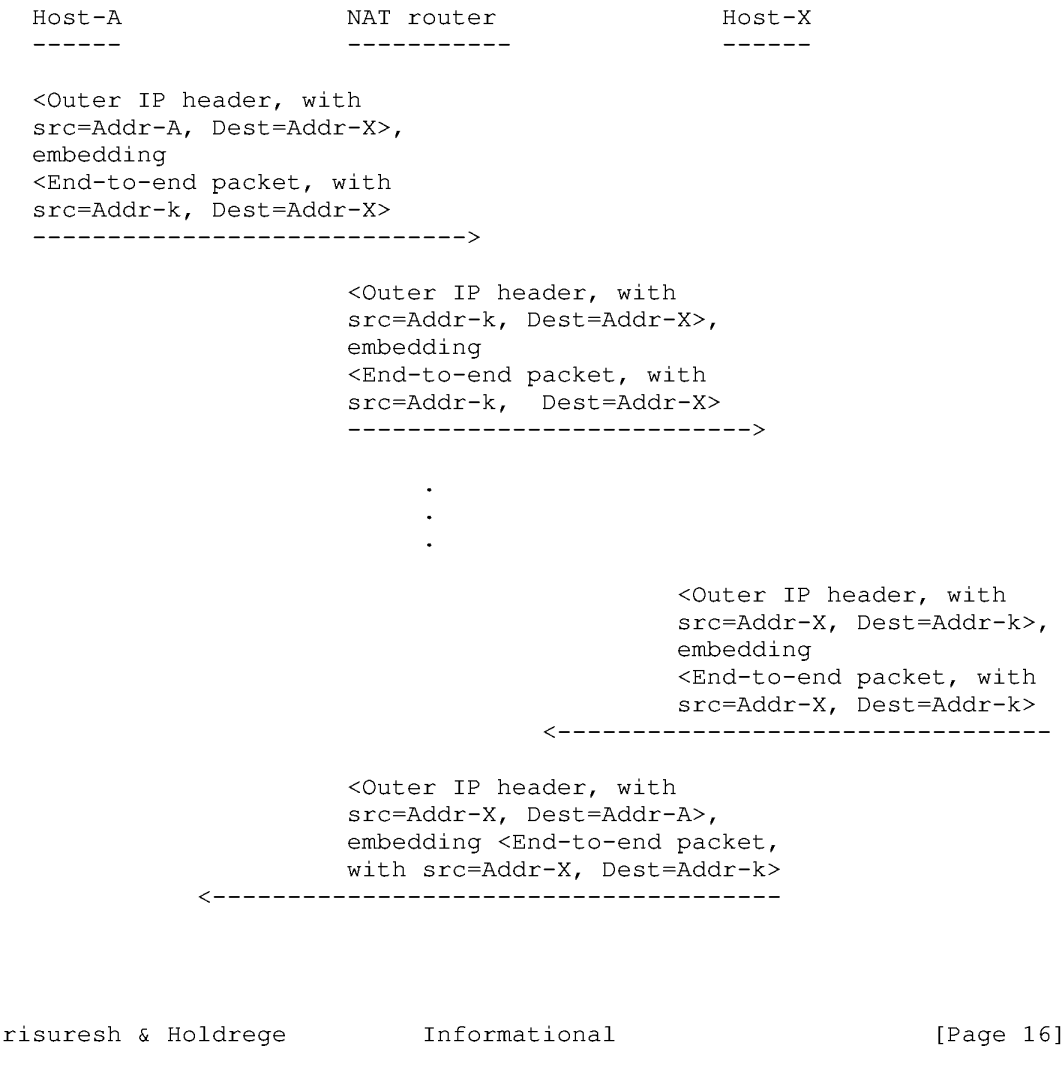
The following is a discussion of routing alternatives that may be pursued for the end-to-end RSA-IP packets within private realm. One approach would be to tunnel the packet to the destination. The outer header can be translated by NAT as normal without affecting the addresses used in the internal header. Another approach would be to set up a bi-directional tunnel between the RSA-IP Client and the border router straddling the two address realms. Packets to and from the client would be tunneled, but packets would be forwarded as

Meta's Exhibit No. 1013
Page 0015

normal between the border router and the remote destination. Note, the tunnel from the client TO the border router may not be necessary. You might be able to just forward the packet directly. This should work so long as your internal network isn't filtering packets based on source addresses (which in this case would be external addresses).

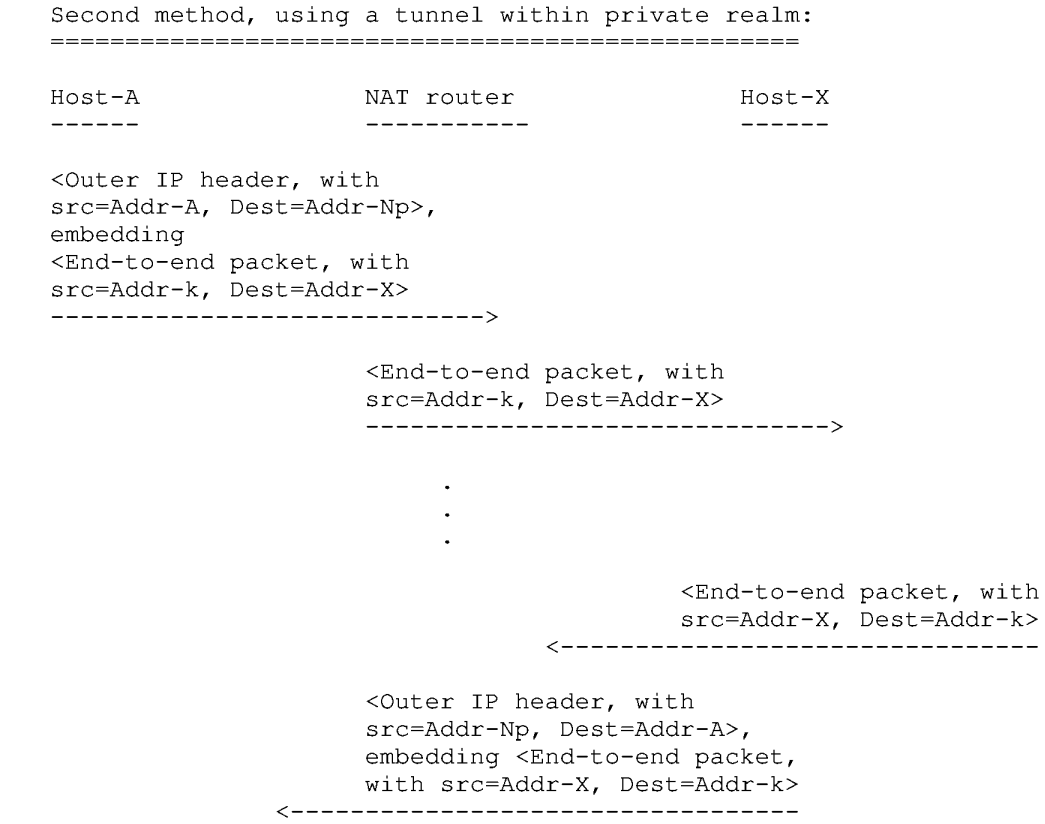
As an example, Host-A in figure 2 above, could assume an address Addr-k from the external realm and act as RSA-IP-Client to allow end-to-end sessions between Addr-k and Addr-X. Traversal of end-to-end packets within private realm may be illustrated as follows:

First method, using NAT router enroute to translate:



Meta's Exhibit No. 1013
Page 0016

Second method, using a tunnel within private realm:



There may be other approaches to pursue.

An RSA-IP-Client has the following characteristics. The collective set of operations performed by an RSA-IP-Client may be termed "RSA-IP".

1. Aware of the realm to which its peer nodes belong.

2. Assumes an address from external realm when communicating with hosts in that realm. Such an address may be assigned statically or obtained dynamically (through a yet-to-be-defined protocol) from a node capable of assigning addresses from external realm. RSA-IP-Server could be the node coordinating external realm address assignment.

Meta's Exhibit No. 1013
Page 0017

3. Route packets to external hosts using an approach amenable to RSA-IP-Server. In all cases, RSA-IP-Client will likely need to act as a tunnel end-point, capable of encapsulating end-to-end packets while forwarding and decapsulating in the return path.

"Realm Specific Address IP Server" (RSA-IP server) is a node resident on both private and external realms, that facilitates routing of external realm packets specific to RSA-IP clients inside a private realm. An RSA-IP-Server may be described as having the following characteristics.

1. May be configured to assign addresses from external realm to RSA-IP-Clients, either statically or dynamically.

2. Must be a router resident on both the private and external address realms.

3. Must be able to provide a mechanism to route external realm packets within private realm. Of the two approaches described, the first approach requires RSA-IP-Server to be a NAT router providing transparent routing for the outer header. This approach requires the external peer to be a tunnel end-point.

   With the second approach, an RSA-IP-Server could be any router (including a NAT router) that can be a tunnel end-point with RSA-IP-Clients. It would detunnel end-to-end packets outbound from RSA-IP-Clients and forward to external hosts. On the return path, it would locate RSA-IP-Client tunnel, based on the destination address of the end-to-end packet and encapsulate the packet in a tunnel to forward to RSA-IP-Client.

RSA-IP-Clients may pursue any of the IPsec techniques, namely transport or tunnel mode Authentication and confidentiality using AH and ESP headers on the embedded packets. Any of the tunneling techniques may be adapted for encapsulation between RSA-IP-Client and RSA-IP-Server or between RSA-IP-Client and external host. For example, IPsec tunnel mode encapsulation is a valid type of encapsulation that ensures IPsec authentication and confidentiality for the embedded end-to-end packets.

5.2 Realm Specific Address and port IP (RSAP-IP)

Realm Specific Address and port IP (RSAP-IP) is a variation of RSIP in that multiple private hosts use a single external address, multiplexing on transport IDentifiers (i.e., TCP/UDP port numbers and ICMP Query IDs).

Meta's Exhibit No. 1013
Page 0018

"RSAP-IP-Client" may be defined similar to RSA-IP-Client with the variation that RSAP-IP-Client assumes a tuple of (external address, transport Identifier) when connecting to hosts in external realm to pursue end-to-end communication. As such, communication with external nodes for an RSAP-IP-Client may be limited to TCP, UDP and ICMP sessions.

"RSAP-IP-Server" is similar to RSA-IP-Server in that it facilitates routing of external realm packets specific to RSAP-IP clients inside a private realm. Typically, an RSAP-IP-Server would also be the one to assign transport tuples to RSAP-IP-Clients.

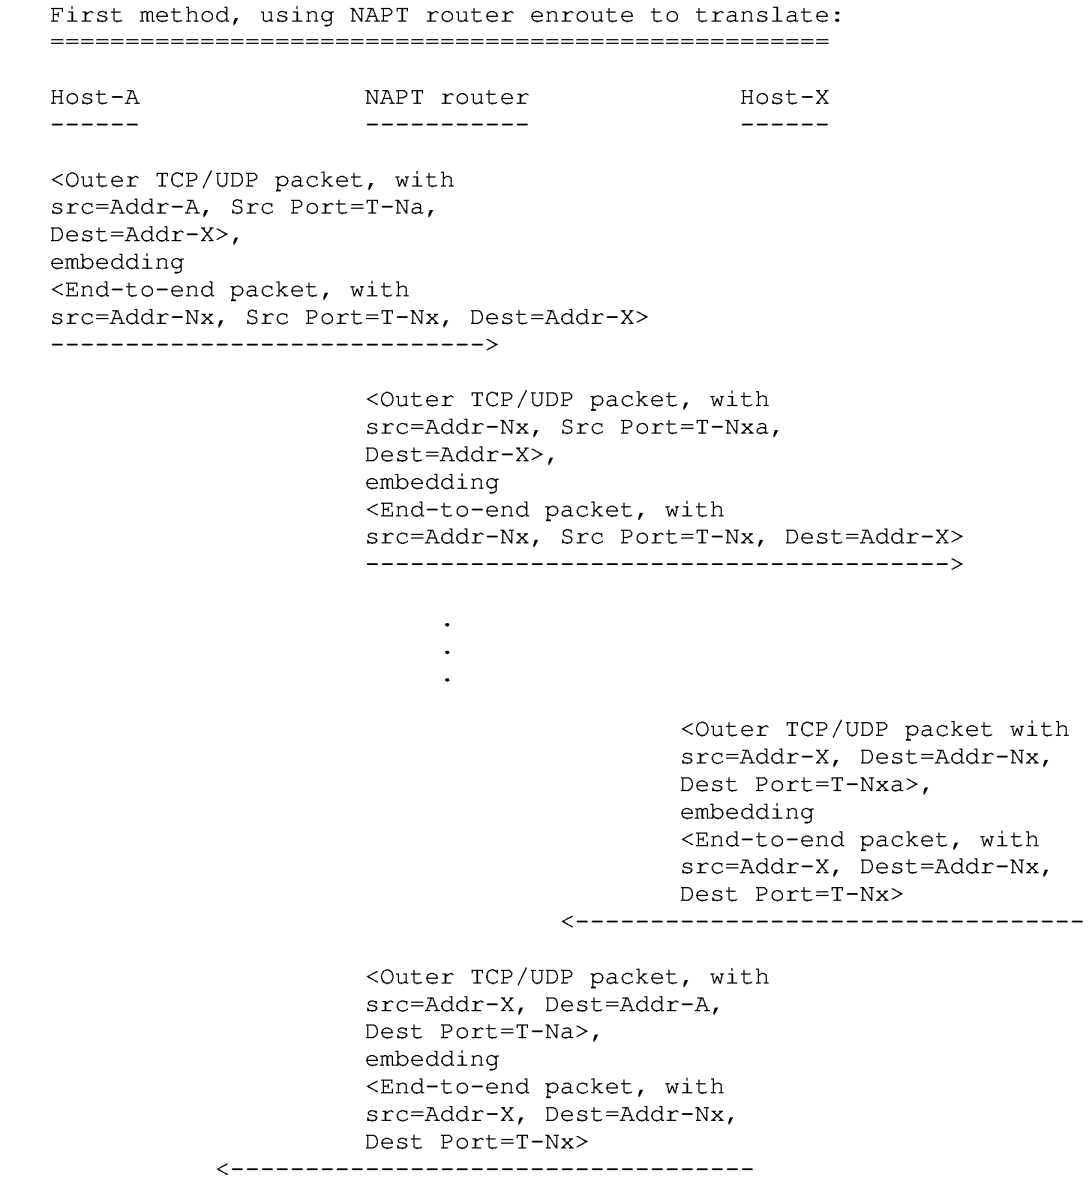
A NAPT router enroute could serve as RSAP-IP-Server, when the outer encapsulation is TCP/UDP based and is addressed between the RSAP-IP-Client and external peer. This approach requires the external peer to be the end-point of TCP/UDP based tunnel. Using this approach, RSAP-IP-Clients may pursue any of the IPsec techniques, namely transport or tunnel mode authentication and confidentiality using AH and ESP headers on the embedded packets. Note however, IPsec tunnel mode is not a valid type of encapsulation, as a NAPT router cannot provide routing transparency to AH and ESP protocols.

Alternately, packets may be tunneled between RSAP-IP-Client and RSAP-IP-Server such that RSAP-IP-Server would detunnel packets outbound from RSAP-IP-Clients and forward to external hosts. On the return path, RSAP-IP-Server would locate RSAP-IP-Client tunnel, based on the tuple of (destination address, transport Identifier) and encapsulate the original packet within a tunnel to forward to RSAP-IP-Client. With this approach, there is no limitation on the tunneling technique employed between RSAP-IP-Client and RSAP-IP-Server. However, there are limitations to applying IPsec based security on end-to-end packets. Transport mode based authentication and integrity may be attained. But, confidentiality cannot be permitted because RSAP-IP-Server must be able to examine the destination transport Identifier in order to identify the RSAP-IP-tunnel to forward inbound packets to. For this reason, only the transport mode TCP, UDP and ICMP packets protected by AH and ESP-authentication can traverse a RSAP-IP-Server using this approach.

As an example, say Host-A in figure 2 above, obtains a tuple of (Addr-Nx, TCP port T-Nx) from NAPT router to act as RSAP-IP-Client to initiate end-to-end TCP sessions with Host-X. Traversal of end-to-end packets within private realm may be illustrated as follows. In the first method, outer layer of the outgoing packet from Host-A uses (private address Addr-A, source port T-Na) as source tuple to communicate with Host-X. NAPT router enroute translates this tuple into (Addr-Nx, Port T-Nxa). This translation is independent of RSAP-IP-Client tuple parameters used in the embedded packet.

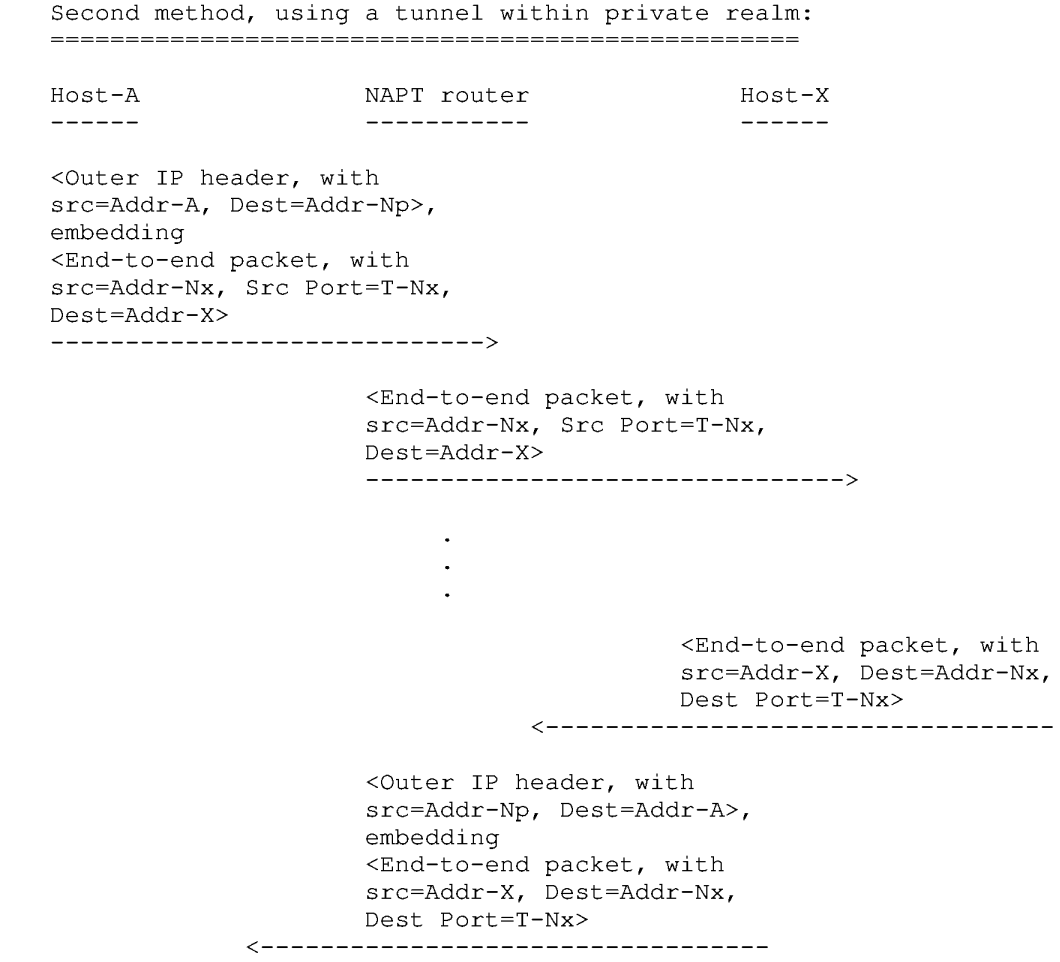
Meta's Exhibit No. 1013
Page 0019

First method, using NAPT router enroute to translate:



Meta's Exhibit No. 1013
Page 0020

Second method, using a tunnel within private realm:



6.0. Private Networks and Tunnels

Let us consider the case where your private network is connected to the external world via tunnels. In such a case, tunnel encapsulated traffic may or may not contain translated packets depending upon the characteristics of address realms a tunnel is bridging.

The following subsections discuss two scenarios where tunnels are used (a) in conjunction with Address translation, and (b) without translation.

Meta's Exhibit No. 1013
Page 0021

6.1. Tunneling translated packets

All variations of address translations discussed in the previous section can be applicable to direct connected links as well as tunnels and virtual private networks (VPNs).

For example, a private network connected to a business partner through a VPN could employ traditional NAT to communicate with the partner. Likewise, it is possible to employ twice NAT, if the partner's address space overlapped with the private network. There could be a NAT device on one end of the tunnel or on both ends of the tunnel. In all cases, traffic across the VPN can be encrypted for security purposes. Security here refers to security for traffic across VPNs alone. End-to-end security requires trusting NAT devices within private network.

6.2. Backbone partitioned private Networks

There are many instances where a private network (such as a corporate network) is spread over different locations and use public backbone for communications between those locations. In such cases, it is not desirable to do address translation, both because large numbers of hosts may want to communicate across the backbone, thus requiring large address tables, and because there will be more applications that depend on configured addresses, as opposed to going to a name server. We call such a private network a backbone-partitioned private network.

Backbone-partitioned stubs should behave as though they were a non-partitioned stub. That is, the routers in all partitions should maintain routes to the local address spaces of all partitions. Of course, the (public) backbones do not maintain routes to any local addresses. Therefore, the border routers must tunnel (using VPNs) through the backbones using encapsulation. To do this, each NAT box will set aside a global address for tunneling.

When a NAT box x in stub partition X wishes to deliver a packet to stub partition Y, it will encapsulate the packet in an IP header with destination address set to the global address of NAT box y that has been reserved for encapsulation. When NAT box y receives a packet with that destination address, it decapsulates the IP header and routes the packet internally. Note, there is no address translation in the process; merely transfer of private network packets over an external network tunnel backbone.

Meta's Exhibit No. 1013
Page 0022

7.0. NAT operational characteristics

NAT devices are application unaware in that the translations are limited to IP/TCP/UDP/ICMP headers and ICMP error messages only. NAT devices do not change the payload of the packets, as payloads tend to be application specific.

NAT devices (without the inclusion of ALGs) do not examine or modify transport payload. For this reason, NAT devices are transparent to applications in many cases. There are two areas, however, where NAT devices often cause difficulties: 1) when an application payload includes an IP address, and 2) when end-to-end security is needed. Note, this is not a comprehensive list.

Application layer security techniques that do not make use of or depend on IP addresses will work correctly in the presence of NAT (e.g., TLS, SSL and ssh). In contrast, transport layer techniques such as IPSec transport mode or the TCP MD5 Signature Option RFC 2385 [Ref 17] do not.

In IPSec transport mode, both AH and ESP have an integrity check covering the entire payload. When the payload is TCP or UDP, the TCP/UDP checksum is covered by the integrity check. When a NAT device modifies an address the checksum is no longer valid with respect to the new address. Normally, NAT also updates the checksum, but this is ineffective when AH and ESP are used. Consequently, receivers will discard a packet either because it fails the IPSec integrity check (if the NAT device updates the checksum), or because the checksum is invalid (if the NAT device leaves the checksum unmodified).

Note that IPsec tunnel mode ESP is permissible so long as the embedded packet contents are unaffected by the outer IP header translation. Although this technique does not work in traditional NAT deployments (i.e., where hosts are unaware that NATs are present), the technique is applicable to Realm-Specific IP as described in Section 5.0.

Note also that end-to-end ESP based transport mode authentication and confidentiality are permissible for packets such as ICMP, whose IP payload content is unaffected by the outer IP header translation.

NAT devices also break fundamental assumptions by public key distribution infrastructures such as Secure DNS RFC 2535 [Ref 18] and X.509 certificates with signed public keys. In the case of Secure

Meta's Exhibit No. 1013
Page 0023

DNS, each DNS RRset is signed with a key from within the zone. Moreover, the authenticity of a specific key is verified by following a chain of trust that goes all the way to the DNS root. When a DNS-ALG modifies addresses (e.g., as in the case of Twice-NAT), verification of signatures fails.

It may be of interest to note that IKE (Session key negotiation protocol) is a UDP based session layer protocol and is not protected by network based IPsec security. Only a portion of the individual payloads within IKE are protected. As a result, IKE sessions are permissible across NAT, so long as IKE payload does not contain addresses and/or transport IDs specific to one realm and not the other. Given that IKE is used to setup IPSec associations, and there are at present no known ways of making IPSec work through a NAT function, it is a future work item to take advantage of IKE through a NAT box.

One of the most popular internet applications "FTP" would not work with the definition of NAT as described. The following sub-section is devoted to describing how FTP is supported on NAT devices. FTP ALG is an integral part of most NAT implementations. Some vendors may choose to include additional ALGs to custom support other applications on the NAT device.

7.1. FTP support

"PORT" command and "PASV" response in FTP control session payload identify the IP address and TCP port that must be used for the data session it supports. The arguments to the PORT command and PASV response are an IP address and a TCP port in ASCII. An FTP ALG is required to monitor and update the FTP control session payload so that information contained in the payload is relevant to end nodes. The ALG must also update NAT with appropriate data session tuples and session orientation so that NAT could set up state information for the FTP data sessions.

Because the address and TCP port are encoded in ASCII, this may result in a change in the size of packet. For instance, 10,18,177,42,64,87 is 18 ASCII characters, whereas 193,45,228,137,64,87 is 20 ASCII characters. If the new size is same as the previous, only the TCP checksum needs adjustment as a result of change of data. If the new size is less than or greater than the previous, TCP sequence numbers must also be changed to reflect the change in length of FTP control data portion. A special table may be used by the ALG to correct the TCP sequence and acknowledge numbers. The sequence number and acknowledgement correction will need to be performed on all future packet of the connection.

Meta's Exhibit No. 1013
Page 0024

8.0. NAT limitations

8.1. Applications with IP-address Content

Not All applications lend themselves easily to address translation by NAT devices. Especially, the applications that carry IP address (and TU port, in case of NAPT) inside the payload. Application Level Gateways, or ALGs must be used to perform translations on packets pertaining to such applications. ALGs may optionally utilize address (and TU port) assignments made by NAT and perform translations specific to the application. The combination of NAT functionality and ALGs will not provide end-to-end security assured by IPsec. However, tunnel mode IPsec can be accomplished with NAT router serving as tunnel end point.

SNMP is one such application with address content in payload. NAT routers would not translate IP addresses within SNMP payloads. It is not uncommon for an SNMP specific ALG to reside on a NAT router to perform SNMP MIB translations proprietary to the private network.

8.2. Applications with inter-dependent control and data sessions

NAT devices operate on the assumption that each session is independent. Session characteristics like session orientation, source and destination IP addresses, session protocol, and source and destination transport level identifiers are determined independently at the start of each new session.

However, there are applications such as H.323 that use one or more control sessions to set the characteristics of the follow-on sessions in their control session payload. Such applications require use of application specific ALGs that can interpret and translate the payload, if necessary. Payload interpretation would help NAT be prepared for the follow-on data sessions.

8.3. Debugging Considerations

NAT increases the probability of mis-addressing. For example, same local address may be bound to different global address at different times and vice versa. As a result, any traffic flow study based purely on global addresses and TU ports could be confused and might misinterpret the results.

If a host is abusing the Internet in some way (such as trying to attack another machine or even sending large amounts of junk mail or something) it is more difficult to pinpoint the source of the trouble because the IP address of the host is hidden in a NAT router.

Meta's Exhibit No. 1013
Page 0025

8.4. Translation of fragmented FTP control packets

Translation of fragmented FTP control packets is tricky when the packets contain "PORT" command or response to "PASV" command. Clearly, this is a pathological case. NAT router would need to assemble the fragments together first and then translate prior to forwarding.

Yet another case would be when each character of packets containing "PORT" command or response to "PASV" is sent in a separate datagram, unfragmented. In this case, NAT would simply have to let the packets through, without translating the TCP payload. Of course, the application will fail if the payload needed to be altered. The application could still work in a few cases, where the payload contents can be valid in both realms, without modifications enroute. For example, FTP originated from a private host would still work while traversing a traditional NAT or bi-directional NAT device, so long as the FTP control session employed PASV command to establish data sessions. The reason being that the address and port number specified by FTP server in the PASV response (sent as multiple unfragmented packets) is valid to the private host, as is. The NAT device will simply view the ensuing data session (also originating from private host) as an independent TCP session.

8.5. Compute intensive

NAT is compute intensive even with the help of a clever checksum adjustment algorithm, as each data packet is subject to NAT lookup and modifications. As a result, router forwarding throughput could be slowed considerably. However, so long as the processing capacity of the NAT device exceeds line processing rate, this should not be a problem.

9.0. Security Considerations

Many people view traditional NAT router as a one-way (session) traffic filter, restricting sessions from external hosts into their machines. In addition, when address assignment in NAT router is done dynamically, that makes it harder for an attacker to point to any specific host in the NAT domain. NAT routers may be used in conjunction with firewalls to filter unwanted traffic.

If NAT devices and ALGs are not in a trusted boundary, that is a major security problem, as ALGs could snoop end user traffic payload. Session level payload could be encrypted end to end, so long as the payload does not contain IP addresses and/or transport identifiers that are valid in only one of the realms. With the exception of RSIP, end-to-end IP network level security assured by current IPsec

Meta's Exhibit No. 1013
Page 0026

techniques is not attainable with NAT devices in between. One of the ends must be a NAT box. Refer section 7.0 for a discussion on why end-to-end IPsec security cannot be assured with NAT devices along the route.

The combination of NAT functionality, ALGs and firewalls will provide a transparent working environment for a private networking domain. With the exception of RSIP, end-to-end network security assured by IPsec cannot be attained for end-hosts within the private network (Refer section 5.0 for RSIP operation). In all other cases, if you want to use end-to-end IPsec, there cannot be a NAT device in the path. If we make the assumption that NAT devices are part of a trusted boundary, tunnel mode IPsec can be accomplished with NAT router (or a combination of NAT, ALGs and firewall) serving as tunnel end point.

NAT devices, when combined with ALGs, can ensure that the datagrams injected into Internet have no private addresses in headers or payload. Applications that do not meet these requirements may be dropped using firewall filters. For this reason, it is not uncommon to find NAT, ALG and firewall functions co-exist to provide security at the borders of a private network. NAT gateways can be used as tunnel end points to provide secure VPN transport of packet data across an external network domain.

Below are some additional security considerations associated with NAT routers.

1. UDP sessions are inherently unsafe. Responses to a datagram could come from an address different from the target address used by sender ([Ref 4]). As a result, an incoming UDP packet might match the outbound session of a traditional NAT router only in part (the destination address and UDP port number of the packet match, but the source address and port number may not). In such a case, there is a potential security compromise for the NAT device in permitting inbound packets with partial match. This UDP security issue is also inherent to firewalls.

   Traditional NAT implementations that do not track datagrams on a per-session basis but lump states of multiple UDP sessions using the same address binding into a single unified session could compromise the security even further. This is because, the granularity of packet matching would be further limited to just the destination address of the inbound UDP packets.

2. Multicast sessions (UDP based) are another source for security weakness for traditional-NAT routers. Once again, firewalls face the same security dilemma as the NAT routers.

Meta's Exhibit No. 1013
Page 0027

Say, a host on private network initiated a multicast session. Datagram sent by the private host could trigger responses in the reverse direction from multiple external hosts. Traditional-NAT implementations that use a single state to track a multicast session cannot determine for certain if the incoming UDP packet is in response to an existing multicast session or the start of new UDP session initiated by an attacker.

3. NAT devices can be a target for attacks.

   Since NAT devices are Internet hosts they can be the target of a number of different attacks, such as SYN flood and ping flood attacks. NAT devices should employ the same sort of protection techniques as Internet-based servers do.

REFERENCES

[1] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot,G. and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[2] Reynolds, J. and J. Postel, "Assigned Numbers", STD 2, RFC 1700, October, 1994.

[3] Braden, R., "Requirements for Internet Hosts -- Communication Layers", STD 3, RFC 1122, October 1989.

[4] Braden, R., "Requirements for Internet Hosts -- Application and Support", STD 3, RFC 1123, October 1989.

[5] Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.

[6] Postel, J. and J. Reynolds, "File Transfer Protocol (FTP)", STD 9, RFC 959, October 1985.

[7] Postel, J., "Transmission Control Protocol (TCP) Specification", STD 7, RFC 793, September 1981.

[8] Postel, J., "Internet Control Message Protocol Specification" STD 5, RFC 792, September 1981.

[9] Postel, J., "User Datagram Protocol (UDP)", STD 6, RFC 768, August 1980.

[10] Mogul, J. and J. Postel, "Internet Standard Subnetting Procedure", STD 5, RFC 950, August 1985.

Meta's Exhibit No. 1013
Page 0028

[11] Carpenter, B., Crowcroft, J. and Y. Rekhter, "IPv4 Address Behavior Today", RFC 2101, February 1997.

[12] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.

[13] Kent, S. and R. Atkinson, "IP Encapsulating Security Payload (ESP)", RFC 2406, November 1998.

[14] Kent, S. and R. Atkinson, "IP Authentication Header", RFC 2402, November 1998.

[15] Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)", RFC 2409, November 1998.

[16] Piper, D., "The Internet IP Security Domain of Interpretation for ISAKMP", RFC 2407, November 1998.

[17] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC 2385, August 1998.

[18] Eastlake, D., "Domain Name System Security Extensions", RFC 2535, March 1999.

Authors' Addresses

Pyda Srisuresh
Lucent Technologies
4464 Willow Road
Pleasanton, CA 94588-8519
U.S.A.

Phone: (925) 737-2153
Fax: (925) 737-2110
EMail: srisuresh@lucent.com

Matt Holdrege
Lucent Technologies
1701 Harbor Bay Parkway
Alameda, CA 94502

Phone: (510) 769-6001
EMail: holdrege@lucent.com

Meta's Exhibit No. 1013
Page 0029

Full Copyright Statement

Copyright (C) The Internet Society (1999). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

Meta's Exhibit No. 1013
Page 0030

RFC: 793

TRANSMISSION CONTROL PROTOCOL

DARPA INTERNET PROGRAM

PROTOCOL SPECIFICATION

September 1981

prepared for

Defense Advanced Research Projects Agency
Information Processing Techniques Office
1400 Wilson Boulevard
Arlington, Virginia 22209

by

Information Sciences Institute
University of Southern California
4676 Admiralty Way
Marina del Rey, California 90291

Meta's Exhibit No. 1014
Page 001

TABLE OF CONTENTS

PREFACE iii

1. INTRODUCTION 1

1.1 Motivation 1
1.2 Scope 2
1.3 About This Document 2
1.4 Interfaces 3
1.5 Operation 3

2. PHILOSOPHY 7

2.1 Elements of the Internetwork System 7
2.2 Model of Operation 7
2.3 The Host Environment 8
2.4 Interfaces 9
2.5 Relation to Other Protocols 9
2.6 Reliable Communication 9
2.7 Connection Establishment and Clearing 10
2.8 Data Communication 12
2.9 Precedence and Security 13
2.10 Robustness Principle 13

3. FUNCTIONAL SPECIFICATION 15

3.1 Header Format 15
3.2 Terminology 19
3.3 Sequence Numbers 24
3.4 Establishing a connection 30
3.5 Closing a Connection 37
3.6 Precedence and Security 40
3.7 Data Communication 40
3.8 Interfaces 44
3.9 Event Processing 52

GLOSSARY 79

REFERENCES 85

Meta's Exhibit No. 1014
Page 003

Meta's Exhibit No. 1014
Page 004

PREFACE

This document describes the DoD Standard Transmission Control Protocol (TCP). There have been nine earlier editions of the ARPA TCP specification on which this standard is based, and the present text draws heavily from them. There have been many contributors to this work both in terms of concepts and in terms of text. This edition clarifies several details and removes the end-of-letter buffer-size adjustments, and redescribes the letter mechanism as a push function.

Jon Postel

Editor

Meta's Exhibit No. 1014
Page 005

RFC: 793
Replaces: RFC 761
IENs: 129, 124, 112, 81,
55, 44, 40, 27, 21, 5

TRANSMISSION CONTROL PROTOCOL

DARPA INTERNET PROGRAM
PROTOCOL SPECIFICATION

1. INTRODUCTION

The Transmission Control Protocol (TCP) is intended for use as a highly reliable host-to-host protocol between hosts in packet-switched computer communication networks, and in interconnected systems of such networks.

This document describes the functions to be performed by the Transmission Control Protocol, the program that implements it, and its interface to programs or users that require its services.

1.1. Motivation

Computer communication systems are playing an increasingly important role in military, government, and civilian environments. This document focuses its attention primarily on military computer communication requirements, especially robustness in the presence of communication unreliability and availability in the presence of congestion, but many of these problems are found in the civilian and government sector as well.

As strategic and tactical computer communication networks are developed and deployed, it is essential to provide means of interconnecting them and to provide standard interprocess communication protocols which can support a broad range of applications. In anticipation of the need for such standards, the Deputy Undersecretary of Defense for Research and Engineering has declared the Transmission Control Protocol (TCP) described herein to be a basis for DoD-wide inter-process communication protocol standardization.

TCP is a connection-oriented, end-to-end reliable protocol designed to fit into a layered hierarchy of protocols which support multi-network applications. The TCP provides for reliable inter-process communication between pairs of processes in host computers attached to distinct but interconnected computer communication networks. Very few assumptions are made as to the reliability of the communication protocols below the TCP layer. TCP assumes it can obtain a simple, potentially unreliable datagram service from the lower level protocols. In principle, the TCP should be able to operate above a wide spectrum of communication systems ranging from hard-wired connections to packet-switched or circuit-switched networks.

Meta's Exhibit No. 1014
Page 007

TCP is based on concepts first described by Cerf and Kahn in [1]. The TCP fits into a layered protocol architecture just above a basic Internet Protocol [2] which provides a way for the TCP to send and receive variable-length segments of information enclosed in internet datagram "envelopes". The internet datagram provides a means for addressing source and destination TCPs in different networks. The internet protocol also deals with any fragmentation or reassembly of the TCP segments required to achieve transport and delivery through multiple networks and interconnecting gateways. The internet protocol also carries information on the precedence, security classification and compartmentation of the TCP segments, so this information can be communicated end-to-end across multiple networks.

Protocol Layering

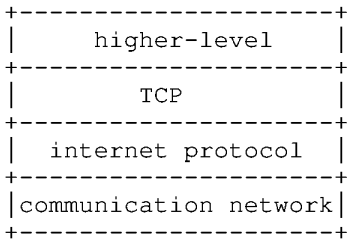


Figure 1

Much of this document is written in the context of TCP implementations which are co-resident with higher level protocols in the host computer. Some computer systems will be connected to networks via front-end computers which house the TCP and internet protocol layers, as well as network specific software. The TCP specification describes an interface to the higher level protocols which appears to be implementable even for the front-end case, as long as a suitable host-to-front end protocol is implemented.

1.2. Scope

The TCP is intended to provide a reliable process-to-process communication service in a multinetwork environment. The TCP is intended to be a host-to-host protocol in common use in multiple networks.

1.3. About this Document

This document represents a specification of the behavior required of any TCP implementation, both in its interactions with higher level protocols and in its interactions with other TCPs. The rest of this

Meta's Exhibit No. 1014
Page 008

section offers a very brief view of the protocol interfaces and operation. Section 2 summarizes the philosophical basis for the TCP design. Section 3 offers both a detailed description of the actions required of TCP when various events occur (arrival of new segments, user calls, errors, etc.) and the details of the formats of TCP segments.

1.4. Interfaces

The TCP interfaces on one side to user or application processes and on the other side to a lower level protocol such as Internet Protocol.

The interface between an application process and the TCP is illustrated in reasonable detail. This interface consists of a set of calls much like the calls an operating system provides to an application process for manipulating files. For example, there are calls to open and close connections and to send and receive data on established connections. It is also expected that the TCP can asynchronously communicate with application programs. Although considerable freedom is permitted to TCP implementors to design interfaces which are appropriate to a particular operating system environment, a minimum functionality is required at the TCP/user interface for any valid implementation.

The interface between TCP and lower level protocol is essentially unspecified except that it is assumed there is a mechanism whereby the two levels can asynchronously pass information to each other. Typically, one expects the lower level protocol to specify this interface. TCP is designed to work in a very general environment of interconnected networks. The lower level protocol which is assumed throughout this document is the Internet Protocol [2].

1.5. Operation

As noted above, the primary purpose of the TCP is to provide reliable, securable logical circuit or connection service between pairs of processes. To provide this service on top of a less reliable internet communication system requires facilities in the following areas:

- Basic Data Transfer
- Reliability
- Flow Control
- Multiplexing
- Connections
- Precedence and Security

The basic operation of the TCP in each of these areas is described in the following paragraphs.

Meta's Exhibit No. 1014
Page 009

Basic Data Transfer:

The TCP is able to transfer a continuous stream of octets in each direction between its users by packaging some number of octets into segments for transmission through the internet system. In general, the TCPs decide when to block and forward data at their own convenience.

Sometimes users need to be sure that all the data they have submitted to the TCP has been transmitted. For this purpose a push function is defined. To assure that data submitted to a TCP is actually transmitted the sending user indicates that it should be pushed through to the receiving user. A push causes the TCPs to promptly forward and deliver data up to that point to the receiver. The exact push point might not be visible to the receiving user and the push function does not supply a record boundary marker.

Reliability:

The TCP must recover from data that is damaged, lost, duplicated, or delivered out of order by the internet communication system. This is achieved by assigning a sequence number to each octet transmitted, and requiring a positive acknowledgment (ACK) from the receiving TCP. If the ACK is not received within a timeout interval, the data is retransmitted. At the receiver, the sequence numbers are used to correctly order segments that may be received out of order and to eliminate duplicates. Damage is handled by adding a checksum to each segment transmitted, checking it at the receiver, and discarding damaged segments.

As long as the TCPs continue to function properly and the internet system does not become completely partitioned, no transmission errors will affect the correct delivery of data. TCP recovers from internet communication system errors.

Flow Control:

TCP provides a means for the receiver to govern the amount of data sent by the sender. This is achieved by returning a "window" with every ACK indicating a range of acceptable sequence numbers beyond the last segment successfully received. The window indicates an allowed number of octets that the sender may transmit before receiving further permission.

Meta's Exhibit No. 1014
Page 0010

Multiplexing:

To allow for many processes within a single Host to use TCP communication facilities simultaneously, the TCP provides a set of addresses or ports within each host. Concatenated with the network and host addresses from the internet communication layer, this forms a socket. A pair of sockets uniquely identifies each connection. That is, a socket may be simultaneously used in multiple connections.

The binding of ports to processes is handled independently by each Host. However, it proves useful to attach frequently used processes (e.g., a "logger" or timesharing service) to fixed sockets which are made known to the public. These services can then be accessed through the known addresses. Establishing and learning the port addresses of other processes may involve more dynamic mechanisms.

Connections:

The reliability and flow control mechanisms described above require that TCPs initialize and maintain certain status information for each data stream. The combination of this information, including sockets, sequence numbers, and window sizes, is called a connection. Each connection is uniquely specified by a pair of sockets identifying its two sides.

When two processes wish to communicate, their TCP's must first establish a connection (initialize the status information on each side). When their communication is complete, the connection is terminated or closed to free the resources for other uses.

Since connections must be established between unreliable hosts and over the unreliable internet communication system, a handshake mechanism with clock-based sequence numbers is used to avoid erroneous initialization of connections.

Precedence and Security:

The users of TCP may indicate the security and precedence of their communication. Provision is made for default values to be used when these features are not needed.

Meta's Exhibit No. 1014
Page 0011

Meta's Exhibit No. 1014
Page 0012

2. PHILOSOPHY

2.1. Elements of the Internetwork System

The internetwork environment consists of hosts connected to networks which are in turn interconnected via gateways. It is assumed here that the networks may be either local networks (e.g., the ETHERNET) or large networks (e.g., the ARPANET), but in any case are based on packet switching technology. The active agents that produce and consume messages are processes. Various levels of protocols in the networks, the gateways, and the hosts support an interprocess communication system that provides two-way data flow on logical connections between process ports.

The term packet is used generically here to mean the data of one transaction between a host and its network. The format of data blocks exchanged within the a network will generally not be of concern to us.

Hosts are computers attached to a network, and from the communication network's point of view, are the sources and destinations of packets. Processes are viewed as the active elements in host computers (in accordance with the fairly common definition of a process as a program in execution). Even terminals and files or other I/O devices are viewed as communicating with each other through the use of processes. Thus, all communication is viewed as inter-process communication.

Since a process may need to distinguish among several communication streams between itself and another process (or processes), we imagine that each process may have a number of ports through which it communicates with the ports of other processes.

2.2. Model of Operation

Processes transmit data by calling on the TCP and passing buffers of data as arguments. The TCP packages the data from these buffers into segments and calls on the internet module to transmit each segment to the destination TCP. The receiving TCP places the data from a segment into the receiving user's buffer and notifies the receiving user. The TCPs include control information in the segments which they use to ensure reliable ordered data transmission.

The model of internet communication is that there is an internet protocol module associated with each TCP which provides an interface to the local network. This internet module packages TCP segments inside internet datagrams and routes these datagrams to a destination internet module or intermediate gateway. To transmit the datagram through the local network, it is embedded in a local network packet.

The packet switches may perform further packaging, fragmentation, or

Meta's Exhibit No. 1014
Page 0013

other operations to achieve the delivery of the local packet to the destination internet module.

At a gateway between networks, the internet datagram is "unwrapped" from its local packet and examined to determine through which network the internet datagram should travel next. The internet datagram is then "wrapped" in a local packet suitable to the next network and routed to the next gateway, or to the final destination.

A gateway is permitted to break up an internet datagram into smaller internet datagram fragments if this is necessary for transmission through the next network. To do this, the gateway produces a set of internet datagrams; each carrying a fragment. Fragments may be further broken into smaller fragments at subsequent gateways. The internet datagram fragment format is designed so that the destination internet module can reassemble fragments into internet datagrams.

A destination internet module unwraps the segment from the datagram (after reassembling the datagram, if necessary) and passes it to the destination TCP.

This simple model of the operation glosses over many details. One important feature is the type of service. This provides information to the gateway (or internet module) to guide it in selecting the service parameters to be used in traversing the next network. Included in the type of service information is the precedence of the datagram. Datagrams may also carry security information to permit host and gateways that operate in multilevel secure environments to properly segregate datagrams for security considerations.

2.3. The Host Environment

The TCP is assumed to be a module in an operating system. The users access the TCP much like they would access the file system. The TCP may call on other operating system functions, for example, to manage data structures. The actual interface to the network is assumed to be controlled by a device driver module. The TCP does not call on the network device driver directly, but rather calls on the internet datagram protocol module which may in turn call on the device driver.

The mechanisms of TCP do not preclude implementation of the TCP in a front-end processor. However, in such an implementation, a host-to-front-end protocol must provide the functionality to support the type of TCP-user interface described in this document.

Meta's Exhibit No. 1014
Page 0014

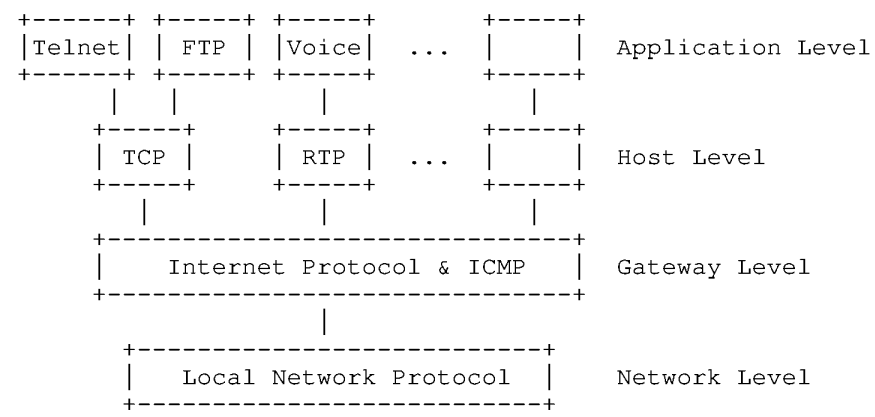
2.4. Interfaces

The TCP/user interface provides for calls made by the user on the TCP to OPEN or CLOSE a connection, to SEND or RECEIVE data, or to obtain STATUS about a connection. These calls are like other calls from user programs on the operating system, for example, the calls to open, read from, and close a file.

The TCP/internet interface provides calls to send and receive datagrams addressed to TCP modules in hosts anywhere in the internet system. These calls have parameters for passing the address, type of service, precedence, security, and other control information.

2.5. Relation to Other Protocols

The following diagram illustrates the place of the TCP in the protocol hierarchy:



Protocol Relationships

Figure 2.

It is expected that the TCP will be able to support higher level protocols efficiently. It should be easy to interface higher level protocols like the ARPANET Telnet or AUTODIN II THP to the TCP.

2.6. Reliable Communication

A stream of data sent on a TCP connection is delivered reliably and in order at the destination.

Meta's Exhibit No. 1014
Page 0015

Transmission is made reliable via the use of sequence numbers and acknowledgments. Conceptually, each octet of data is assigned a sequence number. The sequence number of the first octet of data in a segment is transmitted with that segment and is called the segment sequence number. Segments also carry an acknowledgment number which is the sequence number of the next expected data octet of transmissions in the reverse direction. When the TCP transmits a segment containing data, it puts a copy on a retransmission queue and starts a timer; when the acknowledgment for that data is received, the segment is deleted from the queue. If the acknowledgment is not received before the timer runs out, the segment is retransmitted.

An acknowledgment by TCP does not guarantee that the data has been delivered to the end user, but only that the receiving TCP has taken the responsibility to do so.

To govern the flow of data between TCPs, a flow control mechanism is employed. The receiving TCP reports a "window" to the sending TCP. This window specifies the number of octets, starting with the acknowledgment number, that the receiving TCP is currently prepared to receive.

2.7. Connection Establishment and Clearing

To identify the separate data streams that a TCP may handle, the TCP provides a port identifier. Since port identifiers are selected independently by each TCP they might not be unique. To provide for unique addresses within each TCP, we concatenate an internet address identifying the TCP with a port identifier to create a socket which will be unique throughout all networks connected together.

A connection is fully specified by the pair of sockets at the ends. A local socket may participate in many connections to different foreign sockets. A connection can be used to carry data in both directions, that is, it is "full duplex".

TCPs are free to associate ports with processes however they choose. However, several basic concepts are necessary in any implementation. There must be well-known sockets which the TCP associates only with the "appropriate" processes by some means. We envision that processes may "own" ports, and that processes can initiate connections only on the ports they own. (Means for implementing ownership is a local issue, but we envision a Request Port user command, or a method of uniquely allocating a group of ports to a given process, e.g., by associating the high order bits of a port name with a given process.)

A connection is specified in the OPEN call by the local port and foreign socket arguments. In return, the TCP supplies a (short) local

Meta's Exhibit No. 1014
Page 0016

connection name by which the user refers to the connection in subsequent calls. There are several things that must be remembered about a connection. To store this information we imagine that there is a data structure called a Transmission Control Block (TCB). One implementation strategy would have the local connection name be a pointer to the TCB for this connection. The OPEN call also specifies whether the connection establishment is to be actively pursued, or to be passively waited for.

A passive OPEN request means that the process wants to accept incoming connection requests rather than attempting to initiate a connection. Often the process requesting a passive OPEN will accept a connection request from any caller. In this case a foreign socket of all zeros is used to denote an unspecified socket. Unspecified foreign sockets are allowed only on passive OPENs.

A service process that wished to provide services for unknown other processes would issue a passive OPEN request with an unspecified foreign socket. Then a connection could be made with any process that requested a connection to this local socket. It would help if this local socket were known to be associated with this service.

Well-known sockets are a convenient mechanism for a priori associating a socket address with a standard service. For instance, the "Telnet-Server" process is permanently assigned to a particular socket, and other sockets are reserved for File Transfer, Remote Job Entry, Text Generator, Echoer, and Sink processes (the last three being for test purposes). A socket address might be reserved for access to a "Look-Up" service which would return the specific socket at which a newly created service would be provided. The concept of a well-known socket is part of the TCP specification, but the assignment of sockets to services is outside this specification. (See [4].)

Processes can issue passive OPENs and wait for matching active OPENs from other processes and be informed by the TCP when connections have been established. Two processes which issue active OPENs to each other at the same time will be correctly connected. This flexibility is critical for the support of distributed computing in which components act asynchronously with respect to each other.

There are two principal cases for matching the sockets in the local passive OPENs and an foreign active OPENs. In the first case, the local passive OPENs has fully specified the foreign socket. In this case, the match must be exact. In the second case, the local passive OPENs has left the foreign socket unspecified. In this case, any foreign socket is acceptable as long as the local sockets match. Other possibilities include partially restricted matches.

Meta's Exhibit No. 1014
Page 0017

If there are several pending passive OPENs (recorded in TCBs) with the same local socket, an foreign active OPEN will be matched to a TCB with the specific foreign socket in the foreign active OPEN, if such a TCB exists, before selecting a TCB with an unspecified foreign socket.

The procedures to establish connections utilize the synchronize (SYN) control flag and involves an exchange of three messages. This exchange has been termed a three-way hand shake [3].

A connection is initiated by the rendezvous of an arriving segment containing a SYN and a waiting TCB entry each created by a user OPEN command. The matching of local and foreign sockets determines when a connection has been initiated. The connection becomes "established" when sequence numbers have been synchronized in both directions.

The clearing of a connection also involves the exchange of segments, in this case carrying the FIN control flag.

2.8. Data Communication

The data that flows on a connection may be thought of as a stream of octets. The sending user indicates in each SEND call whether the data in that call (and any preceeding calls) should be immediately pushed through to the receiving user by the setting of the PUSH flag.

A sending TCP is allowed to collect data from the sending user and to send that data in segments at its own convenience, until the push function is signaled, then it must send all unsent data. When a receiving TCP sees the PUSH flag, it must not wait for more data from the sending TCP before passing the data to the receiving process.

There is no necessary relationship between push functions and segment boundaries. The data in any particular segment may be the result of a single SEND call, in whole or part, or of multiple SEND calls.

The purpose of push function and the PUSH flag is to push data through from the sending user to the receiving user. It does not provide a record service.

There is a coupling between the push function and the use of buffers of data that cross the TCP/user interface. Each time a PUSH flag is associated with data placed into the receiving user's buffer, the buffer is returned to the user for processing even if the buffer is not filled. If data arrives that fills the user's buffer before a PUSH is seen, the data is passed to the user in buffer size units.

TCP also provides a means to communicate to the receiver of data that at some point further along in the data stream than the receiver is

Meta's Exhibit No. 1014
Page 0018

currently reading there is urgent data. TCP does not attempt to define what the user specifically does upon being notified of pending urgent data, but the general notion is that the receiving process will take action to process the urgent data quickly.

2.9. Precedence and Security

The TCP makes use of the internet protocol type of service field and security option to provide precedence and security on a per connection basis to TCP users. Not all TCP modules will necessarily function in a multilevel secure environment; some may be limited to unclassified use only, and others may operate at only one security level and compartment. Consequently, some TCP implementations and services to users may be limited to a subset of the multilevel secure case.

TCP modules which operate in a multilevel secure environment must properly mark outgoing segments with the security, compartment, and precedence. Such TCP modules must also provide to their users or higher level protocols such as Telnet or THP an interface to allow them to specify the desired security level, compartment, and precedence of connections.

2.10. Robustness Principle

TCP implementations will follow a general principle of robustness: be conservative in what you do, be liberal in what you accept from others.

Meta's Exhibit No. 1014
Page 0019

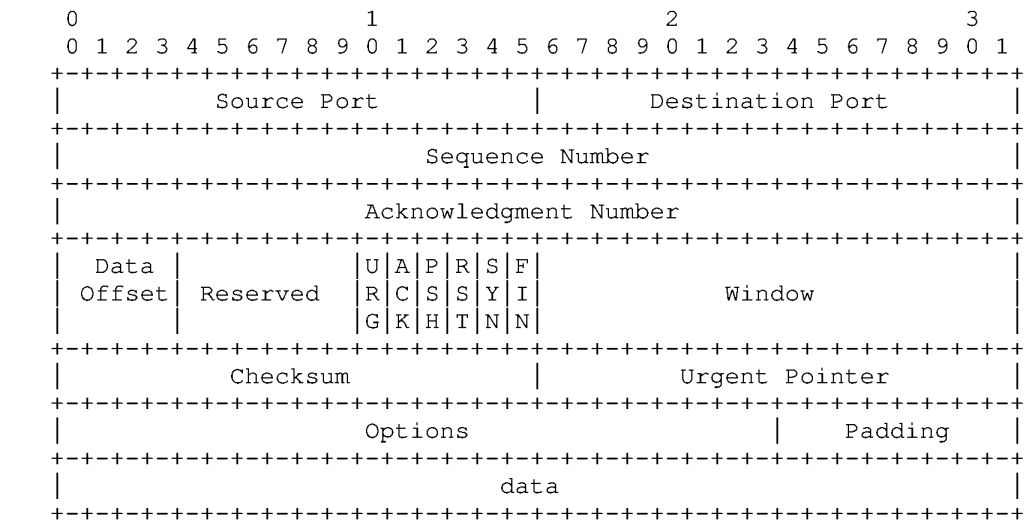
Meta's Exhibit No. 1014
Page 0020

3. FUNCTIONAL SPECIFICATION

3.1. Header Format

TCP segments are sent as internet datagrams. The Internet Protocol header carries several information fields, including the source and destination host addresses [2]. A TCP header follows the internet header, supplying information specific to the TCP protocol. This division allows for the existence of host level protocols other than TCP.

TCP Header Format



TCP Header Format

Note that one tick mark represents one bit position.

Figure 3.

Source Port: 16 bits

The source port number.

Destination Port: 16 bits

The destination port number.

Meta's Exhibit No. 1014
Page 0021

Sequence Number: 32 bits

The sequence number of the first data octet in this segment (except when SYN is present). If SYN is present the sequence number is the initial sequence number (ISN) and the first data octet is ISN+1.

Acknowledgment Number: 32 bits

If the ACK control bit is set this field contains the value of the next sequence number the sender of the segment is expecting to receive. Once a connection is established this is always sent.

Data Offset: 4 bits

The number of 32 bit words in the TCP Header. This indicates where the data begins. The TCP header (even one including options) is an integral number of 32 bits long.

Reserved: 6 bits

Reserved for future use. Must be zero.

Control Bits: 6 bits (from left to right):

URG: Urgent Pointer field significant
ACK: Acknowledgment field significant
PSH: Push Function
RST: Reset the connection
SYN: Synchronize sequence numbers
FIN: No more data from sender

Window: 16 bits

The number of data octets beginning with the one indicated in the acknowledgment field which the sender of this segment is willing to accept.

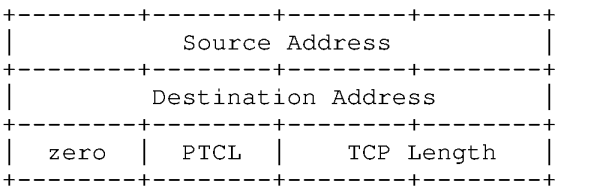
Checksum: 16 bits

The checksum field is the 16 bit one's complement of the one's complement sum of all 16 bit words in the header and text. If a segment contains an odd number of header and text octets to be checksummed, the last octet is padded on the right with zeros to form a 16 bit word for checksum purposes. The pad is not transmitted as part of the segment. While computing the checksum, the checksum field itself is replaced with zeros.

The checksum also covers a 96 bit pseudo header conceptually

Meta's Exhibit No. 1014
Page 0022

prefixed to the TCP header. This pseudo header contains the Source Address, the Destination Address, the Protocol, and TCP length. This gives the TCP protection against misrouted segments. This information is carried in the Internet Protocol and is transferred across the TCP/Network interface in the arguments or results of calls by the TCP on the IP.



The TCP Length is the TCP header length plus the data length in octets (this is not an explicitly transmitted quantity, but is computed), and it does not count the 12 octets of the pseudo header.

Urgent Pointer: 16 bits

This field communicates the current value of the urgent pointer as a positive offset from the sequence number in this segment. The urgent pointer points to the sequence number of the octet following the urgent data. This field is only be interpreted in segments with the URG control bit set.

Options: variable

Options may occupy space at the end of the TCP header and are a multiple of 8 bits in length. All options are included in the checksum. An option may begin on any octet boundary. There are two cases for the format of an option:

Case 1: A single octet of option-kind.

Case 2: An octet of option-kind, an octet of option-length, and the actual option-data octets.

The option-length counts the two octets of option-kind and option-length as well as the option-data octets.

Note that the list of options may be shorter than the data offset field might imply. The content of the header beyond the End-of-Option option must be header padding (i.e., zero).

A TCP must implement all options.

Meta's Exhibit No. 1014
Page 0023

Currently defined options include (kind indicated in octal):

Kind	Length	Meaning
0	-	End of option list.
1	-	No-Operation.
2	4	Maximum Segment Size.

Specific Option Definitions

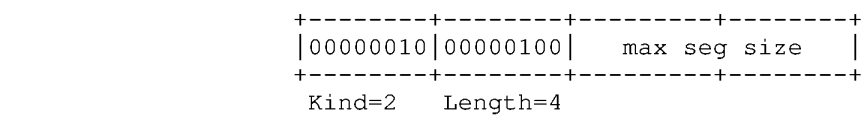
End of Option List

This option code indicates the end of the option list. This might not coincide with the end of the TCP header according to the Data Offset field. This is used at the end of all options, not the end of each option, and need only be used if the end of the options would not otherwise coincide with the end of the TCP header.

No-Operation

This option code may be used between options, for example, to align the beginning of a subsequent option on a word boundary. There is no guarantee that senders will use this option, so receivers must be prepared to process options even if they do not begin on a word boundary.

Maximum Segment Size



Meta's Exhibit No. 1014
Page 0024

Maximum Segment Size Option Data: 16 bits

If this option is present, then it communicates the maximum receive segment size at the TCP which sends this segment. This field must only be sent in the initial connection request (i.e., in segments with the SYN control bit set). If this option is not used, any segment size is allowed.

Padding: variable

The TCP header padding is used to ensure that the TCP header ends and data begins on a 32 bit boundary. The padding is composed of zeros.

3.2. Terminology

Before we can discuss very much about the operation of the TCP we need to introduce some detailed terminology. The maintenance of a TCP connection requires the remembering of several variables. We conceive of these variables being stored in a connection record called a Transmission Control Block or TCB. Among the variables stored in the TCB are the local and remote socket numbers, the security and precedence of the connection, pointers to the user's send and receive buffers, pointers to the retransmit queue and to the current segment. In addition several variables relating to the send and receive sequence numbers are stored in the TCB.

Send Sequence Variables

```
SND.UNA - send unacknowledged
SND.NXT - send next
SND.WND - send window
SND.UP  - send urgent pointer
SND.WL1 - segment sequence number used for last window update
SND.WL2 - segment acknowledgment number used for last window
          update
ISS     - initial send sequence number
```

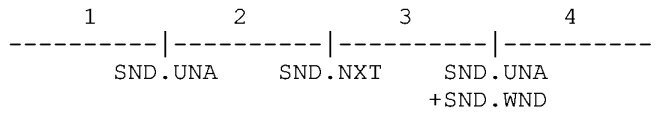
Receive Sequence Variables

```
RCV.NXT - receive next
RCV.WND - receive window
RCV.UP  - receive urgent pointer
IRS     - initial receive sequence number
```


Meta's Exhibit No. 1014
Page 0025

The following diagrams may help to relate some of these variables to the sequence space.

Send Sequence Space



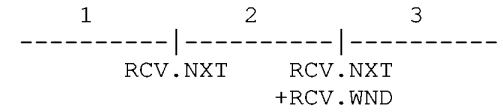
1 - old sequence numbers which have been acknowledged
2 - sequence numbers of unacknowledged data
3 - sequence numbers allowed for new data transmission
4 - future sequence numbers which are not yet allowed

Send Sequence Space

Figure 4.

The send window is the portion of the sequence space labeled 3 in figure 4.

Receive Sequence Space



1 - old sequence numbers which have been acknowledged
2 - sequence numbers allowed for new reception
3 - future sequence numbers which are not yet allowed

Receive Sequence Space

Figure 5.

The receive window is the portion of the sequence space labeled 2 in figure 5.

There are also some variables used frequently in the discussion that take their values from the fields of the current segment.

Meta's Exhibit No. 1014
Page 0026

Current Segment Variables

SEG.SEQ - segment sequence number
SEG.ACK - segment acknowledgment number
SEG.LEN - segment length
SEG.WND - segment window
SEG.UP - segment urgent pointer
SEG.PRC - segment precedence value

A connection progresses through a series of states during its lifetime. The states are: LISTEN, SYN-SENT, SYN-RECEIVED, ESTABLISHED, FIN-WAIT-1, FIN-WAIT-2, CLOSE-WAIT, CLOSING, LAST-ACK, TIME-WAIT, and the fictional state CLOSED. CLOSED is fictional because it represents the state when there is no TCB, and therefore, no connection. Briefly the meanings of the states are:

LISTEN - represents waiting for a connection request from any remote TCP and port.

SYN-SENT - represents waiting for a matching connection request after having sent a connection request.

SYN-RECEIVED - represents waiting for a confirming connection request acknowledgment after having both received and sent a connection request.

ESTABLISHED - represents an open connection, data received can be delivered to the user. The normal state for the data transfer phase of the connection.

FIN-WAIT-1 - represents waiting for a connection termination request from the remote TCP, or an acknowledgment of the connection termination request previously sent.

FIN-WAIT-2 - represents waiting for a connection termination request from the remote TCP.

CLOSE-WAIT - represents waiting for a connection termination request from the local user.

CLOSING - represents waiting for a connection termination request acknowledgment from the remote TCP.

LAST-ACK - represents waiting for an acknowledgment of the connection termination request previously sent to the remote TCP (which includes an acknowledgment of its connection termination request).

Meta's Exhibit No. 1014
Page 0027

TIME-WAIT - represents waiting for enough time to pass to be sure the remote TCP received the acknowledgment of its connection termination request.

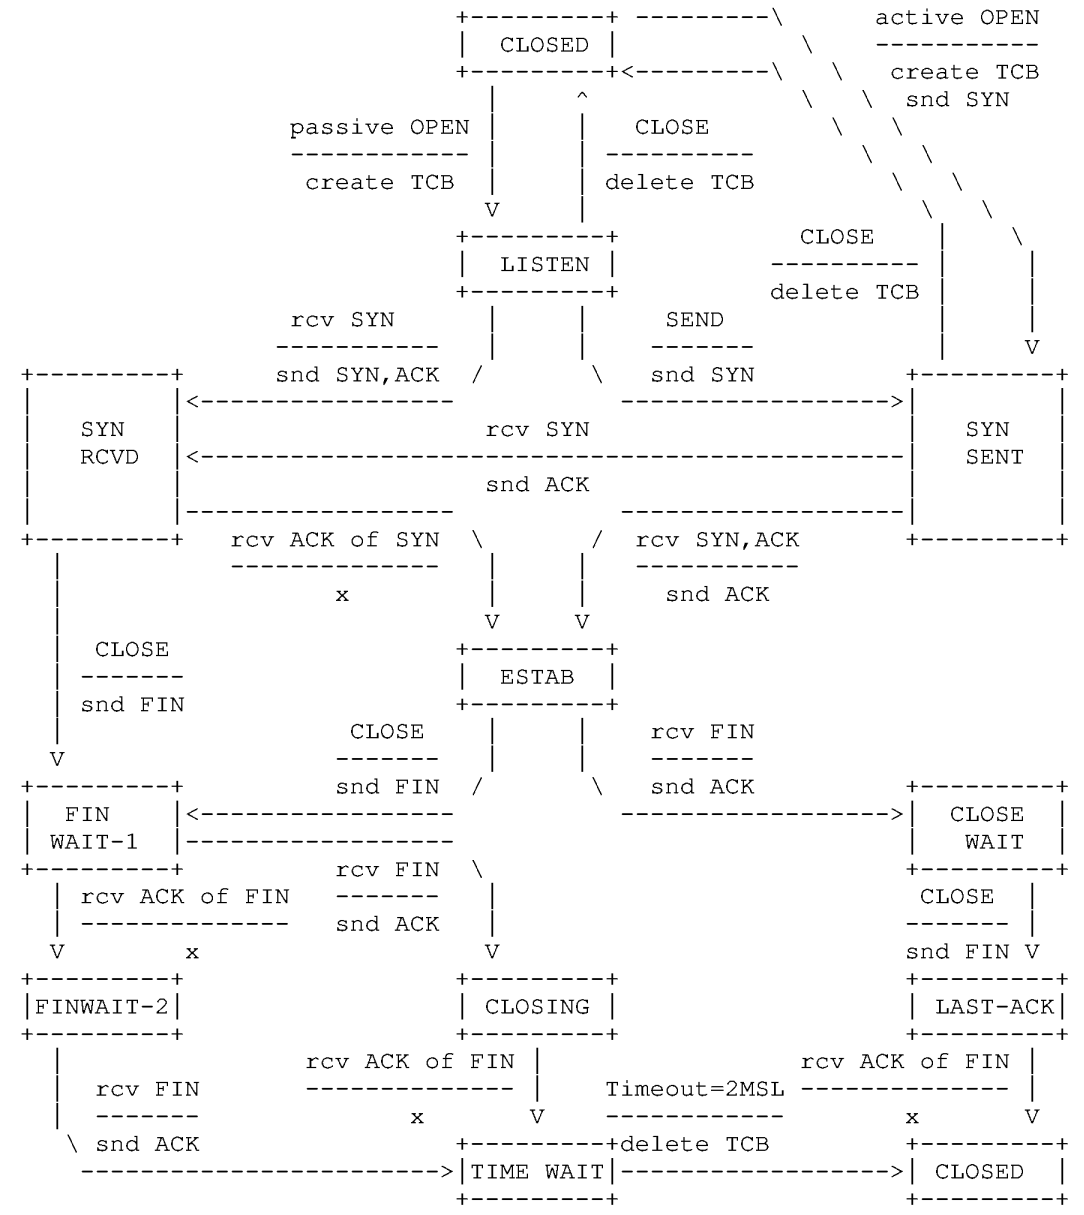
CLOSED - represents no connection state at all.

A TCP connection progresses from one state to another in response to events. The events are the user calls, OPEN, SEND, RECEIVE, CLOSE, ABORT, and STATUS; the incoming segments, particularly those containing the SYN, ACK, RST and FIN flags; and timeouts.

The state diagram in figure 6 illustrates only state changes, together with the causing events and resulting actions, but addresses neither error conditions nor actions which are not connected with state changes. In a later section, more detail is offered with respect to the reaction of the TCP to events.

NOTE BENE: this diagram is only a summary and must not be taken as the total specification.

Meta's Exhibit No. 1014
Page 0028



TCP Connection State Diagram
Figure 6.

3.3. Sequence Numbers

A fundamental notion in the design is that every octet of data sent over a TCP connection has a sequence number. Since every octet is sequenced, each of them can be acknowledged. The acknowledgment mechanism employed is cumulative so that an acknowledgment of sequence number X indicates that all octets up to but not including X have been received. This mechanism allows for straight-forward duplicate detection in the presence of retransmission. Numbering of octets within a segment is that the first data octet immediately following the header is the lowest numbered, and the following octets are numbered consecutively.

It is essential to remember that the actual sequence number space is finite, though very large. This space ranges from 0 to 2**32 - 1. Since the space is finite, all arithmetic dealing with sequence numbers must be performed modulo 2**32. This unsigned arithmetic preserves the relationship of sequence numbers as they cycle from 2**32 - 1 to 0 again. There are some subtleties to computer modulo arithmetic, so great care should be taken in programming the comparison of such values. The symbol "=<" means "less than or equal" (modulo 2**32).

The typical kinds of sequence number comparisons which the TCP must perform include:

(a) Determining that an acknowledgment refers to some sequence number sent but not yet acknowledged.

(b) Determining that all sequence numbers occupied by a segment have been acknowledged (e.g., to remove the segment from a retransmission queue).

(c) Determining that an incoming segment contains sequence numbers which are expected (i.e., that the segment "overlaps" the receive window).

Meta's Exhibit No. 1014
Page 0030

In response to sending data the TCP will receive acknowledgments. The following comparisons are needed to process the acknowledgments.

SND.UNA = oldest unacknowledged sequence number

SND.NXT = next sequence number to be sent

SEG.ACK = acknowledgment from the receiving TCP (next sequence number expected by the receiving TCP)

SEG.SEQ = first sequence number of a segment

SEG.LEN = the number of octets occupied by the data in the segment (counting SYN and FIN)

SEG.SEQ+SEG.LEN-1 = last sequence number of a segment

A new acknowledgment (called an "acceptable ack"), is one for which the inequality below holds:

SND.UNA < SEG.ACK =< SND.NXT

A segment on the retransmission queue is fully acknowledged if the sum of its sequence number and length is less or equal than the acknowledgment value in the incoming segment.

When data is received the following comparisons are needed:

RCV.NXT = next sequence number expected on an incoming segments, and is the left or lower edge of the receive window

RCV.NXT+RCV.WND-1 = last sequence number expected on an incoming segment, and is the right or upper edge of the receive window

SEG.SEQ = first sequence number occupied by the incoming segment

SEG.SEQ+SEG.LEN-1 = last sequence number occupied by the incoming segment

A segment is judged to occupy a portion of valid receive sequence space if

RCV.NXT =< SEG.SEQ < RCV.NXT+RCV.WND

or

RCV.NXT =< SEG.SEQ+SEG.LEN-1 < RCV.NXT+RCV.WND

Meta's Exhibit No. 1014
Page 0031

The first part of this test checks to see if the beginning of the segment falls in the window, the second part of the test checks to see if the end of the segment falls in the window; if the segment passes either part of the test it contains data in the window.

Actually, it is a little more complicated than this. Due to zero windows and zero length segments, we have four cases for the acceptability of an incoming segment:

Segment Length	Receive Window	Test
0	0	SEG.SEQ = RCV.NXT
0	>0	RCV.NXT =< SEG.SEQ < RCV.NXT+RCV.WND
>0	0	not acceptable
>0	>0	RCV.NXT =< SEG.SEQ < RCV.NXT+RCV.WND or RCV.NXT =< SEG.SEQ+SEG.LEN-1 < RCV.NXT+RCV.WND

Note that when the receive window is zero no segments should be acceptable except ACK segments. Thus, it is be possible for a TCP to maintain a zero receive window while transmitting data and receiving ACKs. However, even when the receive window is zero, a TCP must process the RST and URG fields of all incoming segments.

We have taken advantage of the numbering scheme to protect certain control information as well. This is achieved by implicitly including some control flags in the sequence space so they can be retransmitted and acknowledged without confusion (i.e., one and only one copy of the control will be acted upon). Control information is not physically carried in the segment data space. Consequently, we must adopt rules for implicitly assigning sequence numbers to control. The SYN and FIN are the only controls requiring this protection, and these controls are used only at connection opening and closing. For sequence number purposes, the SYN is considered to occur before the first actual data octet of the segment in which it occurs, while the FIN is considered to occur after the last actual data octet in a segment in which it occurs. The segment length (SEG.LEN) includes both data and sequence space occupying controls. When a SYN is present then SEG.SEQ is the sequence number of the SYN.

Meta's Exhibit No. 1014
Page 0032

Initial Sequence Number Selection

The protocol places no restriction on a particular connection being used over and over again. A connection is defined by a pair of sockets. New instances of a connection will be referred to as incarnations of the connection. The problem that arises from this is -- "how does the TCP identify duplicate segments from previous incarnations of the connection?" This problem becomes apparent if the connection is being opened and closed in quick succession, or if the connection breaks with loss of memory and is then reestablished.

To avoid confusion we must prevent segments from one incarnation of a connection from being used while the same sequence numbers may still be present in the network from an earlier incarnation. We want to assure this, even if a TCP crashes and loses all knowledge of the sequence numbers it has been using. When new connections are created, an initial sequence number (ISN) generator is employed which selects a new 32 bit ISN. The generator is bound to a (possibly fictitious) 32 bit clock whose low order bit is incremented roughly every 4 microseconds. Thus, the ISN cycles approximately every 4.55 hours. Since we assume that segments will stay in the network no more than the Maximum Segment Lifetime (MSL) and that the MSL is less than 4.55 hours we can reasonably assume that ISN's will be unique.

For each connection there is a send sequence number and a receive sequence number. The initial send sequence number (ISS) is chosen by the data sending TCP, and the initial receive sequence number (IRS) is learned during the connection establishing procedure.

For a connection to be established or initialized, the two TCPs must synchronize on each other's initial sequence numbers. This is done in an exchange of connection establishing segments carrying a control bit called "SYN" (for synchronize) and the initial sequence numbers. As a shorthand, segments carrying the SYN bit are also called "SYNs". Hence, the solution requires a suitable mechanism for picking an initial sequence number and a slightly involved handshake to exchange the ISN's.

The synchronization requires each side to send it's own initial sequence number and to receive a confirmation of it in acknowledgment from the other side. Each side must also receive the other side's initial sequence number and send a confirming acknowledgment.

1) A --> B SYN my sequence number is X
2) A <-- B ACK your sequence number is X
3) A <-- B SYN my sequence number is Y
4) A --> B ACK your sequence number is Y

Meta's Exhibit No. 1014
Page 0033

Because steps 2 and 3 can be combined in a single message this is called the three way (or three message) handshake.

A three way handshake is necessary because sequence numbers are not tied to a global clock in the network, and TCPs may have different mechanisms for picking the ISN's. The receiver of the first SYN has no way of knowing whether the segment was an old delayed one or not, unless it remembers the last sequence number used on the connection (which is not always possible), and so it must ask the sender to verify this SYN. The three way handshake and the advantages of a clock-driven scheme are discussed in [3].

Knowing When to Keep Quiet

To be sure that a TCP does not create a segment that carries a sequence number which may be duplicated by an old segment remaining in the network, the TCP must keep quiet for a maximum segment lifetime (MSL) before assigning any sequence numbers upon starting up or recovering from a crash in which memory of sequence numbers in use was lost. For this specification the MSL is taken to be 2 minutes. This is an engineering choice, and may be changed if experience indicates it is desirable to do so. Note that if a TCP is reinitialized in some sense, yet retains its memory of sequence numbers in use, then it need not wait at all; it must only be sure to use sequence numbers larger than those recently used.

The TCP Quiet Time Concept

This specification provides that hosts which "crash" without retaining any knowledge of the last sequence numbers transmitted on each active (i.e., not closed) connection shall delay emitting any TCP segments for at least the agreed Maximum Segment Lifetime (MSL) in the internet system of which the host is a part. In the paragraphs below, an explanation for this specification is given. TCP implementors may violate the "quiet time" restriction, but only at the risk of causing some old data to be accepted as new or new data rejected as old duplicated by some receivers in the internet system.

TCPs consume sequence number space each time a segment is formed and entered into the network output queue at a source host. The duplicate detection and sequencing algorithm in the TCP protocol relies on the unique binding of segment data to sequence space to the extent that sequence numbers will not cycle through all 2**32 values before the segment data bound to those sequence numbers has been delivered and acknowledged by the receiver and all duplicate copies of the segments have "drained" from the internet. Without such an assumption, two distinct TCP segments could conceivably be

Meta's Exhibit No. 1014
Page 0034

assigned the same or overlapping sequence numbers, causing confusion at the receiver as to which data is new and which is old. Remember that each segment is bound to as many consecutive sequence numbers as there are octets of data in the segment.

Under normal conditions, TCPs keep track of the next sequence number to emit and the oldest awaiting acknowledgment so as to avoid mistakenly using a sequence number over before its first use has been acknowledged. This alone does not guarantee that old duplicate data is drained from the net, so the sequence space has been made very large to reduce the probability that a wandering duplicate will cause trouble upon arrival. At 2 megabits/sec. it takes 4.5 hours to use up 2**32 octets of sequence space. Since the maximum segment lifetime in the net is not likely to exceed a few tens of seconds, this is deemed ample protection for foreseeable nets, even if data rates escalate to 10's of megabits/sec. At 100 megabits/sec, the cycle time is 5.4 minutes which may be a little short, but still within reason.

The basic duplicate detection and sequencing algorithm in TCP can be defeated, however, if a source TCP does not have any memory of the sequence numbers it last used on a given connection. For example, if the TCP were to start all connections with sequence number 0, then upon crashing and restarting, a TCP might re-form an earlier connection (possibly after half-open connection resolution) and emit packets with sequence numbers identical to or overlapping with packets still in the network which were emitted on an earlier incarnation of the same connection. In the absence of knowledge about the sequence numbers used on a particular connection, the TCP specification recommends that the source delay for MSL seconds before emitting segments on the connection, to allow time for segments from the earlier connection incarnation to drain from the system.

Even hosts which can remember the time of day and used it to select initial sequence number values are not immune from this problem (i.e., even if time of day is used to select an initial sequence number for each new connection incarnation).

Suppose, for example, that a connection is opened starting with sequence number S. Suppose that this connection is not used much and that eventually the initial sequence number function (ISN(t)) takes on a value equal to the sequence number, say S1, of the last segment sent by this TCP on a particular connection. Now suppose, at this instant, the host crashes, recovers, and establishes a new incarnation of the connection. The initial sequence number chosen is S1 = ISN(t) -- last used sequence number on old incarnation of connection! If the recovery occurs quickly enough, any old

Meta's Exhibit No. 1014
Page 0035

duplicates in the net bearing sequence numbers in the neighborhood of S1 may arrive and be treated as new packets by the receiver of the new incarnation of the connection.

The problem is that the recovering host may not know for how long it crashed nor does it know whether there are still old duplicates in the system from earlier connection incarnations.

One way to deal with this problem is to deliberately delay emitting segments for one MSL after recovery from a crash- this is the "quite time" specification. Hosts which prefer to avoid waiting are willing to risk possible confusion of old and new packets at a given destination may choose not to wait for the "quite time". Implementors may provide TCP users with the ability to select on a connection by connection basis whether to wait after a crash, or may informally implement the "quite time" for all connections. Obviously, even where a user selects to "wait," this is not necessary after the host has been "up" for at least MSL seconds.

To summarize: every segment emitted occupies one or more sequence numbers in the sequence space, the numbers occupied by a segment are "busy" or "in use" until MSL seconds have passed, upon crashing a block of space-time is occupied by the octets of the last emitted segment, if a new connection is started too soon and uses any of the sequence numbers in the space-time footprint of the last segment of the previous connection incarnation, there is a potential sequence number overlap area which could cause confusion at the receiver.

3.4. Establishing a connection

The "three-way handshake" is the procedure used to establish a connection. This procedure normally is initiated by one TCP and responded to by another TCP. The procedure also works if two TCP simultaneously initiate the procedure. When simultaneous attempt occurs, each TCP receives a "SYN" segment which carries no acknowledgment after it has sent a "SYN". Of course, the arrival of an old duplicate "SYN" segment can potentially make it appear, to the recipient, that a simultaneous connection initiation is in progress. Proper use of "reset" segments can disambiguate these cases.

Several examples of connection initiation follow. Although these examples do not show connection synchronization using data-carrying segments, this is perfectly legitimate, so long as the receiving TCP doesn't deliver the data to the user until it is clear the data is valid (i.e., the data must be buffered at the receiver until the connection reaches the ESTABLISHED state). The three-way handshake reduces the possibility of false connections. It is the

Meta's Exhibit No. 1014
Page 0036

implementation of a trade-off between memory and messages to provide information for this checking.

The simplest three-way handshake is shown in figure 7 below. The figures should be interpreted in the following way. Each line is numbered for reference purposes. Right arrows (-->) indicate departure of a TCP segment from TCP A to TCP B, or arrival of a segment at B from A. Left arrows (<--), indicate the reverse. Ellipsis (...) indicates a segment which is still in the network (delayed). An "XXX" indicates a segment which is lost or rejected. Comments appear in parentheses. TCP states represent the state AFTER the departure or arrival of the segment (whose contents are shown in the center of each line). Segment contents are shown in abbreviated form, with sequence number, control flags, and ACK field. Other fields such as window, addresses, lengths, and text have been left out in the interest of clarity.

```
    TCP A                                                TCP B

1.  CLOSED                                               LISTEN

2.  SYN-SENT    --> <SEQ=100><CTL=SYN>               --> SYN-RECEIVED

3.  ESTABLISHED <-- <SEQ=300><ACK=101><CTL=SYN,ACK>  <-- SYN-RECEIVED

4.  ESTABLISHED --> <SEQ=101><ACK=301><CTL=ACK>       --> ESTABLISHED

5.  ESTABLISHED --> <SEQ=101><ACK=301><CTL=ACK><DATA> --> ESTABLISHED
```

Basic 3-Way Handshake for Connection Synchronization

Figure 7.

In line 2 of figure 7, TCP A begins by sending a SYN segment indicating that it will use sequence numbers starting with sequence number 100. In line 3, TCP B sends a SYN and acknowledges the SYN it received from TCP A. Note that the acknowledgment field indicates TCP B is now expecting to hear sequence 101, acknowledging the SYN which occupied sequence 100.

At line 4, TCP A responds with an empty segment containing an ACK for TCP B's SYN; and in line 5, TCP A sends some data. Note that the sequence number of the segment in line 5 is the same as in line 4 because the ACK does not occupy sequence number space (if it did, we would wind up ACKing ACK's!).

Meta's Exhibit No. 1014
Page 0037

Simultaneous initiation is only slightly more complex, as is shown in figure 8. Each TCP cycles from CLOSED to SYN-SENT to SYN-RECEIVED to ESTABLISHED.

```
    TCP A                                                TCP B

1.  CLOSED                                               CLOSED

2.  SYN-SENT     --> <SEQ=100><CTL=SYN>              ...

3.  SYN-RECEIVED <-- <SEQ=300><CTL=SYN>              <-- SYN-SENT

4.               ... <SEQ=100><CTL=SYN>              --> SYN-RECEIVED

5.  SYN-RECEIVED --> <SEQ=100><ACK=301><CTL=SYN,ACK> ...

6.  ESTABLISHED  <-- <SEQ=300><ACK=101><CTL=SYN,ACK> <-- SYN-RECEIVED

7.               ... <SEQ=101><ACK=301><CTL=ACK>     --> ESTABLISHED
```

Simultaneous Connection Synchronization

Figure 8.

The principle reason for the three-way handshake is to prevent old duplicate connection initiations from causing confusion. To deal with this, a special control message, reset, has been devised. If the receiving TCP is in a non-synchronized state (i.e., SYN-SENT, SYN-RECEIVED), it returns to LISTEN on receiving an acceptable reset. If the TCP is in one of the synchronized states (ESTABLISHED, FIN-WAIT-1, FIN-WAIT-2, CLOSE-WAIT, CLOSING, LAST-ACK, TIME-WAIT), it aborts the connection and informs its user. We discuss this latter case under "half-open" connections below.

Meta's Exhibit No. 1014
Page 0038

```
    TCP A                                                TCP B

1.  CLOSED                                               LISTEN

2.  SYN-SENT    --> <SEQ=100><CTL=SYN>               ...

3.  (duplicate) ... <SEQ=90><CTL=SYN>               --> SYN-RECEIVED

4.  SYN-SENT    <-- <SEQ=300><ACK=91><CTL=SYN,ACK>  <-- SYN-RECEIVED

5.  SYN-SENT    --> <SEQ=91><CTL=RST>               --> LISTEN

6.              ... <SEQ=100><CTL=SYN>               --> SYN-RECEIVED

7.  SYN-SENT    <-- <SEQ=400><ACK=101><CTL=SYN,ACK>  <-- SYN-RECEIVED

8.  ESTABLISHED --> <SEQ=101><ACK=401><CTL=ACK>      --> ESTABLISHED
```

Recovery from Old Duplicate SYN

Figure 9.

As a simple example of recovery from old duplicates, consider figure 9. At line 3, an old duplicate SYN arrives at TCP B. TCP B cannot tell that this is an old duplicate, so it responds normally (line 4). TCP A detects that the ACK field is incorrect and returns a RST (reset) with its SEQ field selected to make the segment believable. TCP B, on receiving the RST, returns to the LISTEN state. When the original SYN (pun intended) finally arrives at line 6, the synchronization proceeds normally. If the SYN at line 6 had arrived before the RST, a more complex exchange might have occurred with RST's sent in both directions.

Half-Open Connections and Other Anomalies

An established connection is said to be "half-open" if one of the TCPs has closed or aborted the connection at its end without the knowledge of the other, or if the two ends of the connection have become desynchronized owing to a crash that resulted in loss of memory. Such connections will automatically become reset if an attempt is made to send data in either direction. However, half-open connections are expected to be unusual, and the recovery procedure is mildly involved.

If at site A the connection no longer exists, then an attempt by the

Meta's Exhibit No. 1014
Page 0039

user at site B to send any data on it will result in the site B TCP receiving a reset control message. Such a message indicates to the site B TCP that something is wrong, and it is expected to abort the connection.

Assume that two user processes A and B are communicating with one another when a crash occurs causing loss of memory to A's TCP. Depending on the operating system supporting A's TCP, it is likely that some error recovery mechanism exists. When the TCP is up again, A is likely to start again from the beginning or from a recovery point. As a result, A will probably try to OPEN the connection again or try to SEND on the connection it believes open. In the latter case, it receives the error message "connection not open" from the local (A's) TCP. In an attempt to establish the connection, A's TCP will send a segment containing SYN. This scenario leads to the example shown in figure 10. After TCP A crashes, the user attempts to re-open the connection. TCP B, in the meantime, thinks the connection is open.

```
    TCP A                                                TCP B

1.  (CRASH)                                  (send 300,receive 100)

2.  CLOSED                                               ESTABLISHED

3.  SYN-SENT --> <SEQ=400><CTL=SYN>              --> (??)

4.  (!!)     <-- <SEQ=300><ACK=100><CTL=ACK>     <-- ESTABLISHED

5.  SYN-SENT --> <SEQ=100><CTL=RST>              --> (Abort!!)

6.  SYN-SENT                                             CLOSED

7.  SYN-SENT --> <SEQ=400><CTL=SYN>              -->
```

Half-Open Connection Discovery

Figure 10.

When the SYN arrives at line 3, TCP B, being in a synchronized state, and the incoming segment outside the window, responds with an acknowledgment indicating what sequence it next expects to hear (ACK 100). TCP A sees that this segment does not acknowledge anything it sent and, being unsynchronized, sends a reset (RST) because it has detected a half-open connection. TCP B aborts at line 5. TCP A will

Meta's Exhibit No. 1014
Page 0040

continue to try to establish the connection; the problem is now reduced to the basic 3-way handshake of figure 7.

An interesting alternative case occurs when TCP A crashes and TCP B tries to send data on what it thinks is a synchronized connection. This is illustrated in figure 11. In this case, the data arriving at TCP A from TCP B (line 2) is unacceptable because no such connection exists, so TCP A sends a RST. The RST is acceptable so TCP B processes it and aborts the connection.

```
      TCP A                                                TCP B

1.  (CRASH)                                      (send 300,receive 100)

2.  (??)    <-- <SEQ=300><ACK=100><DATA=10><CTL=ACK> <-- ESTABLISHED

3.          --> <SEQ=100><CTL=RST>                   --> (ABORT!!)
```

Active Side Causes Half-Open Connection Discovery

Figure 11.

In figure 12, we find the two TCPs A and B with passive connections waiting for SYN. An old duplicate arriving at TCP B (line 2) stirs B into action. A SYN-ACK is returned (line 3) and causes TCP A to generate a RST (the ACK in line 3 is not acceptable). TCP B accepts the reset and returns to its passive LISTEN state.

```
    TCP A                                          TCP B

1.  LISTEN                                         LISTEN

2.       ... <SEQ=Z><CTL=SYN>                -->  SYN-RECEIVED

3.  (??) <-- <SEQ=X><ACK=Z+1><CTL=SYN,ACK>   <--  SYN-RECEIVED

4.       --> <SEQ=Z+1><CTL=RST>              -->  (return to LISTEN!)

5.  LISTEN                                         LISTEN
```

Old Duplicate SYN Initiates a Reset on two Passive Sockets

Figure 12.

Meta's Exhibit No. 1014
Page 0041

A variety of other cases are possible, all of which are accounted for by the following rules for RST generation and processing.

Reset Generation

As a general rule, reset (RST) must be sent whenever a segment arrives which apparently is not intended for the current connection. A reset must not be sent if it is not clear that this is the case.

There are three groups of states:

1. If the connection does not exist (CLOSED) then a reset is sent in response to any incoming segment except another reset. In particular, SYNs addressed to a non-existent connection are rejected by this means.

   If the incoming segment has an ACK field, the reset takes its sequence number from the ACK field of the segment, otherwise the reset has sequence number zero and the ACK field is set to the sum of the sequence number and segment length of the incoming segment. The connection remains in the CLOSED state.

2. If the connection is in any non-synchronized state (LISTEN, SYN-SENT, SYN-RECEIVED), and the incoming segment acknowledges something not yet sent (the segment carries an unacceptable ACK), or if an incoming segment has a security level or compartment which does not exactly match the level and compartment requested for the connection, a reset is sent.

   If our SYN has not been acknowledged and the precedence level of the incoming segment is higher than the precedence level requested then either raise the local precedence level (if allowed by the user and the system) or send a reset; or if the precedence level of the incoming segment is lower than the precedence level requested then continue as if the precedence matched exactly (if the remote TCP cannot raise the precedence level to match ours this will be detected in the next segment it sends, and the connection will be terminated then). If our SYN has been acknowledged (perhaps in this incoming segment) the precedence level of the incoming segment must match the local precedence level exactly, if it does not a reset must be sent.

   If the incoming segment has an ACK field, the reset takes its sequence number from the ACK field of the segment, otherwise the reset has sequence number zero and the ACK field is set to the sum of the sequence number and segment length of the incoming segment. The connection remains in the same state.

Meta's Exhibit No. 1014
Page 0042

3. If the connection is in a synchronized state (ESTABLISHED, FIN-WAIT-1, FIN-WAIT-2, CLOSE-WAIT, CLOSING, LAST-ACK, TIME-WAIT), any unacceptable segment (out of window sequence number or unacceptible acknowledgment number) must elicit only an empty acknowledgment segment containing the current send-sequence number and an acknowledgment indicating the next sequence number expected to be received, and the connection remains in the same state.

If an incoming segment has a security level, or compartment, or precedence which does not exactly match the level, and compartment, and precedence requested for the connection,a reset is sent and connection goes to the CLOSED state. The reset takes its sequence number from the ACK field of the incoming segment.

Reset Processing

In all states except SYN-SENT, all reset (RST) segments are validated by checking their SEQ-fields. A reset is valid if its sequence number is in the window. In the SYN-SENT state (a RST received in response to an initial SYN), the RST is acceptable if the ACK field acknowledges the SYN.

The receiver of a RST first validates it, then changes state. If the receiver was in the LISTEN state, it ignores it. If the receiver was in SYN-RECEIVED state and had previously been in the LISTEN state, then the receiver returns to the LISTEN state, otherwise the receiver aborts the connection and goes to the CLOSED state. If the receiver was in any other state, it aborts the connection and advises the user and goes to the CLOSED state.

3.5. Closing a Connection

CLOSE is an operation meaning "I have no more data to send." The notion of closing a full-duplex connection is subject to ambiguous interpretation, of course, since it may not be obvious how to treat the receiving side of the connection. We have chosen to treat CLOSE in a simplex fashion. The user who CLOSEs may continue to RECEIVE until he is told that the other side has CLOSED also. Thus, a program could initiate several SENDs followed by a CLOSE, and then continue to RECEIVE until signaled that a RECEIVE failed because the other side has CLOSED. We assume that the TCP will signal a user, even if no RECEIVEs are outstanding, that the other side has closed, so the user can terminate his side gracefully. A TCP will reliably deliver all buffers SENT before the connection was CLOSED so a user who expects no data in return need only wait to hear the connection was CLOSED successfully to know that all his data was received at the destination TCP. Users must keep reading connections they close for sending until the TCP says no more data.

Meta's Exhibit No. 1014
Page 0043

There are essentially three cases:

1) The user initiates by telling the TCP to CLOSE the connection

2) The remote TCP initiates by sending a FIN control signal

3) Both users CLOSE simultaneously

Case 1: Local user initiates the close

In this case, a FIN segment can be constructed and placed on the outgoing segment queue. No further SENDs from the user will be accepted by the TCP, and it enters the FIN-WAIT-1 state. RECEIVEs are allowed in this state. All segments preceding and including FIN will be retransmitted until acknowledged. When the other TCP has both acknowledged the FIN and sent a FIN of its own, the first TCP can ACK this FIN. Note that a TCP receiving a FIN will ACK but not send its own FIN until its user has CLOSED the connection also.

Case 2: TCP receives a FIN from the network

If an unsolicited FIN arrives from the network, the receiving TCP can ACK it and tell the user that the connection is closing. The user will respond with a CLOSE, upon which the TCP can send a FIN to the other TCP after sending any remaining data. The TCP then waits until its own FIN is acknowledged whereupon it deletes the connection. If an ACK is not forthcoming, after the user timeout the connection is aborted and the user is told.

Case 3: both users close simultaneously

A simultaneous CLOSE by users at both ends of a connection causes FIN segments to be exchanged. When all segments preceding the FINs have been processed and acknowledged, each TCP can ACK the FIN it has received. Both will, upon receiving these ACKs, delete the connection.

Meta's Exhibit No. 1014
Page 0044

```
    TCP A                                                TCP B

1.  ESTABLISHED                                          ESTABLISHED

2.  (Close)
    FIN-WAIT-1  --> <SEQ=100><ACK=300><CTL=FIN,ACK>  --> CLOSE-WAIT

3.  FIN-WAIT-2  <-- <SEQ=300><ACK=101><CTL=ACK>      <-- CLOSE-WAIT

4.                                                       (Close)
    TIME-WAIT   <-- <SEQ=300><ACK=101><CTL=FIN,ACK>  <-- LAST-ACK

5.  TIME-WAIT   --> <SEQ=101><ACK=301><CTL=ACK>      --> CLOSED

6.  (2 MSL)
    CLOSED
```

Normal Close Sequence

Figure 13.

```
    TCP A                                                TCP B

1.  ESTABLISHED                                          ESTABLISHED

2.  (Close)                                              (Close)
    FIN-WAIT-1  --> <SEQ=100><ACK=300><CTL=FIN,ACK>  ... FIN-WAIT-1
                <-- <SEQ=300><ACK=100><CTL=FIN,ACK>  <--
                ... <SEQ=100><ACK=300><CTL=FIN,ACK>  -->

3.  CLOSING     --> <SEQ=101><ACK=301><CTL=ACK>      ... CLOSING
                <-- <SEQ=301><ACK=101><CTL=ACK>      <--
                ... <SEQ=101><ACK=301><CTL=ACK>      -->

4.  TIME-WAIT                                            TIME-WAIT
    (2 MSL)                                              (2 MSL)
    CLOSED                                               CLOSED
```

Simultaneous Close Sequence

Figure 14.

Meta's Exhibit No. 1014
Page 0045

3.6. Precedence and Security

The intent is that connection be allowed only between ports operating with exactly the same security and compartment values and at the higher of the precedence level requested by the two ports.

The precedence and security parameters used in TCP are exactly those defined in the Internet Protocol (IP) [2]. Throughout this TCP specification the term "security/compartment" is intended to indicate the security parameters used in IP including security, compartment, user group, and handling restriction.

A connection attempt with mismatched security/compartment values or a lower precedence value must be rejected by sending a reset. Rejecting a connection due to too low a precedence only occurs after an acknowledgment of the SYN has been received.

Note that TCP modules which operate only at the default value of precedence will still have to check the precedence of incoming segments and possibly raise the precedence level they use on the connection.

The security paramaters may be used even in a non-secure environment (the values would indicate unclassified data), thus hosts in non-secure environments must be prepared to receive the security parameters, though they need not send them.

3.7. Data Communication

Once the connection is established data is communicated by the exchange of segments. Because segments may be lost due to errors (checksum test failure), or network congestion, TCP uses retransmission (after a timeout) to ensure delivery of every segment. Duplicate segments may arrive due to network or TCP retransmission. As discussed in the section on sequence numbers the TCP performs certain tests on the sequence and acknowledgment numbers in the segments to verify their acceptability.

The sender of data keeps track of the next sequence number to use in the variable SND.NXT. The receiver of data keeps track of the next sequence number to expect in the variable RCV.NXT. The sender of data keeps track of the oldest unacknowledged sequence number in the variable SND.UNA. If the data flow is momentarily idle and all data sent has been acknowledged then the three variables will be equal.

When the sender creates a segment and transmits it the sender advances SND.NXT. When the receiver accepts a segment it advances RCV.NXT and sends an acknowledgment. When the data sender receives an

Meta's Exhibit No. 1014
Page 0046

acknowledgment it advances SND.UNA. The extent to which the values of these variables differ is a measure of the delay in the communication. The amount by which the variables are advanced is the length of the data in the segment. Note that once in the ESTABLISHED state all segments must carry current acknowledgment information.

The CLOSE user call implies a push function, as does the FIN control flag in an incoming segment.

Retransmission Timeout

Because of the variability of the networks that compose an internetwork system and the wide range of uses of TCP connections the retransmission timeout must be dynamically determined. One procedure for determining a retransmission time out is given here as an illustration.

An Example Retransmission Timeout Procedure

Measure the elapsed time between sending a data octet with a particular sequence number and receiving an acknowledgment that covers that sequence number (segments sent do not have to match segments received). This measured elapsed time is the Round Trip Time (RTT). Next compute a Smoothed Round Trip Time (SRTT) as:

```
SRTT = ( ALPHA * SRTT ) + ((1-ALPHA) * RTT)
```

and based on this, compute the retransmission timeout (RTO) as:

```
RTO = min[UBOUND,max[LBOUND,(BETA*SRTT)]]
```

where UBOUND is an upper bound on the timeout (e.g., 1 minute), LBOUND is a lower bound on the timeout (e.g., 1 second), ALPHA is a smoothing factor (e.g., .8 to .9), and BETA is a delay variance factor (e.g., 1.3 to 2.0).

The Communication of Urgent Information

The objective of the TCP urgent mechanism is to allow the sending user to stimulate the receiving user to accept some urgent data and to permit the receiving TCP to indicate to the receiving user when all the currently known urgent data has been received by the user.

This mechanism permits a point in the data stream to be designated as the end of urgent information. Whenever this point is in advance of the receive sequence number (RCV.NXT) at the receiving TCP, that TCP must tell the user to go into "urgent mode"; when the receive sequence number catches up to the urgent pointer, the TCP must tell user to go

Meta's Exhibit No. 1014
Page 0047

into "normal mode". If the urgent pointer is updated while the user is in "urgent mode", the update will be invisible to the user.

The method employs a urgent field which is carried in all segments transmitted. The URG control flag indicates that the urgent field is meaningful and must be added to the segment sequence number to yield the urgent pointer. The absence of this flag indicates that there is no urgent data outstanding.

To send an urgent indication the user must also send at least one data octet. If the sending user also indicates a push, timely delivery of the urgent information to the destination process is enhanced.

Managing the Window

The window sent in each segment indicates the range of sequence numbers the sender of the window (the data receiver) is currently prepared to accept. There is an assumption that this is related to the currently available data buffer space available for this connection.

Indicating a large window encourages transmissions. If more data arrives than can be accepted, it will be discarded. This will result in excessive retransmissions, adding unnecessarily to the load on the network and the TCPs. Indicating a small window may restrict the transmission of data to the point of introducing a round trip delay between each new segment transmitted.

The mechanisms provided allow a TCP to advertise a large window and to subsequently advertise a much smaller window without having accepted that much data. This, so called "shrinking the window," is strongly discouraged. The robustness principle dictates that TCPs will not shrink the window themselves, but will be prepared for such behavior on the part of other TCPs.

The sending TCP must be prepared to accept from the user and send at least one octet of new data even if the send window is zero. The sending TCP must regularly retransmit to the receiving TCP even when the window is zero. Two minutes is recommended for the retransmission interval when the window is zero. This retransmission is essential to guarantee that when either TCP has a zero window the re-opening of the window will be reliably reported to the other.

When the receiving TCP has a zero window and a segment arrives it must still send an acknowledgment showing its next expected sequence number and current window (zero).

The sending TCP packages the data to be transmitted into segments

Meta's Exhibit No. 1014
Page 0048

which fit the current window, and may repackage segments on the retransmission queue. Such repackaging is not required, but may be helpful.

In a connection with a one-way data flow, the window information will be carried in acknowledgment segments that all have the same sequence number so there will be no way to reorder them if they arrive out of order. This is not a serious problem, but it will allow the window information to be on occasion temporarily based on old reports from the data receiver. A refinement to avoid this problem is to act on the window information from segments that carry the highest acknowledgment number (that is segments with acknowledgment number equal or greater than the highest previously received).

The window management procedure has significant influence on the communication performance. The following comments are suggestions to implementers.

Window Management Suggestions

Allocating a very small window causes data to be transmitted in many small segments when better performance is achieved using fewer large segments.

One suggestion for avoiding small windows is for the receiver to defer updating a window until the additional allocation is at least X percent of the maximum allocation possible for the connection (where X might be 20 to 40).

Another suggestion is for the sender to avoid sending small segments by waiting until the window is large enough before sending data. If the the user signals a push function then the data must be sent even if it is a small segment.

Note that the acknowledgments should not be delayed or unnecessary retransmissions will result. One strategy would be to send an acknowledgment when a small segment arrives (with out updating the window information), and then to send another acknowledgment with new window information when the window is larger.

The segment sent to probe a zero window may also begin a break up of transmitted data into smaller and smaller segments. If a segment containing a single data octet sent to probe a zero window is accepted, it consumes one octet of the window now available. If the sending TCP simply sends as much as it can whenever the window is non zero, the transmitted data will be broken into alternating big and small segments. As time goes on, occasional pauses in the receiver making window allocation available will

Meta's Exhibit No. 1014
Page 0049

result in breaking the big segments into a small and not quite so big pair. And after a while the data transmission will be in mostly small segments.

The suggestion here is that the TCP implementations need to actively attempt to combine small window allocations into larger windows, since the mechanisms for managing the window tend to lead to many small windows in the simplest minded implementations.

3.8. Interfaces

There are of course two interfaces of concern: the user/TCP interface and the TCP/lower-level interface. We have a fairly elaborate model of the user/TCP interface, but the interface to the lower level protocol module is left unspecified here, since it will be specified in detail by the specification of the lowel level protocol. For the case that the lower level is IP we note some of the parameter values that TCPs might use.

User/TCP Interface

The following functional description of user commands to the TCP is, at best, fictional, since every operating system will have different facilities. Consequently, we must warn readers that different TCP implementations may have different user interfaces. However, all TCPs must provide a certain minimum set of services to guarantee that all TCP implementations can support the same protocol hierarchy. This section specifies the functional interfaces required of all TCP implementations.

TCP User Commands

The following sections functionally characterize a USER/TCP interface. The notation used is similar to most procedure or function calls in high level languages, but this usage is not meant to rule out trap type service calls (e.g., SVCs, UUOs, EMTs).

The user commands described below specify the basic functions the TCP must perform to support interprocess communication. Individual implementations must define their own exact format, and may provide combinations or subsets of the basic functions in single calls. In particular, some implementations may wish to automatically OPEN a connection on the first SEND or RECEIVE issued by the user for a given connection.

Meta's Exhibit No. 1014
Page 0050

In providing interprocess communication facilities, the TCP must not only accept commands, but must also return information to the processes it serves. The latter consists of:

(a) general information about a connection (e.g., interrupts, remote close, binding of unspecified foreign socket).

(b) replies to specific user commands indicating success or various types of failure.

Open

Format: OPEN (local port, foreign socket, active/passive [, timeout] [, precedence] [, security/compartment] [, options]) -> local connection name

We assume that the local TCP is aware of the identity of the processes it serves and will check the authority of the process to use the connection specified. Depending upon the implementation of the TCP, the local network and TCP identifiers for the source address will either be supplied by the TCP or the lower level protocol (e.g., IP). These considerations are the result of concern about security, to the extent that no TCP be able to masquerade as another one, and so on. Similarly, no process can masquerade as another without the collusion of the TCP.

If the active/passive flag is set to passive, then this is a call to LISTEN for an incoming connection. A passive open may have either a fully specified foreign socket to wait for a particular connection or an unspecified foreign socket to wait for any call. A fully specified passive call can be made active by the subsequent execution of a SEND.

A transmission control block (TCB) is created and partially filled in with data from the OPEN command parameters.

On an active OPEN command, the TCP will begin the procedure to synchronize (i.e., establish) the connection at once.

The timeout, if present, permits the caller to set up a timeout for all data submitted to TCP. If data is not successfully delivered to the destination within the timeout period, the TCP will abort the connection. The present global default is five minutes.

The TCP or some component of the operating system will verify the users authority to open a connection with the specified

Meta's Exhibit No. 1014
Page 0051

precedence or security/compartment. The absence of precedence or security/compartment specification in the OPEN call indicates the default values must be used.

TCP will accept incoming requests as matching only if the security/compartment information is exactly the same and only if the precedence is equal to or higher than the precedence requested in the OPEN call.

The precedence for the connection is the higher of the values requested in the OPEN call and received from the incoming request, and fixed at that value for the life of the connection.Implementers may want to give the user control of this precedence negotiation. For example, the user might be allowed to specify that the precedence must be exactly matched, or that any attempt to raise the precedence be confirmed by the user.

A local connection name will be returned to the user by the TCP. The local connection name can then be used as a short hand term for the connection defined by the <local socket, foreign socket> pair.

Send

Format: SEND (local connection name, buffer address, byte count, PUSH flag, URGENT flag [,timeout])

This call causes the data contained in the indicated user buffer to be sent on the indicated connection. If the connection has not been opened, the SEND is considered an error. Some implementations may allow users to SEND first; in which case, an automatic OPEN would be done. If the calling process is not authorized to use this connection, an error is returned.

If the PUSH flag is set, the data must be transmitted promptly to the receiver, and the PUSH bit will be set in the last TCP segment created from the buffer. If the PUSH flag is not set, the data may be combined with data from subsequent SENDs for transmission efficiency.

If the URGENT flag is set, segments sent to the destination TCP will have the urgent pointer set. The receiving TCP will signal the urgent condition to the receiving process if the urgent pointer indicates that data preceding the urgent pointer has not been consumed by the receiving process. The purpose of urgent is to stimulate the receiver to process the urgent data and to indicate to the receiver when all the currently known urgent

Meta's Exhibit No. 1014
Page 0052

data has been received. The number of times the sending user's TCP signals urgent will not necessarily be equal to the number of times the receiving user will be notified of the presence of urgent data.

If no foreign socket was specified in the OPEN, but the connection is established (e.g., because a LISTENing connection has become specific due to a foreign segment arriving for the local socket), then the designated buffer is sent to the implied foreign socket. Users who make use of OPEN with an unspecified foreign socket can make use of SEND without ever explicitly knowing the foreign socket address.

However, if a SEND is attempted before the foreign socket becomes specified, an error will be returned. Users can use the STATUS call to determine the status of the connection. In some implementations the TCP may notify the user when an unspecified socket is bound.

If a timeout is specified, the current user timeout for this connection is changed to the new one.

In the simplest implementation, SEND would not return control to the sending process until either the transmission was complete or the timeout had been exceeded. However, this simple method is both subject to deadlocks (for example, both sides of the connection might try to do SENDs before doing any RECEIVEs) and offers poor performance, so it is not recommended. A more sophisticated implementation would return immediately to allow the process to run concurrently with network I/O, and, furthermore, to allow multiple SENDs to be in progress. Multiple SENDs are served in first come, first served order, so the TCP will queue those it cannot service immediately.

We have implicitly assumed an asynchronous user interface in which a SEND later elicits some kind of SIGNAL or pseudo-interrupt from the serving TCP. An alternative is to return a response immediately. For instance, SENDs might return immediate local acknowledgment, even if the segment sent had not been acknowledged by the distant TCP. We could optimistically assume eventual success. If we are wrong, the connection will close anyway due to the timeout. In implementations of this kind (synchronous), there will still be some asynchronous signals, but these will deal with the connection itself, and not with specific segments or buffers.

In order for the process to distinguish among error or success indications for different SENDs, it might be appropriate for the

Meta's Exhibit No. 1014
Page 0053

buffer address to be returned along with the coded response to the SEND request. TCP-to-user signals are discussed below, indicating the information which should be returned to the calling process.

Receive

Format: RECEIVE (local connection name, buffer address, byte count) -> byte count, urgent flag, push flag

This command allocates a receiving buffer associated with the specified connection. If no OPEN precedes this command or the calling process is not authorized to use this connection, an error is returned.

In the simplest implementation, control would not return to the calling program until either the buffer was filled, or some error occurred, but this scheme is highly subject to deadlocks. A more sophisticated implementation would permit several RECEIVEs to be outstanding at once. These would be filled as segments arrive. This strategy permits increased throughput at the cost of a more elaborate scheme (possibly asynchronous) to notify the calling program that a PUSH has been seen or a buffer filled.

If enough data arrive to fill the buffer before a PUSH is seen, the PUSH flag will not be set in the response to the RECEIVE. The buffer will be filled with as much data as it can hold. If a PUSH is seen before the buffer is filled the buffer will be returned partially filled and PUSH indicated.

If there is urgent data the user will have been informed as soon as it arrived via a TCP-to-user signal. The receiving user should thus be in "urgent mode". If the URGENT flag is on, additional urgent data remains. If the URGENT flag is off, this call to RECEIVE has returned all the urgent data, and the user may now leave "urgent mode". Note that data following the urgent pointer (non-urgent data) cannot be delivered to the user in the same buffer with preceeding urgent data unless the boundary is clearly marked for the user.

To distinguish among several outstanding RECEIVEs and to take care of the case that a buffer is not completely filled, the return code is accompanied by both a buffer pointer and a byte count indicating the actual length of the data received.

Alternative implementations of RECEIVE might have the TCP

Meta's Exhibit No. 1014
Page 0054

allocate buffer storage, or the TCP might share a ring buffer with the user.

Close

Format: CLOSE (local connection name)

This command causes the connection specified to be closed. If the connection is not open or the calling process is not authorized to use this connection, an error is returned. Closing connections is intended to be a graceful operation in the sense that outstanding SENDs will be transmitted (and retransmitted), as flow control permits, until all have been serviced. Thus, it should be acceptable to make several SEND calls, followed by a CLOSE, and expect all the data to be sent to the destination. It should also be clear that users should continue to RECEIVE on CLOSING connections, since the other side may be trying to transmit the last of its data. Thus, CLOSE means "I have no more to send" but does not mean "I will not receive any more." It may happen (if the user level protocol is not well thought out) that the closing side is unable to get rid of all its data before timing out. In this event, CLOSE turns into ABORT, and the closing TCP gives up.

The user may CLOSE the connection at any time on his own initiative, or in response to various prompts from the TCP (e.g., remote close executed, transmission timeout exceeded, destination inaccessible).

Because closing a connection requires communication with the foreign TCP, connections may remain in the closing state for a short time. Attempts to reopen the connection before the TCP replies to the CLOSE command will result in error responses.

Close also implies push function.

Status

Format: STATUS (local connection name) -> status data

This is an implementation dependent user command and could be excluded without adverse effect. Information returned would typically come from the TCB associated with the connection.

This command returns a data block containing the following information:

local socket,

Meta's Exhibit No. 1014
Page 0055

foreign socket,
local connection name,
receive window,
send window,
connection state,
number of buffers awaiting acknowledgment,
number of buffers pending receipt,
urgent state,
precedence,
security/compartment,
and transmission timeout.

Depending on the state of the connection, or on the implementation itself, some of this information may not be available or meaningful. If the calling process is not authorized to use this connection, an error is returned. This prevents unauthorized processes from gaining information about a connection.

Abort

Format: ABORT (local connection name)

This command causes all pending SENDs and RECEIVES to be aborted, the TCB to be removed, and a special RESET message to be sent to the TCP on the other side of the connection. Depending on the implementation, users may receive abort indications for each outstanding SEND or RECEIVE, or may simply receive an ABORT-acknowledgment.

TCP-to-User Messages

It is assumed that the operating system environment provides a means for the TCP to asynchronously signal the user program. When the TCP does signal a user program, certain information is passed to the user. Often in the specification the information will be an error message. In other cases there will be information relating to the completion of processing a SEND or RECEIVE or other user call.

The following information is provided:

Local Connection Name	Always
Response String	Always
Buffer Address	Send & Receive
Byte count (counts bytes received)	Receive
Push flag	Receive
Urgent flag	Receive

Meta's Exhibit No. 1014
Page 0056

TCP/Lower-Level Interface

The TCP calls on a lower level protocol module to actually send and receive information over a network. One case is that of the ARPA internetwork system where the lower level module is the Internet Protocol (IP) [2].

If the lower level protocol is IP it provides arguments for a type of service and for a time to live. TCP uses the following settings for these parameters:

Type of Service = Precedence: routine, Delay: normal, Throughput: normal, Reliability: normal; or 00000000.

Time to Live = one minute, or 00111100.

Note that the assumed maximum segment lifetime is two minutes. Here we explicitly ask that a segment be destroyed if it cannot be delivered by the internet system within one minute.

If the lower level is IP (or other protocol that provides this feature) and source routing is used, the interface must allow the route information to be communicated. This is especially important so that the source and destination addresses used in the TCP checksum be the originating source and ultimate destination. It is also important to preserve the return route to answer connection requests.

Any lower level protocol will have to provide the source address, destination address, and protocol fields, and some way to determine the "TCP length", both to provide the functional equivlent service of IP and to be used in the TCP checksum.

Meta's Exhibit No. 1014
Page 0057

3.9. Event Processing

The processing depicted in this section is an example of one possible implementation. Other implementations may have slightly different processing sequences, but they should differ from those in this section only in detail, not in substance.

The activity of the TCP can be characterized as responding to events. The events that occur can be cast into three categories: user calls, arriving segments, and timeouts. This section describes the processing the TCP does in response to each of the events. In many cases the processing required depends on the state of the connection.

Events that occur:

User Calls

OPEN
SEND
RECEIVE
CLOSE
ABORT
STATUS

Arriving Segments

SEGMENT ARRIVES

Timeouts

USER TIMEOUT
RETRANSMISSION TIMEOUT
TIME-WAIT TIMEOUT

The model of the TCP/user interface is that user commands receive an immediate return and possibly a delayed response via an event or pseudo interrupt. In the following descriptions, the term "signal" means cause a delayed response.

Error responses are given as character strings. For example, user commands referencing connections that do not exist receive "error: connection not open".

Please note in the following that all arithmetic on sequence numbers, acknowledgment numbers, windows, et cetera, is modulo 2**32 the size of the sequence number space. Also note that "=<" means less than or equal to (modulo 2**32).

Meta's Exhibit No. 1014
Page 0058

A natural way to think about processing incoming segments is to imagine that they are first tested for proper sequence number (i.e., that their contents lie in the range of the expected "receive window" in the sequence number space) and then that they are generally queued and processed in sequence number order.

When a segment overlaps other already received segments we reconstruct the segment to contain just the new data, and adjust the header fields to be consistent.

Note that if no state change is mentioned the TCP stays in the same state.

Meta's Exhibit No. 1014
Page 0059

OPEN Call

CLOSED STATE (i.e., TCB does not exist)

Create a new transmission control block (TCB) to hold connection state information. Fill in local socket identifier, foreign socket, precedence, security/compartment, and user timeout information. Note that some parts of the foreign socket may be unspecified in a passive OPEN and are to be filled in by the parameters of the incoming SYN segment. Verify the security and precedence requested are allowed for this user, if not return "error: precedence not allowed" or "error: security/compartment not allowed." If passive enter the LISTEN state and return. If active and the foreign socket is unspecified, return "error: foreign socket unspecified"; if active and the foreign socket is specified, issue a SYN segment. An initial send sequence number (ISS) is selected. A SYN segment of the form <SEQ=ISS><CTL=SYN> is sent. Set SND.UNA to ISS, SND.NXT to ISS+1, enter SYN-SENT state, and return.

If the caller does not have access to the local socket specified, return "error: connection illegal for this process". If there is no room to create a new connection, return "error: insufficient resources".

LISTEN STATE

If active and the foreign socket is specified, then change the connection from passive to active, select an ISS. Send a SYN segment, set SND.UNA to ISS, SND.NXT to ISS+1. Enter SYN-SENT state. Data associated with SEND may be sent with SYN segment or queued for transmission after entering ESTABLISHED state. The urgent bit if requested in the command must be sent with the data segments sent as a result of this command. If there is no room to queue the request, respond with "error: insufficient resources". If Foreign socket was not specified, then return "error: foreign socket unspecified".

Meta's Exhibit No. 1014
Page 0060

OPEN Call

SYN-SENT STATE
SYN-RECEIVED STATE
ESTABLISHED STATE
FIN-WAIT-1 STATE
FIN-WAIT-2 STATE
CLOSE-WAIT STATE
CLOSING STATE
LAST-ACK STATE
TIME-WAIT STATE

Return "error: connection already exists".

Meta's Exhibit No. 1014
Page 0061

SEND Call

CLOSED STATE (i.e., TCB does not exist)

If the user does not have access to such a connection, then return "error: connection illegal for this process".

Otherwise, return "error: connection does not exist".

LISTEN STATE

If the foreign socket is specified, then change the connection from passive to active, select an ISS. Send a SYN segment, set SND.UNA to ISS, SND.NXT to ISS+1. Enter SYN-SENT state. Data associated with SEND may be sent with SYN segment or queued for transmission after entering ESTABLISHED state. The urgent bit if requested in the command must be sent with the data segments sent as a result of this command. If there is no room to queue the request, respond with "error: insufficient resources". If Foreign socket was not specified, then return "error: foreign socket unspecified".

SYN-SENT STATE
SYN-RECEIVED STATE

Queue the data for transmission after entering ESTABLISHED state. If no space to queue, respond with "error: insufficient resources".

ESTABLISHED STATE
CLOSE-WAIT STATE

Segmentize the buffer and send it with a piggybacked acknowledgment (acknowledgment value = RCV.NXT). If there is insufficient space to remember this buffer, simply return "error: insufficient resources".

If the urgent flag is set, then SND.UP <- SND.NXT-1 and set the urgent pointer in the outgoing segments.

Meta's Exhibit No. 1014
Page 0062

SEND Call

FIN-WAIT-1 STATE
FIN-WAIT-2 STATE
CLOSING STATE
LAST-ACK STATE
TIME-WAIT STATE

Return "error: connection closing" and do not service request.

Meta's Exhibit No. 1014
Page 0063

RECEIVE Call

CLOSED STATE (i.e., TCB does not exist)

If the user does not have access to such a connection, return "error: connection illegal for this process".

Otherwise return "error: connection does not exist".

LISTEN STATE
SYN-SENT STATE
SYN-RECEIVED STATE

Queue for processing after entering ESTABLISHED state. If there is no room to queue this request, respond with "error: insufficient resources".

ESTABLISHED STATE
FIN-WAIT-1 STATE
FIN-WAIT-2 STATE

If insufficient incoming segments are queued to satisfy the request, queue the request. If there is no queue space to remember the RECEIVE, respond with "error: insufficient resources".

Reassemble queued incoming segments into receive buffer and return to user. Mark "push seen" (PUSH) if this is the case.

If RCV.UP is in advance of the data currently being passed to the user notify the user of the presence of urgent data.

When the TCP takes responsibility for delivering data to the user that fact must be communicated to the sender via an acknowledgment. The formation of such an acknowledgment is described below in the discussion of processing an incoming segment.

Meta's Exhibit No. 1014
Page 0064

CLOSE-WAIT STATE

Since the remote side has already sent FIN, RECEIVEs must be satisfied by text already on hand, but not yet delivered to the user. If no text is awaiting delivery, the RECEIVE will get a "error: connection closing" response. Otherwise, any remaining text can be used to satisfy the RECEIVE.

CLOSING STATE
LAST-ACK STATE
TIME-WAIT STATE

Return "error: connection closing".

Meta's Exhibit No. 1014
Page 0065

CLOSE Call

CLOSED STATE (i.e., TCB does not exist)

If the user does not have access to such a connection, return "error: connection illegal for this process".

Otherwise, return "error: connection does not exist".

LISTEN STATE

Any outstanding RECEIVEs are returned with "error: closing" responses. Delete TCB, enter CLOSED state, and return.

SYN-SENT STATE

Delete the TCB and return "error: closing" responses to any queued SENDs, or RECEIVEs.

SYN-RECEIVED STATE

If no SENDs have been issued and there is no pending data to send, then form a FIN segment and send it, and enter FIN-WAIT-1 state; otherwise queue for processing after entering ESTABLISHED state.

ESTABLISHED STATE

Queue this until all preceding SENDs have been segmentized, then form a FIN segment and send it. In any case, enter FIN-WAIT-1 state.

FIN-WAIT-1 STATE
FIN-WAIT-2 STATE

Strictly speaking, this is an error and should receive a "error: connection closing" response. An "ok" response would be acceptable, too, as long as a second FIN is not emitted (the first FIN may be retransmitted though).

Meta's Exhibit No. 1014
Page 0066

CLOSE Call

CLOSE-WAIT STATE

Queue this request until all preceding SENDs have been segmentized; then send a FIN segment, enter CLOSING state.

CLOSING STATE
LAST-ACK STATE
TIME-WAIT STATE

Respond with "error: connection closing".

Meta's Exhibit No. 1014
Page 0067

ABORT Call

CLOSED STATE (i.e., TCB does not exist)

If the user should not have access to such a connection, return "error: connection illegal for this process".

Otherwise return "error: connection does not exist".

LISTEN STATE

Any outstanding RECEIVEs should be returned with "error: connection reset" responses. Delete TCB, enter CLOSED state, and return.

SYN-SENT STATE

All queued SENDs and RECEIVEs should be given "connection reset" notification, delete the TCB, enter CLOSED state, and return.

SYN-RECEIVED STATE
ESTABLISHED STATE
FIN-WAIT-1 STATE
FIN-WAIT-2 STATE
CLOSE-WAIT STATE

Send a reset segment:

<SEQ=SND.NXT><CTL=RST>

All queued SENDs and RECEIVEs should be given "connection reset" notification; all segments queued for transmission (except for the RST formed above) or retransmission should be flushed, delete the TCB, enter CLOSED state, and return.

CLOSING STATE
LAST-ACK STATE
TIME-WAIT STATE

Respond with "ok" and delete the TCB, enter CLOSED state, and return.

Meta's Exhibit No. 1014
Page 0068

STATUS Call

CLOSED STATE (i.e., TCB does not exist)

If the user should not have access to such a connection, return "error: connection illegal for this process".

Otherwise return "error: connection does not exist".

LISTEN STATE

Return "state = LISTEN", and the TCB pointer.

SYN-SENT STATE

Return "state = SYN-SENT", and the TCB pointer.

SYN-RECEIVED STATE

Return "state = SYN-RECEIVED", and the TCB pointer.

ESTABLISHED STATE

Return "state = ESTABLISHED", and the TCB pointer.

FIN-WAIT-1 STATE

Return "state = FIN-WAIT-1", and the TCB pointer.

FIN-WAIT-2 STATE

Return "state = FIN-WAIT-2", and the TCB pointer.

CLOSE-WAIT STATE

Return "state = CLOSE-WAIT", and the TCB pointer.

CLOSING STATE

Return "state = CLOSING", and the TCB pointer.

LAST-ACK STATE

Return "state = LAST-ACK", and the TCB pointer.

Meta's Exhibit No. 1014
Page 0069

TIME-WAIT STATE

Return "state = TIME-WAIT", and the TCB pointer.

Meta's Exhibit No. 1014
Page 0070

SEGMENT ARRIVES

If the state is CLOSED (i.e., TCB does not exist) then

all data in the incoming segment is discarded. An incoming segment containing a RST is discarded. An incoming segment not containing a RST causes a RST to be sent in response. The acknowledgment and sequence field values are selected to make the reset sequence acceptable to the TCP that sent the offending segment.

If the ACK bit is off, sequence number zero is used,

<SEQ=0><ACK=SEG.SEQ+SEG.LEN><CTL=RST,ACK>

If the ACK bit is on,

<SEQ=SEG.ACK><CTL=RST>

Return.

If the state is LISTEN then

first check for an RST

An incoming RST should be ignored. Return.

second check for an ACK

Any acknowledgment is bad if it arrives on a connection still in the LISTEN state. An acceptable reset segment should be formed for any arriving ACK-bearing segment. The RST should be formatted as follows:

<SEQ=SEG.ACK><CTL=RST>

Return.

third check for a SYN

If the SYN bit is set, check the security. If the security/compartment on the incoming segment does not exactly match the security/compartment in the TCB then send a reset and return.

<SEQ=SEG.ACK><CTL=RST>

Meta's Exhibit No. 1014
Page 0071

If the SEG.PRC is greater than the TCB.PRC then if allowed by the user and the system set TCB.PRC<-SEG.PRC, if not allowed send a reset and return.

<SEQ=SEG.ACK><CTL=RST>

If the SEG.PRC is less than the TCB.PRC then continue.

Set RCV.NXT to SEG.SEQ+1, IRS is set to SEG.SEQ and any other control or text should be queued for processing later. ISS should be selected and a SYN segment sent of the form:

<SEQ=ISS><ACK=RCV.NXT><CTL=SYN,ACK>

SND.NXT is set to ISS+1 and SND.UNA to ISS. The connection state should be changed to SYN-RECEIVED. Note that any other incoming control or data (combined with SYN) will be processed in the SYN-RECEIVED state, but processing of SYN and ACK should not be repeated. If the listen was not fully specified (i.e., the foreign socket was not fully specified), then the unspecified fields should be filled in now.

fourth other text or control

Any other control or text-bearing segment (not containing SYN) must have an ACK and thus would be discarded by the ACK processing. An incoming RST segment could not be valid, since it could not have been sent in response to anything sent by this incarnation of the connection. So you are unlikely to get here, but if you do, drop the segment, and return.

If the state is SYN-SENT then

first check the ACK bit

If the ACK bit is set

If SEG.ACK =< ISS, or SEG.ACK > SND.NXT, send a reset (unless the RST bit is set, if so drop the segment and return)

<SEQ=SEG.ACK><CTL=RST>

and discard the segment. Return.

If SND.UNA =< SEG.ACK =< SND.NXT then the ACK is acceptable.

second check the RST bit

Meta's Exhibit No. 1014
Page 0072

SEGMENT ARRIVES

If the RST bit is set

If the ACK was acceptable then signal the user "error: connection reset", drop the segment, enter CLOSED state, delete TCB, and return. Otherwise (no ACK) drop the segment and return.

third check the security and precedence

If the security/compartment in the segment does not exactly match the security/compartment in the TCB, send a reset

If there is an ACK

<SEQ=SEG.ACK><CTL=RST>

Otherwise

<SEQ=0><ACK=SEG.SEQ+SEG.LEN><CTL=RST,ACK>

If there is an ACK

The precedence in the segment must match the precedence in the TCB, if not, send a reset

<SEQ=SEG.ACK><CTL=RST>

If there is no ACK

If the precedence in the segment is higher than the precedence in the TCB then if allowed by the user and the system raise the precedence in the TCB to that in the segment, if not allowed to raise the prec then send a reset.

<SEQ=0><ACK=SEG.SEQ+SEG.LEN><CTL=RST,ACK>

If the precedence in the segment is lower than the precedence in the TCB continue.

If a reset was sent, discard the segment and return.

fourth check the SYN bit

This step should be reached only if the ACK is ok, or there is no ACK, and it the segment did not contain a RST.

If the SYN bit is on and the security/compartment and precedence

Meta's Exhibit No. 1014
Page 0073

are acceptable then, RCV.NXT is set to SEG.SEQ+1, IRS is set to SEG.SEQ. SND.UNA should be advanced to equal SEG.ACK (if there is an ACK), and any segments on the retransmission queue which are thereby acknowledged should be removed.

If SND.UNA > ISS (our SYN has been ACKed), change the connection state to ESTABLISHED, form an ACK segment

<SEQ=SND.NXT><ACK=RCV.NXT><CTL=ACK>

and send it. Data or controls which were queued for transmission may be included. If there are other controls or text in the segment then continue processing at the sixth step below where the URG bit is checked, otherwise return.

Otherwise enter SYN-RECEIVED, form a SYN,ACK segment

<SEQ=ISS><ACK=RCV.NXT><CTL=SYN,ACK>

and send it. If there are other controls or text in the segment, queue them for processing after the ESTABLISHED state has been reached, return.

fifth, if neither of the SYN or RST bits is set then drop the segment and return.

Meta's Exhibit No. 1014
Page 0074

SEGMENT ARRIVES

Otherwise,

first check sequence number

SYN-RECEIVED STATE
ESTABLISHED STATE
FIN-WAIT-1 STATE
FIN-WAIT-2 STATE
CLOSE-WAIT STATE
CLOSING STATE
LAST-ACK STATE
TIME-WAIT STATE

Segments are processed in sequence. Initial tests on arrival are used to discard old duplicates, but further processing is done in SEG.SEQ order. If a segment's contents straddle the boundary between old and new, only the new parts should be processed.

There are four cases for the acceptability test for an incoming segment:

Segment Length	Receive Window	Test
0	0	SEG.SEQ = RCV.NXT
0	>0	RCV.NXT =< SEG.SEQ < RCV.NXT+RCV.WND
>0	0	not acceptable
>0	>0	RCV.NXT =< SEG.SEQ < RCV.NXT+RCV.WND or RCV.NXT =< SEG.SEQ+SEG.LEN-1 < RCV.NXT+RCV.WND

If the RCV.WND is zero, no segments will be acceptable, but special allowance should be made to accept valid ACKs, URGs and RSTs.

If an incoming segment is not acceptable, an acknowledgment should be sent in reply (unless the RST bit is set, if so drop the segment and return):

<SEQ=SND.NXT><ACK=RCV.NXT><CTL=ACK>

After sending the acknowledgment, drop the unacceptable segment and return.

Meta's Exhibit No. 1014
Page 0075

In the following it is assumed that the segment is the idealized segment that begins at RCV.NXT and does not exceed the window. One could tailor actual segments to fit this assumption by trimming off any portions that lie outside the window (including SYN and FIN), and only processing further if the segment then begins at RCV.NXT. Segments with higher begining sequence numbers may be held for later processing.

second check the RST bit,

SYN-RECEIVED STATE

If the RST bit is set

If this connection was initiated with a passive OPEN (i.e., came from the LISTEN state), then return this connection to LISTEN state and return. The user need not be informed. If this connection was initiated with an active OPEN (i.e., came from SYN-SENT state) then the connection was refused, signal the user "connection refused". In either case, all segments on the retransmission queue should be removed. And in the active OPEN case, enter the CLOSED state and delete the TCB, and return.

ESTABLISHED
FIN-WAIT-1
FIN-WAIT-2
CLOSE-WAIT

If the RST bit is set then, any outstanding RECEIVEs and SEND should receive "reset" responses. All segment queues should be flushed. Users should also receive an unsolicited general "connection reset" signal. Enter the CLOSED state, delete the TCB, and return.

CLOSING STATE
LAST-ACK STATE
TIME-WAIT

If the RST bit is set then, enter the CLOSED state, delete the TCB, and return.

Meta's Exhibit No. 1014
Page 0076

SEGMENT ARRIVES

third check security and precedence

SYN-RECEIVED

If the security/compartment and precedence in the segment do not exactly match the security/compartment and precedence in the TCB then send a reset, and return.

ESTABLISHED STATE

If the security/compartment and precedence in the segment do not exactly match the security/compartment and precedence in the TCB then send a reset, any outstanding RECEIVEs and SEND should receive "reset" responses. All segment queues should be flushed. Users should also receive an unsolicited general "connection reset" signal. Enter the CLOSED state, delete the TCB, and return.

Note this check is placed following the sequence check to prevent a segment from an old connection between these ports with a different security or precedence from causing an abort of the current connection.

fourth, check the SYN bit,

SYN-RECEIVED
ESTABLISHED STATE
FIN-WAIT STATE-1
FIN-WAIT STATE-2
CLOSE-WAIT STATE
CLOSING STATE
LAST-ACK STATE
TIME-WAIT STATE

If the SYN is in the window it is an error, send a reset, any outstanding RECEIVEs and SEND should receive "reset" responses, all segment queues should be flushed, the user should also receive an unsolicited general "connection reset" signal, enter the CLOSED state, delete the TCB, and return.

If the SYN is not in the window this step would not be reached and an ack would have been sent in the first step (sequence number check).

Meta's Exhibit No. 1014
Page 0077

fifth check the ACK field,

if the ACK bit is off drop the segment and return

if the ACK bit is on

SYN-RECEIVED STATE

If SND.UNA =< SEG.ACK =< SND.NXT then enter ESTABLISHED state and continue processing.

If the segment acknowledgment is not acceptable, form a reset segment,

<SEQ=SEG.ACK><CTL=RST>

and send it.

ESTABLISHED STATE

If SND.UNA < SEG.ACK =< SND.NXT then, set SND.UNA <- SEG.ACK. Any segments on the retransmission queue which are thereby entirely acknowledged are removed. Users should receive positive acknowledgments for buffers which have been SENT and fully acknowledged (i.e., SEND buffer should be returned with "ok" response). If the ACK is a duplicate (SEG.ACK < SND.UNA), it can be ignored. If the ACK acks something not yet sent (SEG.ACK > SND.NXT) then send an ACK, drop the segment, and return.

If SND.UNA < SEG.ACK =< SND.NXT, the send window should be updated. If (SND.WL1 < SEG.SEQ or (SND.WL1 = SEG.SEQ and SND.WL2 =< SEG.ACK)), set SND.WND <- SEG.WND, set SND.WL1 <- SEG.SEQ, and set SND.WL2 <- SEG.ACK.

Note that SND.WND is an offset from SND.UNA, that SND.WL1 records the sequence number of the last segment used to update SND.WND, and that SND.WL2 records the acknowledgment number of the last segment used to update SND.WND. The check here prevents using old segments to update the window.

Meta's Exhibit No. 1014
Page 0078

SEGMENT ARRIVES

FIN-WAIT-1 STATE

In addition to the processing for the ESTABLISHED state, if our FIN is now acknowledged then enter FIN-WAIT-2 and continue processing in that state.

FIN-WAIT-2 STATE

In addition to the processing for the ESTABLISHED state, if the retransmission queue is empty, the user's CLOSE can be acknowledged ("ok") but do not delete the TCB.

CLOSE-WAIT STATE

Do the same processing as for the ESTABLISHED state.

CLOSING STATE

In addition to the processing for the ESTABLISHED state, if the ACK acknowledges our FIN then enter the TIME-WAIT state, otherwise ignore the segment.

LAST-ACK STATE

The only thing that can arrive in this state is an acknowledgment of our FIN. If our FIN is now acknowledged, delete the TCB, enter the CLOSED state, and return.

TIME-WAIT STATE

The only thing that can arrive in this state is a retransmission of the remote FIN. Acknowledge it, and restart the 2 MSL timeout.

sixth, check the URG bit,

ESTABLISHED STATE
FIN-WAIT-1 STATE
FIN-WAIT-2 STATE

If the URG bit is set, RCV.UP <- max(RCV.UP,SEG.UP), and signal the user that the remote side has urgent data if the urgent pointer (RCV.UP) is in advance of the data consumed. If the user has already been signaled (or is still in the "urgent mode") for this continuous sequence of urgent data, do not signal the user again.

Meta's Exhibit No. 1014
Page 0079

CLOSE-WAIT STATE
CLOSING STATE
LAST-ACK STATE
TIME-WAIT

This should not occur, since a FIN has been received from the remote side. Ignore the URG.

seventh, process the segment text,

ESTABLISHED STATE
FIN-WAIT-1 STATE
FIN-WAIT-2 STATE

Once in the ESTABLISHED state, it is possible to deliver segment text to user RECEIVE buffers. Text from segments can be moved into buffers until either the buffer is full or the segment is empty. If the segment empties and carries an PUSH flag, then the user is informed, when the buffer is returned, that a PUSH has been received.

When the TCP takes responsibility for delivering the data to the user it must also acknowledge the receipt of the data.

Once the TCP takes responsibility for the data it advances RCV.NXT over the data accepted, and adjusts RCV.WND as apporopriate to the current buffer availability. The total of RCV.NXT and RCV.WND should not be reduced.

Please note the window management suggestions in section 3.7.

Send an acknowledgment of the form:

<SEQ=SND.NXT><ACK=RCV.NXT><CTL=ACK>

This acknowledgment should be piggybacked on a segment being transmitted if possible without incurring undue delay.

Meta's Exhibit No. 1014
Page 0080

SEGMENT ARRIVES

CLOSE-WAIT STATE
CLOSING STATE
LAST-ACK STATE
TIME-WAIT STATE

This should not occur, since a FIN has been received from the remote side. Ignore the segment text.

eighth, check the FIN bit,

Do not process the FIN if the state is CLOSED, LISTEN or SYN-SENT since the SEG.SEQ cannot be validated; drop the segment and return.

If the FIN bit is set, signal the user "connection closing" and return any pending RECEIVEs with same message, advance RCV.NXT over the FIN, and send an acknowledgment for the FIN. Note that FIN implies PUSH for any segment text not yet delivered to the user.

SYN-RECEIVED STATE
ESTABLISHED STATE

Enter the CLOSE-WAIT state.

FIN-WAIT-1 STATE

If our FIN has been ACKed (perhaps in this segment), then enter TIME-WAIT, start the time-wait timer, turn off the other timers; otherwise enter the CLOSING state.

FIN-WAIT-2 STATE

Enter the TIME-WAIT state. Start the time-wait timer, turn off the other timers.

CLOSE-WAIT STATE

Remain in the CLOSE-WAIT state.

CLOSING STATE

Remain in the CLOSING state.

LAST-ACK STATE

Remain in the LAST-ACK state.

Meta's Exhibit No. 1014
Page 0081

SEGMENT ARRIVES

TIME-WAIT STATE

Remain in the TIME-WAIT state. Restart the 2 MSL time-wait timeout.

and return.

Meta's Exhibit No. 1014
Page 0082

USER TIMEOUT

For any state if the user timeout expires, flush all queues, signal the user "error: connection aborted due to user timeout" in general and for any outstanding calls, delete the TCB, enter the CLOSED state and return.

RETRANSMISSION TIMEOUT

For any state if the retransmission timeout expires on a segment in the retransmission queue, send the segment at the front of the retransmission queue again, reinitialize the retransmission timer, and return.

TIME-WAIT TIMEOUT

If the time-wait timeout expires on a connection delete the TCB, enter the CLOSED state and return.

Meta's Exhibit No. 1014
Page 0083

Meta's Exhibit No. 1014
Page 0084

GLOSSARY

1822
: BBN Report 1822, "The Specification of the Interconnection of a Host and an IMP". The specification of interface between a host and the ARPANET.

ACK
: A control bit (acknowledge) occupying no sequence space, which indicates that the acknowledgment field of this segment specifies the next sequence number the sender of this segment is expecting to receive, hence acknowledging receipt of all previous sequence numbers.

ARPANET message
: The unit of transmission between a host and an IMP in the ARPANET. The maximum size is about 1012 octets (8096 bits).

ARPANET packet
: A unit of transmission used internally in the ARPANET between IMPs. The maximum size is about 126 octets (1008 bits).

connection
: A logical communication path identified by a pair of sockets.

datagram
: A message sent in a packet switched computer communications network.

Destination Address
: The destination address, usually the network and host identifiers.

FIN
: A control bit (finis) occupying one sequence number, which indicates that the sender will send no more data or control occupying sequence space.

fragment
: A portion of a logical unit of data, in particular an internet fragment is a portion of an internet datagram.

FTP
: A file transfer protocol.

Meta's Exhibit No. 1014
Page 0085

header
: Control information at the beginning of a message, segment, fragment, packet or block of data.

host
: A computer. In particular a source or destination of messages from the point of view of the communication network.

Identification
: An Internet Protocol field. This identifying value assigned by the sender aids in assembling the fragments of a datagram.

IMP
: The Interface Message Processor, the packet switch of the ARPANET.

internet address
: A source or destination address specific to the host level.

internet datagram
: The unit of data exchanged between an internet module and the higher level protocol together with the internet header.

internet fragment
: A portion of the data of an internet datagram with an internet header.

IP
: Internet Protocol.

IRS
: The Initial Receive Sequence number. The first sequence number used by the sender on a connection.

ISN
: The Initial Sequence Number. The first sequence number used on a connection, (either ISS or IRS). Selected on a clock based procedure.

ISS
: The Initial Send Sequence number. The first sequence number used by the sender on a connection.

leader
: Control information at the beginning of a message or block of data. In particular, in the ARPANET, the control information on an ARPANET message at the host-IMP interface.

Meta's Exhibit No. 1014
Page 0086

left sequence
: This is the next sequence number to be acknowledged by the data receiving TCP (or the lowest currently unacknowledged sequence number) and is sometimes referred to as the left edge of the send window.

local packet
: The unit of transmission within a local network.

module
: An implementation, usually in software, of a protocol or other procedure.

MSL
: Maximum Segment Lifetime, the time a TCP segment can exist in the internetwork system. Arbitrarily defined to be 2 minutes.

octet
: An eight bit byte.

Options
: An Option field may contain several options, and each option may be several octets in length. The options are used primarily in testing situations; for example, to carry timestamps. Both the Internet Protocol and TCP provide for options fields.

packet
: A package of data with a header which may or may not be logically complete. More often a physical packaging than a logical packaging of data.

port
: The portion of a socket that specifies which logical input or output channel of a process is associated with the data.

process
: A program in execution. A source or destination of data from the point of view of the TCP or other host-to-host protocol.

PUSH
: A control bit occupying no sequence space, indicating that this segment contains data that must be pushed through to the receiving user.

RCV.NXT
: receive next sequence number

Meta's Exhibit No. 1014
Page 0087

RCV.UP
receive urgent pointer

RCV.WND
receive window

receive next sequence number
This is the next sequence number the local TCP is expecting to receive.

receive window
This represents the sequence numbers the local (receiving) TCP is willing to receive. Thus, the local TCP considers that segments overlapping the range RCV.NXT to RCV.NXT + RCV.WND - 1 carry acceptable data or control. Segments containing sequence numbers entirely outside of this range are considered duplicates and discarded.

RST
A control bit (reset), occupying no sequence space, indicating that the receiver should delete the connection without further interaction. The receiver can determine, based on the sequence number and acknowledgment fields of the incoming segment, whether it should honor the reset command or ignore it. In no case does receipt of a segment containing RST give rise to a RST in response.

RTP
Real Time Protocol: A host-to-host protocol for communication of time critical information.

SEG.ACK
segment acknowledgment

SEG.LEN
segment length

SEG.PRC
segment precedence value

SEG.SEQ
segment sequence

SEG.UP
segment urgent pointer field

Meta's Exhibit No. 1014
Page 0088

SEG.WND
segment window field

segment
A logical unit of data, in particular a TCP segment is the unit of data transfered between a pair of TCP modules.

segment acknowledgment
The sequence number in the acknowledgment field of the arriving segment.

segment length
The amount of sequence number space occupied by a segment, including any controls which occupy sequence space.

segment sequence
The number in the sequence field of the arriving segment.

send sequence
This is the next sequence number the local (sending) TCP will use on the connection. It is initially selected from an initial sequence number curve (ISN) and is incremented for each octet of data or sequenced control transmitted.

send window
This represents the sequence numbers which the remote (receiving) TCP is willing to receive. It is the value of the window field specified in segments from the remote (data receiving) TCP. The range of new sequence numbers which may be emitted by a TCP lies between SND.NXT and SND.UNA + SND.WND - 1. (Retransmissions of sequence numbers between SND.UNA and SND.NXT are expected, of course.)

SND.NXT
send sequence

SND.UNA
left sequence

SND.UP
send urgent pointer

SND.WL1
segment sequence number at last window update

SND.WL2
segment acknowledgment number at last window update

Meta's Exhibit No. 1014
Page 0089

SND.WND
: send window

socket
: An address which specifically includes a port identifier, that is, the concatenation of an Internet Address with a TCP port.

Source Address
: The source address, usually the network and host identifiers.

SYN
: A control bit in the incoming segment, occupying one sequence number, used at the initiation of a connection, to indicate where the sequence numbering will start.

TCB
: Transmission control block, the data structure that records the state of a connection.

TCB.PRC
: The precedence of the connection.

TCP
: Transmission Control Protocol: A host-to-host protocol for reliable communication in internetwork environments.

TOS
: Type of Service, an Internet Protocol field.

Type of Service
: An Internet Protocol field which indicates the type of service for this internet fragment.

URG
: A control bit (urgent), occupying no sequence space, used to indicate that the receiving user should be notified to do urgent processing as long as there is data to be consumed with sequence numbers less than the value indicated in the urgent pointer.

urgent pointer
: A control field meaningful only when the URG bit is on. This field communicates the value of the urgent pointer which indicates the data octet associated with the sending user's urgent call.

Meta's Exhibit No. 1014
Page 0090

REFERENCES

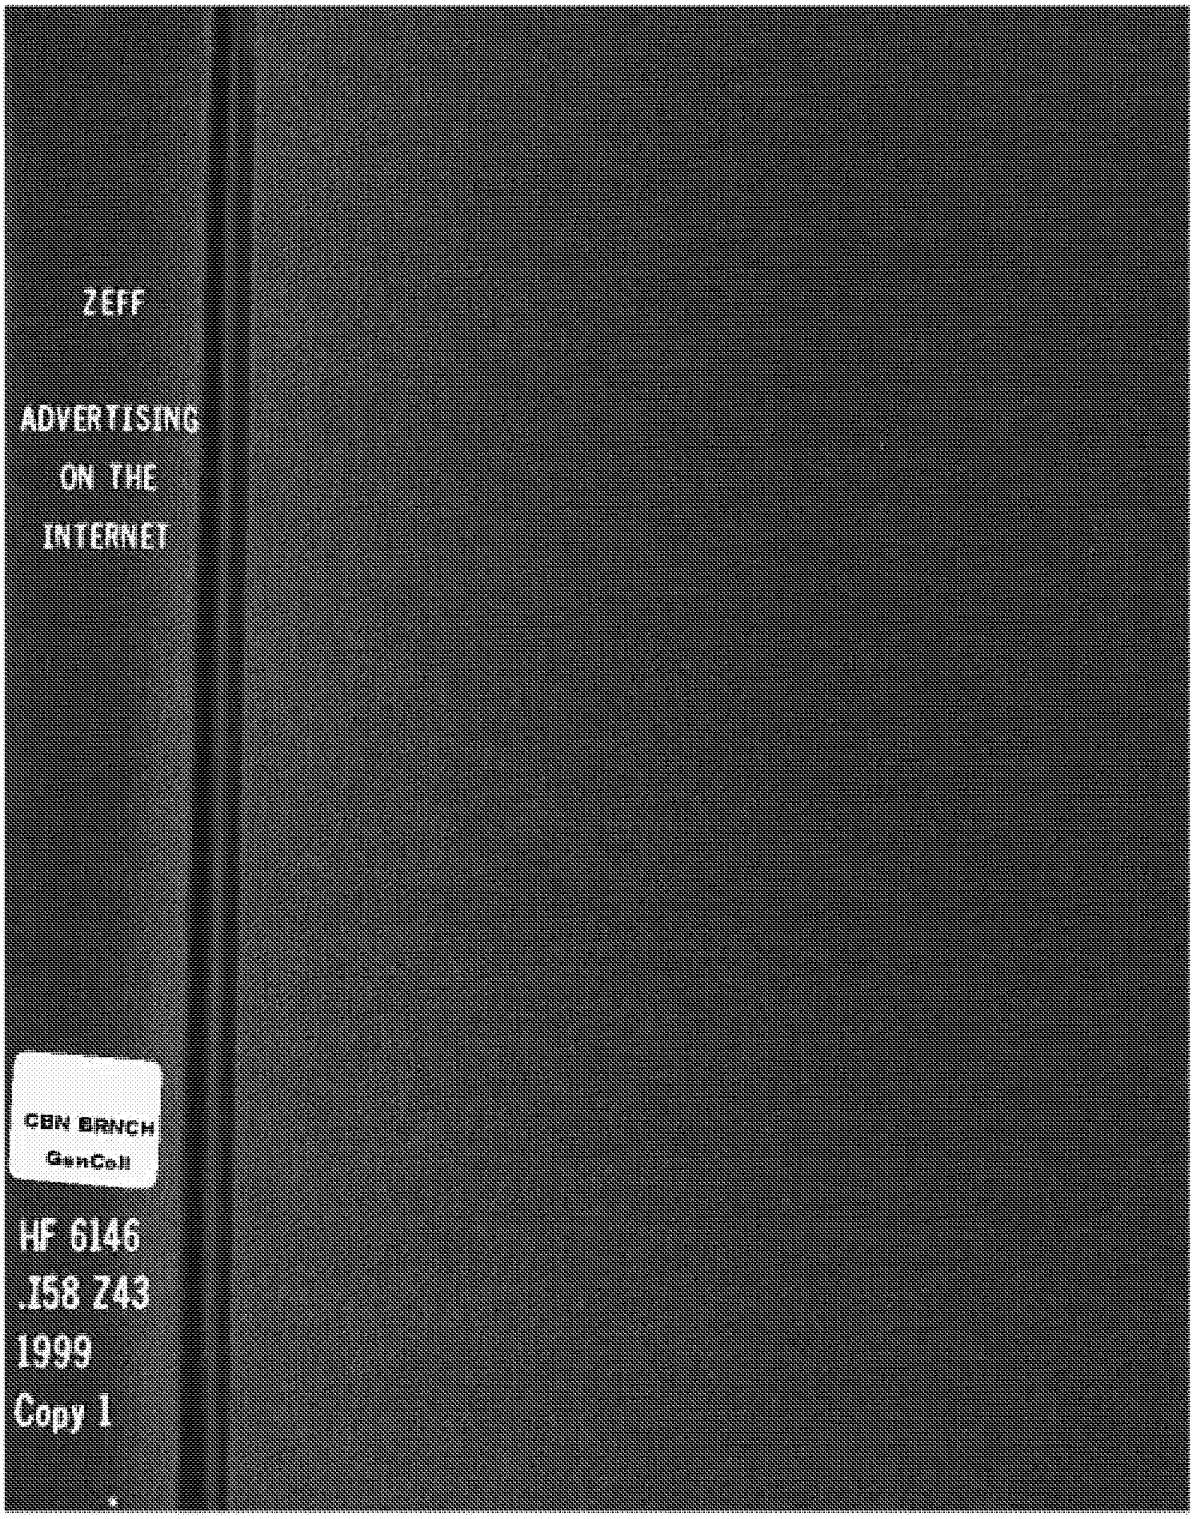
[1] Cerf, V., and R. Kahn, "A Protocol for Packet Network Intercommunication", IEEE Transactions on Communications, Vol. COM-22, No. 5, pp 637-648, May 1974.

[2] Postel, J. (ed.), "Internet Protocol - DARPA Internet Program Protocol Specification", RFC 791, USC/Information Sciences Institute, September 1981.

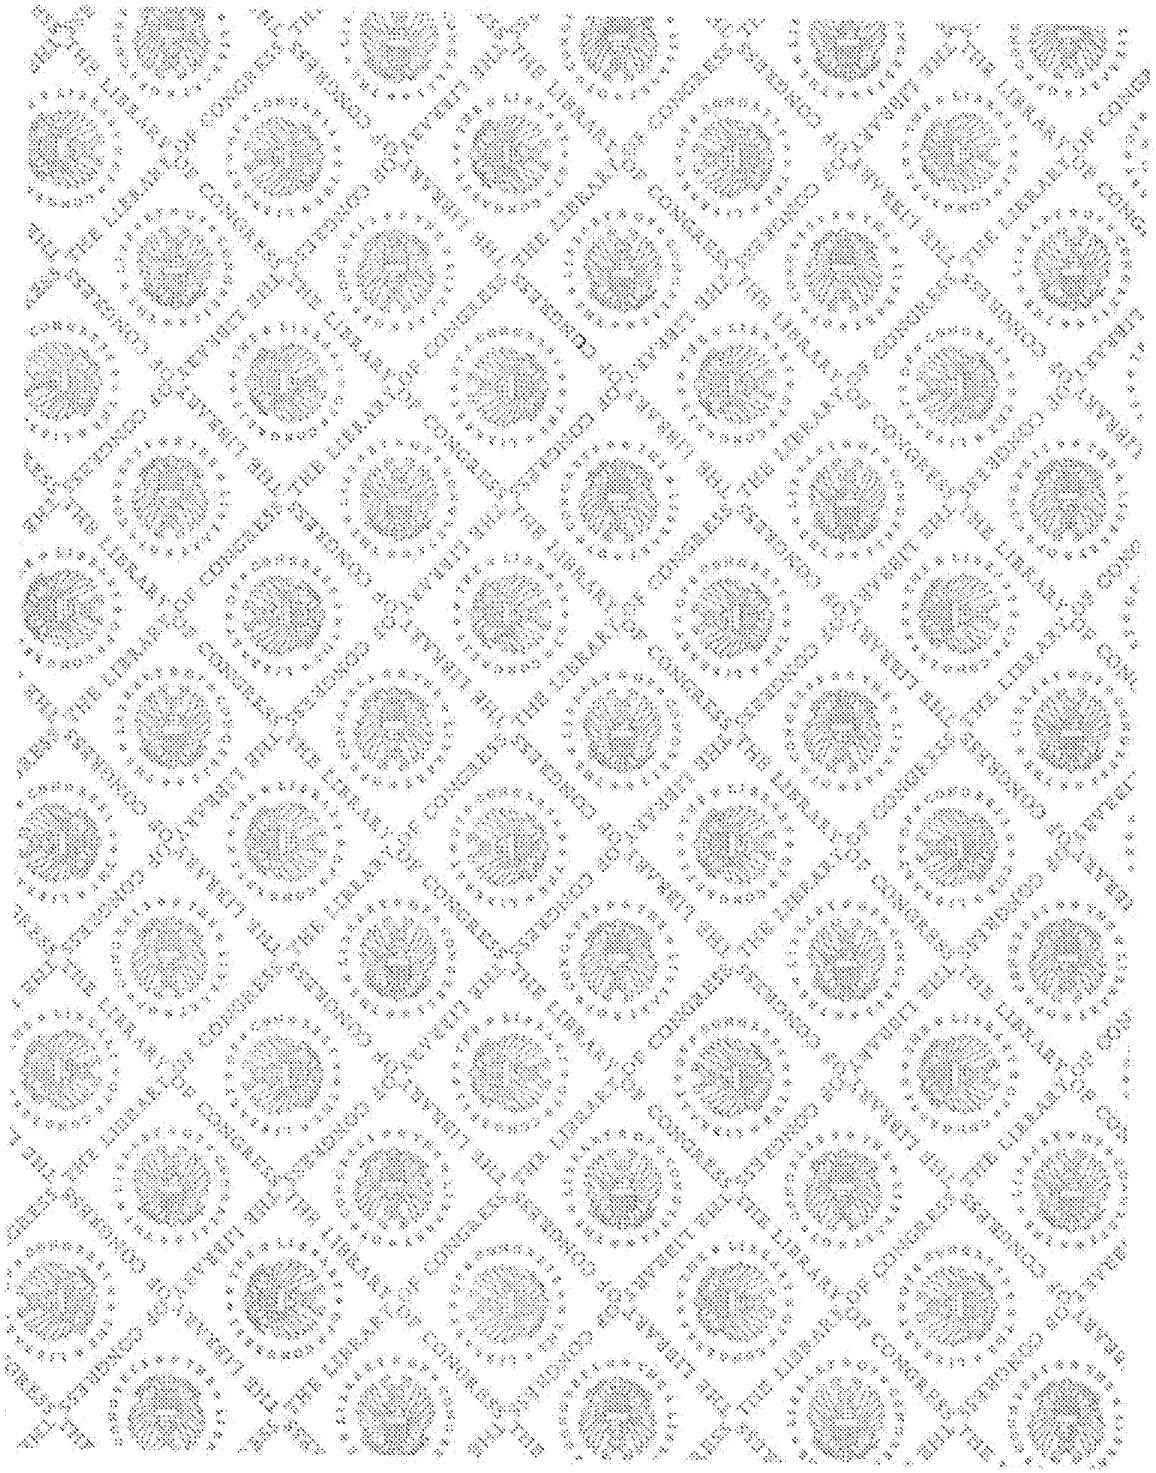
[3] Dalal, Y. and C. Sunshine, "Connection Management in Transport Protocols", Computer Networks, Vol. 2, No. 6, pp. 454-473, December 1978.

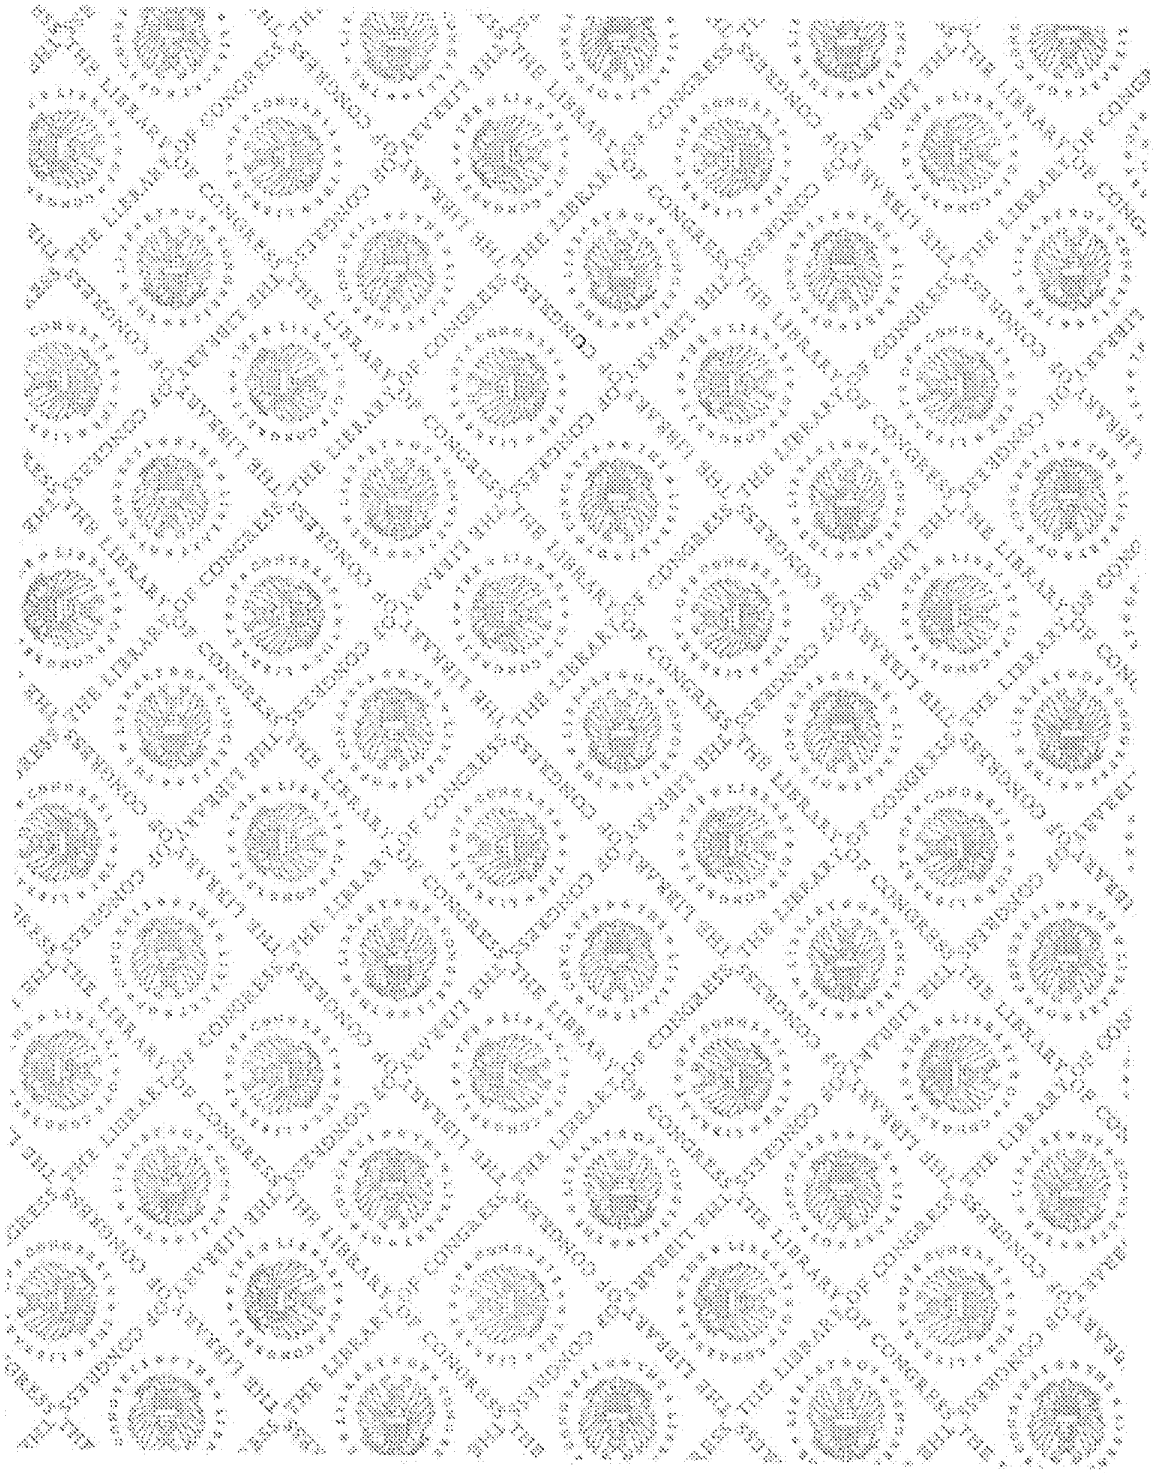
[4] Postel, J., "Assigned Numbers", RFC 790, USC/Information Sciences Institute, September 1981.

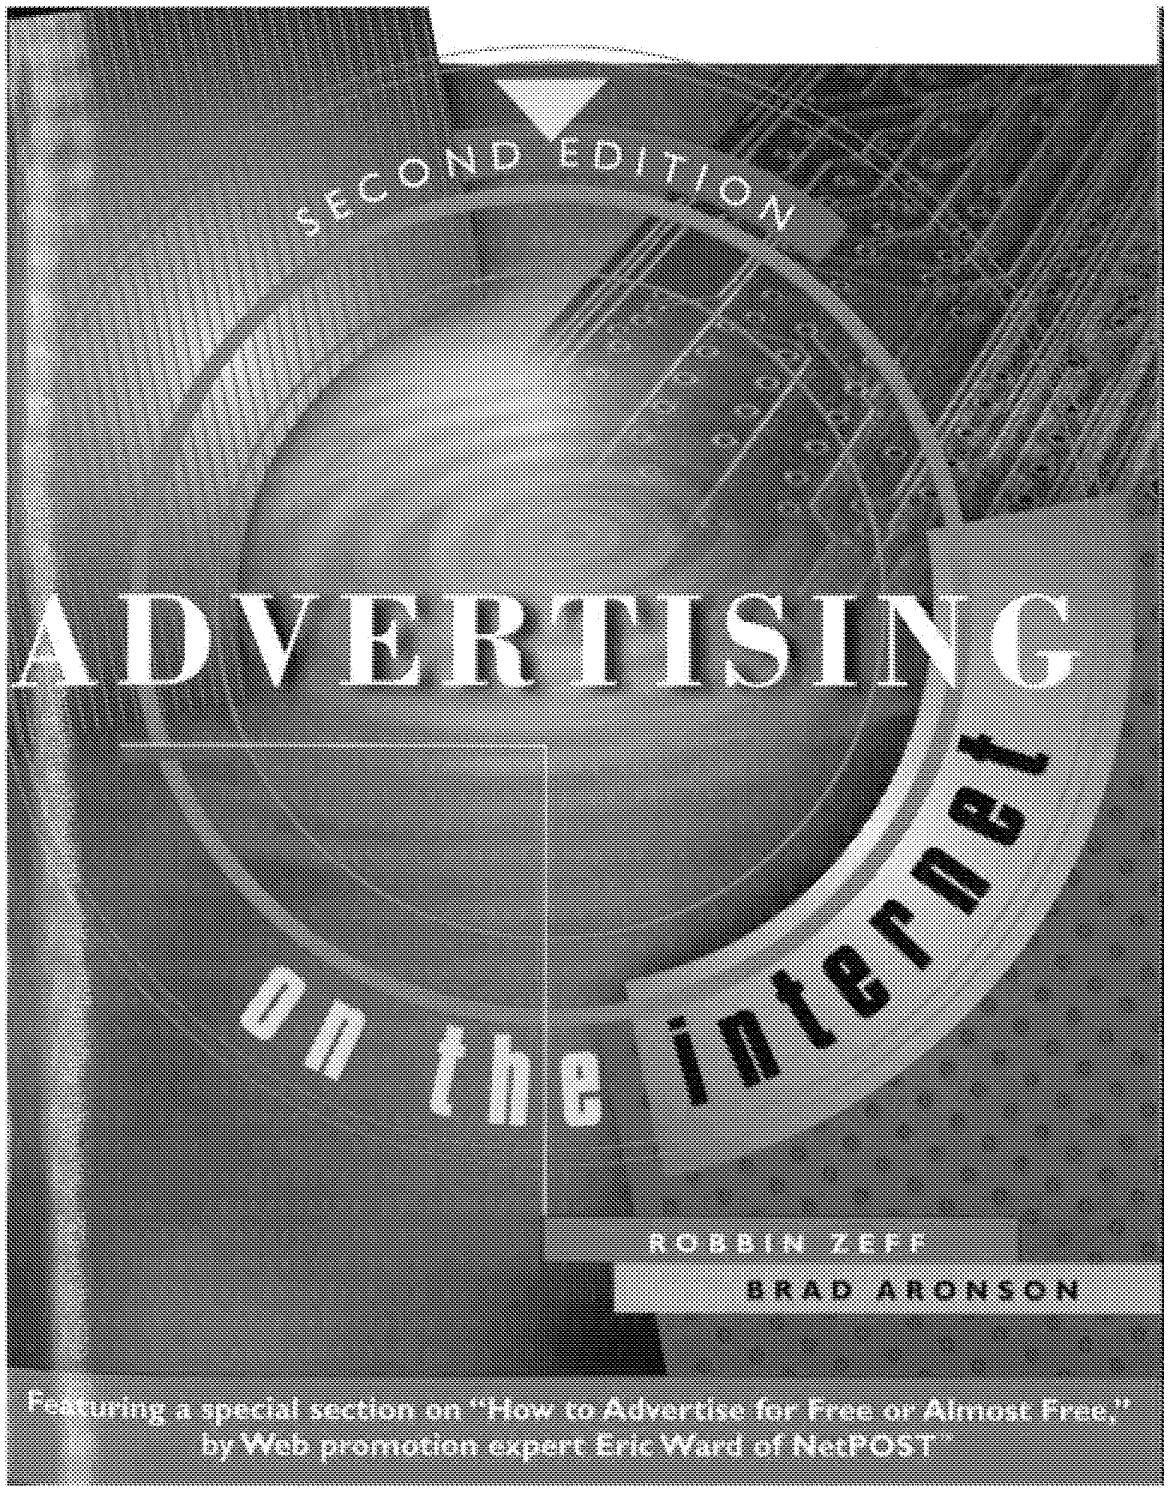
Meta's Exhibit No. 1014
Page 0091



Meta's Exhibit No. 1004
Page 001







Meta's Exhibit No. 1004
Page 004

Advertising on the Internet

Second Edition

Robbin Zeff
Brad Aronson

Wiley Computer Publishing



John Wiley & Sons, Inc.
NEW YORK • CHICHESTER • WEINHEIM • BRISBANE • SINGAPORE • TORONTO

Meta's Exhibit No. 1004
Page 005

HF6146
I58Z43
1999
Copy 1

Publisher: Robert Ipsen
Editor: Cary Sullivan
Managing Editor: Brian Snapp
Associate New Media Editor: Mike Sosa
Text Design & Composition: Benchmark Productions, Inc.

Designations used by companies to distinguish their products are often claimed as trademarks. In all instances where John Wiley & Sons, Inc., is aware of a claim, the product names appear in initial capital or ALL CAPITAL LETTERS. Readers, however, should contact the appropriate companies for more complete information regarding trademarks and registration.

This book is printed on acid-free paper. ♾

Copyright © 1999 by Robbin Zeff and Bradley Aronson. All rights reserved.

Published by John Wiley & Sons, Inc.

Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4744. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 605 Third Avenue, New York, NY 10158-0012, (212) 850-6011, fax (212) 850-6008, e-mail: permreq@wiley.com.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold with the understanding that the publisher is not engaged in professional services. If professional advice or other expert assistance is required, the services of a competent professional person should be sought.

Library of Congress Cataloging-in-Publication Data:

Zeff, Robbin Lee.
Advertising on the Internet/Robbin Zeff, Bradley Aronson. 2nd ed.
p. cm.
"Wiley Computer Publishing."
Includes index.
ISBN 0-471-34404-4 (pbk./website: alk. paper)
1. Internet advertising. 2. Internet advertising Directories. 3. Internet advertising—Computer network resources. 4. Internet advertising—Rates. I. Aronson, Brad. II. Title.
HF6146.I58Z43 1999
658.8'0025'4678--dc21 99-21722
CIP

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

CONTENTS

Acknowledgments xi

Chapter 1 The Internet Advertising Landscape 1

Setting the Stage 1

What Is Online Advertising? 11
- Online Advertising and Traditional Advertising Vehicles 11
- Internet Advertising—A Definition 12
- The Online Advertising Advantage 13

A Short History of Online Advertising 14
- In the Beginning 14
- Setbacks: The Spam Wars 14
- Acceptance 15

Wake Up, It's a Mass Market 16
- Critics 17

Players in the Internet Advertising Industry 18
- The Sellers 18
- The Buyers 19
- Advertising Infrastructure 21

Chapter 2 Online Ad Models 23

Web Site as Ad 23

Ad Models 26
- Advertising via E-mail 26
- Advertising via the Web 33
- Testing 45
- Buttons 46
- Text Links 47
- Sponsorships 50
- Advertorial 51
- Push Technology 54
- Interstitials 55

A Few More Ad Models outside the Box 59
- New Screen Real Estate 59
- Screensavers 60

Meta's Exhibit No. 1004
Page 007

	Bookmarks and Toolbars	60
	Cursors	60
	Undervalued Web Space	60
	ISPs	62
	Other Formats	63
	Guarding Your Domain	63
	Resources for Reviewing Ads and Ad Campaigns	64
Chapter 3	**Internet Direct Marketing**	**69**
	Direct E-mail	70
	E-mail Lists	71
	Banner Advertising	81
	Banner to Web Site or Mini-Page	85
	Banner as Direct Response Piece	86
	Loyalty Programs	86
	Coupons	88
	Free Samples and Trials	89
	Contests and Games	91
	Internet Direct Marketing Agencies	93
Chapter 4	**The Art and Science of Web Measurement**	**95**
	Building Blocks	95
	Key Terminology in Web Measurement	97
	What Can Be Learned from Log Analysis	99
	Most Frequently Requested Pages	99
	Number of Visits	99
	Pages Viewed Per Visit and Length of Visit	100
	Distinction Between First-Time and Repeat Visitors	100
	Most Popular Usage Times	100
	Popular Paths Followed through the Site	100
	Click-through Rate	100
	Browser Used	100
	Referring URL	101
	Domain and Host of the User	101
	User Computer Platform	101
	User Entrance Page	101
	User Exit Page	101
	Web Measurement Tools	101
	Third-Party Outsourcing	102
	In-House Management of Software	103
	The Problems with Web Measurement	104
	Technological Hurdles	106

	Auditing of Web Site Traffic	110
	Rating the Web One Click at a Time	112
Chapter 5	**Ad Management**	**121**
	Ad Management for Publishers	122
	In-House Ad Management	122
	Ad Management Service Bureaus	123
	Ad Management for Advertisers	125
	Third-Party Ad Management	127
	Products for Both Publishers and Advertisers	129
Chapter 6	**Targeting—The Online Advertising Advantage**	**133**
	Targeting Based on Content and Context	134
	Targeting Based on Registration Information	135
	Using Cookies for Targeting	139
	Targeting through Database Mining	140
	Targeting through Profiling and Personalization	141
	Recommendations as Ads	141
	Personalization	142
	Collaborative Filtering	145
	Behavioral Targeting	145
	Putting It All Together: How Ad Management Solutions Facilitate Targeting	146
	Privacy	147
	Keeping Customers Is a Continuous Process	150
Chapter 7	**Pricing Online Ads**	**153**
	Pricing Models	154
	CPM	154
	Click-Through	157
	Flat Fee	159
	Pay Per Viewer	162
	Pay Per Purchase	162
	Affiliate Programs	163
	Auctions	172
	Real-Time Ad Buying Through Networks	174
	Hybrid	175
	Trends in Pricing	176
	Problems with Comparing Value	176
	Flexibility Is the Key	177

Chapter 8	**Market Research—Information Is Your Competitive Advantage**	**179**
	Building Your Body of Knowledge	180
	Industry Events: Where the Hunters and Gatherers Meet	180
	Trade Associations	180
	Market Research Firms	181
	Stats, Facts, and Figures	182
	Academic Studies	188
	Psychographic Studies	188
	Age	190
	Gender	192
	Race	193
	Language	195
	Household Income	195
	Usage Patterns	196
	General Online Advertising Resources	198
Chapter 9	**Buying Online Ads**	**203**
	Determining Campaign Goals	203
	Traffic Building	204
	Selling Products	205
	Generating Leads	206
	Building Brand	207
	The Site Selection Process	208
	Audience	208
	Reporting	209
	Ability to Quickly Rotate Creatives	209
	Banner Size and Placement on the Page	209
	Amount of Ad Inventory Available for Purchase	211
	Performance Guarantees	212
	Technology Capabilities	213
	Targeting Capabilities	213
	Comparing Audience	215
	Types of Sites	215
	Paying for Media Buys	217
	CPM	217
	Cost Per Action	217
	Hybrid	218
	Pricing for Buys	218
	Search Engines and Directories	218
	Content Sites	218
	Network Buys	219
	Free E-mail Services	219
	Spot Buys	220
	Allocating Your Campaign Budget	220
	Testing	223

Chapter 10 Selling Online Advertising 225

Identifying and Defining What you Have to Sell 226

Preparing Your Site's Infrastructure 226
- Monitoring and Measuring Site Traffic 226
- Ad Models 226
- Ad Management 227
- Auditing 227

Know Your Audience 227
- The Credibility Issue 228
- Pricing Ads—Developing the Rate Structure 228
- Agency and Sales Commissions 230
- Discounts and Special Deals 230
- Selling Day Parts 231
- Competitive Analysis 231

Media Kit 232
- Site Overview and Features 232
- Contact Information 234
- Advertising and Sponsorship Programs 234
- Advertising Rates 235
- Site Traffic 235
- Audience Demographics 237
- Production Specifications 237
- Delivery Specifications 237
- Reporting 237
- Deciding the Format of Your Media Kit 237
- Sign on the Dotted Line 238

Selling Strategies 239
- Site Representation Firm 239
- Ad Networks 245
- Differences Between Ad Networks and Representation Firms 248
- Auctions 249

Sales Staff 250
- Sales Training 250
- The Never-Ending Quest to Keep Good Sales Staff 251

Spread the Word about Your Advertising Opportunities 252
- One Last Word on Client Care and Feeding 253

Chapter 11 Going Local 257

Profile of a Local Information User 258

Targeting by Content 263
- Online Versions of Local Media 263
- Informational Sites with Regional Focuses 269
- Ad Networks Go Local 270
- Local Organizations 274

Ranking Local Sites 275
Targeting Through Registration Data 275
IP Targeting 275

Chapter 12 International Online Advertising 281
Why Go Global? 281
Are You a Candidate for Going Global? 284
Obstacles to Adoption 284
Connection 284
Internet Service Provider 285
Competing Systems 285
Turning Obstacles into Opportunities 285
Translation Issues 286
Translating Your Web Site 286
Translating Your Ad 289
Localization 290
Communicating with International Customers 290
Locating Appropriate Sites in Other Countries 294
Portals and Directories 294
Ad Networks 296
Testing the Market 297
Targeting by Language, Not Country 299
Educating the International Market 300

Chapter 13 Legal Issues Every Internet Advertiser Should Know 305
The European Experience 305
The U.S. Situation 308
Unsolicited Commercial E-mail (Spam) 308
Personal Identifying Information 309
Consumers' Personal Identifying Information 310
Industry Guidelines 314

Chapter 14 The Internet Advertising Convergence 317
It's Everywhere, Used by Everyone 319
Watch Out E-mail, Here Comes Instant Messaging 322
Internet, Intranet, and Extranet . . . Oh My! 323
Moving Full Speed Ahead 324

How to Advertise for Free or Almost Free (by Eric Ward) 327
First Things First: Preparing to Announce and Submit Your Site 327
The Practical, Tactical Approach to Search Engine and Directory Submissions 328

You're Ready to Submit Your Site, But Is Your Site Ready to Be Submitted? 328
Getting Organized: The Web Site Questionnaire 331
Finding the Most Appropriate Resources to Announce and Submit Your Site 336
Keeping a History of the Submissions You Make 337
General Search Engines 337
General Directories 338
"What's New" Announcement Sites 340
Best of the Web and What's Cool Outlets 340
Usenet Newsgroups 342
Learn Appropriate Internet Marketing Practices 342
Sites That Help You Announce Your Site 342
Content Rating, Reciprocal Links, Offline Promotions, and Other Techniques 345
Rating Your Web Site's Content 345
Offline Promotional Products 346
Reciprocal Links 347
Inbound Links 347
Conclusion 349

Appendix Resource Directory

Ad Management Technologies 352
Ad Networks 357
Alternative Ad Models 365
Associations 366
Auction 369
Auditing Services 369
Banner Exchange Programs 370
Conferences and Events 372
Contests 373
Directories of E-mail Discussion Lists and Newsletters 374
E-mail List Brokers/Services 375
E-mail List Software 378
Loyalty Programs 379
Market Research Firms 380
Measurement Technologies 390
Personalization Technologies 393
Promotion Services 395

Rating Firms 398
Representation Firms 398
Resources 400
E-mail Discussion Lists and Newsletters 400
Informational Sites 404
Online Publications 409

Glossary 411

About the Authors 417

Index 419

Meta's Exhibit No. 1004
Page 0014

ACKNOWLEDGMENTS

Writing a book is an intimate experience between a person and the words they choose, but being able to gather the creative ideas and energy to produce the book takes a great deal of support. I have many to thank for their assistance and support.

I want to begin by thanking all those who contributed directly to this book by either writing an entire chapter or one of the editorial sidebars found throughout the book: Rochelle Blaustein, Andy Bourland, Rex Briggs, Ann Burgraff, Melinda Gipson, Ann Handley, Richard Hopple, Bob Ivins, Leslie Laredo, Tom Lix, Liddy Manson, Evan Neufeld, Ali Partovi, Katherine Randolph, Todd Tweedy, David Wamsley, Eric Ward, Tony Winders, and David Zinman.

For always providing me with a bottomless cup of coffee and a creative and supportive environment, I want to thank the regulars and staff at Metro 29 Diner and the Lyon's Village Starbucks.

I couldn't have finished this book and kept my business afloat without the help of my staff at The Zeff Group who assisted in researching this book: Jee Hyung Kim and Mimi Talamonti. As always, they produced A+ work.

I owe a great deal of thanks to the members of the online advertising and marketing community. This book would not have been possible without the generosity they shared in terms of information and ideas. I want to point out a few who went above and beyond the call of duty: ClickZ, FrankelBiz, the Internet Advertising Discussion List, Jupiter Communications, the Media Center, and the Online Advertising Discussion List.

I also need to thank the hundreds of students who have taken my Internet advertising class at ZDU (www.zdu.com). Their perspective from the four corners of the globe continually opens my eyes to the breadth and depth of this industry.

I must thank my publisher, John Wiley & Sons, Inc., for continually encouraging my publishing projects, especially my editor Cary Sullivan.

Writing a book is an exercise in fortitude, and I wouldn't have the stamina without the support of my friends and family. In writing this book Toby Dershowitz, Virginia Diamond, Matt Edelstein, Laura Kragie, Ronnie Levin, and Leon MacMullen were dear and patient friends. My family (my parents Jack and Debbie; my sister Tova, my brother Brick, and the entire Nadler clan—Jami, Zanni, Dustin, and Devon) deserve gratitude for always being there and never doubting that I would resurface when the book was done.

Meta's Exhibit No. 1004
Page 0015

And finally, I lost a good friend during the writing of this book. I want to dedicate this book to Sammy; he never knew he was a cat.

—*Robbin Zeff*

I'm very fortunate in that I have always managed to have great friends, family, and colleagues.

The first thank-you is owed to my wife Mia. She did the majority of the work on the resources section of the book, and she was understanding of the hectic schedule as we brought the second edition of this book to a close. In general my schedule isn't usually normal, and she is always supportive.

I'd like to mention my parents Joe and Joan, my brother Rob (the investment wizard) and my grandparents Alva, Al, and Lucille. There is no substitute for a supportive family. All the articles I constantly receive from my grandparents are also better than any clipping service to which I could subscribe.

My uncle, Steven Krupnick, has always been a great advisor and the best landlord a company could want. Andy Ehrlich, who has an office in our building, has also been an excellent advisor. Both Steve and Andy have given their time generously to help me personally and to help in the growth of i-frontier.

I wrote a big portion of this edition while I was in Seattle, and I am thankful for the hospitality of my wife's family: Erik, Carol, Ike, and Erik. Even when I'm working in Seattle, if I'm visiting them, it feels like a vacation.

Folks from my office who contributed to the research and graphics include Paul Walcutt, Adam Embick, and Manlio Lo Conte. I'd also like to thank the rest of the hard-working people who have helped move i-frontier forward. Without a great team, I would never have had time to write the second edition. Those people include: Jeremy Lockhorn, Vincent Moos, Collie Andrews, Jodin Trocheck, Dan Hughes, Rob Aronson, Jennie Leigh, Frank Chiachiere, Raquel Torres Coates, Dan DaCosta. Others include Jon Hurwitz, Arvin Casas, Nathan Solomon, Jay Valinis.

Of course, I'd also like to thank Robbin Zeff, the co-author of this book. Her attention to detail is why no one will find a more comprehensive book about Web advertising. I hope we'll write a publication together again soon . . . after I have a little time to catch up on my sleep.

—*Brad Aronson*

CHAPTER 1

The Internet Advertising Landscape

Since the emergence of modern advertising more than 70 years ago, members of the profession have recognized that it's a direct reflection of society at the time. "Day by day a picture of our time is recorded completely and vividly in the advertising in American newspapers and magazines," said a pamphlet produced in 1926 by the N.W. Ayer & Sons advertising agency. Today, the picture of our time would also be presented on a computer screen, with the sound of a modem connecting in the background.

Setting the Stage

The Internet, the largest global network of networked computers, was founded in the spirit of free access to information. But with its transition from a platform used primarily by research, government, and educational institutions to a commercial entity, established business and entrepreneurial hopefuls have put up Web sites and begun using the Internet to make their mark—and their fortune—on the digital frontier. From established media companies like Time Warner to media companies that began as online business ventures such as Yahoo!, the Internet is following the American tradition of turning uncharted territory into opportunity. And, in the tradition of all great communication media, advertising has moved quickly to the forefront as a source for supporting online commercial ventures.

The first format of online promotion was the Web site itself. Then came advertisements—banners and buttons—whose goal it was to bring users from publishers' Web sites to the advertisers' Web sites. These ads followed the print model's standard of clearly defined borders. The limited screen real estate allotted for banners and buttons (less than 10 percent), as well as the advertisers' desire to reach more consumers without requiring users

Meta's Exhibit No. 1004
Page 0017

to come to their Web sites, inspired the development of new models. From sponsorships to interstitials to push technology, new online advertising models are blurring the lines between advertising and editorial, and providing advertisers with more obtrusive ads similar to television advertising.

By the end of 1998, online advertising had grown to a $1.8 billion industry, with revenue projections of almost $8 billion by the year 2002 (see Figure 1.1). Not bad for an industry that got its start just four years earlier in 1994. The Aberdeen Group (www.aberdeen.com) predicts that Internet advertising spending will reach $5.1 billion by the year 2000. It should be noted when evaluating online ad revenue figures that barter is a common practice in the North American online ad market and often miscounted as actual spending (see Figure 1.2).

In a February 1998 study conducted by the American Association of Advertising Agencies, 45 percent of those surveyed said Internet advertising already was or was expected to become a significant segment of their advertising budget. Of these, 23 percent said Internet advertising was already a major part of their advertising strategy, 50 percent projected that it would become so within two to five years, and the remaining 27 percent expected it to be within the next 5 to 10 years.

Indeed, with the growth in ad spending, we're also seeing a growth in the number of Web sites seeking advertising. According to AdKnowledge, in December 1998 there were 1,430 brand named sites seeking advertising that listed rate cards in AdKnowledge's database (see Figure 1.3). Since the AdKnowledge database primarily reflects major brands, the actual number is much higher. In reality, many more sites are accepting advertising. For example, LinkExchange alone has 400,000 members making the real number of sites seeking advertising more than a half million, with most of them small- to midsized sites.

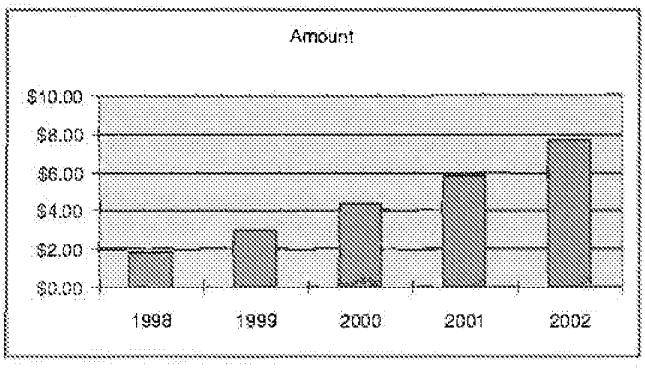


Figure 1.1 Web ad spending projections 1998–2002.
Source: 1998 Jupiter Communications (www.jup.com)

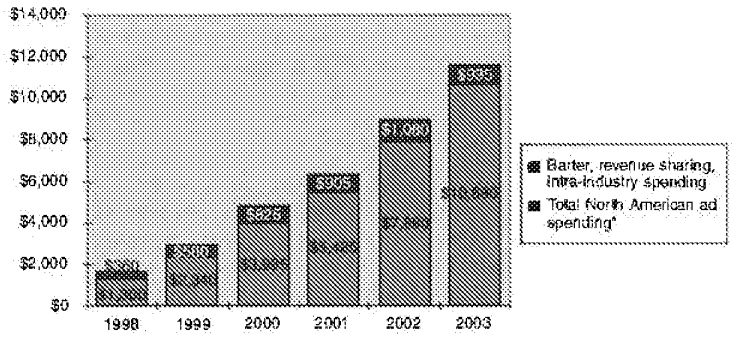


Figure 1.2 The North American online advertising market.
Source: 1998 Forrester Research and Zenith Media (www.forrester.com)

Why is ad spending increasing? Because online buying is increasing. During the 1998 holiday shopping season, AOL members alone spent US $1.2 billion with online retailers. AOL sees great similarities between the online world and the offline world when it comes to reaching consumers. "Marketers either have to create their own shopping mall or buy space in somebody else's," says Bob Pittman, AOL president. "Traditional marketers know that building your own mall is harder." Indeed, on the Web it is easier and more efficient to advertise on someone else's site.

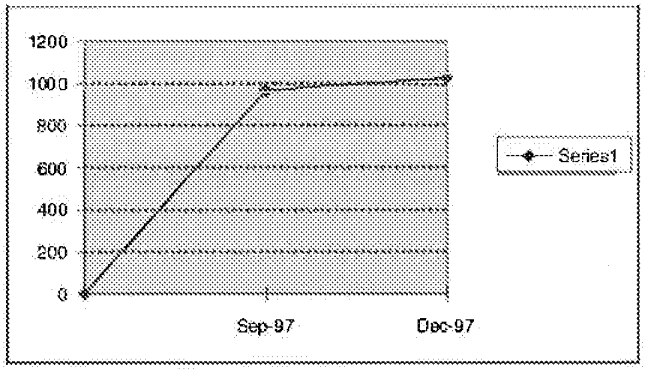


Figure 1.3 Total number of Web sites seeking advertising.
Source: AdKnowledge, Inc. (www.adknowledge.com)
Note: Data reflects the total number of unique sites tracked and included in the AdKnowledge Planner database.

Of the hundreds of thousands of Web publishers, only a handful are generating significant revenues from advertising. This imbalance is not unique to the online world. It occurs in broadcast TV, magazines, and newspapers, where the majority of ad dollars go to a well-defined group. For example, of the 6,000 magazines in operation, only 31 receive in excess of $200 million in advertising revenues. Time Inc. alone commands one-third of the entire advertising market.

In many respects, the Web is merely following an age-old tradition whereby media buyers look for the biggest and most convenient bang for their buck. On the Web, the top 10 sites have attracted the lion's share of advertising revenue simply because they are the sites with the highest traffic. Traditional advertisers are accustomed to working with a mass-market medium, and high volume fits that profile. Today, the Internet does reach millions. According to Nua Internet Surveys (www.nua.ie), there were 87 million people in the United States and Canada online as of November 1998, for a worldwide total of 150 million users (see Table 1.1).

Will advertising be a viable revenue source for all Web sites big and small? Smaller sites will never be able to deliver the high-volume audience of major publishing sites, but they can deliver a targeted and loyal audience. It is this targetability and loyalty that will interest businesses in advertising on smaller sites representing niche markets.

And small sites are making money from advertising. Ray Owens, president of the "Joke a Day" humor mailing list (www.jokeaday.com), is proud to talk about how he turned telling old jokes via e-mail into "a $150,000 a year job, enabling me to spend my day sitting around in my bathrobe, surfing the Net, quitting my 'real job,' and generally living the 'American Dream.'" The "Joke A Day" site sells out its ad inventory six months in advance and is the largest daily humor mailing list in the world, reaching 115,000 people in 137 countries, with an average of 400 new people signing up for the list every day.

Table 1.1 Number of People on the Internet

REGION	INTERNET USERS
Africa	0.80 million
Asia/Pacific	24.33 million
Europe	32.76 million
Middle East	0.78 million
Canada and U.S.A.	87 million
South America	4.5 million
World Total	**150 million**

Source: Nua Internet Surveys (www.nua.ie)

Table 1.2 Internet Growth Rate in Relation to Other Media

MEDIA	YEARS TO REACH 50 MILLION USERS
Radio	38
TV	13
Cable	10
Internet (estimate)	5

Source: Morgan Stanley Technology Research

Many smaller sites are extremely attractive to major advertisers because of their ability to provide access and exposure to a defined online community, but the small publishers must generate enough traffic so that they have enough advertising inventory to sell. Most do not have the inventory to earn significant sums from selling advertising on a CPM (cost per one thousand times an ad is seen). The average CPM is around $30. This means that if a publisher sells all of his ad inventory he will be paid three cents for every time he can deliver an ad. At that rate, a fairly significant amount of traffic is needed to create a solid revenue stream. This is not to say that advertising is the only revenue model available to Web publishers. The revenue stream can include a combination of advertising, subscriptions, licensing content, and sales and merchandising.

The Internet is entering U.S. homes and spreading across the globe faster than any other medium in history (see Table 1.2). According to Geoffrey Ramse, Statsmaster at eMarketer (a Web site that tracks Internet statistics, [www.estats.com]), "The net is growing at a rapid pace. In America, 36 new, first-time users jump online every minute of the day—one new net user every 1.67 seconds" (see Figure 1.4). By the first quarter of 1999, AOL reported 15 million users on its service alone. Indeed, the Internet is substantially influencing how Americans spend their time.

At the end of 1998, America Online commissioned a study to help better understand the impact of the interactive medium. The America Online/Roper Starch Cyberstudy 1998 found that online communication is now permanently imbedded into how we communicate at work and at home. The study found that two-thirds of Internet consumers have been online three or more years, with three-quarters believing Internet services have bettered their lives. Nearly 90 percent responded that they would miss the Internet if it were no longer available. How valuable is the Net? When asked what they would prefer if stranded on a desert island, over two-thirds said a computer with Internet access, only 23 percent said a telephone, and barely 9 percent said a television. "My, how times have changed!" exclaimed Steve Case, AOL CEO in his Jan 8, 1999 letter to AOL members in reviewing the results of the study.

Other findings include these:

- Approximately 90 percent of online consumers regularly or occasionally go online to communicate with friends and family.

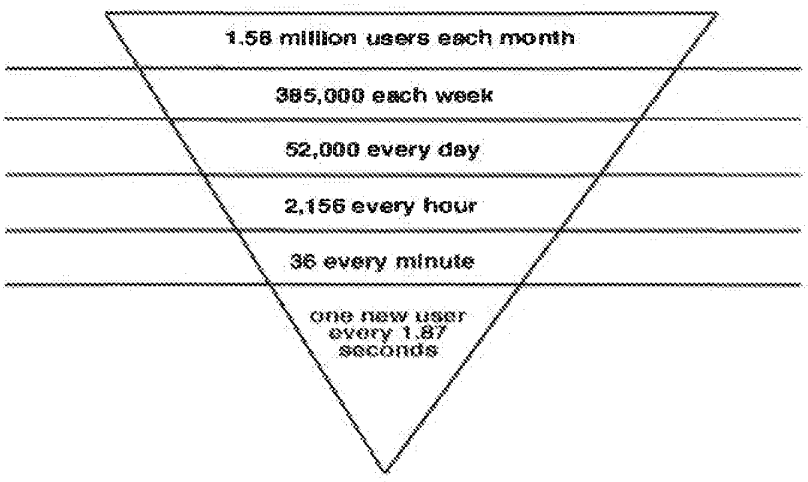


Figure 1.4 Rate of growth of new Internet users.
Source: 1998 estats (www.estats.com)

- In the workplace, more than 50 percent of online consumers claim e-mail is their preferred mode of communicating with colleagues, clients, and customers.
- The majority of online consumers feel that the areas the Internet will most greatly affect are education, work, media, and entertainment.

The impact of the Internet today is remarkable—even more so when we recognize that only one-quarter of American households are online, with nearly two-thirds of these getting online for the first time within the past two years.

Indeed, advertisers are jumping online, but the industry has a long way to go before it will be a staple in the media budgets of big consumer brands. Many of the major advertisers are also the major publishers that advertise their sites to increase traffic to entice more advertisers; compare the top 50 online advertisers in June 1998 (see Table 1.3) with the top 50 online publishers (see Table 1.4).

Indeed, the only constant in the Internet has been its phenomenal growth, development, adoption, and acceptance. As the newest advertising medium on the block, the Web presents challenges from measurability to auditing to customized ad delivery, but what was promise and potential just months ago is now practiced with quantifiable results today. Every year brings significant technological developments. Every quarter sees rapid growth in the online population. Every month, new market leaders emerge in industry segments that did not exist before. And consumers are spending more and more time online. Internet usage has skyrocketed. The AOL members' daily usage has

Table 1.3 Top 50 Advertisers Online (January–June 1998)

RANK	ADVERTISER	YTD SPENDING
1	Microsoft Corp.	$ 17,230,670
2	IBM Corp.	$ 14,864,940
3	Digital Equipment Corp.	$ 5,565,804
4	Excite Inc.	$ 4,918,606
5	General Motors Corp.	$ 4,725,033
6	InfoSeek Corp.	$ 4,515,489
7	Netscape Communications Corp.	$ 3,767,448
8	Lycos Inc.	$ 3,672,075
9	3Com Corp.	$ 3,493,988
10	Datek Securities	$ 3,489,768
11	AT&T Corp.	$ 3,438,059
12	Intel Corp.	$ 3,349,221
13	Walt Disney Co.	$ 3,330,960
14	Compaq Computer Corp.	$ 3,314,658
15	Yahoo Corp.	$ 3,148,160
16	Barnes & Noble Inc.	$ 3,111,673
17	Hewlett Packard Co.	$ 2,983,649
18	CBS Sportsline USA Inc.	$ 2,563,586
19	Honda Motor Co. Inc.	$ 2,511,553
20	Sony Corp.	$ 2,442,539
21	CDnow	$ 2,438,828
22	i-Village	$ 2,402,210
23	Toyota Motor Sales USA, Inc.	$ 2,284,458
24	VISA International	$ 2,236,762
25	Fidelity Investments	$ 2,099,978
26	Sun Microsystems Inc.	$ 2,046,611
27	Amazon.com Books	$ 1,997,383
28	Ford Motor Co.	$ 1,969,524
29	Charles Schwab Corp.	$ 1,907,353
30	SPRINT Communications Co.	$ 1,903,749

continues

Table 1.3 Top 50 Advertisers Online (January–June 1998) *(Continued)*

RANK	ADVERTISER	YTD $ SPENDING
31	Preview Travel Inc.	$ 1,773,367
32	OnSale	$ 1,698,380
33	Procter & Gamble Co.	$ 1,637,977
34	First USA Inc.	$ 1,592,040
35	Egghead Inc.	$ 1,534,144
36	American Express Co.	$ 1,531,462
37	Bell Atlantic Corp.	$ 1,524,291
38	Lotus Development Corp.	$ 1,489,550
39	Ziff Davis Inc.	$ 1,476,981
40	Hearst Corp.	$ 1,475,911
41	Donaldson Lufkin & Jenrette	$ 1,465,012
42	N2K Inc.	$ 1,452,203
43	SAP AG	$ 1,420,033
44	United Parcel Service of America	$ 1,416,423
45	Consumerinfo.com Inc	$ 1,410,167
46	Webnet Marketing	$ 1,390,208
47	CNET Inc	$ 1,349,890
48	Eastman Kodak Co.	$ 1,348,568
49	Verio Corp.	$ 1,335,851
50	Volvo North America Corp.	$ 1,310,878

Source: Jupiter Communications, 1998 (www.jup.com)

Table 1.4 Top 50 Publishers, Ranked by Advertising Revenue (January–June 1998)

RANK	PUBLISHER TITLE	YTD REVENUE
1	Yahoo	$ 54,161,508
2	Excite Network–Excite.com and Webcrawler	$ 47,114,500
3	InfoSeek	$ 33,640,584
4	Lycos	$ 24,775,500
5	Microsoft Network	$ 19,885,509
6	Netscape	$ 18,962,508
7	CNET	$ 18,584,000

Table 1.4 *(Continued)*

RANK	PUBLISHER TITLE	YTD REVENUE
8	ZD Net	$ 17,948,239
9	Altavista	$ 14,804,610
10	CMP Net	$ 9,461,446
11	CNN Interactive	$ 9,349,000
12	Wired Digital	$ 9,092,875
13	Sportsline	$ 8,547,500
14	Wall Street Journal Interactive	$ 8,491,023
15	IDG	$ 8,017,303
16	Pathfinder	$ 6,688,000
17	ESPN Sportszone	$ 6,117,331
18	Weather Channel	$ 6,010,000
19	Homearts Network	$ 5,760,000
20	USA Today	$ 5,107,314
21	Newspage (Individual Inc)	$ 5,060,000
22	AudioNet	$ 5,034,943
23	GEOCITIES	$ 4,268,750
24	New York Times	$ 3,251,000
25	RealAudio	$ 3,246,000
26	Disney Online	$ 3,130,300
27	iVillage	$ 2,870,750
28	Windows95.com	$ 2,717,395
29	NBC	$ 2,635,000
30	Warner Brothers	$ 2,430,000
31	Quote.com	$ 2,312,500
32	Tripod	$ 2,130,468
33	Playboy	$ 2,074,544
34	Discovery Channel Online	$ 2,015,000
35	The Globe	$ 2,015,000
36	Women.com	$ 1,967,050
37	Riddler	$ 1,532,500
38	Car and Driver	$ 1,509,115

continues

Table 1.4 Top 50 Publishers, Ranked by Advertising Revenue (January–June 1998) *(Continued)*

RANK	PUBLISHER TITLE	YTD REVENUE
39	Sony Station	$ 1,475,000
40	Internet.com	$ 1,473,150
41	Morningstar	$ 1,400,000
42	MapQuest	$ 1,371,772
43	Epicurious	$ 1,334,500
44	E! Online	$ 1,307,500
45	Los Angeles Times	$ 1,235,250
46	Jumbo! The Official Web Shareware Site	$ 1,178,699
47	Attitude Network	$ 1,113,016
48	Stockmaster	$ 1,091,000
49	PC Quote	$ 1,081,666
50	Gamesville - Home of Bingo Zone	$ 1,055,000

Source: Jupiter Communications, 1998 (www.jup.com)

increased from 14 minutes per day in September 1996 to 45 minutes per day in July 1998 (see Figure 1.5).

The online mantra is supported in story after story in the press, with Web addresses plastered on everything and everywhere from billboards to beer bottles. Every movie, sports team, and beverage is marketing online. There is no question that Web advertising is starting to take center stage.

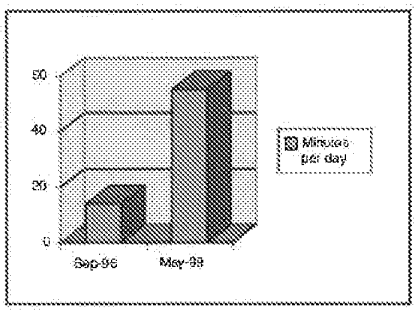


Figure 1.5 AOL member's daily usage 1996–1998.
Source: AOL Company Reports, July 1998

What Is Online Advertising?

Online advertising, like all advertising, attempts to disseminate information in order to affect a buyer-seller transaction. But Web advertising differs from other media by enabling consumers to interact with the advertisement. Consumers can click on the ad for more information, or take the next step and purchase the product in the same online session. Consumers can even purchase products from within Web advertisements.

On the Web, every element of the interactive canvas can be made into an ad, from product placement (a product name or logo used in text or in a display) to wallpaper designs (background images appearing behind editorial content) made up of a corporate logo or any other brand imagery. For example, Disney used a black and white spotted background on Yahoo! to promote the release of the remake of *101 Dalmatians* (the ad concept was developed by i-traffic [www.i-traffic.com]).

Web advertising gives advertisers the opportunity to precisely target an audience, enabling them to deliver advertisements that are customized to each user's particular interest and tastes. Other media offer targeting capabilities, but as you'll see later in this book, the Web offers a level of targeting that advertisers can't find anywhere else.

Online Advertising and Traditional Advertising Vehicles

Can we compare online advertising to a television commercial or an ad in a magazine? Should a Web advertisement be compared to a direct mail piece or a roadside billboard? Should we look to the cable TV or direct response industry for historic precedence? Looking at online advertising in terms of traditional ad models will reveal the breadth and depth of this medium as an advertising vehicle. To that end, we'll discuss billboards, print, direct marketing, and television.

Outdoor advertising: The Billboard. Web banner ads are often likened to roadside billboards. Like a billboard, which can refer passersby to an upcoming attraction or location, a banner ad may point Internet users to a Web site. But roadside billboards are static, whereas Web users can interact with a banner ad by clicking on the ad to get to other sites and more information. Some banners even allow consumers to take an action (like playing a game) without leaving the ad—a huge difference between banners and billboards. With a Web ad, an advertiser can measure exactly how many times a banner was clicked. Unless the billboard is a direct response campaign with a special phone number, there usually isn't an accurate method for determining how many people respond to it.

Print advertising. In a newspaper or a magazine, a quarter-, half-, or full-page advertisement is clearly identifiable. Advertisements are formatted in a way that allows readers to easily separate them from the editorial content. If a magazine advertisement resembles editorial text, the magazine will place a banner across the top identifying the text as a "special advertising section." On the Web, banner advertising corresponds to the print model with its clearly recognizable borders.

Current screen real estate for banner advertising is limited—usually less than 10 percent of the viewable area. However, as you'll see in Chapter 2, Web advertising is not merely banners, and it is not always clearly identified as advertising; it can actually be quite difficult for a consumer to differentiate between advertising and editorial.

Television advertising. The 30-second television commercial is one of the premiere advertising formats. Television commercials allow an advertiser to "take over" the television screen completely, delivering a full-motion advertisement that can target a consumer's interests, needs, and even emotions.

Should a Web ad be more like a TV commercial? Some think so. The idea of an ad commanding the entire screen has inspired the development of ad models that are more dynamic and intrusive, as will be discussed in Chapter 2.

Cable television advertising. Cable television traded the reach and frequency of broadcast TV for better access to niche audiences. The cable audience is more upscale and extremely attractive to advertisers. Increased targeting has lead to specialized commercials with targeted and provocative messaging. Like cable, the Web offers many targeted ad buys, in addition to the buys that allow for massive reach.

Direct marketing. From direct mail to telemarketing, direct marketers can accurately track the number of people who respond to a specific promotion and use this number to determine the profitability of each campaign. The Internet's innate accountability has led advertisers to take a direct marketing approach to the Web—specifically, basing their campaigns on their cost per response.

Internet Advertising—A Definition

Defining Internet advertising has almost become a sport within the industry. At first people defined it in terms of whatever industry they came from. If a person came from cable, then Internet advertising was like cable. If a person came from print, then Internet advertising was like print. Then the debate shifted to whether Internet advertising was more like traditional advertising, which focused on building brand awareness, or direct marketing, which sought to generate leads or execute sales.

The days of always having to compare the Internet with TV, cable, or print for credibility are over. This either/or definition has shifted to one of overlap and inclusion. Internet advertising is the convergence of traditional advertising and direct response marketing (see Figure 1.6).

The AOL advertising team takes the definition one step further. "We don't look at Internet advertising as just brand building or direct marketing," says Steve Keenan, VP of marketing and operations at AOL. "You can have branding like TV, radio, print. You can also get vast amounts of information to people like print or direct marketing. And you can also make a transaction happen; it's just like going to the store or calling an 800 telephone number." For AOL, Internet advertising is the convergence of branding, information dissemination, and sales transactions all in one place.

Internet Advertising

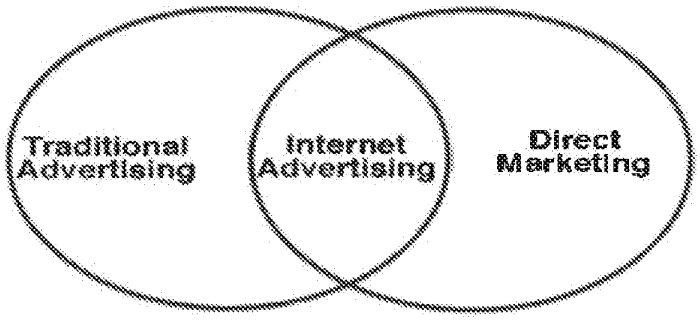


Figure 1.6 Internet advertising is the convergence of traditional advertising and direct response marketing.
Source: 1999 The Zefi Group (www.zefi.com)

The Online Advertising Advantage

Online advertising holds four distinct advantages.

Targetability

Online advertisers have an entirely new range of targeting capabilities They can focus on users from specific companies, SIC codes, or geographical regions and nations, as well as by time of day, computer platform, and browser. They can target using the databases that serve as the backbone of direct marketing. They can even target based on a person's personal preferences and actual behavior.

Tracking

Marketers can track how users interact with their brands and learn what is of interest to their current and prospective customers. For example, a car manufacturer can track how a user progresses through its site to determine whether more users are interested in the safety information or the "extras" that come with a particular model.

Advertisers can also measure the response to an advertisement (through the number of times an ad is clicked on, the number of purchases or leads an ad generated, etc.), which is difficult to do with traditional television, print, and billboard advertising.

Deliverability and Flexibility

Online, an ad is delivered in real time 24 hours a day, 7 days a week, 365 days a year. Furthermore, an ad campaign can be launched, updated, or canceled immediately. An advertiser can follow a campaign's progress daily, notice that a campaign is generating very little response in the first week, and replace it by week 2. This is a big difference from print, where an ad cannot be changed until a new edition of the publication is published; or TV, where the high costs of ad development make frequent changes prohibitive.

Interactivity

An advertiser's goal is to engage the prospect with a brand or product. This can be done more effectively online, where consumers can interact with the product, test the product, and, if they choose, buy the product. For example, an advertisement for software can take a user directly to the location where a demo can be downloaded and tested immediately. If the consumer likes the software, he or she can purchase it right then and there. No other medium moves the consumer from information seeker to purchaser friction free.

A Short History of Online Advertising

The history of online advertising spans a very short period of time, from the launching of the commercial online service Prodigy, through the so-called spam wars, and to the coming of age with the launch of HotWired, originally the online sister company of *Wired* magazine.

In the Beginning

Long before the commercialization of the Web, Prodigy was testing the online advertising waters. When the service went live in 1990, it adopted advertising as a revenue source. Prodigy was the lone experiment in this area; the other commercial online services shied away from incorporating advertising into their content mix. CompuServe went online in 1979 but did not sell its first sponsorship until 1995. Likewise, America Online (AOL) did not start addressing advertising seriously until the fall of 1995. Today, many members of Prodigy's original online advertising sales team head up advertising sales efforts at the top sites.

Setbacks: The Spam Wars

As Prodigy was testing advertising within the subscriber-only environment, others were exploring the commercial implications of Internet applications. In 1994, the law firm of Canter and Siegel thought they had discovered the silver bullet for tapping the Internet as a cheap marketing medium and posted an advertisement for green card

assistance to over 7,000 newsgroups (Internet discussion areas). Internet etiquette, known as *netiquette*, dictates that discussion forums not be used to post blatantly commercial materials that have nothing to do with the subject of the discussion group. Such unrelated and commercial posts, as well as unsolicited commercial e-mails, became known as *spam*. Spam was disruptive to the discussion group, polluting the topic of conversation, and often resulted in large volumes of e-mail complaints, known as *flames* sent to the spammer and his Internet Service Provider (ISP).

Although Canter and Siegel did receive requests for green card assistance as a result of the posts, the resulting outrage hurt the firm's reputation within the online community. Canter and Siegel's ISP crashed 15 times under the flood of 30,000 flames in only 18 hours. Their ISP subsequently canceled their account.

Canter and Siegel next set up accounts at two other ISPs, insisting they would place even more advertisements. They received tremendous press coverage in print, but were blacklisted from the Internet community—ISPs dropped their accounts rather than face thousands of angry e-mails.

Acceptance

The Canter and Siegel spam wars occurred just before *Wired* magazine launched HotWired (www.hotwired.com), a Web property with an advertising business model. Fearing backlash from the cyberspace community, HotWired reduced the on-screen dimensions it had originally planned to use for its first advertising unit, the banner ad.

HotWired signed its first ad contract with AT&T on April 15, 1994, and launched its site on October 27, 1994 (see Figure 1.7). Its inaugural sponsors represented a cross-section of the advertising community, from such blue chip advertisers as AT&T and IBM to Pepsi's alcoholic beverage Zima.

To HotWired's pleasant surprise, the only negative comments received were in response to portions of the Web site that were still under construction; no one criticized the advertisements.

Although HotWired hoped to profit from its advertising sales, it did not begin with a dedicated sales force. Instead, it used the national sales force for its print magazine *Wired*. This proved satisfactory because there were only two agencies buying the vast majority of online advertising: Modem Media (now Modem Media Poppe Tyson [www.modemmedia.com]) and Messner Vetere Berge McNamee Schmetterer (www.mvbms.com). Together, the clients of these two agencies accounted for almost half of HotWired's charter advertisers. "Our charter advertisers were visionaries," recalls Rick Boyce, vice president of advertising at HotWired. The story goes that the folks at Modem Media were so convinced of the significance of being a charter sponsor of HotWired that they bought ads before they had final approval from AT&T, since no one at AT&T was ready to take responsibility for online ad buying. Modem Media's gamble paid off: the charter sponsors received tremendous press coverage as part of the site's launch.

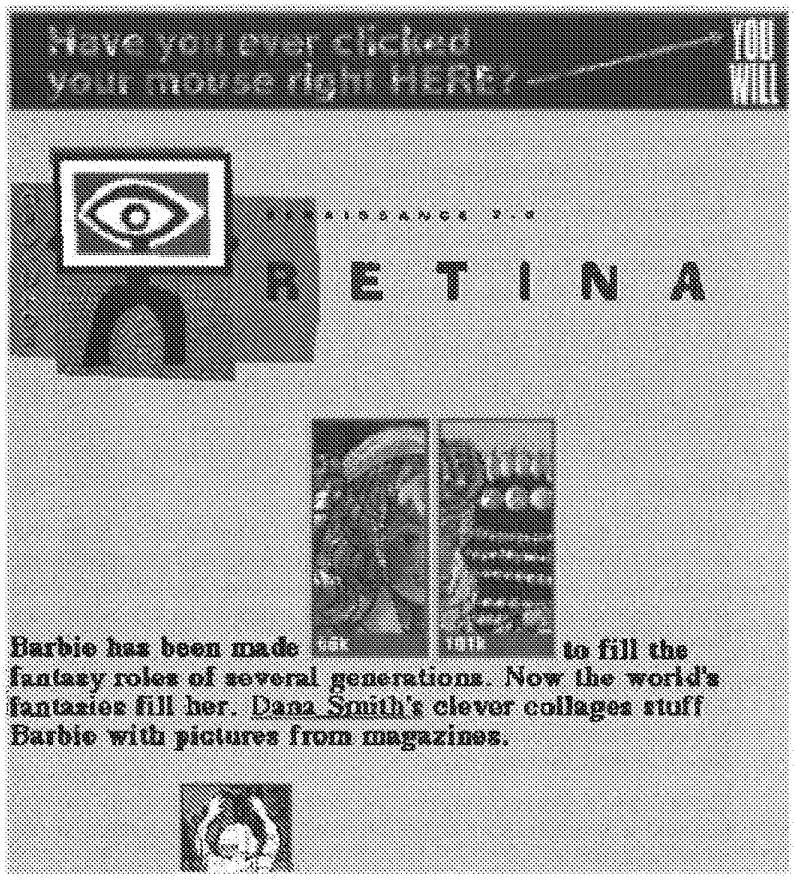


Figure 1.7 First banner ad on HotWired.
Source: HotWired (www.hotwired.com)

Word (www.word.com) was the second "zine" (Web magazine) to sign sponsors. It too received its first sponsorships—Zima and MasterCard—from the Internet advertising visionaries at Modem Media. Soon, other sites were scrambling to accept advertising, and today there are hundreds of thousands of sites that accept advertising.

Wake Up, It's a Mass Market

From 1994 to 1996, some of the major advertising agencies on Madison Avenue maintained a wait-and-see attitude toward the importance and viability of Internet adver-

tising. The joke used to be that the mantra on Madison Avenue was "wake me when it's a mass market." Madison Avenue is awake now.

Over time there have been pivotal events that have cemented a medium's place to a particular decade in the history of communication. For the 1920s, it was the radio delivering Calvin Coolidge's State of the Union addresses into the home. In the 1960s, it was the televising of the Nixon-Kennedy debate that completely changed politics and campaigning. And for the Information Age of the 1990s, the pivotal event will surely be the release of the Starr Report, where over a two-day period 24.7 million unique visitors used the Internet to view the text of the report. And the Start Report even became an advertising issue, because sites sold ad space next to this high-traffic-generating content.

The Internet started out as a medium used primarily by the infamous early adopters, those consumers who will spare no expense to have the latest electronic or computer gizmo. However, in 1995 the number of current users started to extend into mass market turf. Indeed AOL, with its 15 million users and climbing, is in the same league as TV mass marketing heavyweights (see Figure 1.8). Correspondingly, Internet users accept online advertising as they do advertising on all other media.

A report conducted by Opinion Research Corp. International in 1998 found that Internet brands such as America Online, Yahoo!, and MSNBC are recognized by at least half of America's adults, with brands such as Netscape, Amazon.com, and Priceline recognized by at least a quarter of U.S. adults. "The most valuable real estate in the world is screen real estate," says Brewster Kahle, president of Alexa Internet.

Critics

Online advertising is not without its critics (see Table 1.5). There are several software products that block ads and promote themselves as speeding up Web page download time. These products include JunkBuster (www.junkbuster.com), @tGuard

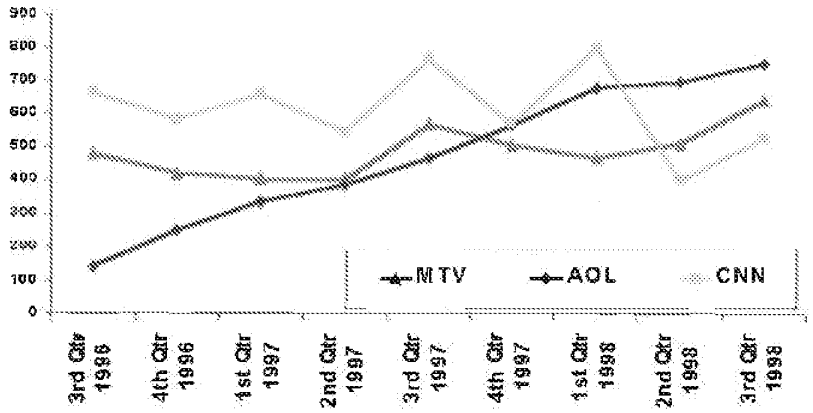


Figure 1.8 AOL is now in the same league as TV mass market heavyweights.

Table 1.5 Web Advertising Opinions

	DISAGREE	AGREE
Advertising is needed to keep content free	11%	53%
Targeted ads are fine if site or ISP doesn't give out my info	18%	50%
Web ads are less obtrusive than direct-mail offers	23%	43%
Web ads are less obtrusive than TV commercials	24%	43%
Would pay extra to subscribe to site without ads	37%	22%
ISPs should be able to sell subscriber lists, like magazines do	61%	16%

Source: 1998 INTECO survey of 2,500 U.S. adults.

(www.atguard.com), and WebWiper (www. webwiper.com), to name a few. Fortunately, these products are not taking off as vigorously as advertising. The online audience understands that there is a price to pay for free access to content, and that price is advertising support. Besides, the average consumer does not find Web advertising irritating enough to warrant taking the time to download and install software to block the ads.

Players in the Internet Advertising Industry

The cast of characters in the Internet advertising industry can be divided into three camps: the sellers, the buyers, and the advertising infrastructure—those companies that make Internet advertising possible by developing the tools to deliver the ads, measure the ads, audit the ad campaigns, and complete the other tasks that make advertising on the Internet work.

It is important to note that in the online advertising world, the line between the buyers and the sellers is narrow. Many publishers buy ads to promote their Web sites, and advertisers develop Web sites to promote their products.

The Sellers

The Web is now made up of hundreds of thousands of sites and the numbers are growing each minute. Young entrepreneurs, many of them college students with free Internet accounts, developed the first sites. Other early players were those folks ready to embrace technology or who looked to the Web as a place to express themselves.

Web Site Publishers

Today, major publishers are media companies. Some sites are produced by a single company, like Washingtonpost.com, which is the online version of the *Washington Post*

newspaper (www.washingtonpost.com). Other sites represent different channels of interest and are owned by one parent company such as Microsoft's Expedia (www.expedia.com) or Carpoint (www.carpoint.com). Still others are joint projects of two or more companies—for example, MSNBC (www.msnbc.com) is a joint project of Microsoft and NBC.

In 1997, AOL was limping into the advertising arena. But in 1998 alone, AOL had signed up more than 400 advertisers, with as many as 55 spending in excess of $1 million annually a piece. "Whether you're 6 or 60, we want you on AOL," says Steve Case, AOL Chairman. AOL hopes to sell advertising that targets a number of demographics—not simply the technically savvy males who were the majority of early Internet users. And though AOL was not the early leader in accepting advertising on online services, it currently dominates the advertising landscape.

Sales Channels

In the seller's camp, the key players are indeed the publishers, who produce the sites with advertising space to sell; but the facilitators of the sales are equally important and include ad networks, representation firms, and ad auctions that sell advertising for the publishers.

Ad networks. To assist sites in generating revenue from advertising, ad networks have formed to sell advertising. The ad network represents a slate of sites serving as a sales force for publishers who participate in their networks. These networks benefit the publishers by providing an aggressive and experienced sales force. Also, a publisher's inventory may be more difficult to sell individually than when it is part of a larger inventory pool. The networks make it possible for advertisers to extend their reach into a variety of sites with one media buy. In addition, advertising agencies are often more willing to meet with a network rather than have dozens of meetings with individual Web site publishers.

Representation firms. Sites with advertising to sell can hire a firm to sell the ad space on their behalf.

Auctions. Publishers can offer soon-to-expire ad space on the auction block for quick sale. There are also ad networks that allow publishers to sell their excess ad inventory in real time.

The Buyers

The buyers consist of the advertisers who have products, services, or Web sites to promote. They are often represented by interactive or traditional agencies that design online advertising campaigns and the media buyers who actually purchase the ad space.

Advertisers

The first advertisers to jump on the Web en masse were the Internet publishers themselves. In fact, the list of the top advertisers for 1996 reads similarly to the list of top

publishers for that year (see Tables 1.3 and 1.4). Indeed, those with the most to gain from the growth of the Web invest the most money.

In 1998, one of the largest consumer advertisers in the world, Procter & Gamble, moved from the sidelines into the headlines. Procter & Gamble hosted a conference in Cincinnati to bring together major global advertisers, online publishers, interactive advertising agencies, leading technology companies, and others to discuss the best means to accelerate the growth and development of online advertising. Titled the Future of Advertising Stakeholders (FAST) Summit (www.fastsummit.com), the two-day meeting focused on three guiding principles:

> FAST will focus on developing action plans that accelerate the development and use of digital advertising models that are both effective and consumer-acceptable in order to enable interactive digital media to become more available, more useful, and more enjoyable for the end user/consumer.
>
> FAST must be driven by committed advertising, media, and technology firms determined to help achieve win/win solutions for the entire industry.
>
> FAST efforts will focus on both broadening acceptance of the current state of the space (through effective knowledge sharing) as well as accelerating the development of breakthroughs that will more rapidly enable its future potential.
>
> (Source: www.fastsummit.com/pages/index.cgi/principles)

The historic event resulted in the formation of the Fast Forward organization (www.fastinfo.org), a broad-based Internet industry coalition formed to accelerate development of online marketing with such committees as consumer acceptance, ad models, measurement, and media buying. Even more importantly, this event brought the major brand advertisers to the table to commit to investing in the Internet as an advertising medium, making advertising on the Internet part of every brand's marketing mix.

Interactive Agencies

The link between the advertiser and the publisher is often the interactive agency. The birth of interactive advertising did not take place on Madison Avenue, the home of the ad agency. Instead, it bloomed in Multimedia Gulch in San Francisco and New York City's Silicon Alley.

In 1994, Ed Artzt, CEO of Procter & Gamble and president of the American Association of Advertising Agencies (a.k.a. 4As), scolded his association for not embracing the possibilities of the new technology. Though at the time Artzt was referring to interactive television, it soon became clear that the Internet was the dominant medium on the horizon. While traditional ad agencies took a wait-and-see approach to the Internet's potential, "boutique shops" pioneered online advertising. These shops did not look at online advertising as new wine in old bottles, but embraced the nuances of cyber culture, to make Web advertising a vehicle that worked.

Following in the tradition established by former P&G CEO Artzt, Denis Beausejour, vice president of advertising at P&G in 1998, proclaimed that in ten years there would be two kinds of agencies: digital agencies and dinosaur agencies.

Today's interactive agencies have their origins in Web site development, direct marketing, and traditional advertising agencies. The agencies that began as Web site development shops had strong technical backgrounds. These shops built the first corporate Web sites and brought in the marketing and advertising expertise to make their companies full-service operations later. Agencies that came from the world of direct response marketing reacted to the Web's remarkable ability to serve as a direct marketing vehicle. And finally, some interactive agencies were built by visionaries from the ad agency world who set up interactive divisions or formed their own boutique shops.

These early roots are quickly becoming a distant memory, as traditional advertising agency giants try to catch up. Some have set up interactive departments within the company. Others have spun off their own interactive shops. And others have merely acquired the expertise. Whichever the case, traditional ad agencies are now in the interactive business, trying to regain their attractiveness with advertisers as full-service providers.

Advertising Infrastructure

A service and software solution infrastructure has formed to provide the tools to help publishers and advertisers deliver results through Internet advertising.

Traffic Measurement

The Internet's keen ability to conduct in-depth measurement of Web site and advertising activity has given birth to an entire measurement industry that includes the counters, the auditors, and the comparative measurement researchers. Today, Web publishers can measure everything that occurs on a Web site and everything that occurs within a Web ad. In other advertising media, measurement is based on sample. On the Web, however, measurement is done by census, every impression is counted.

Targeting and Personalization

Targeting and personalization are key advantages of Internet advertising and essential components of the advertising infrastructure for publisher and advertiser alike. The targeting and personalization software companies are central to the online advertising industry.

Ad Management

Software solutions that coordinate the logistics of ad placement are available for the online advertising industry. The software takes care of the daily management of ad placement, rotation, reporting, and billing—all the required components of a Web advertising campaign.

The technological infrastructure of online advertising sets it apart from traditional advertising. Chapters 4 and 5 explain how these tools work. But before we jump into a discussion about the tools of the trade, we will examine the types of ad models that are in use today and being developed for tomorrow.

It's That Simple. It's That Complex.

Tom Lix, President/CEO (tlix@newmarket.net)
NewMarket Network (www.newmarket.net)

The Internet is just like anything else, like print, like television, like radio, like billboards, like the insides of a book of matches. Except for one thing. People can talk back. They can have a conversation. They can ask questions. They can complain. They can compliment. They can suggest. And, they can demand. Not only "can" they do all these things, but they "will" do all these things and more, much more.

As marketers we've always talked about communicating with our customers, even in recent years when the mantra of "1 to 1 marketing" appeared. But in reality, we never really communicated with our customers. Rather, we talked "to" them—sometimes as a market, and sometimes as an individual—but always talking "to" them, not with them.

The Internet is changing all of that. While most people don't admit it, I'd venture a guess that most people in business are scared to death (by the way, my advice is, get over it real fast).

So here's the deal. Embrace this dialog; use it to your advantage and to your customer's advantage. Be willing to listen to all the problems, face the tough questions, deal with the irate customers, and use it to improve (just do it quickly).

Consider advertising as a means to talk with, not at your customers. Advertising is no longer a means to define or obfuscate your brand. It is no longer a manageable device to communicate "selected" attributes and benefits. Advertising is certainly not a simple tool to carefully portray or reinforce an image that "you" want to deliver. The Internet means we have to talk with our customers. We have to listen to them. We have to respond to them.

The Internet also means our customers not only talk to us, but they talk to each other quickly and in numbers that you never imagined. The old adage about dissatisfied customers telling ten of their friends is now about telling thousands. Traditional advertising just won't be adequate anymore. The only thing that will suffice is an open, honest dialog with our customers. It's that simple. It's that complex.

CHAPTER 2

Online Ad Models

The first form of Web advertising was the Web site itself. But as the Web became cluttered with commercial sites, simply building a Web site wasn't enough to reach Internet consumers. Advertisers needed a tool for driving users to their Web sites, and publishers needed a way to pay for their Web endeavors. The first Web advertising models were banners and buttons that resembled the traditional print advertising format by appearing within clearly defined borders. But, the limited screen real estate used for banners and buttons, along with an average click-through rate that is now under 1 percent (according to NetRatings [www.netratings.com]) inspired the development of new, more obtrusive types of Web advertising. In this chapter we'll examines the many advertising models used on the Internet today.

Web Site as Ad

For many companies, the first foray into Internet advertising was building a Web site. These sites were usually built because the company thought the Web was "cool," having a Web site would make the company look "hip" or the company was too afraid not to have a Web site and appear behind the competition. The first sites ran the gamut from "brochureware," which was nothing more than print material reproduced online, to the opposite extreme—a site that employed every technological and design feature available.

Both tactics were disappointing in failing to deliver results and caused advertisers to rethink their use of corporate Web sites. Brochureware didn't take advantage of the depth of information that could be presented in a user-friendly manner through the Web,

Meta's Exhibit No. 1004
Page 0039

and so-called cool sites became prohibitively expensive to produce and lacked a clear objective other than showcasing new technology. Consequently, advertisers began to produce Web sites that contained product information potential customers actually needed. Web users are looking for information; when a person visits a company's Web site, that person expects product information. Buyers of large-ticket items—especially lifestyle items like cars and homes—usually want as much information as possible. This is not to say that advertisers are no longer trying to make these sites interesting and fun, but they are not investing hundreds of thousands of dollars in Web sites designed only to entertain and not educate the users on their wares. Web sites are required to fulfill clear business objectives.

Companies are learning to develop innovative Web strategies that move a potential customer from browser to buyer. For example, on the Cadillac site (www.cadillac.com), a consumer can plan every aspect of buying or leasing a new Cadillac through an interactive showroom. In the showroom, the consumer chooses the model of the car she is interested in and then the screen presents an image of that car. From that point, the user decides on the options she wants—from exterior color to interior design. Each time she chooses a feature, the sample car is redrawn on the screen to include the newly chosen option. After the user has composed her perfectly outfitted model, she is provided with a suggested retail price as well as financing options. From there the user can locate a dealer in her area. All of this is done in the comfort and privacy of the user's home; there are no set hours and no pushy car dealers.

Sometimes the product itself isn't conducive to a heavy information type site and a humorous approach works well. One of the first to succeed at this strategy was the Ragu spaghetti sauce site (www.eat.com). The site centers around Mama, the woman behind Ragu. The site offers humorous stories from Mama, interesting background information on Italy, and Italian lessons where Web-surfing parents can learn phrases such as "Clean your room, or no spaghetti for you tonight." The Ragu site was an instant hit receiving high traffic and tremendous publicity for blending content with the type of irreverent humor that the Internet audience adores.

The tradition of blending humor with information has progressed and we now have sites that are witty, informative, and successful. A good example of this genre is the Web site for the radio program Car Talk (www.cartalk.com). The Car Talk site uses irreverent humor to discuss the trials and tribulations of persistent car problems (see Figure 2.1a) as well as responding to letters from their audience (see Figure 2.1b). The site was designed and developed by Public Interactive (www.publicinteractive.com).

We're learning that brand credibility is most successful when the information offered is kept in context and on topic. An outstanding example of this is the site for Tide detergent, titled The Tide Clothesline (www.tide.com). One might think that there is little compelling content to place on a Web site for a laundry detergent. However, the designers of Tide kept clear of kitsch and instead came up with site features that have a high value to their customers. They start off by giving laundry tips and time-saving suggestions and then move on to one of the most valuable features of the site: "The Stain Detective" (www.tide.com/stainDet/) (see Figure 2.2). This page is a search

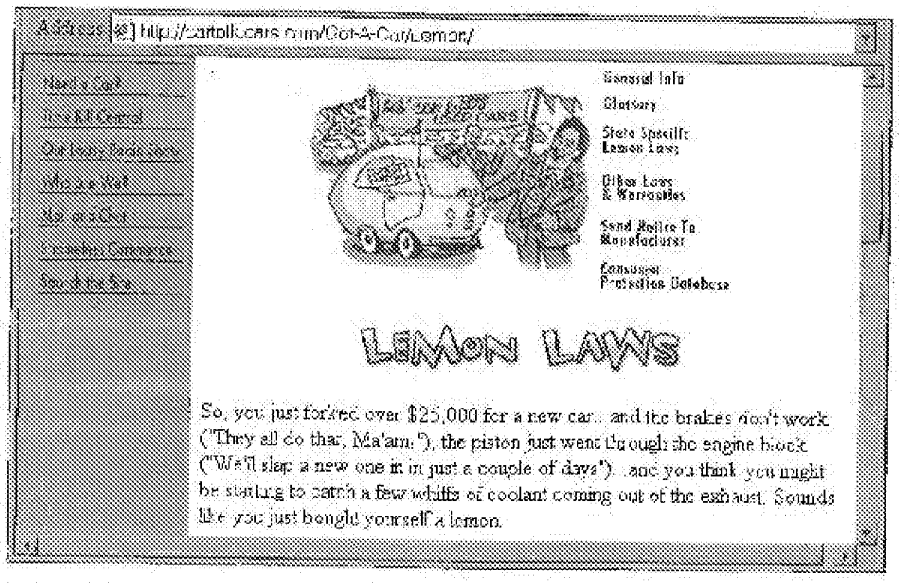


Figure 2.1a The Car Talk site uses humor to provide a wealth of information about "Lemon Laws."

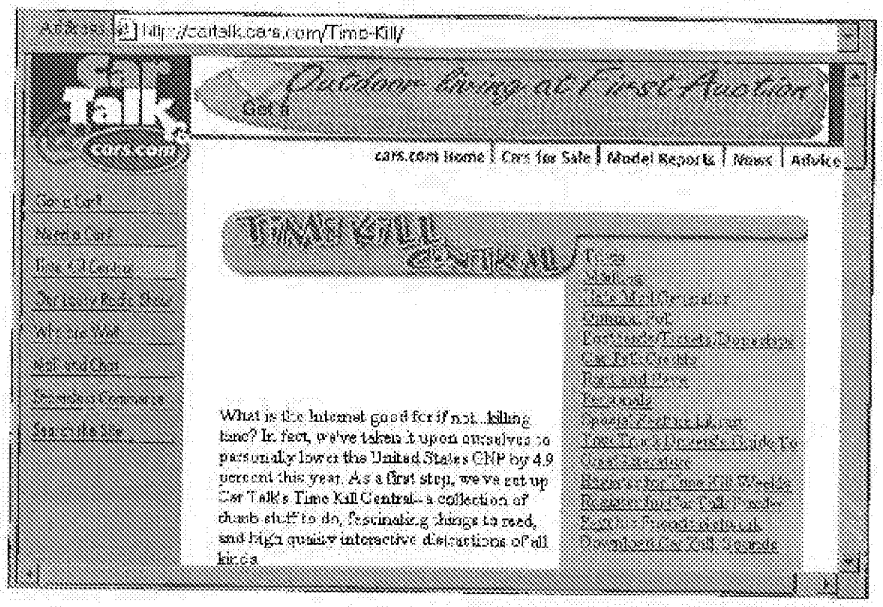


Figure 2.1b Humor goes a long way in branding the Car Talk style of interacting with their customers.

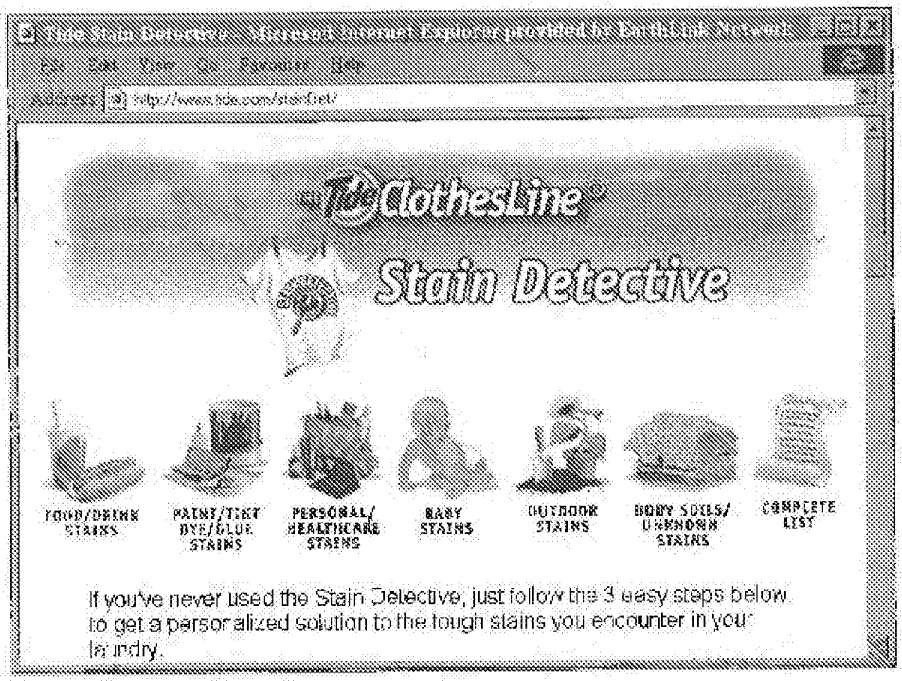


Figure 2.2 Tide adds value to its site through its Stain Detective feature.

engine providing extremely useful stain removal information, making the Tide site a valuable resource based on its own content alone.

Although our discussion has focused on Web site advertising, some brand/product sites fulfill other functions such as customer service. The Federal Express site (www.fedex.com), for example, does an excellent job of servicing its customers by offering package tracking; in that way, the site is an excellent tool for building brand loyalty and creating an image of a highly advanced delivery service—not to mention that the online service represents savings to the company (see Figure 2.3).

Ad Models

Online advertising models make use of the Internet's two most popular applications: e-mail and the Web.

Advertising via E-mail

What is the first thing people do after connecting to the Internet? They check their e-mail. Although e-mail is the most widely used Internet application, and an effective advertising vehicle, it is often overlooked by advertisers. When marketers hear about e-mail marketing, they often scoff at the idea, calling it spam. *Spam*, which is unsolicited

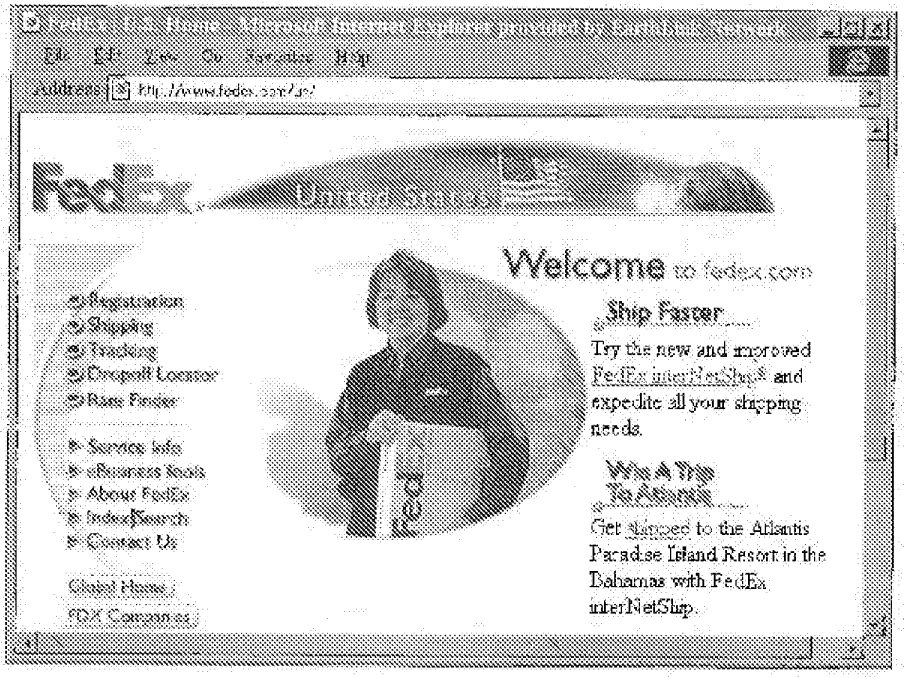


Figure 2.3 Federal Express makes customer service and ease of use top priorities.

e-mail messages and advertisements, is a dangerous marketing tactic that we do not recommend. However, there are many effective ways to use e-mail advertising, which we will discuss now, as well as in the next chapter, which covers Internet direct marketing.

Sponsoring Discussion Lists and E-mail Newsletters

An e-mail newsletter is a publication created by a business or individual and then sent out to people by request. An e-mail discussion group, on the other hand, is made up of "conversations" that subscribers are having on a particular topic. Anyone who wants to contribute to the discussion sends a message to a moderator or an e-mail address that automatically sends the message to all subscribers. Most good e-mail discussion lists have a moderator who reviews all of the messages and decides which are appropriate to be sent to the entire list. This protects the list members from being inundated with irrelevant or unintelligible posts.

In most cases, subscribers to e-mail discussion lists can choose between receiving each message individually or receiving digests (compilations of all messages sent each day or week). There are currently thousands of niche discussion lists and newsletters on every subject imaginable. Table 2.1 lists directories of e-mail discussion lists and newsletters.

E-mail discussion lists and newsletters offer inexpensive and effective advertising opportunities for reaching a niche market. Advertisers can sponsor e-mail discussion lists for a week or month at a time.

Table 2.1 Directories of E-mail Discussion Lists and Newsletters

DIRECTORY	URL
Yahoo List	(http://dir.yahoo.com/Computers_and_Internet/Internet/Mailing_Lists/Web_Directories/)
Liszt	(www.liszt.com)
Reference.com	(www.reference.com)
eGroups.com	(www.egroups.com/index.html)
The List of Lists	(http://tile.net/lists/)
Publicly Accessible Mailing Lists	(www.neosoft.com/internet/paml/)

The benefit of sponsoring a discussion group or an e-mail newsletter is that it reaches a precisely targeted audience. Participants are not casual surfers checking out the content of a particular Web site; they are people who have taken action to subscribe to the list.

E-mail newsletters and discussion groups can be delivered as pure text or as HTML files, which allow messages to look like Web pages. Whether or not the recipient of the newsletter receives an HTML version depends on two things. First, the user needs to be using e-mail software that can read HTML e-mails. Second, the sender has to send HTML e-mails. Figure 2.4 shows an e-mail newsletter with text formatting, and Figure 2.5 shows an e-mail newsletter with HTML formatting.

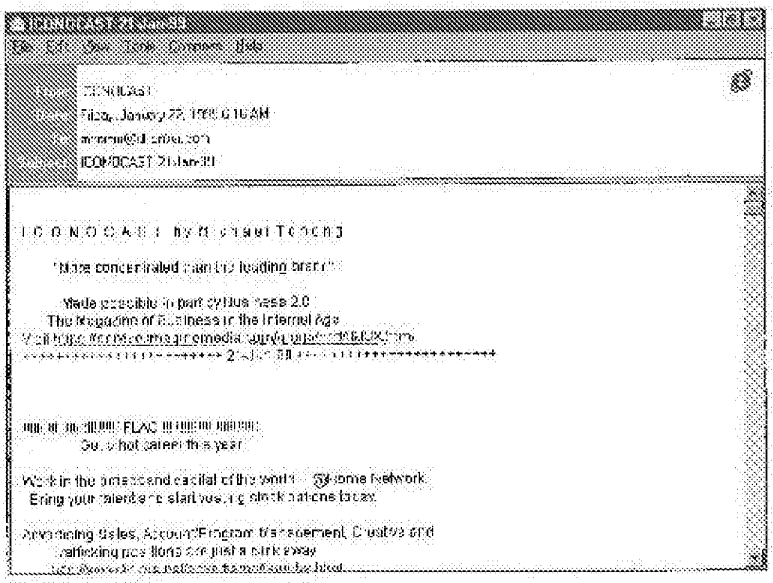


Figure 2.4 E-mail newsletter in text format. This is the beginning of the ICONOCAST e-mail newsletter.

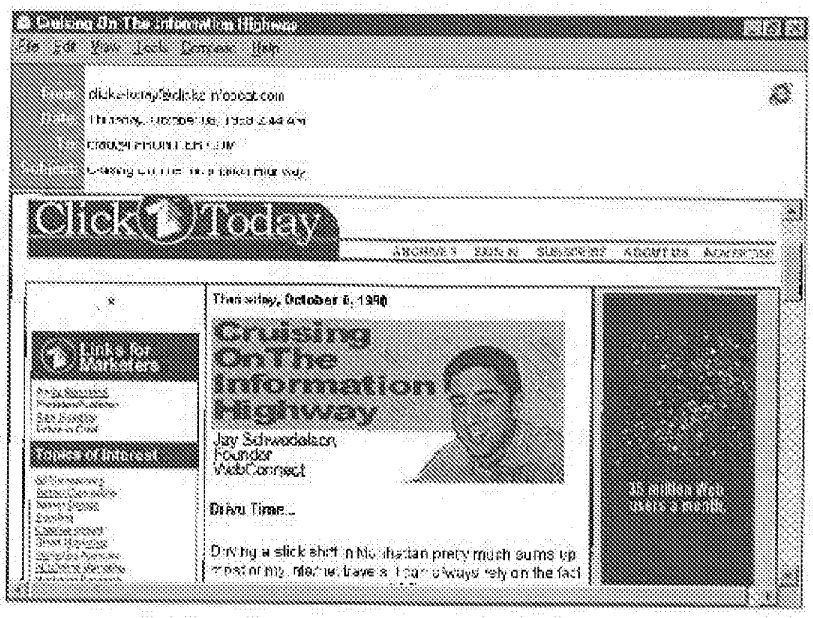


Figure 2.5 E-mail newsletter as HTML, from ClickZ.

Tips for Maximizing the Usefulness of E-mail Sponsorships

Only sponsor moderated lists. This increases the likelihood that your ad will not run alongside (and therefore be associated with) irrelevant posts to the list.

Be introduced. If you're sponsoring a moderated discussion list, ask the moderator to send out a note to the list introducing you as the sponsor. This tip comes from Richard Hoy, moderator of the Online Advertising discussion list (www.o-a.com). Figure 2.6 shows an example of a note that Hoy sends to his list to introduce a new advertiser.

Request multiple placements. In many cases an advertiser is the sole sponsor of a particular publication. If this is the case, the publisher of the list cannot place any other advertisers in the newsletters or discussion list posts that you sponsor. So, why not ask for multiple placements? Have different variations of your ad run in the beginning, middle, and end of the e-mail. You've just tripled your chances of getting a response.

Make an offer. Your ad is surrounding content that people want to read. A strong offer from a discount to a special give away will entice people to take action.

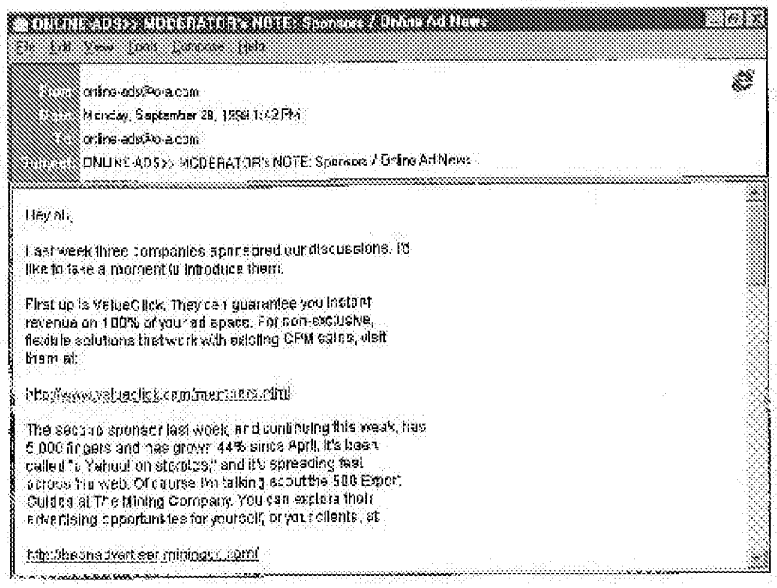


Figure 2.6 A note to the Online Advertising Discussion List sent by moderator Richard Hoy to introduce new advertisers.

When advertising in a text e-mail newsletter, the advertiser can know the total number of people who receive the e-mail message, because the sender can track exactly how many e-mails are sent. However, it is impossible to know how many people actually read the sponsorship text in the message. This is for two reasons. First, some people subscribe to a lot of newsletters and discussion groups that they don't read, because there isn't enough time in the day. And second, many people forward e-mail to others actually increasing the list's circulation. Another consideration is that in a text e-mail newsletter, an advertisement is plain text limiting the advertising to only a certain number of characters to convey its message. See Figure 2.7 for an example of a text sponsorship.

When you're sponsoring HTML e-mail, you can run the same banner and button advertisements that you run on the Web. Figure 2.8 shows an e-mail newsletter from Andover Update. This is an example of a banner ad within an HTML e-mail newsletter. There are benefits to sponsoring an HTML e-mail newsletter rather than a text newsletter. For starters, you can use a graphical advertisement, which will attract more attention than a text advertisement. Also, you'll know how many times the e-mail newsletter with your ad was opened. Every time an HTML e-mail is opened, the graphics are pulled from the publisher's ad server. By looking at the number of times your banner was pulled from the ad server, the publisher can give you an indication of the reach of your advertisement. (Text-based e-mail newsletters cannot provide this information, since they do not pull files from the server when the impressions can be counted.) On the

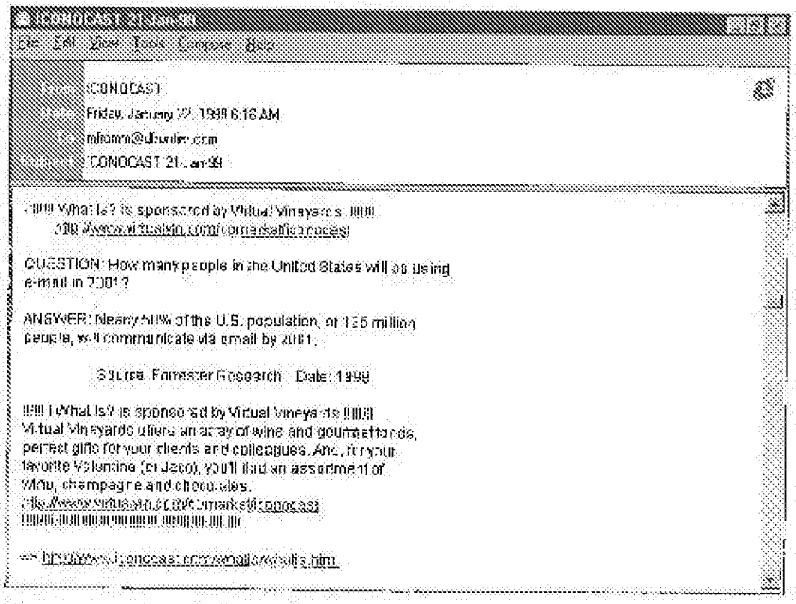


Figure 2.7 Virtual Vineyards e-mail sponsorships of ICONOCAST

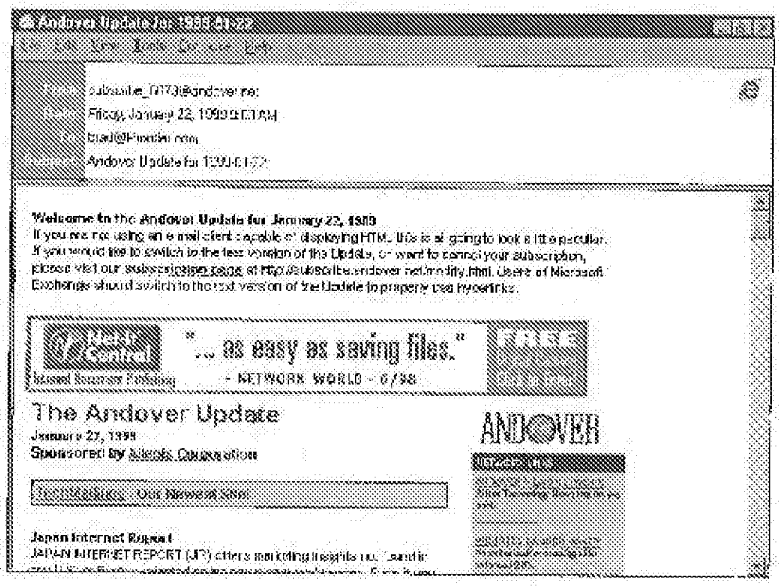


Figure 2.8 Example of an HTML ad in an e-mail newsletter.

other hand, HTML e-mail newsletters can be as cluttered as Web pages, making your ad less noticeable than an ad in a text-only newsletter.

Direct E-mail

In addition to sponsoring discussion lists and e-mail newsletters, which usually give the advertiser only one or two lines of copy, advertisers are using direct e-mail. Early attempts at direct e-mail took the form of spam. Although companies are still using spam, it has proven to deliver extremely poor results (not to mention causing people to hate companies associated with spam). Acceptable direct e-mail advertising is sent to folks who have requested e-mail solicitations on a particular subject. A more thorough discussion of direct e-mail is provided in Chapter 3.

Ad-Supported E-mail

Companies like Hotmail (www.hotmail.com), shown in Figure 2.9, Juno (www.juno.com), and USA.Net (www.usa.net) were some of the first marketers to offer free e-mail access to people who would use their e-mail readers, which display paid advertising. Now, most of the search engines and many additional Web sites offer free e-mail. For a list of free e-mail providers, you can visit:

http://dir.yahoo.com/Business_and_Economy/Companies/Internet_Services/E-mail_Providers/Free_E-mail/

Most of these services require users to visit their Web site (and therefore have Internet access) to use the free e-mail. The exception is Juno. Juno provides complete Internet e-mail access (from the software to the dial-in connections) to anyone with a Windows-

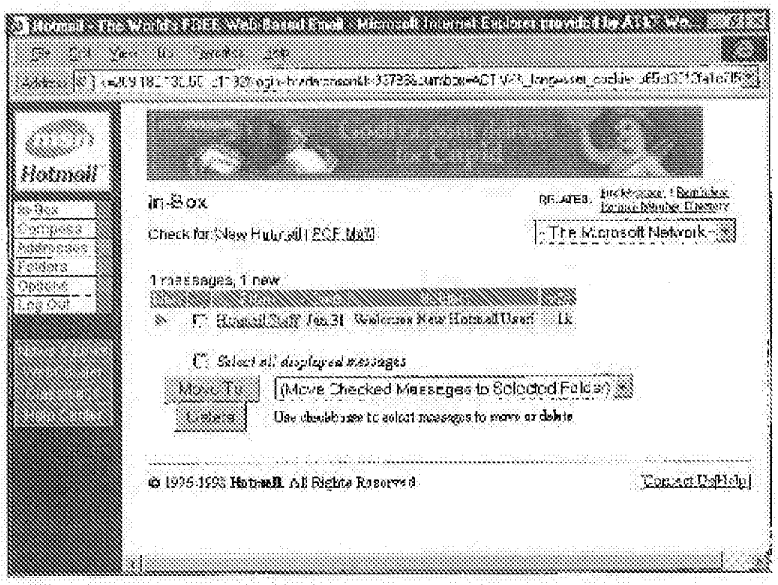


Figure 2.9 Hotmail interface and advertisements.

compatible PC and a modem. No other software or Internet access is needed. To set up a Juno account, all a user must do is complete a demographic survey. And, of course, advertisers can target based on the answers. When the user sends and receives e-mail, Juno advertisements rotate on a timely basis (see Figure 2.10).

Advertising via the Web

The bulk of online advertising is occurring on the Web. The Web, with its audio, video, and interactive capabilities, provides tremendous opportunity for effective and creative ad development.

Banners

As we mentioned in Chapter 1, Web advertising began with banners, those rectangular graphics located on Web pages (Figure 2.11 displays a banner ad on PriceSCAN). Banners come in a range of sizes, but are usually around 7 inches wide by 1 inch deep or 468 × 60 pixels (see Figure 2.12 for a graphical representation of the standard banner ad units). Although advertisers and publishers have been saying that advertising must go "beyond the banner," this format currently receives the vast majority of Web ad spending and will continue to be a Web advertising staple for the foreseeable future.

The industry has now come to realize that we need banners; the key is to learn how to build better banners. There are three primary categories of banners: static, animated, and interactive.

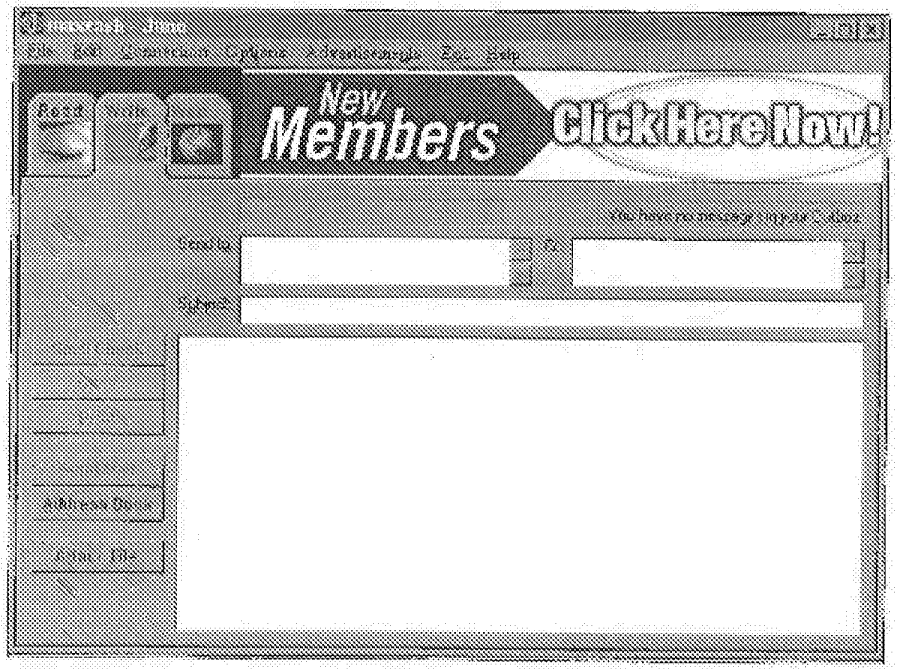


Figure 2.10 Juno interface and advertisement.



Figure 2.11 PriceSCAN displaying a banner advertisement.

Static banners. Static banners are fixed images on a site and were the first type of banner used in the early years of Web advertising. The advantages of a static banner are its ease of production and its universal acceptance by all sites. Its drawback is that with all the new innovations in banner ad technology, a static banner looks stale and boring. Additionally, static banners generate less user response than animated and interactive banners.

Animated banners. These are banners that move or spin or have some form of action. They most often use GIF89 animation, which operates like the old flip books full of consecutive images. Most animated banners have from 2 to 20 frames. Animated banners are extremely popular, and with good reason. Animated banners pull higher click-through rates than static banners. With multiple frames, an animated banner can deliver more information and graphical impact than a static banner. Moreover, this type of banner is relatively inexpensive to produce and small in size, usually under 15 kilobytes.

Interactive banners. When the call rang out for better banners, the interactive banner emerged. Interactive banners serve to engage the user in some way, either by playing a game, inserting information, answering a question, pulling down a menu, filling out a form, or making a purchase. These are ads that require direct interaction, which is more involving than a mere "click." Interactive banners are divided into two types: HTML and Rich Media.

HTML banners. HTML banners are a low tech option that allows the consumer to enter data into the banner or make a choice from a pull-down or radio button menu

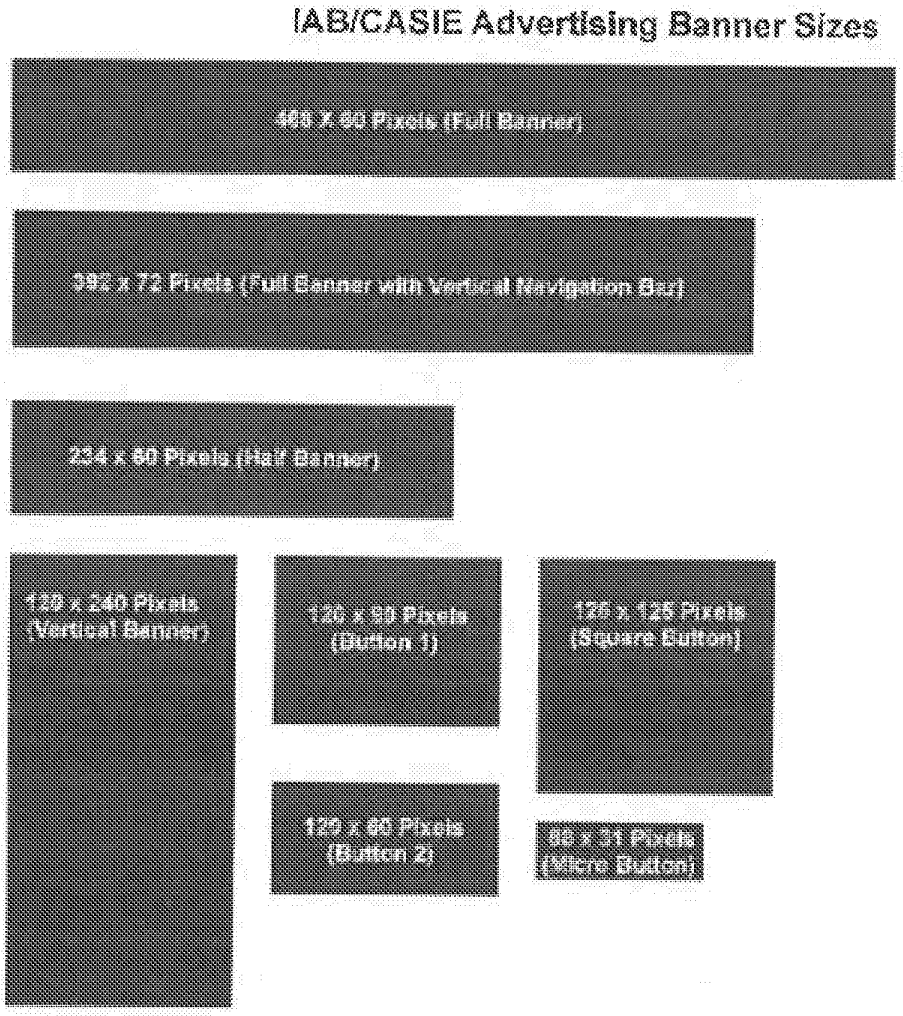


Figure 2.12 Graphical representation of the standard banner ad units.

within a banner. In our experience, HTML banners have had much higher response than animated banners. HTML banners allow a user to get involved: to choose which section of the site to visit, or to ask a question in the form, or to participate in a game (see Figures 2.13 and 2.14 for examples of HTML banners). This type of banner can be seen by a user with a low speed connection and an early version of a browser. The banner requires a few additional steps on the site's part to serve it up, but is quickly receiving broad acceptance.

Rich media. This is the incorporation of advanced technology into the banner ad itself.

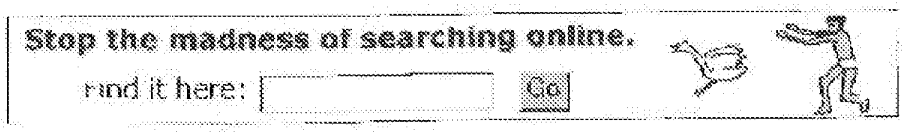


Figure 2.13a This is an HTML banner for the Internet portal site 4anything.com. The banner's text plays on the frustration of searching online and the graphic ties into 4anything's television advertising, which compares searching on most portals to a wild goose chase. In this advertisement, the user can type a search query into the banner, instead of having to click on the banner and start searching after arriving at the site. See Figure 2.13b to see what happens after a search query is entered.

Rich Media Banners

When we wrote the first edition of this book, push (discussed later in this chapter) was the hot topic that everyone was discussing. As of the writing of this book, rich media is the current craze. Tom Hespos, media manager of K2 Design, defines rich media as "any new technology that allows for either a more detailed advertising message to be presented, or for a higher level of interactivity than you might get in a standard GIF banner." Rich media is considered higher bandwidth advertising that delivers more of a brand impact than a GIF-animated banner advertisement.

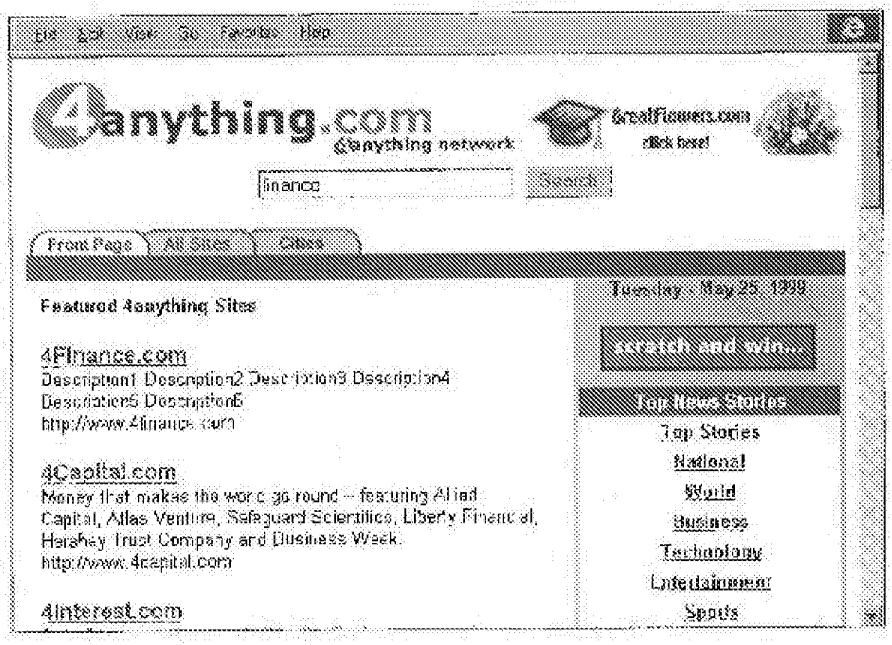


Figure 2.13b After typing in the query, the user goes to a search results page. This example is the search results page for "finance."

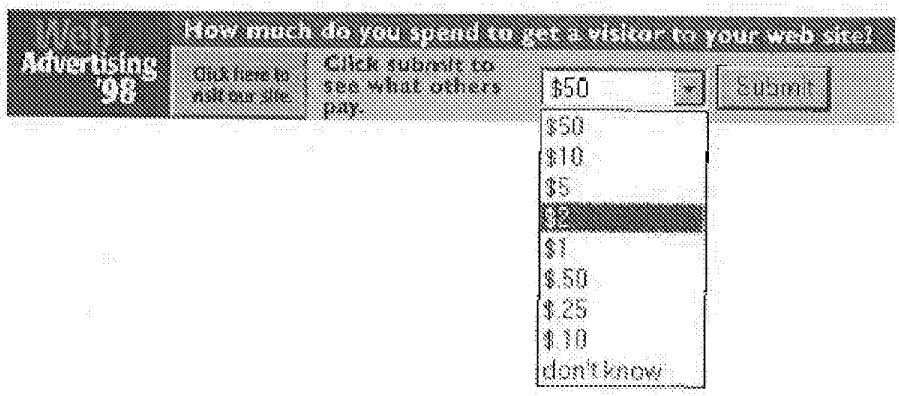


Figure 2.14a This banner asks the question that is on the minds of many media buyers: "How much do you pay to get a visitor to your Web site? After the media buyer chooses an answer from a pull-down menu, she goes to a page that is a real-time tabulation of all the answers so far. Response rates can be high in a banner like this, because it allows for user interactivity.

Rich media banners allow a consumer to complete a transaction within the advertisement without ever leaving the publisher's Web site. A consumer usually visits a Web site to read the content, which makes it difficult to lure the consumer away—even with an enticing banner. However, if an advertiser can make a great offer in the banner, and

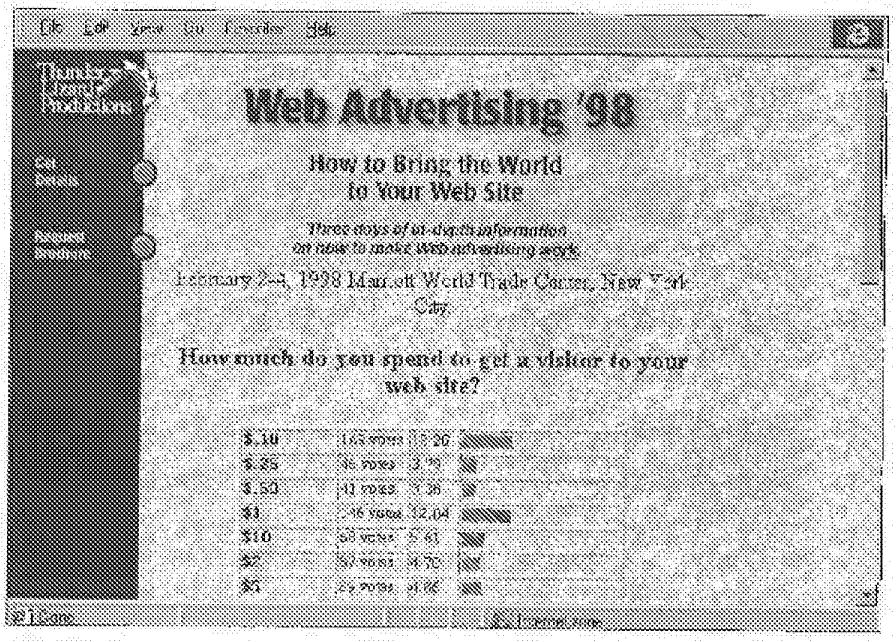


Figure 2.14b This is the results page that a media buyer receives.

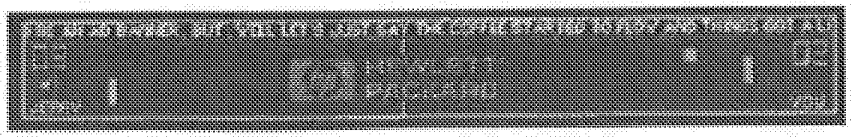


Figure 2.15 The Hewlett Packard Java pong game ad.

the consumer can complete a transaction or receive the entire advertising message through the banner without leaving the Web site, it seems logical that the consumer will be more responsive. One of the early rich media banners used java programing to bring its banners to life and was developed for Hewlett Packard. The goal of the ad was to deliver the message: "HP is a high tech company that is on the cutting edge." To convey this message, HP didn't need to drive traffic to its Web site. HP created a banner advertisement that was actually the game Pong (see Figure 2.15). Consumers could play the game within the banner, which contained the HP logo. To our knowledge, HP's ad was the first banner that contained a game. The game was incredibly popular, and HP's target audience spent a lot of time playing this game and thinking about how "cool" HP was.

Java ad banners can also allow consumers to order a product from within the banner. This feature is ideal for taking advantage of consumers' impulse-buying habits by facilitating the entire buying process from the banner ad itself. Figure 2.16 is the Java banner for 1-800-FLOWERS that allows customers to order Mother's Day flowers directly through the banner ad.

Rich media as a general term is used to denote ads that make use of advanced technology. Let's look at some of the technology used to enhance banner ads.

InterVU

Using InterVU (www.intervu.com) technology, a banner can display a short video clip from within a window in the banner. InterVU calls this its V-Banner. The looping video clips are usually 3 to 5 seconds. Despite the quality of the video, advertisers have found that InterVU banners can increase response rates; probably, this is because a video banner has the potential to stand out more than an animated banner. InterVU hosts these high bandwidth ads (and content) for Web publishers so that the publishers don't have slowdowns on their sites when a lot of people request high-bandwidth, rich media ads. Figure 2.17 shows an InterVU V-Banner for Cobra Golf. The right side of that banner has a video clip of a golfer swinging.

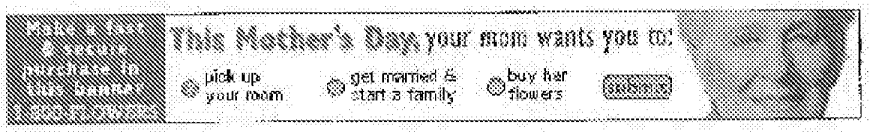


Figure 2.16 1-800-FLOWERS ad that allows consumers to order Mother's Day flowers through the banner.

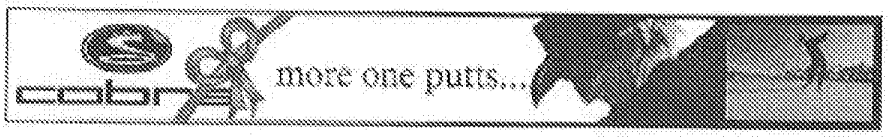


Figure 2.17 InterVU V-Banner ad for Cobra Golf.

Another good example of an InterVU banner is the one that was developed for the movie *Air Force One*. The looping video on the right side of the banner catches the viewer's attention. The video certainly stands out more than any of the animated GIFs that might be on the page competing for the user's attention. It is important to note that InterVU banners are GIF banners that include a video portion. The video takes more time than a GIF to load. This means that even when users leave the site before the video in the ad has loaded, they will still see the GIF ad.

RealAudio and RealVideo (www.real.com)

RealAudio and RealVideo (www.real.com) technology allow vibrant audio and video to be incorporated into the banner. Advertisers can use a variety of combinations of audio and video in these ads. For example, an advertiser can have the banner play audio when the a is downloading or after it is clicked. Figure 2.18 shows a RealAudio/Video ad developed by Modem Media.Poppe Tyson for Reebok. In the middle there is video of Allen Iverson playing basketball (something sure to attract the attention of surfers of NBA.com, which is where this ad ran). There is also audio of him talking about basketball with a commentator. The combination makes it enticing for a user to click through on the ad.

Narrative Communications' Enliven Technology

Narrative Communications' Enliven technology (www.enliven.com) delivers a small Java applet allowing for transactions to be completed within a banner.

Figure 2.19 is a Hewlett Packard Enliven ad. This ad allows users to print out spec sheets for HP's LaserJet 3100. Figure 2.20a-b shows an Enliven ad used by Eddie Bauer. This ad allows a consumer to order from within the banner; the order is tied directly into Eddie Bauer's system.

One of the benefits of an Enliven ad is that surfers can complete a transaction (requesting something or even making a purchase) without leaving the content site, encouraging the consumer to buy on impulse. However, these banners can take a long time to load. Some publishers have even complained that ads like Enliven slow down their site. In any event, as the technology to make the downloads faster improves, this will be a great means of making banners more interactive.

Macromedia's Flash

Macromedia's Flash (www.macromedia.com/software/flash) is a series of audio-visual players that allow a user to view streaming content, which means the user can

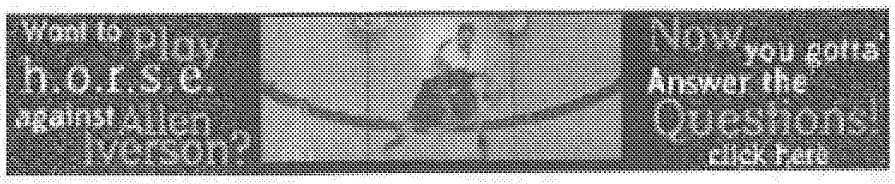


Figure 2.18 Reebok's RealAudio/Video banner ad.

see the ad as it is downloading instead of having to wait for the entire ad to download before seeing anything. Because the content is streamed, it doesn't take long to download the ad, so a banner ad that utilizes Flash will have more animation than an animated GIF that is the same file size. However, a user must have the plug-in to view this ad. As with other rich media ads, the publisher can serve the ad in such a way that only surfers who have the plug-in will receive the rich advertisement, while other folks will receive an advertisement that they can use with their plug-ins and browser.

Thinking Media's ActiveAds

Thinking Media's ActiveAds (www.thethinkingmedia.com/activeads/activeads.html) allow advertisers to create banner ads with which users can have a high level of interactivity and complete transactions. ActiveAds banners can include features such as secure, in-banner credit card transactions, subscription forms, games, animation, and streaming audio—delivered in a basic "player" file of 5 KB. Thinking Media's Java-based ads can be created with a software developer kit (Sonata). Sonata consists of a number of reusable Java components for such features as text animation, image animation, streaming audio, real-time polls, real-time auctions, subscription forms, frequent flyer miles, SSL credit card transactions, decision trees in-banner, clickable news tickers, games, quizzes, and timers. Thinking Media also provides a tracking feature, ActiveTrack, which reports what happens *inside* the banner: how long a banner was visible to the user on the screen, how much time surfers actually spent on average interacting with it, as well as information specific to tracking an individual banner.

Figure 2.21 is an example of an ActiveAds banner for HomePC. The first frame is an offer for a free issue of the magazine. A user is instructed to drag the sticker to the other side of the banner. This makes the offer pop onto the screen, and then an order form appears. The user can order right from within the banner.

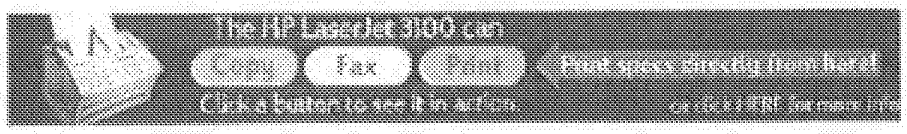


Figure 2.19 HP's Enliven banner, which allows people to print a spec sheet from the banner.



Figure 2.20a Eddie Bauer's Enliven banner presents the consumer with an offer to purchase pants within the banner.

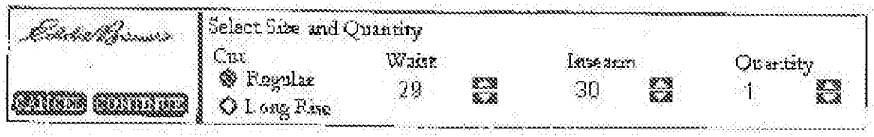


Figure 2.20b After a consumer clicks the banner, he can answer various questions within the banner to complete the purchase.

Figure 2.22 is an example of an ActiveAds banner from CDnow. This ad offers users a chance to win a $10 discount. As you can see from the pictures, users can choose a quiz category and answer the question in the banner. If they get the question right, the banner congratulates them and offers them a discount. If they get it wrong, they are encouraged to "try again." This is a fun banner that encourages people to interact with the advertisement. We think it's great for promoting a brand or service. However, e-commerce advertisers with the goal of selling products should test this carefully. The ad, while it can certainly be great for a branding experience, can create the appearance that there are additional barriers to overcome before simply purchasing a product; users may fear that they have to participate in a game or some other process, when they really want to take advantage of a discount and continue surfing.

Both Sides of Rich Media

Many publishers and advertisers see rich media as the way to get more advertising dollars committed to the Web. These marketers have primarily expressed an interest in using rich media to improve Internet advertising's branding efforts. A banner is usually only the size of six postage stamps, and the file size is usually only around 10 KB. This doesn't leave much space for a message, or much KB size to use video or audio to better convey a message. Marketers often say that TV is much more powerful, because the full motion video and audio allow for more emotional advertisements. Marketers are

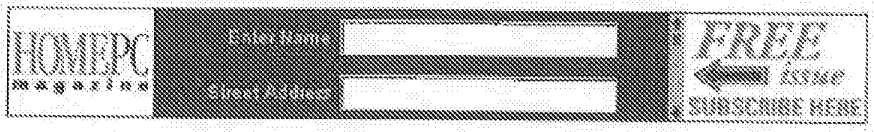


Figure 2.21 ActiveAds banner for HomePC.

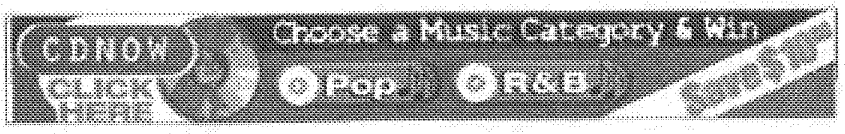


Figure 2.22a ActiveAds banner for CDnow (introductory animation).



Figure 2.22b ActiveAds banner for CDnow (quiz question).

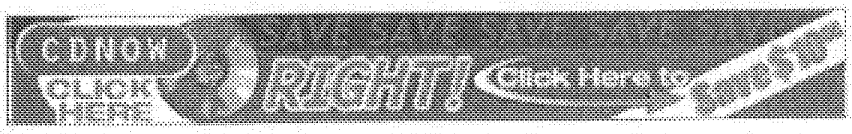


Figure 2.22c ActiveAds banner for CDnow (discount offer).

hoping that rich media will provide a step towards better branding opportunities through Internet advertising. Rich media's market penetration is growing, as presented in Table 2.2.

Rich media is not without its critics. There are many who feel that since most users still connect to the Internet through low bandwidth, the push toward rich media is premature. AOL, which connects more people to the Internet than any service in the world, refuses to use rich media for this very reason. Some interesting arguments for

Table 2.2 Rich Media Penetration

PLAYER/FILE TYPE	PENETRATION	SOURCE
Java (Director, Flash, Emblaze, Enliven, AdController, Microsites, etc.)	76%	InterVU
QuickTime	86%	Apple Computer
MPEG	80%	Microsoft
AVI	78%	Microsoft
RealPlayer	65%	RealNetworks
Shockwave	20%	Macromedia
Flash	35%	Macromedia
Cosmo Player	20%	Cosmo Software

Source: Rich Media Handbook produced by InterVU, 1998

Beware the Cost of Implementation

High technology banners, which use rich media, can be great for increasing response. However, a lot of time must go into technology-enhanced banners, and we're not just talking about the creation of the advertisements. Getting Web sites to display your advertisements correctly and getting their ad serving software to work with these ads is very time consuming. In some cases we've found that the cost of the additional time actually surpasses the value in the increased responsiveness from the ad.

Another note of warning is that many rich media ads require a user to have a particular plug-in. In those cases you want to make sure an alternative ad is delivered to the users who don't have the plug-in. Otherwise a good number of people may not see anything. Also be sure to check the download speeds of technology-enhanced banners. They're often much slower than is desirable. Keep in mind that few people will wait for an ad to download. To determine the percentage of sites that will take your rich media advertisements, look at Figure 2.23. This chart was created by AdKnowledge from their database of Web sites that accept advertising.

Ad Technology	Accepting Sites
• Active Ads	52
• Enliven	215
• HTML	606
• Java	535
• JavaScript	275
• MicroSites	66
• RealAudio	178
• Shockwave	304
• V-Banner	139
• VRML	87

Data from January 1999.

Figure 2.23 The ChannelSeven.com rich media Report Card provides a summary of Web sites and the rich media formats they accept. This information is provided by AdKnowledge, a directory of Web sites that accept advertising. The data is self-reported by the sites.

Building Better Banners

Advertisers can precisely measure the response to their banner advertisements via the aforementioned click-through determinations. (Chapter 4 has more on Web measurement and tracking.) This tracking capability has led advertisers to continually test new banners to see if they will achieve higher response rates. Here are some hints for creating a responsive banner ad:

Pay attention to placement. Some sites will outperform others for an advertiser. The more focused a site is on the target audience, the better the response. In addition to finding a focused site, advertisers should pay particular attention to finding the best pages within a site. Performance on two pages of the same site can differ enormously. As Kevin O'Connor, CEO of DoubleClick, an advertising network, points out: "The best-performing page on a poor-performing site will outperform the worst page on a good-performing site."

Chart frequency—how often a particular user sees a particular ad. According to research performed by DoubleClick and I/Pro, users are most likely to respond the first one to two times they see a banner ad. Response drops significantly after the second time.

Call to action. Why should users click on the banner? Do they win a prize? Receive a discount? Without a reason to click, there is no incentive for a user to leave a publisher's Web site. The incentive should be relevant to the site on which you're advertising. Obviously, a sports offer will work better on ESPN's Web site than on the Epicurious food Web site. "Free" works well, but you won't always get the most qualified audience.

Add "click here" to your banners. Banners with "click here" consistently pull better than those without. There will always be new users of the Web, many of whom won't know where to click. This may also explain why some advertisements have a better response rate when there is hyperlinked text underneath the ad. Hypertext links signal to even the newest users that something is clickable.

Create a sense of urgency. Give users a reason to "click now."

Do not use too much text. Encourage users to look at the banner with a simple offer and design.

Incorporate animation. The banner is a small portion of the Web page, and there are a lot of other graphics and text competing for a user's attention. Animation makes a banner stand out.

Try HTML and rich media banners. Rich media banners tend to have higher click-through rates. So, if the sites on which you advertise accept rich media, give it a try.

Make good use of colors. Use colors that stand out. These will differ depending on the colors used by the site on which you are advertising.

Rotate banners. Using multiple banners throughout a campaign almost always increases response.

and against rich media are presented in the editorials at the end of this chapter by leaders in the online ad industry: Ali Partovi, product manager and original partner of LinkExchange, and Richard Hopple, chairman and CEO of Unicast Communications.

A nice resource for rich media information is ChannelSeven's TurboAds site (www.turboads.com/index.shtml). This site, which is built for the advertising community, has numerous case studies and information about rich media technologies.

The Branding Power of Banners

The role of banner advertisements is changing. Banner ads used to be used for the sole purpose of getting a Web surfer to click and go to an advertiser's Web site. Consumer responsiveness to banners has been dropping. When we wrote the first edition of this book, the average click-through rate (the number of times someone clicks a banner ad divided by the number of times the ad is seen) was 2 percent. According to NetRatings (www.netratings.com), the click-through rate has been steadily dropping, and, as of this writing, it is below 1 percent. To maintain the effectiveness of their advertising, advertisers and agencies must either negotiate lower rates or find Web advertisements that will supplement banners to increase the overall response to a campaign. Advertisers have realized that they need to get a benefit beyond clicks from their banner ads.

Advertisers are now looking to derive a branding value from their ads. For example, most advertisers do (and all advertisers should) keep at least a corner of their banners for a brand image. This allows advertisers to deliver some sort of a message to those people who don't click! Advertisers have also begun to purchase more sponsorships and advertorials, which allows them to integrate their message or call to action within the content, and, in some cases, to deliver their message without forcing the user to leave the site. Studies, such as those conducted by Millward Brown Interactive (www.mbinteractive) and the Internet Advertising Bureau (www.mbinteractive.com/site/iab/study.html) are showing the strong branding ability of banners. And according to a 1999 study by Ipsos-ASI and AOL, consumers are as likely to remember an online banner ad as a television commercial.

Testing

In the early days of banner development (1994 to 1995), very little was known about the dynamics of banner ads on the Web. The same banner would be sent to every site, though admittedly the list of sites was short. Today, the rigorous testing of banners by advertisers and publishers has helped maximize response rates. Rich Boyce, VP of advertising at HotWired, advises, "Be prepared to test, learn, and evolve. Make adjustments as needed. And remember that a Web advertisement is always a work in progress. It's a continual process of refinement. Unlike a television commercial that once produced is done, with Web advertising, the creative is never done."

One of the benefits of advertising on the Web is the ease, accuracy, and speed with which you can conduct a performance test. Ad management software facilitates real-

time reporting for banner testing. Most banner campaigns begin with several ads that are tested over a 24-hour period before they are refined and tested again. Consequently, the resulting banners used in the full campaign have been well researched for maximum click-through effectiveness. Conducting banner ad testing is relatively inexpensive—a banner ad test costs a few thousand dollars, whereas a test of a TV commercial costs tens of thousands of dollars.

Banners only give advertisers limited space with which to promote their brands or sell their products, so it is essential to make the most of that real estate. No doubt we will continue to see major changes in banner creatives as advertisers strive to effectively use a banner's space. Advertisers will also begin to use other forms of advertising, such as sponsorship (which we discuss later in this chapter), that give them the opportunity to use more online real estate to promote their brands.

Buttons

Buttons are small banner-type advertisements that can be placed anywhere on a page and are linked to the button sponsor. Buttons were embraced by consumers and advertisers throughout the Web faster than banner advertisements—probably because they always led to free downloadable software.

One of the first of these campaigns was the "Download Netscape Now" button. Scott Heiferman, CEO of i-traffic (www.i-traffic.com), helped get that campaign started, and brought a grassroots enthusiasm to the program. The i-traffic team surfed the Web for sites that posted text stating that they were "optimized for Netscape." Hundreds of sites with this listing were contacted by i-traffic and offered a T-shirt if they switched to the button. Hundreds of T-shirts later, the Netscape Now button is a standard on Web sites. If imitation is the greatest form of flattery, button campaigns are flattering indeed. Proof positive is that the Microsoft Internet Explorer launch had a strong "download button" component.

The benefit of buttons, especially download buttons, is that they are simple in what they offer. A person clicks on a download button because he or she wants to download the software immediately. And, since the Web community is a strong market for software, buttons have resulted in millions of software downloads.

Buttons can also help advertisers build brand awareness. The ubiquitous Download Netscape Now button not only led to large numbers of software downloads, but also helped familiarize Web users with Netscape and establish it as the first leader in browser market share.

Buttons are commonly used as a way to have a constant presence in specific Web content without paying the high fee of banner advertisements. Figure 2.24 shows a Lycos page with two buttons: one for CDnow and one for barnesandnoble.com. These ad placements work well for the publisher and the advertiser. The publisher can offer advertisers button ads on every page of the site. This builds a nice brand awareness for the advertiser, and with the correct offer, these buttons can generate high response



Figure 2.24 Button ads for CDnow and barnesandnoble.com go down the right column of the page.

rates. The publisher can have multiple buttons on a page (since they don't take up a lot of space) and still have 100 percent of the banner inventory left to sell. To increase the response rate of button advertisements, advertisers often include text specific to the content alongside the button. For example, Figure 2.25 shows a button ad for Amazon on Yahoo! Above the button ad there is always text related to the search performed by the user. In this case, the Yahoo! search was for "goats" and the text above the button reads, "Buy Books on Goats." This type of customization based on the user's search query will significantly increase response rates.

Some button advertisements are large enough that they are the most noticeable advertisement on the page. Figure 2.26 shows a page from PriceSCAN (a price search engine). The button advertisement stands out more than any other element on the page. Not surprisingly, PriceSCAN said that their button advertisements are in higher demand than any of their other ad units. In general button advertisements are good for publishers because it gives them additional inventory to sell—another ad unit on each page. Of course, this has to be balanced with issues of page clutter and number of ads, both of which can reduce advertising effectiveness.

Text Links

Text links are some of the least intrusive, yet most effective advertisements. In the online advertising industry most of the commotion is related to finding new ways to

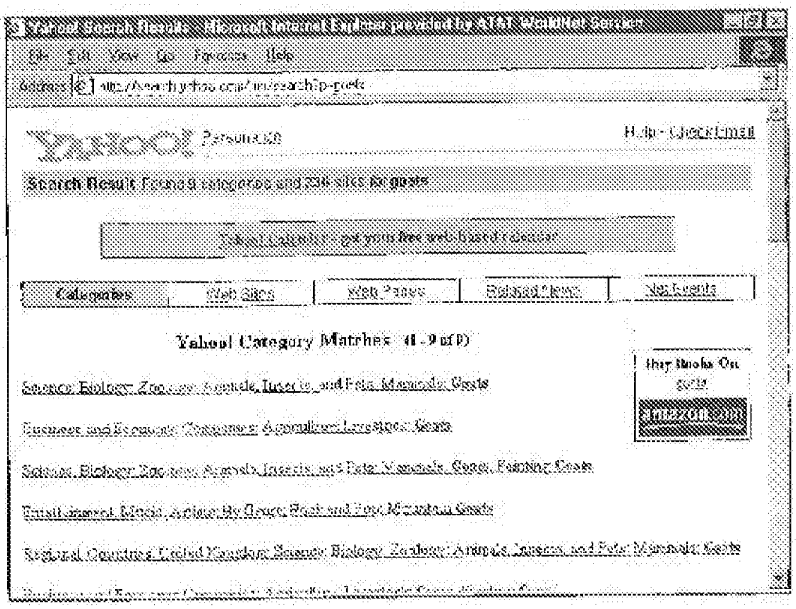


Figure 2.25 Button ad on the right for Amazon.com with customized text on Yahoo!

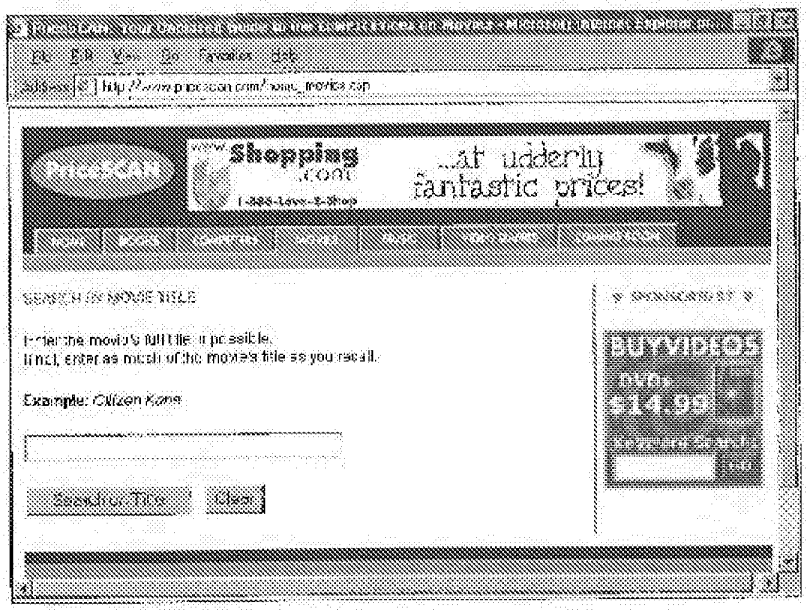


Figure 2.26 PriceSCAN identifies its button ads as sponsors.

advertise at a higher bandwidth. However, sometimes the least bandwidth-intensive and simplest of advertisements can be most effective.

One example of a text link is Figure 2.27. Here you see an AltaVista search that was performed for the keyword "superman." To the right of the search results is a text advertisement paid for by the bookseller Amazon.com. The advertisement tells users to go to "Books about superman." Amazon has the text link change to match whatever search the user performs. The power of these text ads is that they can be tailored to any search a consumer performs on AltaVista; this would be impossible through a banner campaign. As we mentioned in the previous section, a button ad along with a text ad can also be very effective.

Another text link example is from the search engine Lycos. Lycos sells a text link that will appear as the first search result when someone performs a search. In Figure 2.28 we have performed a search for allergy medicine "Claritin." The text ad that pops up says, "Bullseye! The Claritin relief zone." This is a paid advertisement from Claritin even though it looks very much like it's the first text listing that Lycos found, rather than an ad.

Many folks in the online ad industry believe that this is a compromise of editorial integrity. We believe it will become commonplace as Web sites search for more advertising revenue streams and advertisers search for more effective means of reaching consumers.

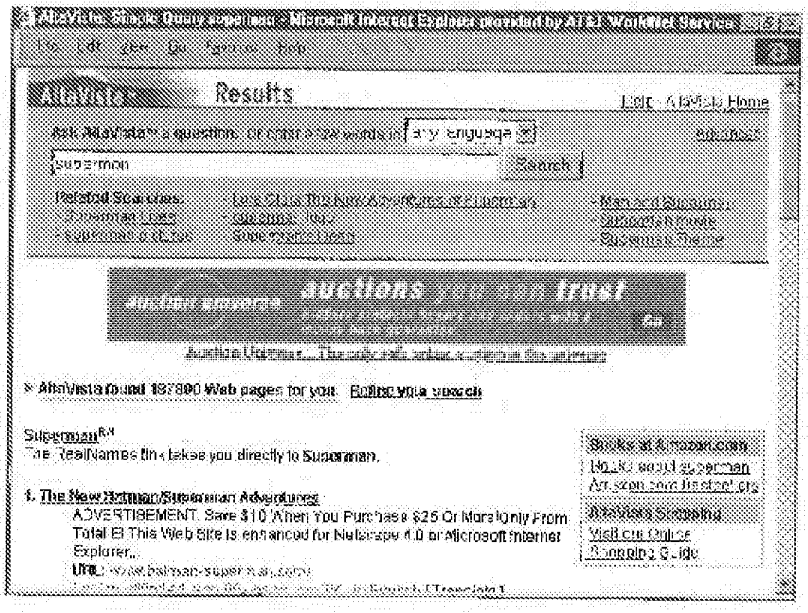


Figure 2.27 A text ad from Amazon.com on AltaVista.

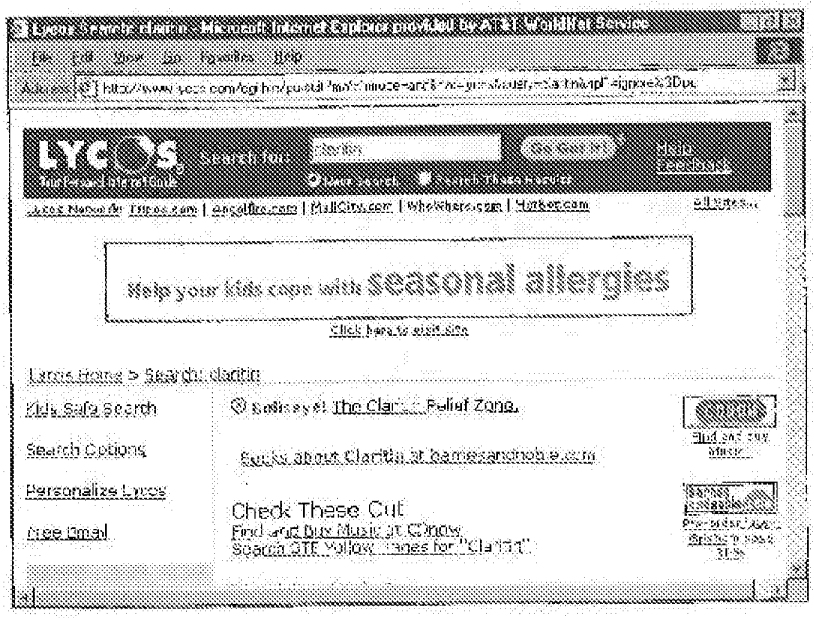


Figure 2.28 The "Bullseye" text ad on Lycos for "The Claritin Relief Zone" looks very much like editorial content, which is why it is so effective.

Some sites sell text listings as one of their main advertising staples. One example is the Year 2000 Web site (www.year2000.com), which is published by the Tenagra Corporation. This site is a comprehensive resource for people who have to deal with the computer problems that will be caused by the Y2K. If you look at Figure 2.29, you'll see the extensive amount of text link advertisements that they sell; all of the listings on the left side of the Web page are paid advertising. These text ads serve as a resource for Web site users; folks visiting the Web site want to find computer companies that address the Y2K problem. If this resource list had been in a graphical format rather than a text format, it would have been much more difficult to use. This is, in effect, a directory like Yahoo!, but all listings are purchased.

Sponsorships

Sponsoring a Web site can take many forms, most of which lead to placements that are beyond typical banner buys. Sponsorships allow an advertiser to have a successful ad campaign without necessarily having to drive traffic to its Web site. Consumers trust the brands that they visit repeatedly to get information. So, delivering your message within such trusted brands can often make that message more powerful. Indeed, sponsorships are on the rise.

An early, very effective sponsorship (an example from the first edition we couldn't possibly remove; it's one of our favorites) was Dockers' sponsorship on HotWired. HotWired was able to smoothly integrate Dockers' content into its Dream Jobs

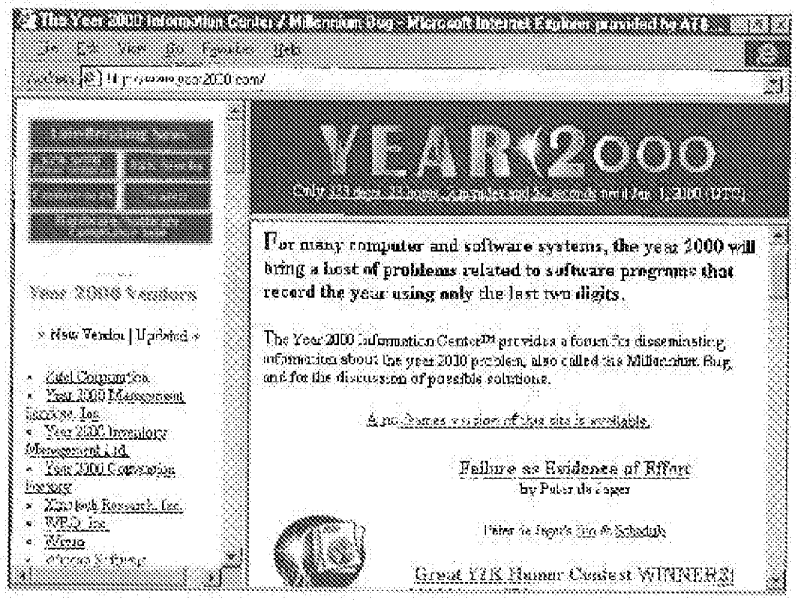


Figure 2.29 The Year 2000 Web site has text advertisements running down the left side of the page.

(www.dreamjobs.com) Web site. (Dream Jobs has since been fully integrated into the HotWired Web site.) Dream Jobs lists employment openings that are considered ideal for its hip, technologically savvy audience. Next to the job descriptions is a meter that rates the jobs based on the dress code. The meter includes the categories "cut off khakis," "stone-washed khakis," and "monkey suit." Dockers clearly conveys a message: Dockers are cool—at cool jobs you can wear them. Levi's doesn't need to drive anyone to the Dockers Web site to succeed with this campaign.

Many sponsorships relate to a specific aspect or feature of a site. Chapstick, a product associated with both cold and hot weather, sponsors the Intellicast Weather site (see Figure 2.30). Charles Schwab sponsors the personal stock tracker on Excite (see Figure 2.31). Excite lets users create their own home pages. Users can choose to have a number of stock quotes appear on their home pages. Those quotes will then appear as stock tickers incorporating advertisements for Schwab. This is a perfect placement; if a user sees that the stocks he is tracking are rapidly moving, he can click the Schwab link to open an account and start trading. Schwab is targeting folks at the point where they will be likely to switch from an imaginary portfolio to actual trading.

Advertorial

An advertorial is a sponsorship that looks more like an editorial than like an advertisement. In print publications, advertisements that resemble editorials have a clear

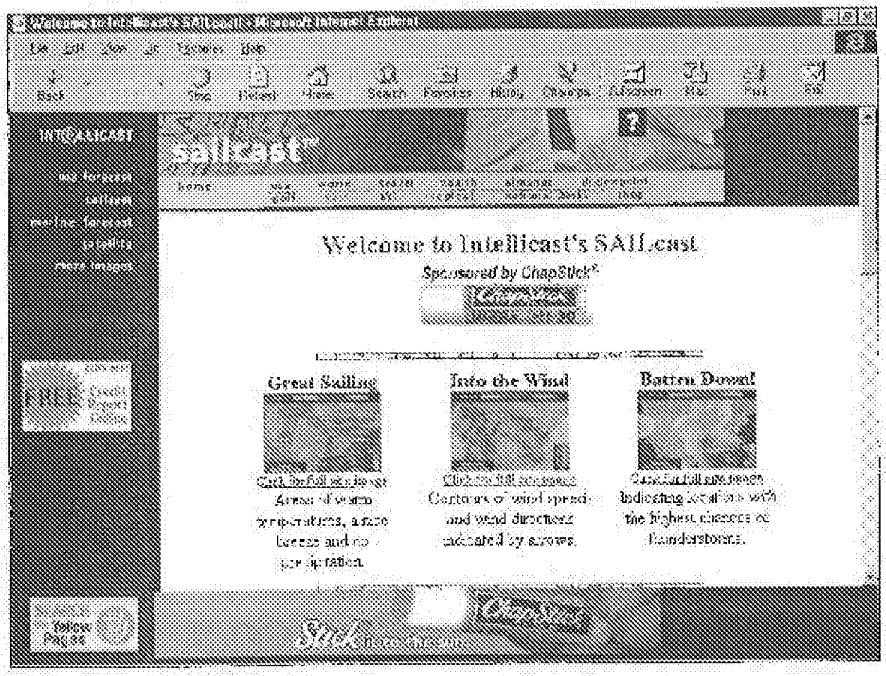


Figure 2.30 Chapstick sponsorship on the Intellicast Weather Site.

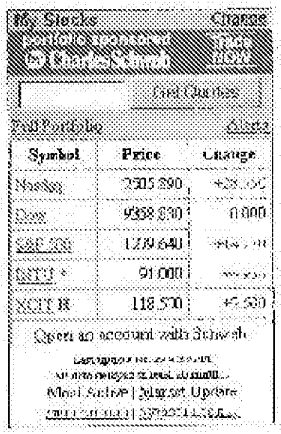


Figure 2.31 Charles Schwab sponsorship of the stock ticker on Excite.

label that usually reads something along the lines of "special advertising section." On the Web, there is usually no clear indication that advertorial areas aren't editorial.

Because consumers trust editorials, advertorials will often have a higher response rate than other Web advertisements. However, advertisers must be careful not to create a negative brand impression by causing consumers to feel deceived; they clicked on something that looked like editorial content, only to discover that it was actually an ad. (Publishers also need to be careful. You don't want to upset your greatest asset: the people who visit your Web site.) Advertorials will be very effective if they deliver the content a consumer expects to receive. On the other hand, if the consumer does not receive that content, the ads will backfire and probably cause a negative brand impression.

A great advertorial example is accessed when a user conducts an Infoseek search for the keyword "health"—and a link to go to the "health news and info" area on Infoseek appears (see Figure 2.32). This area has various text listings, including "Conditions A-Z" and "-in-depth reports" on specific problems. These all look like editorial links, but the links go to OnHealth (www.onhealth.com), a health site that advertises on Infoseek.

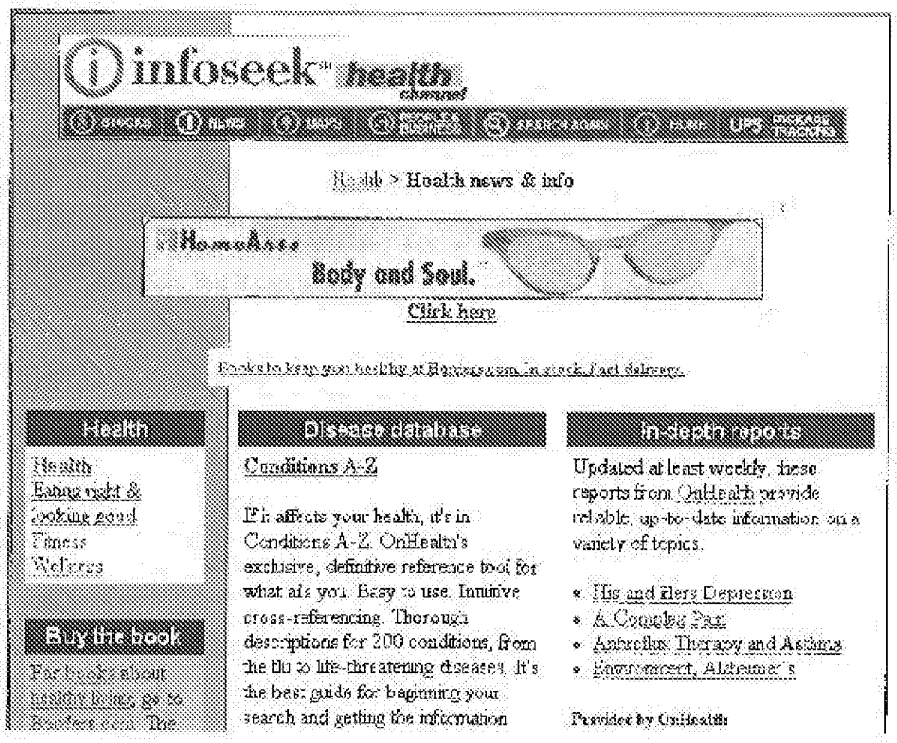


Figure 2.32 Results of a Search on Infoseek using the keyword "health."

Another example comes from Better Health (www.betterhealth.com, see Figures 2.33a and 2.33b). One editorial option on the Better Health site is "Skin care options for aging skin. . . ." Clicking on the link takes the user to a co-branded site that is actually sponsored by Renova, a product to help aging skin. The benefit to Renova is that it gets the credibility of teaming with the respected Better Health site, which already reaches its target market.

Push Technology

Advertisements delivered through *push* technology are sent (usually with content) directly to a user rather than waiting for the user to appear. Push technology includes e-mail, which was already discussed in a previous section of this chapter. The push technology explored here includes other systems that enable Web users to sign up to receive information broadcast to their computers.

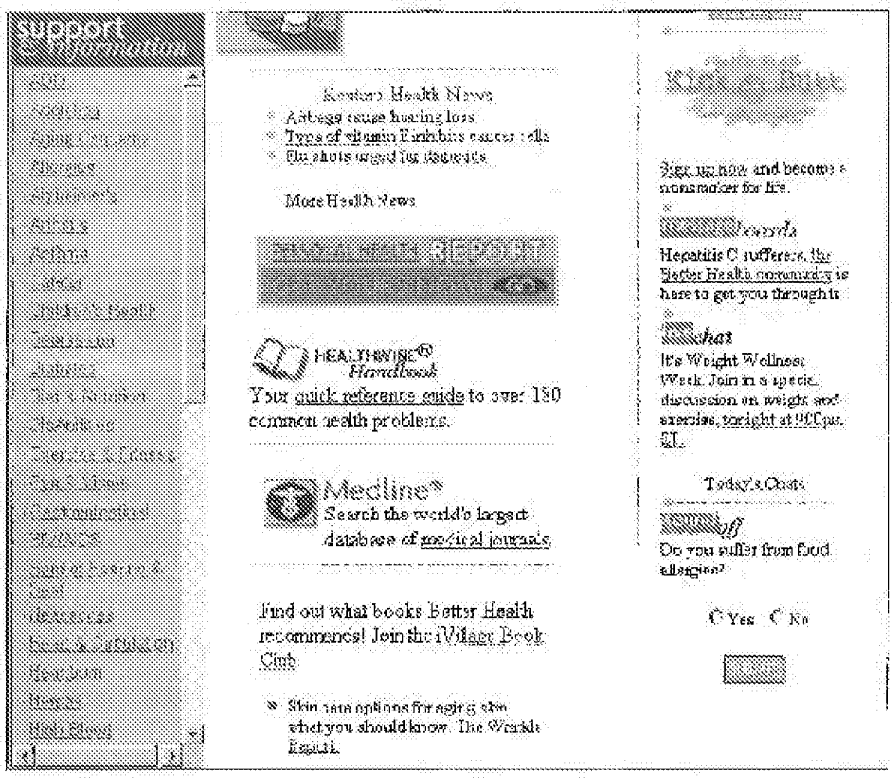


Figure 2.33a Renova text ad on the bottom of Better Health for "The Wrinkle Report."

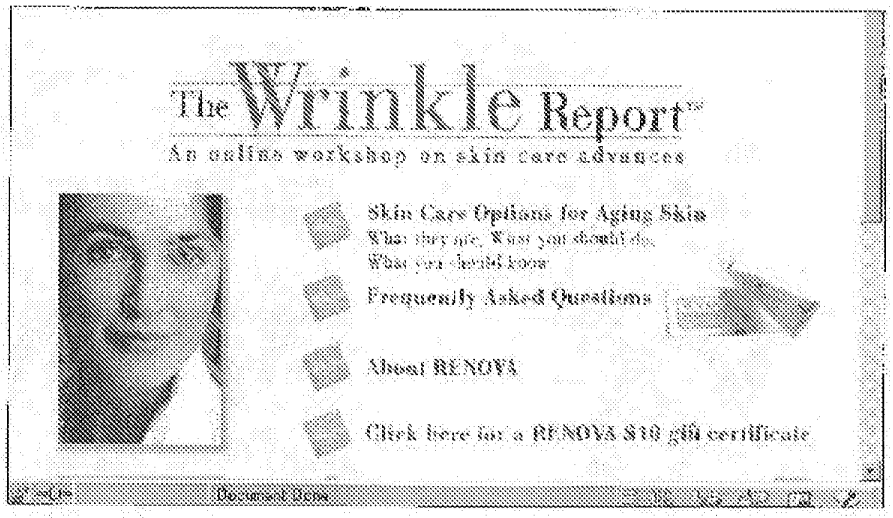


Figure 2.33b "The Wrinkle Report" is a Renova co-branded Web site on Better Health.

PointCast (www.pointcast.com) is one of the many players offering push technology. After downloading the PointCast client (software), users can request that PointCast send them news updates on specific topics. When they register, consumers choose what types of information are of interest to them (stock quotes, news headlines, etc.). Then, when the consumer's computer is idle, a PointCast screensaver appears, containing the information the consumer requested, along with advertisements. This system works best when the computer is always connected to the Internet, such as with a direct connection. But it can work on a dial-in basis.

The main problem some advertisers have with PointCast is that the ads only appear when a computer is idle, and an idle computer often means that the monitor is not being watched. On the other hand, the ad could be seen by many more people in the same office.

Interstitials

Interstitials are advertisements that pop onto the screen and interrupt users. They are sometimes called "pop-ups," "e-mercials, or "intermercials." This ad model is similar to television ads in that the ad interrupts programming. Interstitials come in different sizes (taking up all of the screen or only pieces) and with different levels in interactivity, from static to fully animated productions. Users can sometimes dismiss the interstitial by clicking out of it (which they can't do with television ads), but there is no warning as to when the ads will appear or reappear.

Advertisers like the prospects of using interstitials because these advertisements are sure to be noticed. Indeed, every AOL user can attest to the fact that one can't help but notice the static pop-up ads that greet the user upon opening AOL (see Figure 2.34).

Some interstitials give advertisers an opportunity to use animated or full-motion video (full-motion video is usually limited to interstitials delivered through push vehicles). Advertisers certainly have a better opportunity to get their brand message across when they are not competing with anything else on the screen and can use more images and screen real estate (see Figure 2.33 for an example of an interstitial advertisement).

The downside to interstitials is that publishers and advertisers need to be very cautious not to upset consumers. The Internet began as a medium for the free exchange of information, and at the beginning, there were no advertisements. There's a small minority of people who feel that the commercialization of the Web is bad and that all advertising is bad. Those aren't the people we worry about (unless they're your target audience). We worry about the average consumers who are accustomed to controlling their Web experience; they choose what sites to visit and what to click. These consumers are often upset and bothered that a site and advertiser will force them to see an interstitial. To avoid this, many sites are now using interstitials that are daughter windows (small windows that pop up over the browser and only take up 1/8 of the screen).

Here are some rules that will help ensure that your interstitials don't upset consumers:

Choose sites already using interstitials. Interstitials usually have the best response when they're placed on a site that has run this type of advertisement in the past; con-

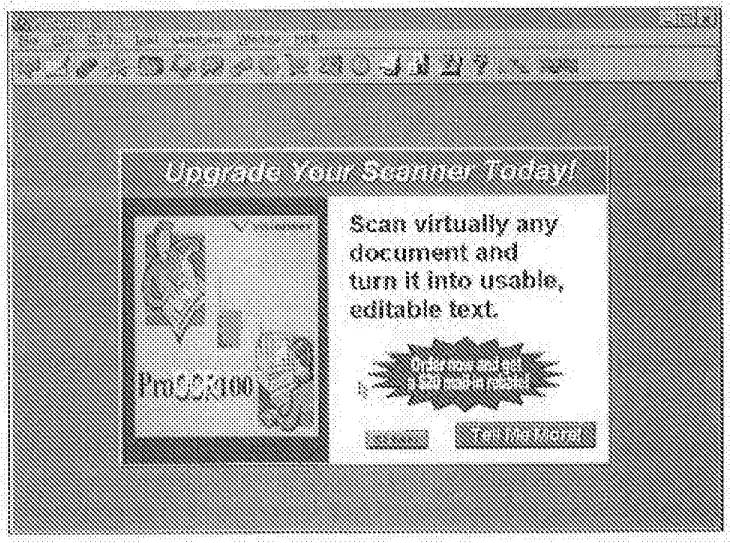


Figure 2.34 Pop-up ad on AOL.

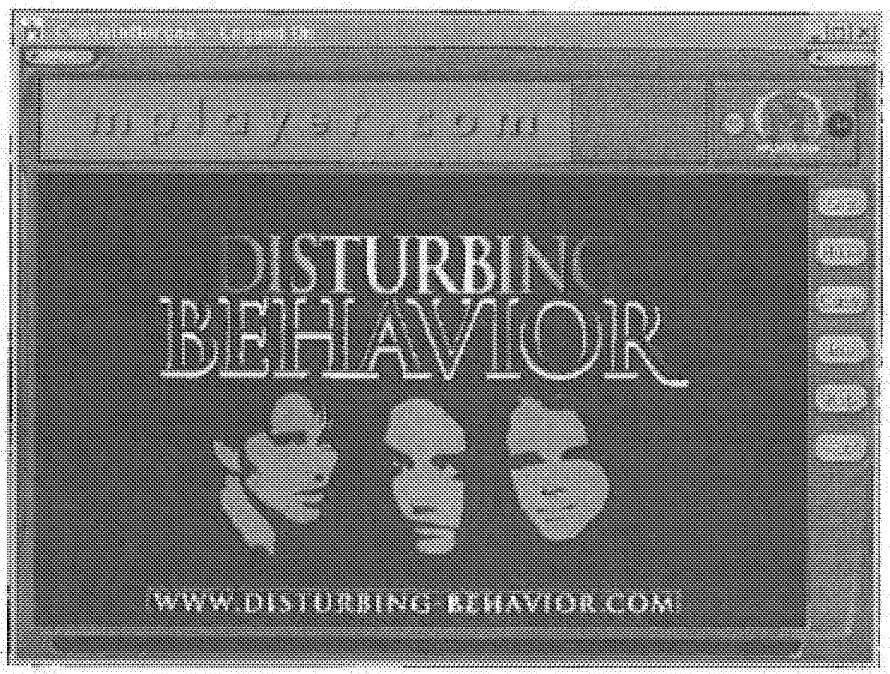


Figure 2.35 Interstitial for the MGM release of the movie *Disturbing Behavior* using MPath technology.

sumers are expecting them. For example, consumers know that Berkley Systems' "You Don't Know Jack" Web site has had interstitial ads since its inception, and interstitials are accepted there.

Try interstitials that are smaller than a full page. These are less obtrusive than a full-page interstitial and seem to be better received by consumers. These small pop-ups take only a quarter of the page and appear in a separate window on top of the browser.

Use interstitials when a user's screen would otherwise be idle. The interstitial shown in Figure 2.36 appears while software is downloading. This eliminated negative feelings about the ad, because it was not preempting any content; it appeared while their screen was idle.

Interstitials can be interactive. Unicast (www.unicast.com) is being used for some exciting interactive campaigns, from sample requests (see Figure 3.12) to other transactional pop-ups.

The interaction can be made into a game, as happens on BoxerJam, a game show site (www.boxerjam.com) where users actually compete against each other in a format

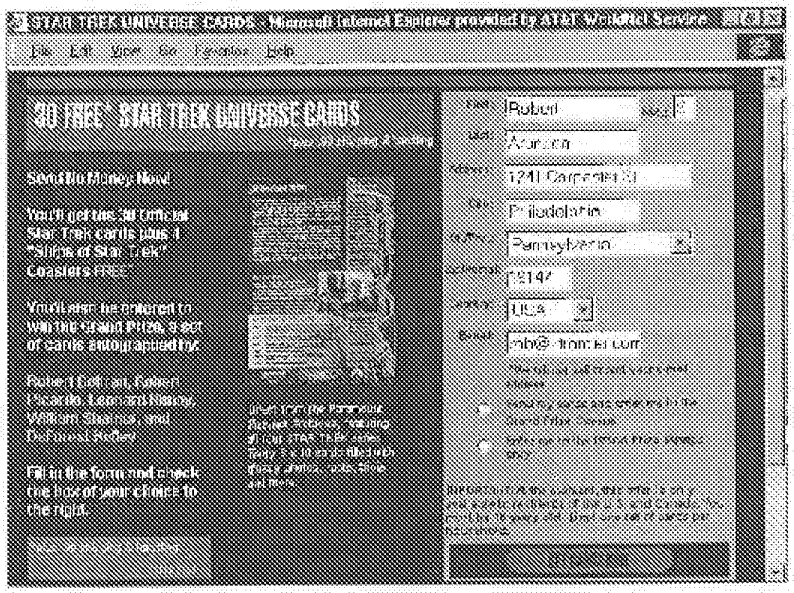


Figure 2.36 Star Trek Universe cards interstitial.

resembling a traditional TV game show. BoxerJam GameShow uses "intermercials" (another term used for interstitials) before the game and between rounds. The ads are used to build brand awareness, drive traffic to a site, and to offer the player an interactive experience through the ad, as demonstrated in Figure 2.37, an ad for Omaha Steaks on the BoxerJam site.

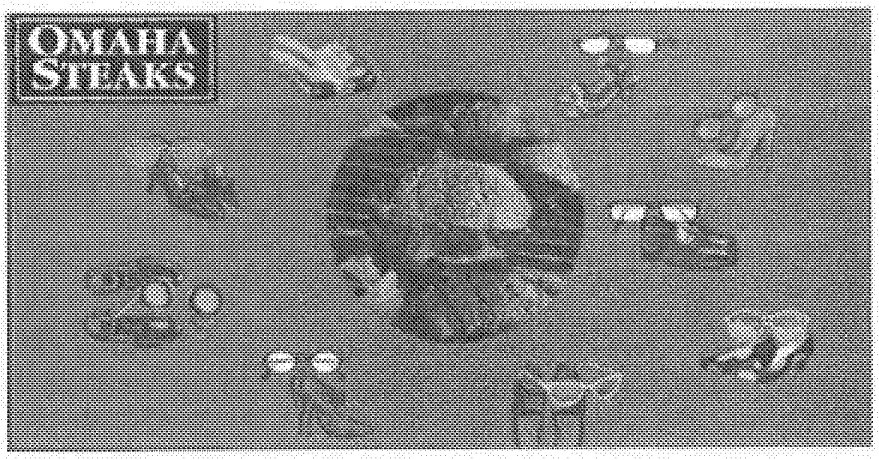


Figure 2.37 Omaha Steaks "intermercial" on the BoxerJam site.

A Few More Ad Models outside the Box

Not all ad models can fit neatly into the categories we've listed in this chapter, so we've written this section to cover those new ad models that that are outside the current category box.

New Screen Real Estate

The first ad model that breaks the mold is delivered through Alexa (www.alexa.com). Alexa is a Web navigation service that works with your browser and accompanies you as you surf, providing useful information about the sites you are viewing and suggesting related sites. Alexa isn't so much a new ad model as a redefining of available ad real estate. The Alexa toolbar travels with users as they surf the Internet. Alexa sells advertising on its toolbar, which is visible to users no matter what site they are visiting, giving advertisers a presence on any site on the Web (see Figure 2.38) In fact, using Alexa you can place your advertising "on" any site on the Web whether that site accepts advertising or not. "Alexa lets you do a competitive stealth play," explains Doug Hansen, Director of Advertising Sales with Alexa. "For example, if you're a major bookseller, you can advertise on your chief competitor's site by advertising on the toolbar

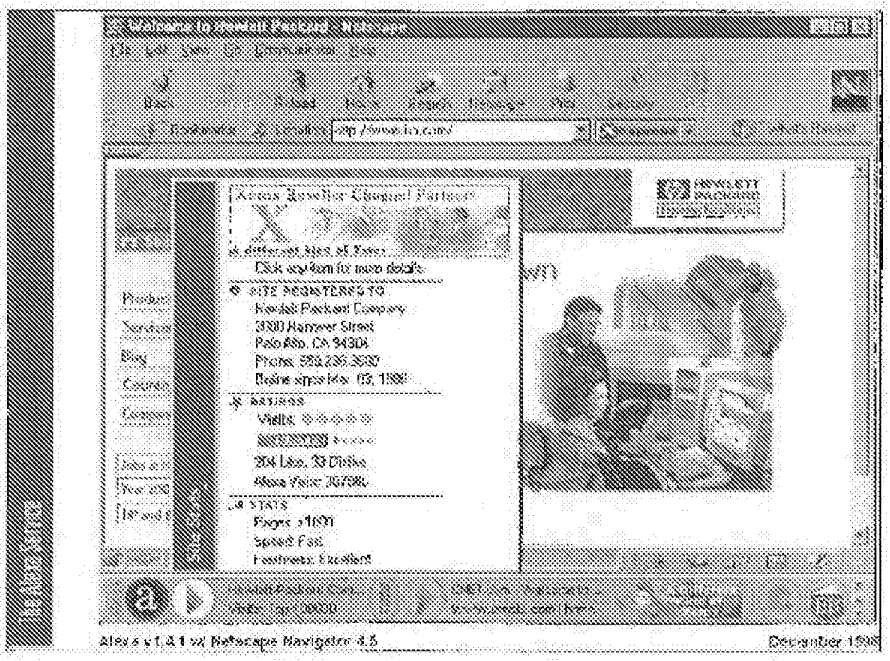


Figure 2.38 The Alexa toolbar lets you place advertising on any site on the Web.

just as the consumer is about to purchase a book. This advertising approach cannot be replicated in traditional site-based ad models." The only limit to the reach of Alexa is the number of people who use their service. As of the writing of this book, Alexa has said that over 1 million people have downloaded their software.

Another exploration of new real estate for advertising space comes from Aureate Media (www.aureatemedia.com), which puts banner ads on software programs themselves. Aureate also hosts the AdSoftware Network, which is composed of a group of software products that have the ability to display electronic advertising.

Screensavers

The entire screensaver itself can be an advertisement. SaveScreen.com (www.savescreen.com) allows consumers to build their own personalized screensavers containing integrated advertising.

Bookmarks and Toolbars

Bookmarks and toolbar buttons on the browser are now available for sale for advertising. Advertisers are paying for premium placement on a user's browser. Advertisers can also be found on browser toolbars. 100hot Surfboard (www.web21.com/surfboard/), a product of Web21 and AdNet Strategies, allows users to download a toolbar that they can add to their browser window (either Netscape or Explorer) which contains the advertiser's logo as a link (see Figure 2.39).

Cursors

Everything on the screen is fair game as an ad vehicle. Even the cursor can be turned into an ad, for a unique form of branding. Comet Systems (www.cometsystems.com) has pioneered a way to customize the mouse cursor, substituting any graphic or animation for the conventional arrow. In this way, you can put your product's image on a user's browser for continual branding. For example, a flower shop can have its cursor be a flower, and when the cursor moves over flower arrangements, it can change to display special discounts. Another example, provided by Comet, is a weather map where the cursor changes as it moves over the map to reflect the weather conditions in each city (see Figure 2.40). The custom cursors will only work for users who have downloaded the Comet Systems software. The download is an easy process, but because users must download this software, it will be a challenge for Comet Systems to get significant reach—unless their software comes bundled with future versions of Netscape and Explorer. Once the software is on your system, you can see the advertising cursor when you go to a site that uses the Comet cursor.

Undervalued Web Space

Nothing is wasted on the Web, not even a page that could be considered a dead end. Columbus Group (www.columbus-group.com), out of Vancouver, British Columbia,

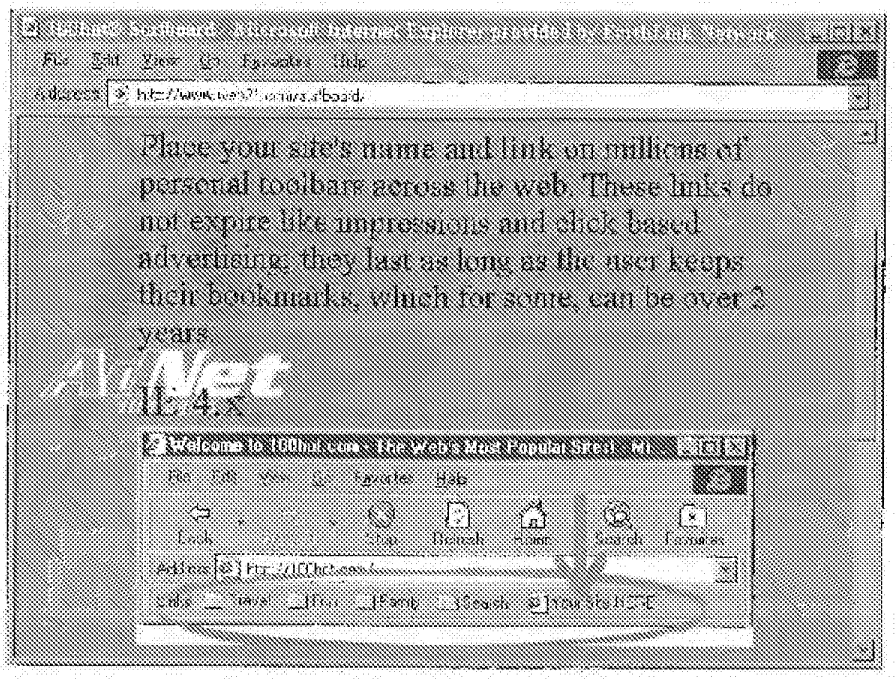


Figure 2.39 100hot SurfBoard lets an advertiser buy a link on a user's browser toolbar.

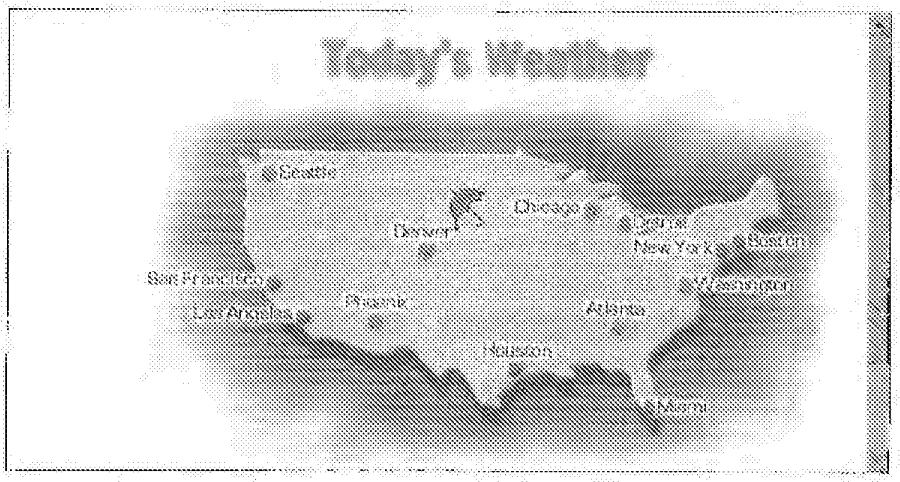


Figure 2.40 A weather technology map using the Comet cursor displays an umbrella by Denver to denote rainy weather.

has a technique that they call "Dead End Advertising," involving the placement of ads on pages that are like dead ends for users navigating the Web, such as thank you pages or receipt-for-payment pages. "The user has just finished a task, and [is] now mentally ready to see something else and go somewhere else," explains Darren Zwack, Director of Online Advertising at Columbus Group. It is on this page that the user is presented with the ad, completely unobstructed by any other content except a link back.

Another variation of redefining the screen real estate is ZDNet's "extramercial," which uses a 3-inch-wide column to the right of a Web page. This screen real estate is available for use with Netscape Navigator 4.0 and IE 4.0 (see Figure 2.41).

ISPs

Even ISPs are getting into the act. NetZero (www.netzero.com) is offering free Internet access for viewing ads. When a user signs up for the service, the demographic data collected will be used for ad targeting. Ads are served in a window that stays open whether the user is surfing the Web or reading e-mail. Figure 2.42 shows an example of an advertisement through the NetZero service.



Figure 2.41 Extramercial on ZDNet adds an extra column after the Web page for advertising space.

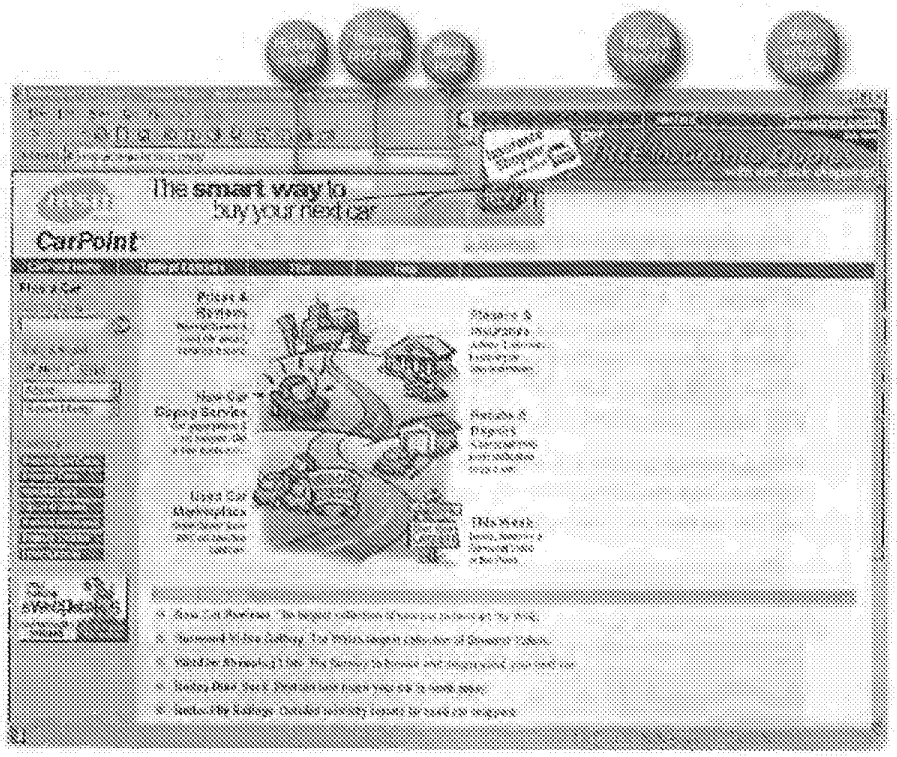


Figure 2.42 NetZero offers free Internet access for the price of viewing ads.

Other Formats

Online ads are moving to even smaller screens, the palmtops. Wired Digital has designed ads for PalmPilots and other Windows CE–compatible hand-held computers. With 13 million units in use by 2001, it makes sense that the ads could subsidize the content.

Guarding Your Domain

"People treat their domain name as their brand," says Douglas Wolford, Network Solutions senior vice president of marketing. Indeed, a good domain is an ad; that's why some domains have been sold for hundreds of thousands of dollars. And, companies are taking seriously the risk that sites will be created to challenge their brand, by buying up possible domains of "anti" sites. For example, Bell Atlantic bought the domains bellatlanticsucks.com and bigyellowsucks.com. One of the earliest "anti" sites was www.aolsucks.com. How AOL would like that domain now!

Resources for Reviewing Ads and Ad Campaigns

There are numerous resources online that highlight successful ads, review ad campaigns, and locate ad development shops. Microscope (www.microscope.com), a ClickZ network site, features a new ad campaign every week and invites leaders in the online ad industry to critique the ads themselves (see Figure 2.43).

ChannelSeven's network of sites (www.channelseven.com) cover a cross-section of the online creative industry. TurboAds (www.turboads.com/index.shtml) is a rich media ad resource. Ad/Insight (www.channelseven.com/adinsight/index.shtml), a section on the ChannelSeven site, provides a collection of ad campaign case studies. AdsGallery (www.adsgallery.com) is a visual yellow pages for creative resources in the online advertising industry and is comprised of a directory of agencies and contractors, as well as an interactive online portfolio of their work (see Figure 2.44).

Four Corners Effective Banners from White Palm (www.whitepalm.com/fourcorners) offers a look at various banners and their response rates. The site also suggests methods for improving responsiveness of banner ads.

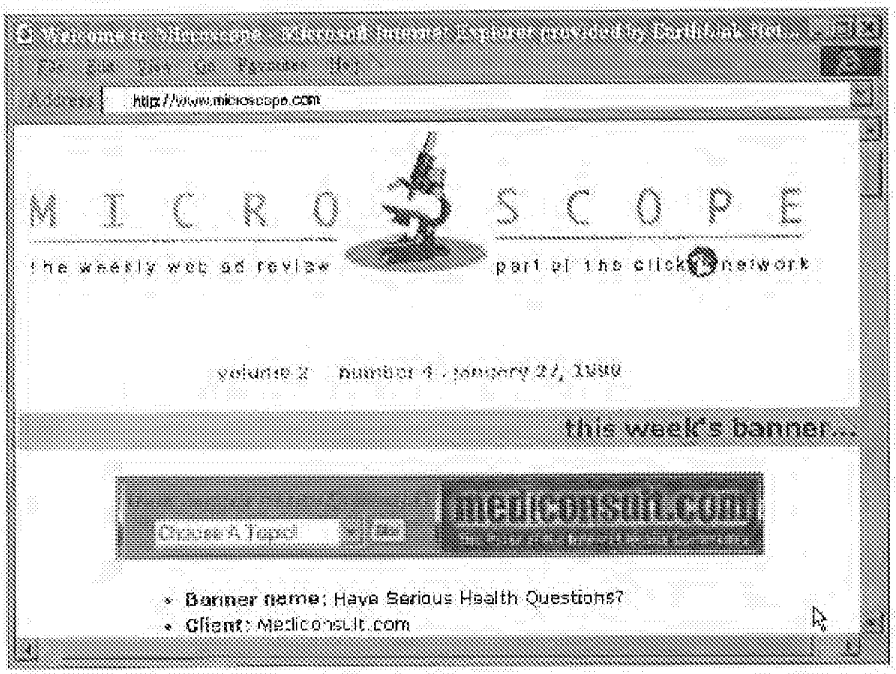


Figure 2.43 Microscope does a weekly review of web advertising campaigns.

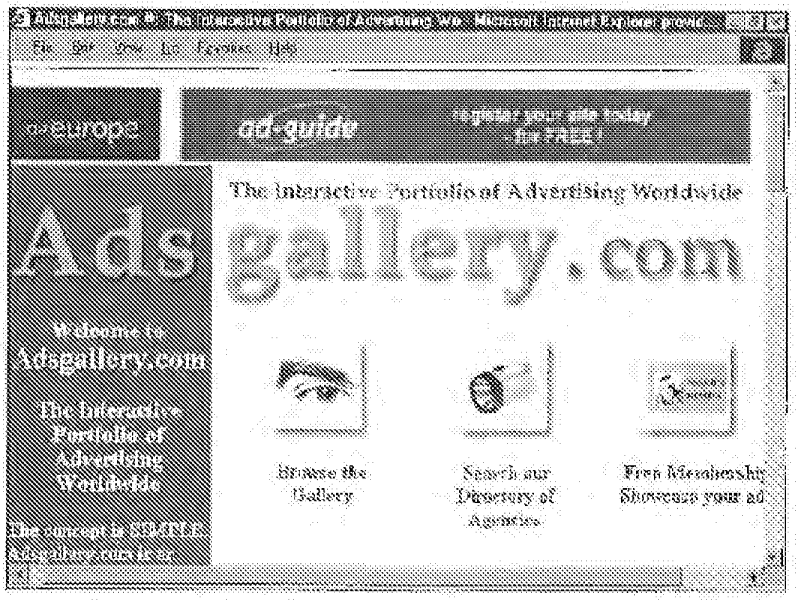


Figure 2.44 AdsGallery is a virtual yellowpages for the digital advertising industry.

In this chapter we've discussed a lot of the different types of ads, from text to banners to sponsorships. Indeed, the types and formats that make up the matrix of online advertising are extremely exciting and dynamic, and constantly pushing the entire industry in new directions. Online users vary as to whether they find ads informative, entertaining, or irritating (see Figure 2.45). However, one thing is certain: online ads are here to stay. In the next chapter, we're going to dive deeper into some of these models and see how they are being used by direct marketers.

	Banner	Standard Real Estate	Banner with DM	Interstitial
Like Very Much/Somewhat	42%	48%	45%	50%
Disliked Somewhat/Very Much	8%	14%	18%	18%
Irritating	7%	9%	15%	15%

Figure 2.45 Percentage of surfers who find ads irritating.
Source: Study conducted by IPSOS-ASI 1998.
Note: Results aggregated from an online copy-testing survey of 180 ads.

Think "Reach," Not "Rich"

Ali Partovi, Product Manager and Original Partner (apartovi@linkexchange.com)
LinkExchange (www.linkexchange.com)

Have you seen a "rich media" ad on your favorite Web site lately? Neither have I! If you're like the majority of users, your connection is not even fast enough for the practical needs of your busy day. You're frustrated and have no patience for an audio or Java ad distracting you from what you want to get done. Your favorite Web site probably knows this, too.

Web designers are finding that graphics-intensive, "beautiful" Web pages are less effective than light, simple ones. Sites like Yahoo! and Amazon that have focused on speed and simplicity have been rewarded with success, and others are now dropping their works of art for more functional designs. The same applies to advertising. Why try to force a multimedia ad on an audience that has clearly voted for simplicity and speed?

Web users are regular people, not technology enthusiasts. They're here to find things, learn things, buy things. They'd prefer to do so quickly and simply—without learning a new interface, without waiting for new technology to download. Every time we put new technology between our users and what they want to do online, we lose some users. So, think about the trade-off between rich media and reach. The smarter path is to aim for reaching the most people, even if we're not reaching them with as "rich" a medium as we'd like.

The vendors of rich media tools and software would like us to believe that rich media ads are the wave of the future. The technology enthusiast in me wishes this were true. However, in the meantime, fortunes are being won and lost on simple text and graphics. I'd rather have a text-only site with zillions of users today, than a "rich media" Web site with the hope of many users tomorrow. And I'd rather advertise with regular text or graphics ads and reach zillions of users, than shrink my advertising audience and impose a frustrating "rich media" experience on them.

People say that brand advertising is handicapped without video and audio, like TV. This might be true (and there are many ad agencies that hope it's true, as their business thrives on producing TV/radio ads). However, it's also true that the Web simply doesn't support rich media well. If your branding strategy is critically dependent on rich media, think about whether you can adapt your strategy for a world where zillions of users watch less TV and spend more and more time on simple text-and-graphics Web sites.

You don't need a rich audiovisual experience to influence consumer purchasing. The best form of advertising is word-of-mouth, and word-of-mouth can come in the form of a simple text-only e-mail. However, the Procter & Gambles of the world are holding their breath and holding their wallets while they wait for the Internet to re-create the rich audiovisual experience that made their TV branding strategy a success. How long can they wait?

It's Just Better Advertising

Richard Hopple, Chairman and CEO, (rhopple@unicast.com)
Unicast Communications (www.unicast.com)

Many argue that the Internet and its users are simply not ready for "rich media," or a more sophisticated form of advertising, and that bandwidth remains . . . a major problem and the solution to today's online advertising needs is "smaller, less intrusive banners."

Smaller, less intrusive banners? Isn't that regressing back to a form of advertising that research has proven to be ineffective? Static Internet banner ads, as a standard, have actually depressed or at least delayed interest in investing from big name advertisers. Since advertising dollars makes the Internet economy go round, this solution needs to be rethought. Without big name advertising funds, the user will ultimately pay.

We've Come Too Far

We've come too far to revert back to an advertiser's worst nightmare: a static banner delivering minimal information that many users choose to ignore. Let's take into consideration the progress that has been made in the world of Internet advertising. After years of skepticism, the Web has revealed itself as a viable medium to risk-averse, traditional advertisers—the advent the online ad industry has been awaiting. With the support of mainstream advertisers like Procter & Gamble, McDonald's, and Coke, the Internet as an advertising medium has the capability to deliver greater returns for both advertisers and publishers and maintain the delivery of free content to the user.

In addition, we've even come so far as to say that the bandwidth bottleneck is no longer an excuse for publishers to dismiss more attractive and effective advertising. Today, there are technologies that begin to solve this ever-vexing problem. In fact, some solutions enable the most innovative advertising concepts to be quickly delivered to users over a 28.8 modem.

The Time Is Now

The time to "sit back and ponder" the medium has passed. It is time to take action and find a better, more effective form of advertising on the Net.

According to research conducted by Grey Interactive and ASI, larger, more complex and sophisticated online ad units are much more effective than standard banner ads. They are more memorable, have higher click-through rates, offer increased communication and the potential for branding.

Which leads us to a solution—the badly misnamed, and in some circles surprisingly controversial, rich media. Essentially, rich media gives the advertiser[s] the opportunity to begin using some of the creative techniques that have been successful for them in other media. It enables animation and sound to capture

It's Just Better Advertising *(Continued)*

the nature of television advertising and the interactivity of the Web. It's not richer advertising. It's just better advertising.

Publishers should be working more aggressively to deploy the existing technologies, understand how they can best be used and how they can be improved going forward. It is in their interest to see that these solutions succeed.

Advertisers should begin paying at least as much attention to the creativity of the advertisements themselves as they have to the buying of the media—the kind of creativity that has powered the advertising industry in other media.

The technologies are here now. It is time to learn how to put them to best use . . . to begin pushing the creative parameters and show ourselves just how effective advertising on the Internet can be. If we don't learn to use them effectively, the big money will be left on the sideline . . . and that would be too bad for everyone.

CHAPTER 4

The Art and Science of Web Measurement

The growth of the Web advertising industry is closely tied to the development of tools for measuring audience size, behavior, and demographics. In traditional advertising, the measurement unit used to determine the cost of reaching an audience is CPM (cost per thousand), based on circulation for print media and projected viewing audience for television. With the Web, the number of times an ad is delivered and clicked is not projected—it is fairly accurately measured. The Web can measure audience behavior by counting the number of times an advertisement is seen, as well as the number of times an ad is clicked.

This chapter will cover the cornerstone of Web advertising, its measurability: what can be measured, how this measurement is conducted, and the products and services available for measurement, as well as the obstacles to effective measurement.

A common myth about the Web is that everything is measurable. As you'll learn in this chapter, the Web is far from the perfectly measurable medium people claim, but it is certainly more accurate than all other media.

Building Blocks

The foundation of Web measurement comes from the information captured by the computer housing a Web site—the server. When you visit a Web site, your browser has a "conversation" with the site's server, and much of that conversation is recorded by the server in a log file. Your browser first requests the HTML page from the server. The code in the HTML page directs the browser as to which graphics it should also request, and from where those graphics should be received. The server records the following user information from each request:

Meta's Exhibit No. 1004
Page 0085

- Date and time of the request.
- What was requested: a specific GIF, HTML page, etc.
- IP number of the computer making the request.
- Browser used by the surfer: Netscape 4, Internet Explorer 3.0, etc.
- Computer platform used by the surfer: Macintosh, PC, etc.
- Referring URL: If you clicked on a hyperlink to get to the Web site, the site from which you came will be recorded in the log files.

A log file records one line for every request. So, if you visit a home page with five graphics, there will be six requests recorded in the log file—one request for the HTML page and one for each of the graphics. For each request the log file records the IP number, specific item requested, browser used, time and date of request, and so on. As you can imagine, since all of this information is recorded for every request, log files become quite large. A busy Web site can generate many megabytes of data in the daily log file alone.

A lot of very useful information is housed within Web site log files. Figure 4.1 shows a raw log file, which should give you an idea of the difficulty of gleaning useful infor-

What Is an IP Number?

An IP (Internet Protocol) address is a numeric code that uniquely identifies a particular computer on the Internet. An IP address is static. An IP number, on the other hand, can be dynamically generated for an Internet user by an ISP (Internet service provider). Every surfer on the Internet has an IP number that is recorded in the log files of the sites that surfer visits. Here's how a surfer's IP number is determined:

- **Surfers using an ISP have IP numbers dynamically assigned to them by the ISP each time they dial in. Suppose your ISP is AOL. You dial into AOL to access the Internet. AOL has a batch of IP numbers that are assigned to customers as they dial into AOL. Now when you surf the Web, you are identified in log files by the IP number you were just assigned. Every time you dial into AOL, you can be assigned a different IP number from AOL's batch of numbers.**
- **If the surfer is accessing the Internet through a direct connection, the surfer's computer will have an IP address. The IP address is static; every time that computer is used to access the Internet, the same IP address is recorded in log files.**
- **If the surfer is accessing the Internet through a server at the office (which is how most business users access the Internet), the office server (like AOL) will dynamically assign an IP number to the office worker. It is important to note that folks accessing the Net through a company connection or an ISP may be going through a proxy server. Depending on how the proxy server is set up, it can assign each user an IP number, and/or all users can be given the same IP.**

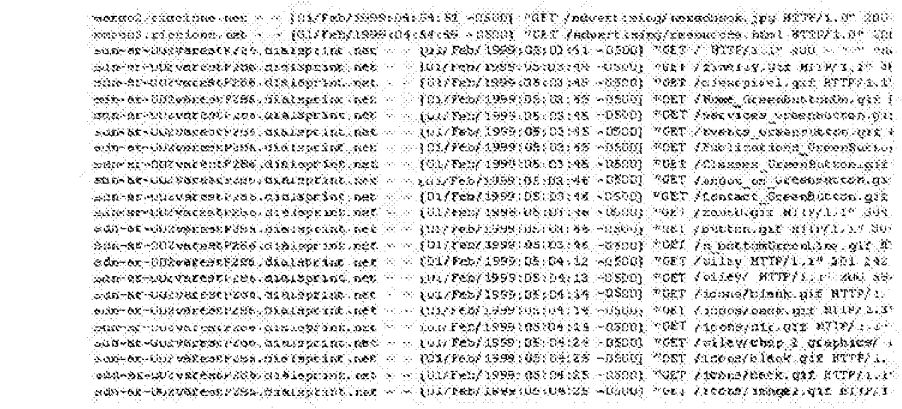


Figure 4.1 The contents of a log file.

mation by just looking through the unformatted text. Fortunately, special analysis tools will crunch all of this data to produce Web site activity reports.

Key Terminology in Web Measurement

Before we dive into the different types of log analysis reports and how they can help you, we will review some of the common terms that you'll see in Web site analysis reports. This new vocabulary blends technology and traditional advertising terms. The following are some of the key terms that will serve as the foundation for understanding the language of Web measurement. For a comprehensive assessment of the concepts and metrics in Web ad measurement, see this book's companion Web site, which includes a reprinting of the Internet Advertising Bureau's "Metrics and Measurement" white paper compiled by the Media Measurement Task Force. This white paper is a fantastic resource when it comes to explaining the intricacies and problems of Internet advertising measurement and reporting.

Ad request. The request of an ad from the server. This happens when someone visits a Web page that has an advertisement and the surfer's browser asks the server to deliver the ad. For a variety of reasons, the ad may not always be successfully served.

Ad view. Technically, the delivery of an ad by the server.

Cache (pronounced *cash*). The temporary storing of a Web page on a user's computer following the page downloading. This speeds up downloading on repeated use, so that the next time you request that page, it is accessed from the cache on your computer instead of from the Web server. Proxy servers and ISPs also cache pages to make access easier for their customers, and to reduce their bandwidth expenses (they don't have to use extra bandwidth every time a page is requested if they have cached that page).

Click rate. Clicks divided by ad requests.

Click-throughs. The number of times an ad is clicked.

Domain name. A naming system that translates the IP number given to computers on the Internet into user-friendly addresses. Domain names are read from left to right, going from specific to general. For example, the domain name for Yahoo! is www.yahoo.com, of which *www* indicates the type of server (World Wide Web), *yahoo* is the general domain name for Yahoo!, and *com* is the top-level domain (commercial). There are two types of top-level domains: categories (com, edu, org, etc.) and countries (uk-United Kingdom, au-Australia, etc.).

Exposures. The number of times that an advertisement is viewed. "Exposure" is commonly used interchangeably with "impression."

Gross exposures, ad views. These are advertising terms that denote the total number of times an ad is seen.

Hit. A hit is a record of each time a file is requested from a server. If a home page consists of eight graphics, nine hits would be recorded each time the home page is requested (one hit for the HTML text and eight for the graphics). Ten people going to that home page would result in 90 hits. Obviously, hits do not reflect the number of people viewing the site, just the number of files requested. Moreover, an increase in hits may be the result of more material on a site (let's say, a redesign of the home page added three new graphic elements) rather than an increase in users visiting the site. Sites still boast about their number of hits, but this basic unit of measurement isn't something that advertisers regard seriously.

Impressions. The number of times an ad is delivered. When an advertiser buys advertising on a CPM basis, the advertiser is paying for every 1,000 impressions that the site can deliver. Different sites measure an impression differently. Some sites count an impression when an ad is requested; others count an impression only when an ad is fully downloaded; and there are other definitions that are used as well. Savvy advertisers will ask the publishers how they define impressions.

Log file. This records all activity on a server. It records the items that are requested (Web pages, graphics, etc.), at what times they are requested, the browser used by the surfer, and other information. Traffic analysis software crunches log files to produce reports that are easy to understand and use as a basis for measuring Web site success.

Page view. This is commonly referred to as viewing all of the elements that comprise a Web page (graphics, text, etc.). It is generally agreed that if someone looks at a Web page that has frames, a page view will be looking at all of the elements of the page. However, some advertisers will count each frame as a separate page view to boost their level of traffic. Number of page views is a common metric for measuring traffic.

Response rate. The number of times an ad is clicked divided by the number of impressions that ad received.

Visitor. A person who goes to a specific Web site. Sites often define their traffic levels in terms of the number of visitors they've had in a given time period.

Visits. This is another term for comparing and measuring Web site traffic. A visit is activity on a Web site from a specific individual. That activity usually counts as a

new visit if the individual has been away from the Web site for a period of thirty minutes or longer.

What Can Be Learned from Log Analysis

Marketers can use Web site analysis software that they either purchase or custom build, or they can use a third-party service to take all of the data in their log files and produce analysis reports. The reports can be used to understand how customers are using the Web site and to help determine action items for improving the effectiveness of the Web site. Successful Web advertising goes hand in hand with an effective Web site. Below are some of the reports that can come from Web site analysis and how we suggest using those reports to improve your Web site.

Most Frequently Requested Pages

This can indicate a few things. First, you may discover that there are certain placements of navigational links that increase traffic to specific sections of your Web site. For example, we have seen that the highest traffic pages are often the ones that are the first navigational options listed on the home page, or those with buttons that are above the fourteen-inch point on the screen. The fourteen-inch point is the place where even material on the screen is visible even on a fourteen-inch monitor (one of the smaller desk top size monitors in use). Buttons above the 14-inch point have worked better than lower placements on the page, because they are seen by the greatest number of people. As we've seen time and time again in focus groups, Web surfers rarely scroll. Almost all Web users (even those with small monitors) can see everything within the first 14 vertical inches of the screen. This can be likened to putting the most important news above the fold on a newspaper.

Assuming that the different sections of your site have equally prominent navigational options, you can use this statistic to learn what is most interesting to your audience. It's probably a good bet that your hot content could generate great ideas for advertisements to which your target audience would respond well. You'll also know what content is worth adding to and where you shouldn't be allocating your resources.

If some pages deep within your site have unusually high traffic, you should consider pushing these pages more towards the front of the site. If they have significant traffic when they're buried, imagine what would happen if they were easier to access.

Number of Visits

This gives you a general idea of your traffic levels; it's a good benchmark of your site's popularity. A visit is most often defined as activity on your Web site from one IP address. If the IP address is not active on your Web site for a period of thirty minutes, then any new activity from that IP address will be counted as a separate visit. Since IP addresses are dynamically assigned by ISPs and proxy servers, and more than one person can have an IP at the same time, activity from one IP can actually be from a variety of folks.

Pages Viewed Per Visit and Length of Visit

This gives you an idea of how interested users are in your content. The number of visitors in a given time period is a good benchmark of overall traffic, but you need to know if you're actually getting good traffic—people who are interested in the content. For example, if 50 percent of your visitors look at one page of your site and then leave, you're in trouble. Watching these numbers will also be a good measure of the effect of new content and redesigns. If redesigns and new content increase the number of pages viewed, then you're succeeding at interesting people in what you're offering.

It's important to note that the length of visit can be deceiving. If your server is slow, visit length may go up because surfers have to wait longer for your content to download. Also, an increased visit length can sometimes be the result of a poor redesign that makes it harder for users to find the content that interests them. Of course, these increases in visit length would be far from an indication that the quality of your site is improving, so it's important to carefully determine why changes in the visit length occur.

Distinction Between First-Time and Repeat Visitors

This measurement is not going to be accurate unless folks have to register at the site or the site employs cookies (both cookies and registration will be discussed later in this chapter).

Most Popular Usage Times

This lets you know at what time a daily update is early enough to reach the most users. It also tells you the best times to perform server maintenance (when you have the least traffic). You can also learn about your consumers. Is your site receiving its heaviest traffic during business hours or on weekends and evenings?

Popular Paths Followed through the Site

This gives you a framework for testing navigational changes that influence how a surfer uses your site. You can answer questions like, "if I move the 'current news' button to the top of the page, will there be a greater number of people clicking that option?" This is great for determining what navigational options and copy will influence the path users take through the site.

Click-through Rate

This provides you with data to give to advertisers. Although the click-through rate isn't as important as actually reaching the right audience, a majority of advertisers use click-through rate to measure success, so this is an important measurement. You'll be able to see how different placements and navigational changes affect your click-through rate. You'll also notice what types of ads work best for your audience.

Browser Used

This will help you to determine what design elements are appropriate. If you want to use Java, you want to make sure that a majority of your Web users are using browsers that work well with Java. You should test-view your Web site with the various browsers that your consumers are using.

Referring URL

This tells you where your users are coming from, which is a good way to track the effectiveness of link trades and advertising. High referrers that are not from advertising are usually sites worth investigating for advertising opportunities. Many marketers think it isn't worthwhile to advertise on sites that deliver them a lot of traffic, because they're already getting that traffic for free. But we've seen that the increase in response from actual advertising usually makes the ad buy well worth the investment.

Domain and Host of the User

Log analysis software can use the IP address that is recorded in your log files to figure out the actual domain and host name. In the case of AOL, the name is www.aol.com. An IP address coming from a direct connection at IBM would be www.ibm.com.

Marketers can obtain good information from this data. For example, for one of our U.S.-based clients we found high amounts of traffic from countries outside the United States. In particular we noticed a lot of traffic from Brazil (.br). To accommodate the Brazilian customers, we recommended modifying the order page to make it easier for taking Brazilian orders. The orders from Brazil rolled in, and the client was glad we had uncovered this traffic pattern from their log files.

User Computer Platform

This provides valuable information for computer or software vendors and manufacturers.

User Entrance Page

This lets you know which pages need to serve a home page function and offer the look and feel of a "welcome." Users will not always enter through your home page, and knowing where they do enter will alert you to design those pages in a manner that draws users deeper into the site.

User Exit Page

It's important to know the last page your users view before leaving. These are pages that you want to redesign to keep users from leaving.

Web Measurement Tools

As we've mentioned, the information contained in log files is presented in a raw format. The ability to organize this information into actionable reports is a function of a variety of software solutions. There are three options available. The first is to build the measurement software yourself. This is what many of the early sites had to do before there were commercial products available. Today, the products are so good that this is relatively uncommon. The second option is to purchase the software and manage the process in-house. And the third option is to outsource the entire process. We will discuss the in-house management and outsourcing options below.

Third-Party Outsourcing

Web publishers with a bank of servers or with servers in different geographic locations can save considerable time and effort by hiring a firm to analyze all of their Web site usage data. Such vendors collect and process the data from a company's server logs to generate analysis reports off-site. The leading company for outsourcing Web measurement is I/PRO (www.ipro.com, see Figure 4.2). Using a third party for the time-consuming usage analysis leaves your company more time to focus on other aspects of the business. And, in addition to traffic reports for your own site, I/PRO provides reports on how your traffic growth and patterns compare to the Internet at large. This is useful for determining if your site's growth is faster or slower than the growth of the Internet as a whole. Moreover, many advertisers like the appearance of objectivity that comes from outsourcing to a third party. Since log files are just text file, one could easily tamper with the file. Having a third party handle the analysis places an extra layer between the publisher and the report.

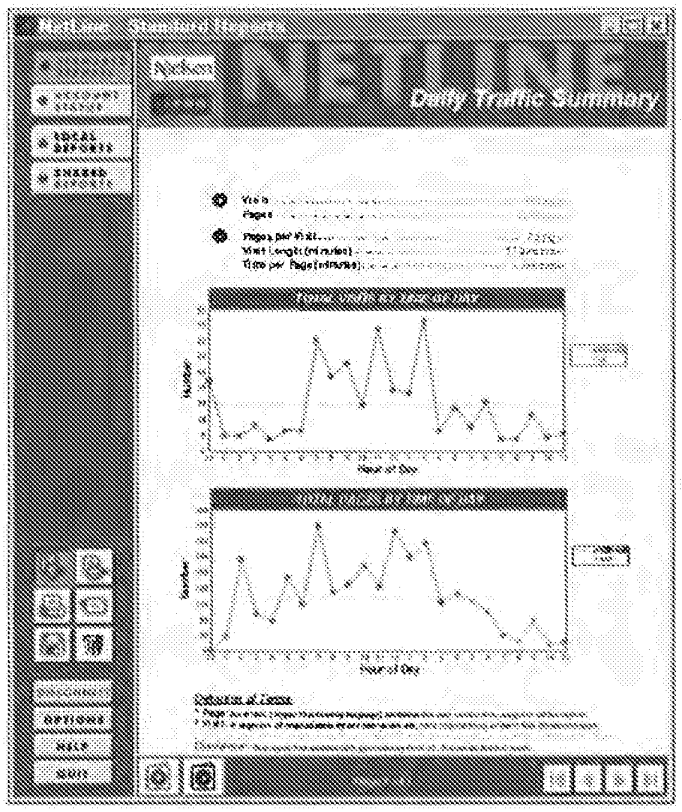


Figure 4.2 I/PRO NetLine Daily Traffic Report

The disadvantages of using such a service are that Web publishers lose control of their usage information. When using a vendor for measurement, a publisher contracts for services on a monthly basis. Payment models vary based on size and frequency of reporting requirements.

In-House Management of Software

Publishers can also purchase their own tracking software. Table 4.1 provides a list of leading Web measurement software solutions. Almost all of these companies will let you download a free trial version of the software from their Web sites, which is a great way to figure out which software will work best for you. Some of these products can service the smallest sites, and others can provide the horsepower to measure traffic on large commercial ventures. Off-the-shelf versions of the software can cost as little as $300. For example, WebTrends Log Analyzer is a popular choice for small businesses on a limited budget (see Figure 4.3). WebTrends Log Analyzer is easy to operate and does not require that the software reside on the server. Instead, the log file is downloaded to the user's computer and the reports are generated there.

Table 4.1 Web Measurement Software

COMPANY	PRODUCT LINE	SAMPLE CLIENTS
Marketwave (www.marketwave.com)	Hit List Pro Hit List Commerce Hit List Enterprise Hit List Live	Microsoft, Intel, IBM, Apple, Disney, AT&T
I/PRO (www.ipro.com)	NetLine	Dell, InfoSeek, NBC, USA Today, Excite, Netscape, Geocities, Kraft
net.Genesis (www.netgenesis.com)	net.Analysis	Lucent, Sun, 3Com, Xerox, Astra
WebTrends (www.webtrends.com)	WebTrends Log Analyzer (popular inexpensive product) WebTrends Enterprise Suite WebTrends Professional Suite WebTrends for Firewalls and VPNs	IBM, Netscape, Microsoft, Compaq, Novell
Andromedia (www.andromedia.com)	Aria 2.5	U.S. Postal Service, 3Com, Sony Online, Earthlink
Accrue Software Inc. (www.accrue.com)	Accrue Insight 2.5	Warner Bros., Hearst New Media, FedEx, Sun, Qualcomm, Motorola

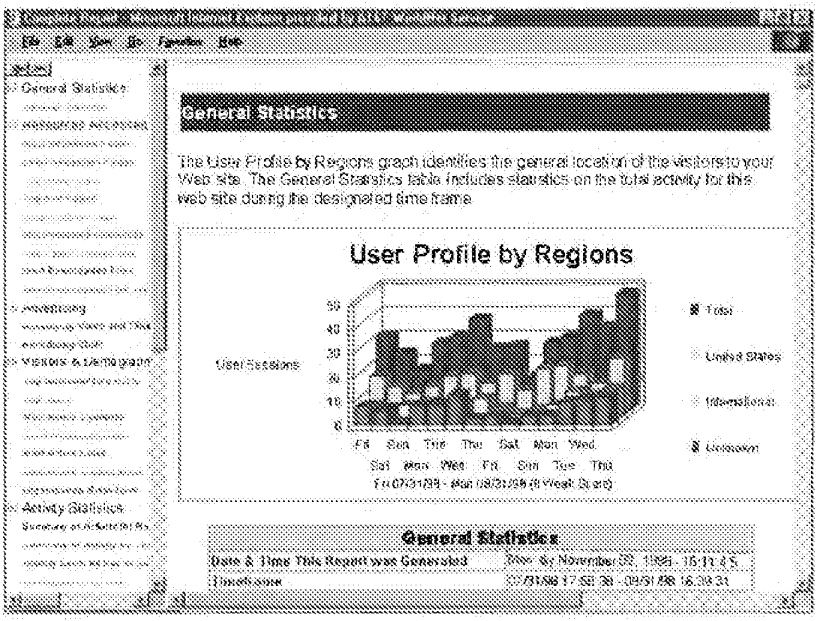


Figure 4.3 WebTrend's report of User Profile by Geographic Location.

For higher traffic sites, the cost can be thousands of dollars. Leaders in this industry include software from Accrue, Andromedia, Marketwave (see Figure 4.4), netGenesis (see Figure 4.5), and WebTrends.

Significant resources are needed for in-house reporting for Web sites that generate a high volume of traffic, including enough disk space to maintain log files, the computer to run the analysis, and someone to manage the software and the process. The process becomes even more complicated when a Web site is spread over several servers across the country.

The Problems with Web Measurement

As a new and rapidly growing industry, it is only natural that there are some problems with Web measurement. One challenge is a technological hurdle—certain developments in this area can skew the accuracy of numbers. Problems also arise as a result of the lack of standardization in the industry, which can cause measurement numbers to be misunderstood and difficult to compare.

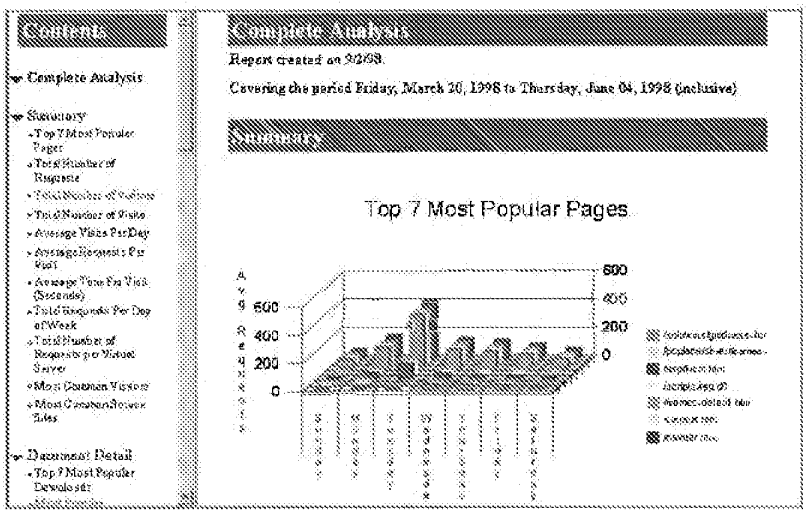


Figure 4.4 Marketwave's Hit List Pro report of most popular pages on a site.

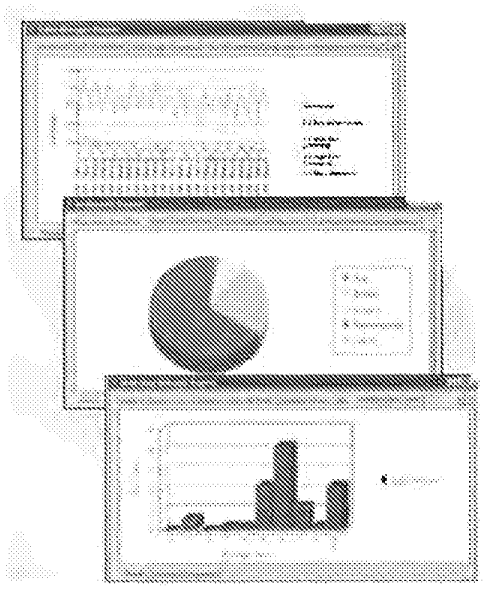


Figure 4.5 net.Genesis's net.Analysis series of Web site usage reports.

Technological Hurdles

Some of the very same technological developments that assist in speeding up Web delivery and facilitating corporate security can skew measurement numbers.

Caching

Caching describes the process of storing Web pages on a hard disk or server to decrease the download time; it eliminates the need to download the same pages each time they are requested. Most browsing software caches recently viewed pages on a user's hard disk, and the larger commercial online services and ISPs, such as AOL, store particularly popular Web pages on their server to expedite the loading of those sites for its members. This also occurs at many large corporations.

Although this process allows users to access information more quickly, publishers cannot track the usage of their Web sites when that usage occurs on someone else's server. For example, thousands of AOL users can surf a Web site, but if it is cached on AOL's servers, the publisher would not know about this activity. This is a problem for two reasons. First, the publisher does not have the correct usage data of the Web site. Second, if the publisher is selling advertising based on impressions, he or she is unable to charge for impressions delivered through cached pages, because no record exists of the impressions in the log files. Note that although impressions aren't recorded, clicks are—when someone clicks a banner it triggers a script that runs through the publisher's server, so sites that are cached will report inflated click-through rates.

The caching server does keep usage records of sites it has cached; this information is recorded in the caching server's log file. Publishers are not against caching, but they would like to have access to the log files of cached sites so they can accurately measure the activity on their Web sites. Help is on the way. There are now software products, such as MatchLogic's TrueCount (www.matchlogic.com), that can measure impressions even when caching occurs.

Proxy Servers and Multiple Users of One Machine

In an effort to maintain security on corporate computer systems, proxy servers act as the gateway from inside a company's firewall to the outside world. Requests for information from the Internet generated inside the firewall must first go to the proxy server, which then gathers the material from the Internet. Proxy servers usually use only one IP address, though larger systems may use several. The result is that a publisher's log file may have identified only one user (the proxy server) when in fact hundreds of computers at a company requested information from the site. This obstacle will be hard to overcome for obvious security reasons.

Another similar situation occurs when several people share a machine. This happens at businesses, academic institutions, and in home environments. For example, at a small company, only one computer may be connected to the Internet, and all employees wishing to surf the Web have to use that machine. This also occurs in an academic

setting where many students may use the same computers in the school's computer lab. And, in the home, every member of a family might use the same PC. The result is that the same IP address will appear several times, giving the impression that one person was returning to the site (or visiting a lot of pages) when in fact it was several. This demonstrates the shortcoming of defining a user by unique IP; the IP identifies the computer, not the user. Also, as we mentioned earlier, proxy servers cache publishers' pages.

Robots

Whereas caching and proxy servers can deflate the true number of requests, robots can inflate the traffic numbers. A robot can be programmed to request a site hundreds or thousands of times a day. A robot may also be a search engine that requests every page of a Web site in order to add those pages to its database. Fortunately, most measurement analysis software is sophisticated enough to identify this measurement aberration for known spiders (the ones used by the major search engines). However, if someone is tracking your site through a spider or an unknown search engine is spidering your site, it can be difficult to detect.

Industry Standards

Without a doubt, one of the growing pains of any industry is that standardization comes on the heels of invention. And in many respects, Web measurement—from terminology to technology—is still developing. With new measurement software coming on the market virtually every month, the lack of industry-wide standards accentuates the confusion. And this confusion extends from terminology to design to pricing structures.

In terms of reporting, standards are needed so advertisers can compare ad performance reports from different publishers. As Richy Glassberg, VP/general manager of Turner Interactive Marketing & Sales and founding member of IAB, puts it, "Standardizing Web measurement in general will make it easier for agencies and clients to buy so that everyone does well."

Any growth industry generates its own terminology, and the Internet is no exception. The Internet has a vocabulary for applications, procedures, and products. Though terms exist to describe the processes and procedures in Web advertising, what is lacking is industry-wide acceptance of the meanings of those terms. Fortunately, organizations like the Internet Advertising Bureau (IAB) (www.iab.net) have developed standards for Web terminology. Unfortunately, not everyone uses their metrics yet.

One example of a term with different meanings is "impression," which is the basis of a majority of ad buys. The methodology for measuring impressions actually differs from site to site. Some sites count an impression as every time a request is made of the server for an advertisement, whereas others count impressions only if the ad is actually delivered. Also, some sites might have ads rotate while the consumer is on the page. In that case, an impression on one site may be a presence on the page for 10 seconds, whereas on other sites an impression is the ad's presence on the page for the user's entire time spent on that page.

Cookies

A discussion about Web measurement isn't complete without a discussion about cookies. This is a topic about which many people fuss, but with which few are actually familiar.

A cookie is text that can be written to a text file on the user's hard drive when visiting a Web site. In most cases, that text will be a unique string, allowing the person to be identified.

Earlier in this chapter we talked about the discussion that a browser has with the server when someone visits a Web site. That discussion can entail the server asking the browser if the user has a cookie from the Web site. If the user has a cookie, the browser will present that information to the server. (Please note that the server can only access cookies that it has placed; it cannot access cookies from other servers.) If the user doesn't have a cookie, the server will try to place one in the user's cookie file (see Figure 4.6 for a look at an actual cookie file).

For example, a site may assign you a cookie with the text string "123." Every time you access a page on the site, the server can request the cookie it placed and recognize that "123" is back again and track where "123" goes on the site. Maximum information can be gained by using cookies in conjunction with a database. For example, the site that

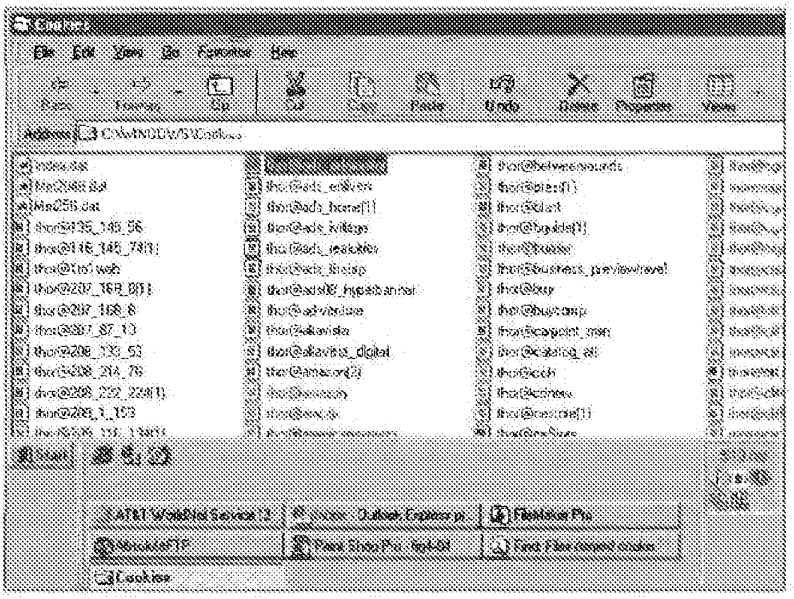


Figure 4.6 A Windows folder of cookie files.

gave you the "123" cookie may have a corresponding "123" record in a database. Every time you do something on the site, your actions can be recorded in the database.

People are often against cookies, because they believe that cookies are an invasion of privacy. In reality, a cookie can't find out anything about you unless you give that information to the Web site. The site knows that "123" is visiting, but not who "123" is or your address unless you provide that information. As we mentioned earlier, there can be a back-end database that is tied to your cookie ID. Then, each time you click on an advertisement or fill out a survey, enter a contest, or otherwise interact with a site, the actions can be recorded in a database as those performed by "123." But unless you fill out a form with your address, the site can't use a cookie to figure out who you are. And, if you fill out a form, the site will know who you are even if it isn't using cookies. The cookie ID tag resides on your hard drive, so the next time you visit the site, that information can be used again.

There are clear benefits for advertisers and publishers to using cookies. For example, if a publisher checks an ID tag, and discovers that according to its database the user visited the auto section of its Web site three times, the publisher can deliver an advertisement to the user for a car manufacturer. To take that scenario a step further, if the user then clicks on the ad and goes to Ford's Web site, Ford can place its own cookie (containing a unique ID) on the user's drive. If Ford enables users to design their dream car on its Web site, and a user chooses to do so, Ford can record that user's

What Determines If a Cookie Is Placed?

Cookies are supported by most browsers, from Netscape 1.1 to more current software. The full specifications on using cookies are provided in a document from Netscape titled "Persistent Client State HTTM Cookies"; it is available at www.netscape.com/newsref/std/cookie_spec.html. Although the server can try to set a cookie, there are situations in which the cookie cannot be placed on the user's hard drive.

The latest browsing software can block cookies. The press has been reporting that cookies invade a user's privacy by clandestinely placing a file on the user's hard drive (we disagree, as we discuss in the text). To address the concerns of users in this regard, both Netscape 3.0 and Internet Explorer 3.0 offer users the option to block cookie files. These browsers also have a feature that alerts users when a cookie is being attached.

Users can manually delete cookie files. As noted, cookies are text files that usually reside on the hard drive in the same directory as the Internet software. All you have to do is delete the text within the file, or delete the entire file, and you remove the cookie.

Proxy servers may preclude cookie use. These computers, which serve as gateways between the company's internal network and the Internet, do not always allow cookies to be placed on individual computers.

dream car in a database record that corresponds to the user's cookie. Then, the next time the user visits, the server can recognize the user based on the ID tag in his or her cookie and deliver an advertisement with the user's dream car. The possibilities for such targeting are obvious.

A publisher using cookies can record when a specific unique ID number saw an advertisement, and use that information to ensure that the same ad is not delivered multiple times to the same browser. This also gives advertisers an opportunity to deliver a series of advertisements to a particular user. After a user sees the first ad in a series, it is recorded in a database. The next time the server recognizes that user, it delivers ad number 2, and so on.

And finally, cookies can also retain other information about users that might be useful. For example, when a user registers at a Web site, the cookie can retain the user name and password so that registration is required only once. When the user returns to that site, the server will automatically check the cookie it gave to that user to determine his or her name and password. A cookie can also remember the items a user placed in a shopping cart. In this way, cookies help create a relationship between the content, the advertiser, and the visitor, without forcing the visitor to repeatedly fill out forms. However, cookies are far from perfect.

When the Cookie Crumbles

Unfortunately, the cookie picture is not all rosy. There are a number of reasons why a server cannot place cookies on a user's hard drive. In fact, only about eight out of ten Internet users can be cookied.

Moreover, cookies are mapped to the browser, not to the individual. The hard drive on which cookies are placed may be used by several people, thus opening the door for profile distortions. For example, at a college computer center where many students may be accessing a particular Web site through the same computer, the first time a student accesses the site from that computer, the computer will receive a cookie file. (In this case we'll assume that the publisher placed a unique ID in the cookie file, to learn more about the user.) Unfortunately, what the publisher thinks is one user is actually 100 different students using the same machine. The same problems (although to a lesser extent) are also caused by family computers that are used by more than one person.

Auditing of Web Site Traffic

Like any other advertising medium, Web advertising needs impartial third-party verification of numbers. Independent third-party audits will provide the foundation for confidence in the Web as a legitimate advertising medium. Web measurement is the timely tracking of all activity on a particular site. At this stage of the Web's evolution, there are a multitude of Web measurement solutions available, and most folks are using them. However, auditing Web measurement reports is mostly being done only by the large commercial sites with blue chip advertisers.

When advertisers are considering a site for ad placements, the advertisers want to know about the site's traffic. As an example, a site might say that it has 30 million page

views each week. How does the advertiser know that this is a true number? How does the advertiser know that the site is actually counting page views correctly? We have seen that sites both purposely and accidentally miscount traffic numbers, which is why audits are important. Audits are also important simply because advertisers are accustomed to receiving them from other advertising vehicles. Currently, most advertisers and their agencies aren't demanding audits, but we think that audits will be expected as the Internet advertising industry matures.

An audit usually tells an advertiser the total traffic of the site as well as the traffic the advertiser's ads receives (see Figure 4.7). An audit will also define the terms as the publisher's site uses them so that an advertiser can be certain what an impression really is. If an advertiser uses third-party ad serving, which we mentioned earlier in this chapter, the advertiser will have a third-party measurement of impressions and clicks, but will not have an audit of total site traffic. Just because something is done by a third-party does not mean it was audited. Also, auditors look for foul play, such as programs that will click on all advertisements to increase click-through rates.

The following are the main audit services currently available for Web sites:

Audit Bureau of Verification Services, Inc. (ABVS) (www.abcinteractiveaudits.com) was established on January 1, 1996. It is a wholly owned subsidiary of the Audit

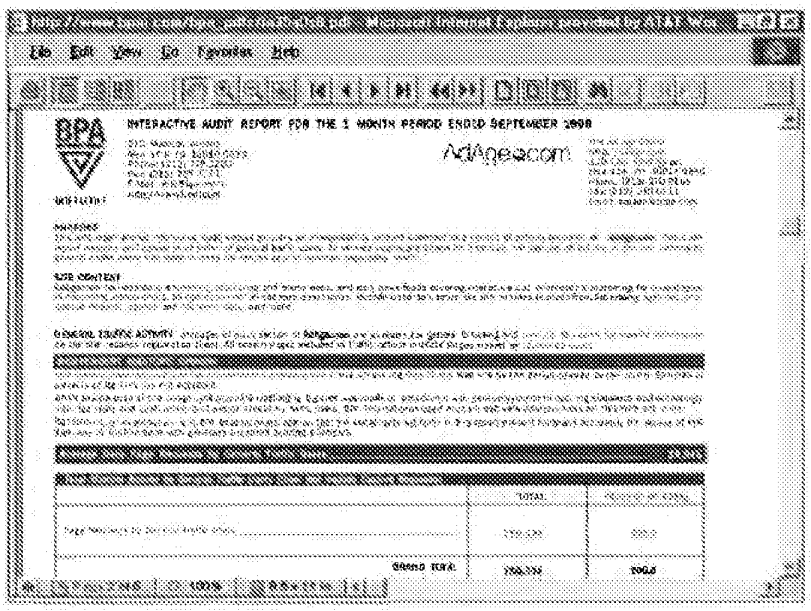


Figure 4.7 An Interactive Audit report by BPA for the AdAge.com Web site.

Bureau of Circulation (ABC). ABC—with over 82 years of circulation auditing experience—formed ABVS "to provide verification requirements of new and nontraditional advertiser-supported media."

BPA International (www.bpai.com), a global leader in magazine auditing since 1931, established an interactive auditing division because "as advertising on the Web grows into a commercially viable medium, advertisers need to know they are buying 'cyberspace' on the basis of accurate numbers that have been verified by an independent third party."

Nielsen, the television rating firm, has teamed with I/PRO (www.ipro.com) to offer a Web site analysis service for auditing and reporting Internet activity.

Rating the Web One Click at a Time

In advertising, broadcasters sell time or space to advertisers to disseminate their message to a particular medium's viewership. And, every medium—print, TV, radio, and so on—has a means to judge this viewership. Viewership is determined through ratings; that is, the number of people exposed to the message. Ratings equal the percentage of a given population group (for example, U.S. households for television) engaged with a medium at an identifiable moment in time.

There is nothing dull about rating Web sites. Whether you are referring to the act of profiling site audiences or to the firms that rate the sites, this is a field full of surprises. In 1998, the rating field shrank from four players to only two teams as the result of mergers. The two biggest players in the field, Media Metrix and RelevantKnowledge, merged into a new bionic Media Metrix ("The Power of Relevant Knowledge") (www.mediametrix.com). The other team consists of the strategic alliance between Nielsen Media Research (www.nielsen.com), the company that owns TV ratings, and NetRatings (www.netratings.com).

How are sites rated? Media Metrix offers coverage of over 15,000 Web sites and online properties, with monthly, weekly, and real-time data reporting, from a representative sample of over 40,000 people measured at home, work, and school. Media Metrix uses a suite of patented software "meters" placed on participants' computers to measure actual Internet and digital media audience usage behavior. The meters capture real-time, actual behavior of Internet users at the time it occurs—minute-by-minute, page-by-page, click-by-click. Tables 4.2–4.6 present the top 25 Web sites in a variety of categories, from Media Metrix. These ratings reflect the integration of the Media Metrix and RelevantKnowledge samples following the merger.

Table 4.2 Top 25 Web Properties
Media Metrix, Inc.
November 1998 Measurement Period
(11/1/98 through 11/30/98)

COMBINED AT HOME AND AT WORK			
RANK	WEB PROPERTY	% REACH	UNIQUE VISITORS (000)
Total Web Universe		100.0	56,781
1	AOL Websites	53.5	30,350
2	Microsoft Sites	48.9	27,751
3	Yahoo Sites	48.0	27,283
4	Lycos	45.2	25,656
5	GeoCities	32.0	18,158
6	The Excite Network	30.0	17,015
7	Netscape	29.9	16,987
8	Disney Online	25.6	14,510
9	Infoseek	20.2	11,488
10	Altavista Search Services	19.2	10,907
11	Time Warner Online	18.6	10,580
12	Bluemountainarts.com	16.0	9,087
13	Xoom Sites	15.4	8,759
14	Amazon	13.8	7,812
15	RealSite Portfolio	12.9	7,326
16	ZDNet Sites	12.6	7,170
17	BroadCast.com Consolidation	11.8	6,691
18	CNET	11.7	6,670
19	LookSmart	10.5	5,982
20	Snap.com Search Services	9.5	5,377
21	Switchboard Network	9.2	5,218
22	Infospace Impressions	8.0	4,554
23	eBay	7.7	4,382
24	Viacom Online	7.7	4,374
25	Go2Net Network	7.5	4,263

Table 4.3 Top 25 Web Sites
Media Metrix, Inc.
November 1998 Measurement Period
(11/1/98 through 11/30/98)

COMBINED AT HOME AND AT WORK			
RANK	WEB SITE	% REACH	UNIQUE VISITORS (000)
Total Web Universe		100.0	56,781
1	aol.com	48.7	27,630
2	yahoo.com	47.0	26,690
3	msn.com	32.7	18,544
4	geocities.com	32.0	18,158
5	netscape.com	29.9	16,987
6	excite.com	25.1	14,237
7	lycos.com	23.2	13,192
8	microsoft.com	22.9	12,991
9	infoseek.com	19.3	10,963
10	Altavista Search Services	19.2	10,907
11	tripod.com	18.5	10,488
12	angelfire.com	16.8	9,533
13	hotmail.com	16.7	9,509
14	bluemountainarts.com	16.0	9,087
15	xoom.com	14.6	8,280
16	amazon.com	13.8	7,812
17	real.com	11.7	6,647
18	zdnet.com	10.8	6,112
19	hotbot.com	10.4	5,898
20	Snap.com Search Services	9.5	5,377
21	switchboard.com	9.1	5,156
22	Disney Online	8.6	4,875
23	icq.com	8.5	4,850
24	simplenet.com	8.5	4,846
25	msnbc.com	8.3	4,693

Meta's Exhibit No. 1004
Page 00103

Table 4.4 Top 25 Web Properties
Media Metrix, Inc.
November 1998 Measurement Period
(11/1/98 through 11/30/98)

TOP 25 WEB PROPERTIES—AT HOME			TOP 25 WEB PROPERTIES—AT WORK		
RANK	WEB PROPERTY	% REACH	RANK	WEB PROPERTY	% REACH
1	AOL Websites	49.2	1	Yahoo Sites	49.6
2	Yahoo Sites	43.7	2	Microsoft Sites	48.7
3	Microsoft Sites	43.1	3	Lycos	39.3
4	Lycos	39.1	4	AOL Websites	36.3
5	GeoCities	27.9	5	Netscape	36.0
6	The Excite Network	25.0	6	The Excite Network	27.5
7	Netscape	24.8	7	GeoCities	24.2
8	The Walt Disney Company Online	21.0	8	The Walt Disney Company Online	23.8
9	Infoseek	15.8	9	Altavista Search Services	21.8
10	Altavista Search Services	15.3	10	Infoseek	20.4
11	Time Warner Online	14.5	11	Time Warner Online	19.5
12	Bluemountainarts.com	14.1	12	Amazon	14.4
13	Xoom Sites	13.3	13	ZDNet Sites	14.0
14	Amazon	10.8	14	RealSite Portfolio	12.4
15	RealSite Portfolio	10.3	15	Bluemountainarts.com	11.5
16	ZDNet Sites	10.1	16	CNET	11.4
17	BroadCast.Com Consolidation	10.1	17	Xoom Sites	10.1
18	CNET	9.7	18	BroadCast.com Consolidation	9.7
19	Snap.com Search Services	8.6	19	Weather Channel	9.3
20	LookSmart	8.5	20	LookSmart	9.3
21	Switchboard Network	7.2	21	Switchboard Network	8.3
22	Viacom Online	7.1	22	Infospace Impressions	7.6
23	eBay	6.7	23	Snap.com Search Services	7.4
24	Go2Net Network	6.3	24	Miningco.com	7.1
25	Infospace Impressions	6.2	25	eBay	7.0

Meta's Exhibit No. 1004
Page 00104

Table 4.5 Top 25 Web Properties
Media Metrix, Inc.
November 1998 Measurement Period
(11/1/98 through 11/30/98)

TOP 25 WEB SITES—AT HOME			TOP 25 WEB SITES—AT WORK		
RANK	WEB SITE	% REACH	RANK	WEB SITE	% REACH
1	aol.com	44.5	1	yahoo.com	48.9
2	yahoo.com	42.8	2	netscape.com	36.0
3	geocities.com	27.9	3	aol.com	34.0
4	msn.com	27.7	4	msn.com	32.0
5	netscape.com	24.8	5	microsoft.com	24.6
6	excite.com	20.8	6	geocities.com	24.2
7	microsoft.com	19.4	7	excite.com	22.8
8	lycos.com	18.6	8	Altavista Search Services	21.8
9	tripod.com	16.3	9	lycos.com	20.9
10	hotmail.com	15.7	10	infoseek.com	19.7
11	Altavista Search Services	15.3	11	amazon.com	14.4
12	angelfire.com	15.0	12	zdnet.com	13.2
13	infoseek.com	15.0	13	tripod.com	12.5
14	bluemountainarts.com	14.1	14	hotmail.com	12.2
15	xoom.com	12.7	15	real.com	11.6
16	amazon.com	10.8	16	bluemountainarts.com	11.5
17	real.com	9.3	17	cnn.com	10.8
18	Snap.com Search & Services	8.6	18	angelfire.com	10.1
19	hotbot.com	8.3	19	msnbc.com	9.5
20	zdnet.com	8.3	20	weather.com	9.3
21	icq.com	8.2	21	hotbot.com	9.3
22	simplenet.com	7.2	22	xoom.com	9.2
23	Disney online	7.2	23	switchboard.com	8.3
24	switchboard.com	7.0	24	infospace.com	7.6
25	ebay.com	6.6	25	espn.com	7.4

Table 4.6 Top Shopping Web Sites
Media Metrix, Inc.
WWW Audience Ratings
November 1998

RANK	DOMAIN NAMES	% REACH	UNIQUE VISITORS (000)
total Web users		100.0	56,781
Shopping		62.7	35,583
1	bluemountainarts.com	16.0	9,087
2	amazon.com	13.8	7,812
3	ebay.com	7.6	4,327
4	Barnes & Noble	6.6	3,769
5	CNET Software Download Services	5.9	3,327
6	etoys.com	5.2	2,957
7	cdnow.com	4.5	2,539
8	columbiahouse.com	4.5	2,536
9	musicblvd.com	4.0	2,262
10	valupage.com	3.7	2,109
11	classifieds2000.com	3.6	2,018
12	coolsavings.com	3.3	1,851
13	spree.com	3.0	1,718
14	bmgmusicservice.com	2.5	1,445
15	onsale.com	2.4	1,359
16	123greetings.com	2.3	1,296
17	beyond.com	2.2	1,277
18	CNET Commerce Services	2.2	1,251
19	netmarket.com	2.1	1,202
20	surplusaction.com	2.1	1,180
21	qvc.com	2.0	1,148
22	wal-mart.com	2.0	1,131
23	carpoint.com	1.9	1,076
24	freeshop.com	1.9	1,075
25	ezspree.com	1.9	1,067
26	egghead.com	1.8	1,003
27	reel.com	1.8	995
28	outpost.com	1.7	939

Table 4.6 *Continued*

RANK	DOMAIN NAMES	% REACH	UNIQUE VISITORS (000)
29	messagemates.com	1.6	913
30	cdw.com	1.6	910
31	priceline.com	1.5	880
32	tucows.com	1.5	851
33	bestbuy.com	1.4	809
34	landsend.com	1.4	809
35	autoweb.com	1.3	763
36	mypoints.com	1.3	758
37	firstauction.com	1.3	747
38	egreetings.com	1.3	739
39	supermarkets.com	1.3	712
40	ubid.com	1.2	691
41	winfiles.com	1.2	660
42	warehouse.com	1.1	628
43	autobytel.com	1.1	625
44	surplusdirect.com	1.0	595
45	netgrocer.com	1.0	575
46	jcrew.com	1.0	554
47	netcart.com	1.0	545
48	ezbid.com	0.9	520
49	worldremote.net	0.9	520
50	filez.com	0.9	516

Chart definitions:

Top 25 Web Properties: The top 25 World Wide Web properties are based on unduplicated audience reach. "Web Properties" include the largest single Web site brands as well as consolidations of multiple Web domains that fall under one brand or common ownership.

Example of "Web Properties": An example of a "Web Property" (i.e., one that offers more than one site) is the Excite Network, which encompasses city.net, excite.com, mckinley.com, and webcrawler.com.

Top 25 Web Sites: The top 25 Web sites are based on unduplicated audience reach. Top sites are from Media Metrix's At Home and At Work samples.

Reach definition and example: The percentage of Web-active individuals (i.e., individuals in the Media Metrix sample) that visited a site once in the given month. Individual visits are unduplicated (not counted twice). For example, if a site has a 2% reach, that means that 2% of Web-active individuals visited that site within the given month.

Unique visitors: The actual number of total users who visited the Web site once in the given month. All Unique Visitors are unduplicated (only counted once) and are in thousands.

Questions to Consider before Choosing a Web Measurement Solution for Your Company

What information do you want to track? Factor in the goals you have set for your Web site, the needs of your advertisers, and the reporting needs of upper management, to determine what information a measurement solution must provide.

Is it more efficient to perform the analysis in-house or should you outsource the process? In-house analysis is less expensive if you have a low to medium traffic site. If, on the other hand, your Web site generates high volume traffic, you need appropriate resources: disk space, adequate computer capabilities, and staff to generate reports. Moreover, the solution must be scalable for continued growth. Note that you need not only enough disk space for your current log file, but enough to archive log files so you can generate long-term trending reports.

What is the volume of traffic at your site? Can the products or service bureau handle it? Specifically, ask software vendors how much disk space and what type of a process you need to generate reports from log files to accommodate the traffic level (usually you will be asked to provide your log file size in megabytes).

How sophisticated does your traffic analysis need to be? Consider your current and future needs in your decision-making process.

How frequently do you need reports? Every day, every week, every month? Daily reports can be expensive when provided by a third-party provider. Also, if you do reports in-house, do you have the staff to put the reports together?

How many different departments need to use the data (marketing, IT, financial, etc.)? Does each department need different report configurations? It is a mistake not to check with all of the departments that will need reports before choosing a solution provider. If they all have different needs, and you

learn what they are, you have a better chance of finding a product that meets all or most of them.

What reports will your advertisers need? If you're a publisher, before you invest in a measurement solution, make sure it meets the needs of the companies footing the bill—your advertisers. Advertisers usually want to know how many times each creative is seen and how many times each is clicked. Advertisers prefer to get these statistics daily so they can tweak their campaigns to get maximum response rates (for example, daily reports will allow an advertiser to drop a poorly performing creative at the beginning of the campaign). Advertisers also want to compare how their ads perform on different pages of your Web site. Please note that unless your site is fairly low in traffic, you will benefit significantly from using adserving software rather than log analysis software to track advertising.

How private do you want to keep your tracking information? Will you allow your traffic statistics to be used in aggregate for industry-wide analysis by the tracking company? If you use a third party to generate traffic usage reports, your data is often used in aggregate for industry-wide analysis, such as I/PRO's analysis of click-through rates.

Will you need to have your reports audited? If so, check with the company that you want to perform the audits. They may recommend a measurement solution that makes the auditing process easier.

After you've considered these questions, you can conduct further research to ensure that you're making a smart purchasing decision. Find out whether any of the measurement solutions you're considering are already serving members of your industry. Ask how the solutions are working. Another alternative is to download a demo or trial version of the software. This gives you a chance to see how products work and how easy it really is to generate the reports you need.

CHAPTER 6

Targeting—The Online Advertising Advantage

For a long time Internet advertisers dreamed of delivering the right ad to the right person at the right time. In the early days of online advertising, the realization of this dream was touted as the true promise and potential of the industry. What was once just promise is now a reality. Various strategies are now readily available to capture the demographics, tastes, and preferences of users to better target content and advertising. In a broadcast television campaign, advertisers can target by geographic location, time of day, and the audience watching a particular show. With a direct mail campaign, marketers can target as finely as Fortune 500 employees who have attended safety and health conferences. Internet advertising can offer an advertiser all the targeting capabilities of other media plus a lot more. Indeed, on the Internet an advertiser can finally target every ad so that it goes out to just the right prospective customer at precisely the right time.

Advertisers use targeting to ensure the efficiency of their campaigns. Online publishers have a variety of targeting options to offer their advertisers (see Figure 6.1). In fact, online targeting moves up a continuum from the most basic means of reaching a customer up to extremely technologically sophisticated methods of learning about customers. The continuum starts with ad placement based on content and context. Next comes targeting gleaned from information received from the users themselves through site registration systems. Then there is targeting based on information available through Web measurement and ad serving tools. Targeting also occurs through database mining, collaborative filtering, behavioral analysis, and personalization tools. Today, these advanced tools are readily available and remarkably accurate. Advertisers are taking advantage of these targeting options to increase the effectiveness of their online ad campaigns.

The industry's intense interest in targeting rests on its benefits. First, targeting eliminates wasted impressions by only delivering ads to interested persons or those in the

Meta's Exhibit No. 1004
Page 00110

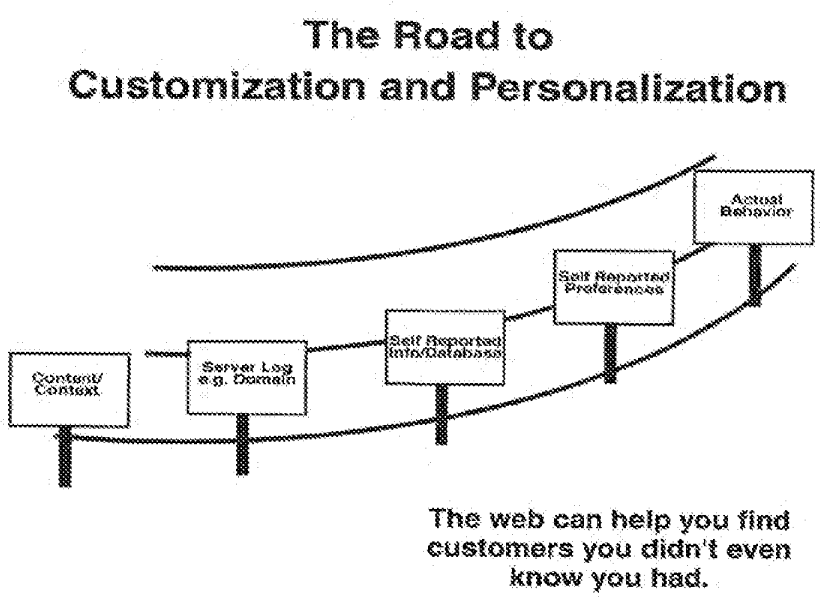


Figure 6.1 Internet advertising targeting continuum.

right target audience. Second, an ad is not considered a nuisance when it is information a person wants. Targeting ads to the right people increases interest and response.

How does the ad targeting actually occur? Targeting is facilitated through ad serving solutions. In the early days of targeting, the ad servers were separate from the targeting technologies, especially those providing personalization and customization. In 1998, these tools began to merge and form alliances to better facilitate ad delivery. In this chapter we're going to discuss the targeting continuum and demonstrate how targeting is the online advertising advantage.

Targeting Based on Content and Context

The most basic form of targeting is through content and context. This is how the bulk of targeting is done in traditional media. In magazines an advertiser can appeal to a particular gender or age group or niche market, by advertising in *Sports Illustrated* or *Teen People* or *Cat Fancy*. Using context to target, you can take your targeting a step further—for example, by advertising a car dealership on the classified ad page listing automobiles for sale, or a sporting goods store in the sports section of the newspaper.

This also holds true for radio, broadcast TV, cable, and even billboards. Targeting through content and context, of course, is also an option for online advertising.

On the Web there are thousands of content sites from which advertisers can choose. And, each of these sites offers a variety of targeting opportunities. There are probably few Web sites as good for reaching sports enthusiasts as ESPN (www.espn.com). And, if advertisers wants to further target within ESPN they can do so. For example, Mercedes-Benz runs advertising in the golf area instead of the basketball area (Figure 6.2 shows a Mercedes ad running in the golf section of ESPN and targeted towards golfers).

Similarly, advertisers can target through keywords on search engines. This is a much more focused buy than simply running an advertisement across a search engine.

Another example of content and context targeting is advertising for Prilosec on the health site InteliHealth (www.intelihealth.com). Although the entire site is about health, it is much more effective for Prilosec, a drug for persistent heartburn, to run advertisements in the heartburn area rather than throughout the site. Figure 6.3 shows the ads running in the targeted heartburn section.

Targeting Based on Registration Information

The simplest and often most accurate way to find out about a user is to ask the user directly. The key here is that users need to provide that information. The publisher



Figure 6.2 Mercedes-Benz ad on ESPN targeting the high-income demographic of golfers.

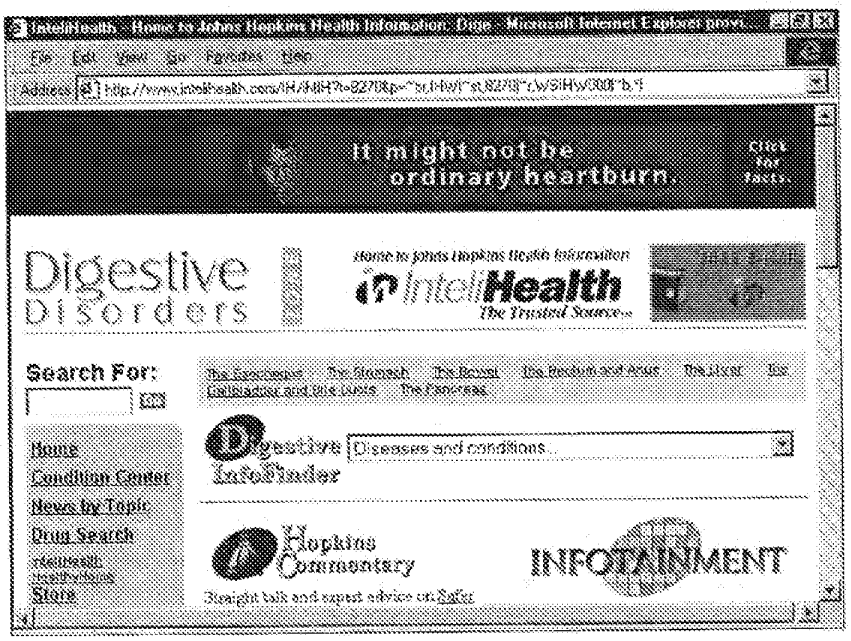


Figure 6.3 Prilosec targets its banner ads for heartburn medication in the "Digestive Disorders" section of InteliHealth.

may be able to get users to provide personal information by offering the opportunity to customize the Web site, enter a contest, receive a discount, or some other means. Basically, the publisher needs to give the user a very good reason to provide demographic data. If the Web surfer provides that data, the ad server can work in conjunction with the database that stores customer profiles and targets ads accordingly. For example, if a pet site asked users to register for a drawing for a year's supply of pet food and the consumers list their type of pet, the dog food makers would know they weren't wasting advertising inventory; they could target only those people who said they had dogs as pets. (Often, if a publisher can get the user to provide a zip code for customized news, weather, or other information, other demographic data can be pulled based on the zip code.)

One of the main misconceptions in the use of ad serving software is that publishers believe the software will make this type of targeting a reality. The software will enable this type of targeting, but the publisher needs to collect the data from the consumer to make it work.

Since targeting can always be improved with user-supplied information, we thought we'd go through some tips for getting your customers to give you that information.

Provide value. In general, users have expressed that they do not want to provide information to Web sites. However, they will provide information, if there is value. In Figure 6.4 users are asked to submit their zip code to The Weather Channel

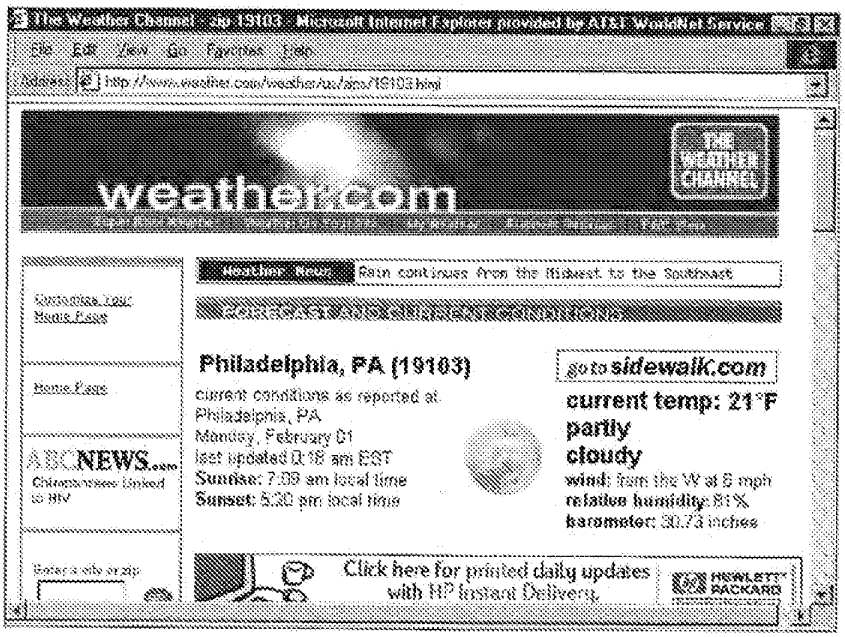


Figure 6.4 Weather.com provides customized weather reports to users who supply their zip code.

(www.weather.com) to get customized weather information. Most visitors to weather.com will probably provide this information, because they want the weather for their geographic area. After getting zip code information, weather.com can use a database to determine the household income, geographic area, and basic demographics about the user.

Tell people how they will benefit from providing you information. MSNBC also requests a zip code for news, sports, and weather. Instead of simply saying "customize this site," MSNBC entices the user by explaining the benefits of customization as the carrot for providing information. If your site asks for multiple pieces of information, it is good to explain how each piece of user-provided information will improve the user's experience on the site.

Only ask users to provide information that will play a part in delivering them value. The more a site asks of a consumer, the less likely the consumer will be to provide information. Therefore, if a question you ask isn't absolutely critical to your business or to providing value for the consumer, don't ask it.

It is interesting to note that there are some exceptions to this rule. We have found specifically that consumers visiting healthcare sites, when given an option, will choose to fill out a long form over a short form, if they think it will help them get information that is more customized to their needs.

Ask consumers for information in a fun way. The Excite home page often asks an interesting poll question. Users can enter their answer and then go to a page to see the cumulative tally from everyone who has answered the poll. In Figure 6.5 you can see a sports poll. This is fun for users, who are voicing their opinions, probably without even realizing that this can add to the user profiles that the site is building based on their behavior and answers to poll questions.

Privacy statement. Consumers are very concerned with what will happen to their information after they provide it. Response rates will improve if you have a clear privacy stance linked from every page on which the user is asked for data. The Direct Marketing Association provides resource information on building a privacy statement (www.dma.org).

Give users a reason to be truthful. This goes back to providing value. If each question doesn't lead to value for consumers, it is very likely that many of the responses you receive will be lies.

In trying to gather consumer data, it's important to keep in mind how the consumer feels. A 1998 Jupiter/NFO Interactive survey of consumers found that only 16 percent of respondents reported that they were "very likely" to register at a site that requires registration. The study went on to ask what factors would encourage them

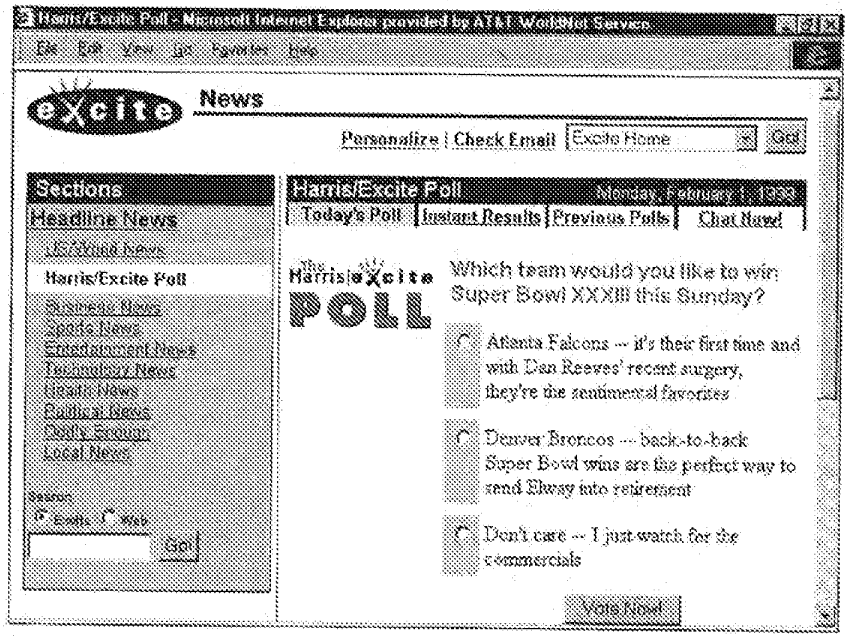


Figure 6.5 Excite uses surveys and polls to gather demographic information.

to fill in a registration form (see Figure 6.6). Over half of the consumers (56 percent) said that their willingness to register would increase greatly with a guarantee that their personal information would not be misused. Other reported important factors included content of higher value (41 percent) and a less time-consuming registration process (32 percent). Incentive programs ranked lower in this survey, with only 23 percent saying affinity points would serve as an enticement, and 21 percent that sweepstakes would.

In addition to putting up a registration page, there are software options like Lead Spinner 2.0 (www.leadspinner.com), which tries to make the registration process more friendly for consumers. The software installs a program so that when the site is loaded, visitors receive a message on the screen that asks them if they wouldn't mind if their visit was logged. If they accept, then the data about the user is sent via e-mail into a database. If the user says no, then nothing is taken and the user can continue through the rest of the site without any further bothers. "As you can see, this is a great tool and service for people to use to find out 'who' is visiting the site as well as turn visitors into potential sales," says Brian Kilcrease, vice president of Glass Spider Designs, Inc.

Using Cookies for Targeting

In the conversation between browsers and servers, cookies can be exchanged, which makes them useful for targeting. As we mentioned in Chapter 4, cookies are a way to identify browsers on a Web site. (A brief refresher: a server can place a cookie on a hard drive. The cookie most often contains a unique string of text that allows the user's browser to be recognized if the user returns to the Web site.) A cookie can track what visitors do on the site and target accordingly. For example, site visitors that go to the sports section numerous times could be "tagged" as sports people and delivered an NFL ad the next time they came back to the Web site.

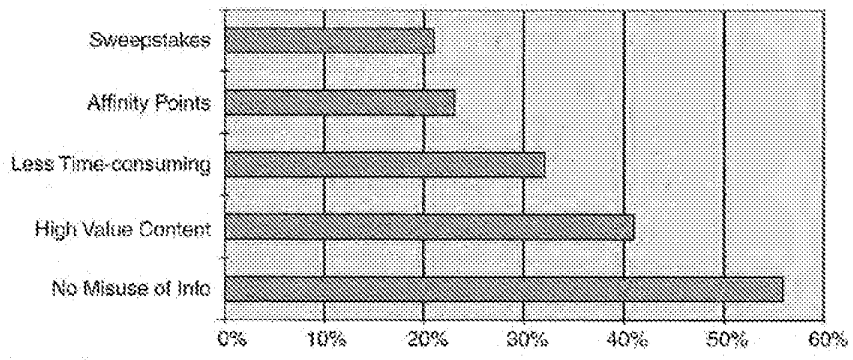


Figure 6.6 Factors that would greatly increase an online user's willingness to complete an online registration form.

Cookies can be used to facilitate user behavioral targeting. Behavioral targeting entails having enough content on your site that you can determine a user's preferences based on his or her actions on your site. For behavioral targeting you can create rules. For example, if a user goes to hiking content, that person should be profiled as an outdoors person. Or behavior-based targeting can be controlled by the software, which gauges a user's likelihood to want certain content or click particular ads based, for example, on the actions of previous users who took similar navigational paths.

An interesting use of cookies for behavior-based targeting is through a service that builds a user profile across those Web sites that use the service. There are a number of ad serving services that cookie users and track them across all participating sites. As users make decisions about what ads to click, or fill out registration forms, the software builds profiles that advertisers can use for targeting. A service like MatchLogic, which tracks users across high-traffic sites like Excite and WebCrawler, can generate a very large database of profiles, which will help advertisers deliver ads to the right people. As of the writing of this book, Engage Technology, another targeting service, had 30 million profiles. When advertisers or publishers deliver an ad through one of these services, they can target based on these profiles. In the future, this type of targeting will allow non-targeted buys to become targeted. For example, an advertiser could run an advertisement on Yahoo! and ask that the ad be delivered to users who request gardening information and who live in zip codes that indicate a certain income level. Later in this chapter we'll mention other companies that provide this type of targeting. It's important to note that a company like MatchLogic can only get registration information if that information is shared by the site on which they're delivering the ads. In the case of Excite, since the company owns MatchLogic, it's clear that the two share data.

Targeting through Database Mining

Combining ad management software solutions with consumer profiling technology is a hot targeting technique that appeared on the scene in late 1998. One example, NetGravity's Global Profile Service (www.netgravity.com), enables advertisers and publishers to target Internet advertising and marketing messages based on comprehensive consumer profiles. NetGravity is working with such industry-leading data providers as MatchLogic (www.matchlogic.com) and Aptex (www.aptex.com) to build a flexible and anonymous consumer profile data source. "By aggregating multiple data providers," says Rick Jackson, vice president of marketing at NetGravity, "NetGravity's Global Profile Service will enable the thousands of Web sites powered by NetGravity to better understand their audience, and target advertisements and direct marketing offers with greater precision and insight." The Global Profile Service is based on using cookies to recognize users and build profiles of their interests. The service allows advertisers to unlock the value of their inventory by analyzing and segmenting audiences more acutely. "Precision Internet marketing and targeting requirements differ greatly from site to site and advertiser to advertiser. Since no two customers require the same data, a flexible solution is required," said Mark Peabody, a senior analyst with Aberdeen Group.

Targeting through Profiling and Personalization

Now we're going to discuss technologies that usually work in conjunction with an ad server to help profile prospects and deliver them customized advertisements or content (see Table 6.1).

Recommendations as Ads

Creating profiles to personalize content and make recommendations to consumers helps sites to secure sales, especially retail sites like CDnow and Amazon.com. The better these sites can service the customer by recommending appropriate products, the more likely the customer is to purchase. Retailers like these can use engines that make recommendations based on previous inquiries and purchases. The engines can also make recommendations by asking users about their tastes. For example, CDnow has incorporated NetPerception's (www.netperceptions.com) recommendation engine into its site. Users are asked to provide a a list of artists they like. The engine then searches its database for people who listed the same artists and recommend their recent purchases. If consumers rate a list of products, they can get recommendations when the engine cross-references their ratings with other users. Engines can also make recommendations by tracking what the user does on the site and comparing it to rules developed by the publisher or to other visitors.

Table 6.1 Personalization Companies

COMPANY	COMPANY TAG LINE	SELECTED CLIENTS
Art Technology Group (ATG) (www.atg.com)	Manages online relationships from advertising to commerce through personalization technology.	BMG Music, Kodak, Newbridge Networks, Sony, Sun Microsystems
BroadVision (www.broadvision.com)	Internet application solutions for one-to-one relationship management across the extended enterprise.	American Airlines, Liberty Financial, Softbank, Sony Communication Network, Sun Microsystems, US West
Net Perceptions (www.netperceptions.com)	Developer and supplier of real-time recommendation solutions.	Amazon, CDnow, Elonline, Ticketmaster
GuestTrack (www.guesttrack.com)	Tailors content "on the fly" based on the interest profile of each member of the client's audience.	

Personalization

Personalization helps a company send relevant messages to different target audiences that are visiting their Web site. For example, age-sensitive ads can be dynamically delivered to customers fitting the age profile. And this is exactly what American Airlines is able to do using Broadvision's (www.broadvision.com) Dynamic Command Center. Once American Airlines knows that a user on its site is over 21 years of age, a beer or wine ad can be shown. Another use of personalization is through action-based targeting to offer purchasing incentives. For example, if an online shopper places an item in a shopping cart and then removes that item, you can offer a discount when she comes back to that item as an incentive to complete the purchase. By knowing the customer's activity history, you can move browsers into buyers.

Art Technology Group's (www.atg.com) Ad Station uses sophisticated behavior tracking to create a profile on each visitor, and uses this profile to deliver targeted and compelling ads (see Figure 6.7). "In the Internet business world, time, ease of use, accessibility, and performance is everything," says Nancy Morrisroe, Director of Public Relations, at Art Technology Group. "We created the Personalization Control Center so that business users don't have to rely on programmers to specify targeting rules for personalization."

Another personalization option comes from Net Perceptions (www.netperceptions.com). Their Ad Targeting technology can be deployed from within the framework of most leading ad server platforms and starts learning about each site visitor's interests and tastes from the first visit and keeps learning with every repeat visit. Tracked variables include recency/frequency, behavioral information, page ID, time of day, demographic information, keywords entered into search engine, and other categorical information. "This product increases Web site advertising revenues by increasing click-throughs for advertisers," says Steve Larsen, Net Perceptions' VP of marketing and business development. "It does this by automatically learning the individual interests and tastes of site visitors, then putting the right ads in front of each visitor in real time." The software deploys different ad targeting technologies at different times, so that it dynamically adjusts between new visitors, about whom it has minimal information, and repeat visitors, about whom it has learned much more. According to Larsen,

Benefits of Personalization and Customization

Segmenting users. You can segment your users into affinity groups for better targeting of promotions.

Matching promotions to users. You can match time-sensitive and topic-specific promotions to users.

Integrating legacy systems with current customer databases. You can customize every contact a user has with your company.

Tailoring the Web site. You can customize a Web site to each user's interests and tastes.

Tailoring the sales copy. You can customize sales copy to each user's interests and tastes.

"Targeting gets increasingly personalized as the software continually learns about each visitor's individual interests and tastes." Although speaking for his own company, Net Perceptions, Larsen is really talking about the goal of all the services we've discussed in this chapter.

Net Perceptions' software learns continually and gets smarter with each customer interaction. Like most software of this kind, NetPerceptions can also integrate user-provided information (as in the cases of Amazon.com and Cdnow who both use Net-Perceptions) to customize their content. Net Perceptions' recommendation engine works like a salesperson, recommending products that match a person's tastes and interests (see Figure 6.8a and 6.8b).

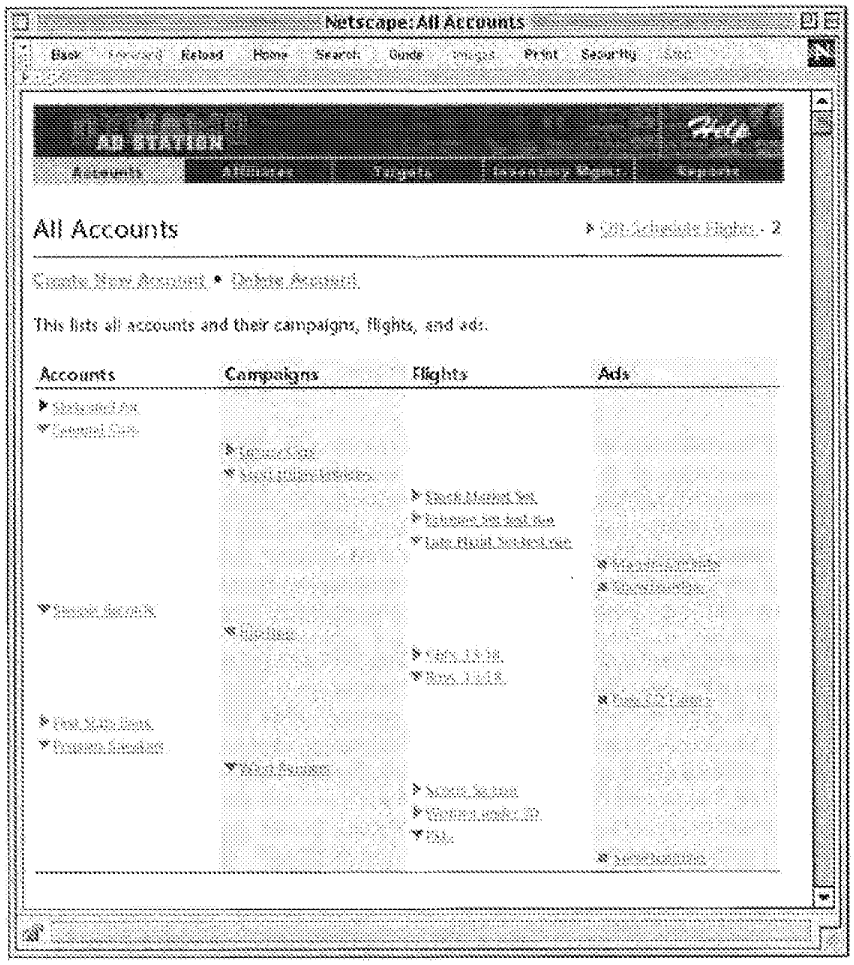


Figure 6.7 ATG's Ad Station.

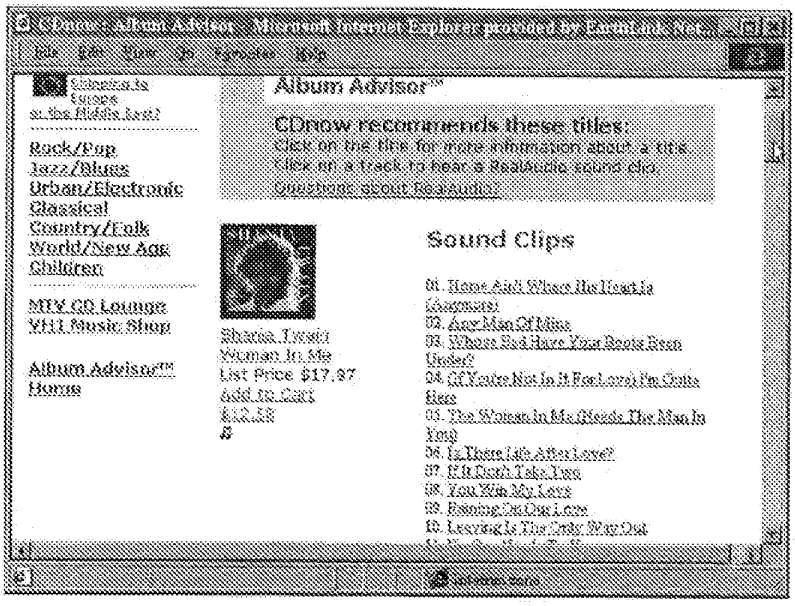


Figure 6.8a CDnow's "Album Adviser" uses the Net Perceptions' recommendation engine. The "Album Adviser" asks a customer to supply names of three of their favorite artists or albums.

Personalization is not language- or cultural-specific. BOL uses Net Perceptions' software extensively as an online bookseller to consumers in the United Kingdom, Germany, France, Netherlands, and Spain (each version of the site in the appropriate language). BOL employs both implicit and explicit data gathering in order to generate highly specific customer ratings. The site offers customers an array of personalization choices, allowing them to get ratings quickly with a *Book Chooser* function, and then continues to develop a long-term relationship using the BOL Recommendation Center.

Personalization doesn't stop on the Web; it is moving off-line as well. Personalization products such as BroadVision's One-to-One (www.broadvision.com) can be integrated with corporate data warehouses that store all kinds of customer information. One-to-One can be accessed by a corporate call center, account executives, or sales staff, whether working with a customer online or over the phone. With this system, the personalization that started online continues no matter how a customer contacts the company. For example, when a customer calls the company, the call center operator can access the customer's record, which lists what that person bought online and the customer's entire contact and purchasing history. This eliminates the need for customers to begin from scratch each time they contact the company. BroadVision sees personalization as a means of integrating all the available customer data to better service customers.

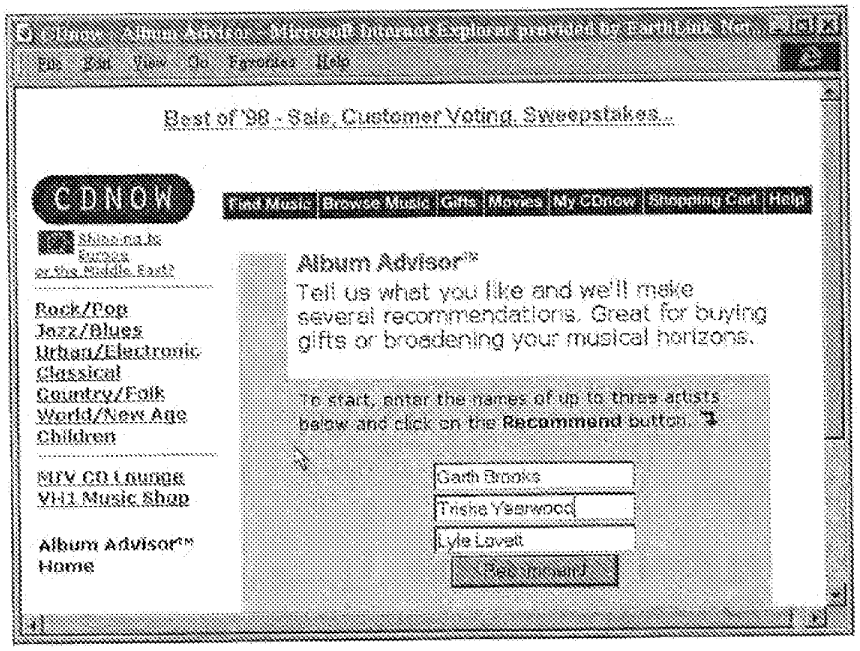


Figure 6.8b The "Album Adviser" then makes recommendations of other artists and albums.

Collaborative Filtering

Another type of personalization is based on collaborative filtering. The company who perfected this technique is Firefly (www.firefly.com). A Firefly-enhanced site works by offering users the option to get a "passport." Users who register for a passport provide personal information building a personal profile. Drawing on each user's profile, Firefly can make recommendations based on its knowledge of passport users with similar profiles. Businesses can then leverage this information to provide each individual with highly personalized service and customized ads. Sites can also use Firefly to build community; users with similar profiles can find and communicate with one another. Users have an incentive to use their passports because it will mean customized content and recommendations. Sites that use Firefly can choose to have a passport exclusively for their site or to be part of the network of passport sites sharing profile information.

Behavioral Targeting

And finally, the last stage on the continuum is targeting based on a person's actual behavior at any moment in time. One of the companies offering this services is Aptex (www.aptex.com) with its SelectCast for AdServers. SelectCast acts as an "intelligent

observer," mining the context and content of all actions—including clicks, queries, page views, and ad impressions. No explicit user feedback or "taste" judgments are ever required. No cookies are used and the information is not captured or saved. SelectCast unobtrusively profiles all user actions, learning the detailed interests of each user, and applying this knowledge to every ad selection. Designated ads can be delivered consistently to selected groups, and new audiences can be automatically identified as they emerge. By correlating updated user and audience profiles against ad inventory, SelectCast continuously fine-tunes ad delivery. Aptex also offers products that help to personalize a Web site by using profiles to deliver custom content and e-commerce promotions. In short, SelectCast compiles demographic and taste information without registration or database mining. Sites that utilize SelectCast have become the secret weapon for successful banner ad campaigns for such noted interactive agencies as Anderson & Lembke and Left Field. Brent Hall, Internet Active Media Planner at Anderson & Lembke says "I use it on the majority of my campaigns." Infoseek uses SelectCast for its UltraMatch service and Brian Monahan of Left Field has found that placing ads on Infoseek's UltraMatch has resulted in better targeting making all their buys as targeted as keyword buys. In short, SelectCast can help you find customers you didn't even know you had.

Putting It All Together: How Ad Management Solutions Facilitate Targeting

As we mentioned in Chapter 4 on Web measurement, information is shared between a site's server and the user's Web browser. Advertisers use this information for ad targeting. As a brief refresher, here are some of the targeting abilities that result from this technology-based conversation:

- IP address: domain and host (once the company name is determined, some services allow for SIC and revenue targeting).
- Geographic location: national, local, etc. (For issues regarding local targeting, see Chapter 11.)
- Browser used: Internet Explorer 4, Internet Explorer 3, Netscape 4, etc.
- Date and time of request.
- Computer platform: Mac or PC, and version of operating system.

Ad management server software now integrates targeting with ad delivery. Below is a list of the ways ad management software integrates targeting into the delivery of ads.

Frequency capping. An ad server can limit the number of times an ad is seen by the same person (usually through the use of cookies). Along those lines, an ad server can also manage the delivery of a sequence of advertisements.

Targeting based on the standard industry code (SIC) of the browser. The SIC represents the sector of the economy in which a company falls. For example, there is an

SIC for ad agencies. When a user goes to a Web site, the browser reveals its IP number. The SIC can be determined by looking up the company at the IP number (let's say it's i-frontier) and then determining that company's SIC from a database (in this case, advertising agency). The same system can be used to target by company size and revenues.

Profiling/personalization. As we mentioned earlier, this allows the advertiser to target based on the user's behavior (what content he chooses to view, etc.) or user-provided registration (or customization) data.

Profiling base wider than site. Some services build profiles of users from all of the sites that use their programs or services. Advertisers can benefit from profiles that have already been developed across other sites.

Research company partner. Many of the services have partnered with database firms to bundle already existing customer databases and database profiles with information gathered online.

Closed loop reporting. A system to provide advertisers with a means for measuring which ads produce specific results.

Auditing. If you're a publisher and you want to provide audited numbers to your advertisers, all of these services can be tied into audits.

Types of ads that can be delivered. We haven't included this information in Table 6.2, because it is constantly changing. However, it is important to make sure that your service can deliver the types of ads you want to use. For example, currently, not all services can deliver all types of rich media ads.

Privacy

As we mentioned earlier in this chapter, a strong privacy policy is important for getting users to provide information. Here is Firefly's privacy policy:

Firefly Network Privacy Policy

Your e-mail address will not be shared outside of the Firefly Network for marketing purposes without your consent.

Profile information will be used to create personalized content, services, and advertising on sites in the Firefly Network. In addition, Firefly and third-party licensees may use your profile to generate aggregate reports and market research.

You may inform Firefly Network, Inc. at (www.firefly.com) at any time if you would like to cancel your account and have your contact information deleted from Firefly's records.

In addition to the Firefly Network Policy, we also take the following steps to ensure your privacy on sites operated by Firefly Network Inc.:

1) Firefly Network Inc. does not require an end-user to provide a name and address to use the service.

Table 6.2 Ad Management Targeting Capabilities

COMPANY, TECHNOLOGY NAME	NAME OF SERVING SOFTWARE THAT CAN BE INSTALLED	NAME OF SERVING SERVICE	PROFILING/ PERSONALIZATION	PROFILE BASE WIDER THAN SITE	RESEARCH COMPANY PARTNER FOR ANALYSIS	CLOSED LOOP REPORTING	COMMENTS
Adknowledge (www.adknowledge.com)	Not offered	SmartBanner	Yes	Yes	Yes	Yes	For advertisers and ad agencies. Also offers functionality that assists an advertiser in planning an ad campaign and executing a buy.
Double Click (Doubleclick.com)	DART	DART	Yes	Yes	Harris Black International	Yes	Offers a service for tracking e-mail advertising.
Engage Technologies (www.engage.com)	Engage Technology	Accipitor AdBureau	Yes	Yes. They currently have 30 million profiles.	VOYAGER Profile from Millward Brown Interactive	No	Over 2000 Web properties currently use Engage, allowing for an extensive profiling base.
AdForce	Not offered	AdForce	Starpoint	Yes	VOYAGER Profile from Millward Brown Interactive	Yes	
MatchLogic (www.matchlogic.com) (owned by Excite)	Not offered	MatchLogic	Yes	Yes	Harris Black International	Yes	Offers unique technology for counting cached impressions.
NetGravity AdCenter (www.netgravity.com)	NetGravity AdServer	NetGravity AdCenter	Yes	Yes	Yes	Yes	Gets their profiling information from MatchLogic and Aptex. Plans to use other vendors as well.

Meta's Exhibit No. 1004
Page 00125

2) All profile data stored by Firefly Network Inc. has been entered solely by the respective end-user.

3) End-users can modify registration data associated with their account, and can control the amount of this information that is disclosed on their member page.

4) Firefly Network Inc. does not provide functionality to perform "reverse searches" which find individual names from address, e-mail, or phone number information.

5) Conversational transcripts of online chat are not retained except for moderated chat sessions.

6) Firefly Mail communications are only accessible by the designated recipients, except that Firefly Network Inc. personnel involved in maintenance of the system may have access to e-mail communications.

7) Recommendations are generated solely based on an end-user's rating preferences and the preferences of other end-users, and are not generated by Firefly Network Inc. or its advertisers.

Some Web sites that use Firefly include barnesandnoble.com, Ziff Davis, and Yahoo! Figure 6.9 shows a page where users of the Barnes and Noble site can sign up to get a passport. After they receive their passport, users can enter titles of books they've enjoyed and books they haven't liked. Based on that information, a user's profile

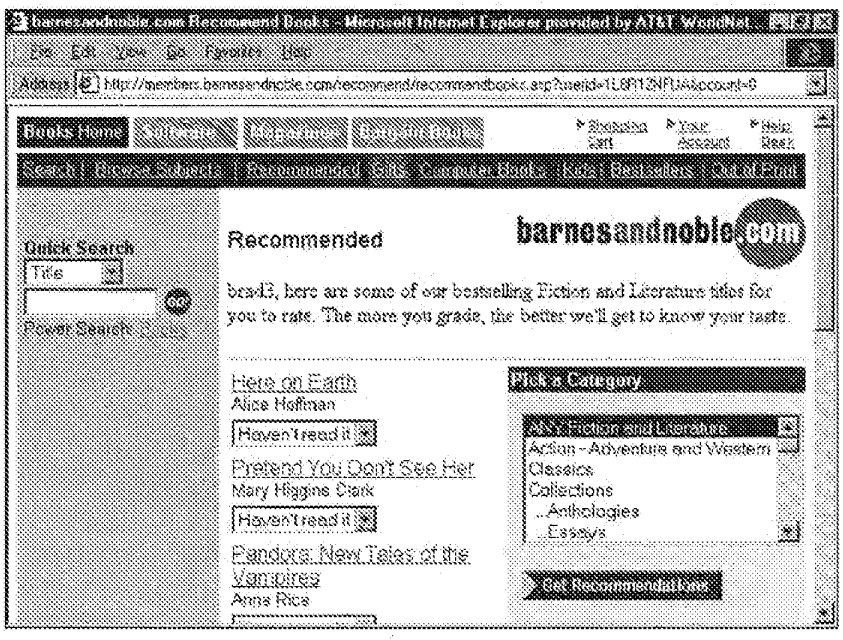


Figure 6.9 Barnesandnoble.com recommendations based on Firefly.

will be compared to those of similar users to determine other books that the user would enjoy.

Privacy is a big concern when it comes to personalization issues. When BOL (the online book-selling venture of Bertelsmann AG) decided to employ Net Perception's personalization technology for its European customer base, it had to address the heightened privacy concerns of European customers. The site developed a unique privacy policy element to allow customers to update their personalized profiles as well as their BOL purchasing history, making the customer master of his profile information. For a full discussion of privacy issues, see Chapter 13.

Keeping Customers Is a Continuous Process

Interestingly enough, companies will develop very sophisticated personalization systems like the ones we mentioned in this chapter, but fail to take the time to respond to customer e-mails. Ken Allard, group director at Jupiter Communications, notes that companies often make sizable investments in personalization technology to derive what users want and then ignore the users when they communicate directly with the company to tell them. In a 1998 study Jupiter found that only 38 percent of major sites respond to customer inquiries within a day (see Figure 6.10). Personalization means being able to reach a person and retrieve an answer, not just fancy profiling systems.

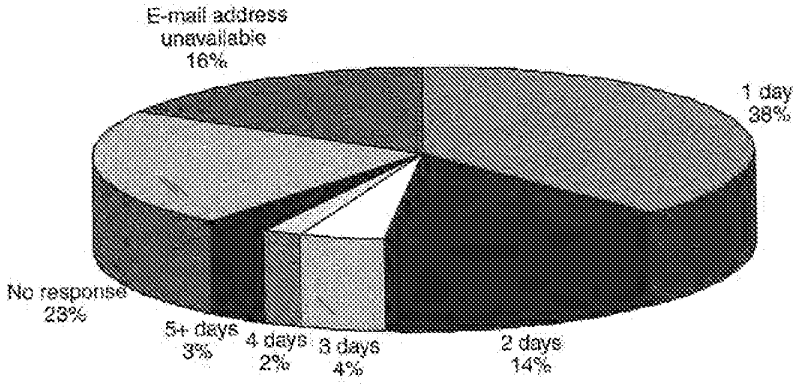


Figure 6.10 Response time to e-mailed questions.

Interactive Marketing: An Unparalleled Opportunity to Reach the Right Customers at the Right Time

Ann Burgraff, Director of Corporate Marketing (ann@netgravity.com)
NetGravity, Inc. (www.netgravity.com)

1998 saw the Internet universe mature dramatically. For online advertising—everywhere, not just in Silicon Valley—there was a revolution. With more and more eyes on the Internet each day, major brands began to recognize the tremendous potential of online advertising to drive online business. Some trends that emerged were partnerships to create standards; initial public offerings; ad agencies taking more control over their online advertising campaigns; and publishers, merchants, and agencies demanding a range of solutions, both service and software, to fit their specific online advertising management needs, and maximize their efforts.

If 1998 was the year of revolution, 1999 will be the year of evolution. The Web has become a 24/7 customer interaction channel capable of delivering a personalized and targeted individual experience for each customer. It is closer to in-person interactivity than any other marketing channel and needs to be treated as such. We see online advertising expanding to tap into this opportunity. Mere online advertising will not be enough. Interactive marketing will be the demand for an even better approach to reach the right people at the right time. Interactive marketing will introduce the need for a whole new set of requirements to meet the needs of publishers, merchants, and ad agencies:

The right fit. It will be key for online marketing solution providers to provide the most appropriate solution for each customer, whether it be a customized solution for a market segment, software, service, or the ability to migrate from software to a service.

Customer intelligence. Offering the means to pinpoint the most profitable audience will be critical in ensuring that our customers' business is successful. It is our job to provide publishers, merchants, and ad agencies with the ability to collect different categories of data, both implicit and explicit, integrate a solution with existing corporate data, and offer the tools to analyze the data and make intelligent decisions based on the results.

The ability to meet the needs of the Global 2000. This means having unlimited scalability, global operations, full range of service, and enterprise-level partners. Interactive marketing will be a critical element of successful Internet strategies.

With interactive marketing maturing, we can begin to see the Internet for what it really is—an unparalleled opportunity to reach the right customers at the right time.

CHAPTER 7

Pricing Online Ads

The Web was born into a dichotomous environment: the Internet culture was founded on free access to information, but its content is extremely expensive to produce, especially if it is expected to rival TV, radio, and print for the public's attention. In the United States, commercial broadcast TV and radio are supported 100 percent by advertising, newspapers 80 percent, magazines 63 percent, and cable TV 50 percent. Naturally, Web publishers are hoping that advertising will cover the lion's share of the costs for their publishing ventures (and ideally even generate some profits). The results look promising—top publishers are beginning to tally significant ad revenues (see Table 1.4), and advertisers are concomitantly increasing their Web advertising budgets each year (see Table 1.3).

When advertising on the Web first began, no one knew how to price. The only known figure was the fee HotWired decided to charge for its first ads in 1994: $30,000 for banner advertising that ran for 12 weeks. "It was a bottom-up approach," recalls Rick Boyce, VP of advertising for HotWired. The figure was derived from the amount of revenue HotWired needed to generate. At that time, there was no pricing model and no precedent, so other sites quoted that same amount.

Modem Media, an interactive agency (now known as Modem Media.Poppe Tyson) based in Norwalk, Connecticut, was one of the first agencies to begin placing banners on Web sites. John Nardone, director of Media and Research Services at Modem Media, was not satisfied with the unaccountable pricing structure, so he developed his own method for calculating an online advertisement's value, which he called the Internet Reach and Involvement Scale. Word spread about Nardone's system, and soon he was getting a dozen calls a day asking him to help determine costs. Nardone remembers, "What I used to tell people was that I'll help you price your ad if you give my client

Meta's Exhibit No. 1004
Page 00129

Zima com an ad for half price. It was a really fun time. We were out on the Internet frontier." Before long, his scale had become the de facto pricing model.

But as more ad models developed, different pricing structures rapidly evolved with them. Currently, pricing models range from flat-fee to paying for the number of times an ad is seen to paying for the number of people who clicked on the ad to paying on a cost-per-lead or cost-per-purchase basis. The purpose of this chapter is to explore all of these models in use today.

Pricing Models

Current pricing models are as varied as the ads themselves. As an industry, Web advertising is looking to other industries for role models; direct parallels cannot always be drawn, but in general, both a mass market and a direct marketing approach apply. The Internet can still be measured based on impression, just as TV, radio, and print can, but many believe that the Internet is actually more applicable to direct marketing, where clicking on an online ad can be likened to opening a direct mail envelope.

CPM

The most dominant pricing model, which has been briefly mentioned throughout this book, is cost per thousand, or CPM. CPM delineates the price charged for displaying an ad one thousand times. This is the unit of measure used in traditional advertising, which is why advertisers and agencies are comfortable with a CPM model.

The calculation to determine CPM is:

- Price ÷ (number of impressions ÷ 1000)

For example, if a site charges $10,000 per banner and guarantees 500,000 impressions, the CPM is $20 ($10,000 divided by 500).

But how do you determine the appropriate CPM for an online ad? Is it $1, $10, $30? Internet advertising CPMs range from $1 all the way up to $200. The price for Web advertising varies across sites, but, in general, similar types of sites are charging rates in the same general range. The types of sites we're going to cover in this chapter include search engines, content sites, online services, and ad networks.

Keep in mind that prices are constantly changing, so the figures you read here should be regarded only as "ballpark" for different sites. Figure 7.1 shows the average CPMs over time from AdKnowledge, a company that maintains a database of sites that accept advertising. As you can see from the chart, ad rates varied significantly during the time period covered by this chart. And, these are gross prices; usually, advertising agencies receive a 15 percent discount off gross prices, paying the *net price*. Publishers are constantly adjusting their rate cards, and prices also fluctuate depending on the number of impressions purchased and over what time period they are delivered. To look at some actual rate cards, refer to information in the accompanying Web site at www.wiley.com/compbooks/zeff.

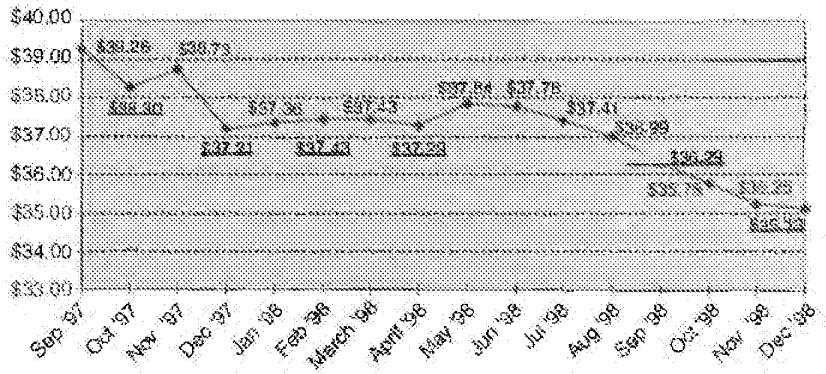


Figure 7.1 Average CPM rate chart from AdKnowledge.
Source: AdKnowledge, Inc.

Search Engines

Search engines and directories are receiving a large portion of advertisers' budgets; in fact, according to Jupiter Communications (see Table 1.4), the top four Web publishers by ad revenue are search sites. It is easy for advertisers to spend a large chunk of money on a high-traffic search engine rather than spreading the ad budget among many smaller sites that will deliver the same audience. The ad buyer only has to deal with one person instead of negotiating with multiple sites, sending ads to each of the sites, and so on. Search engines offer several pricing categories, each distinguished by its level of targetability.

General rotation. A banner will randomly appear during searches. General rotation prices usually range from a $15 to $25 CPM. On buys over $50K advertisers can sometimes get rates in the single digits.

Run of category. Advertising will rotate throughout a category such as arts, entertainment, or sports. Yahoo! (www.yahoo.com) charges a $20 CPM for run of category with a minimum of 500,000 impressions. It is interesting to note that Yahoo! is one of the few sites that does not offer a discount to advertising agencies.

Keyword targeting. Search engines also allow advertisers to "buy keywords"; this means the banner appears only when a user searches using specific keywords. Because the audience is more highly targeted, the CPM is higher. For example, the search engine Excite (www.excite.com) charges a $24 CPM for general rotation, a $60 CPM for advertising in a specific Excite channel, and a $70 CPM for keyword targeting. Keywords are in limited supply (unlike general rotation impressions) and in high demand by advertisers. Therefore, many of the search engines charge an additional fee to have exclusive rights to a keyword. (For example, at the time of this writing, the search engine Snap! (www.snap.com) charges a $1,000 fee for exclusivity for a keyword.) In many instances, the exclusive right to a keyword also means that the advertiser has the first right to purchase future inventory of that keyword.

As we mentioned under "Text Links" in Chapter 2, the search engine Lycos (www.lycos.com) has created additional keyword inventory by charging for the banner impression as well as a text impression; this is an ad placement that they recently developed. Because of the high demand for keywords, we think keyword prices will continue to rise faster than any other ad inventory (Excite is already at a $70 CPM), and we will see more sites following Lycos's lead and developing additional keyword type placements.

Online Services

AOL offers a variety of CPM rates based on ad placement. At AOL, rates for run of service are a $30 CPM (250,000 minimum impressions). AOL sells much of its advertising inventory in packages. One type of package they offer is an industry package—an ad buy that they recommend for a specific industry. For example, AOL has packages for parenting, real estate, shopping, and sports. The price for these packages ranges from a $13 to a $17 CPM, with costs from $9,000 to $24,000 for one month. These buys will include some non-targeted placements. AOL also has targeting of specific site areas like movie reviews, record labels, and game shows online. The CPM for these buys is about $60.

Content Sites

Content site advertising typically costs more than search engine general rotation because the audience is more targeted; users are interested in the specific content that is presented. For instance, ESPN's SportsZone (www.sportszone.com) sells a promotional package that costs $100,000 for 1.2 million impressions per month ($83.33 CPM). A targeted sports package is $250,000 and includes sponsorship of the area of choice and a guaranteed delivery of 1 million impressions ($25 CPM). A general rotation through SportsZone costs $15,000 per month for 500,000 impressions ($30 CPM). Content sites generally charge a $20–$45 CPMs. However, if the content is of a general nature, or it is not high quality, CPMs could drop below that price point. And, ad sites that are extremely targeted can charge more than the $45. An example is ClickZ, a site targeted toward media buyers, which charges up to a $70 CPM.

Ad Networks

DoubleClick (www.doubleclick.net) is an example of an advertising network (see Table 10.2 for a list of ad networks). A media buyer can make one purchase with DoubleClick and have an ad placed on dozens of sites. For example, an advertiser who purchases impressions in DoubleClick's Sports/Leisure/Entertainment affinity would be able to run banners on the various sites that DoubleClick represents in this category. Put another way, a network gives an advertiser much greater reach than most individual sites. The CPMs for the DoubleClick network run from $30 to $70 depending upon the category. Technology is at the high end, with a $70 CPM, and the entertainment affinity is at the low end, with a $30 CPM. For an additional price, advertisers can also target by geography, domain name, SIC code, and company name and size.

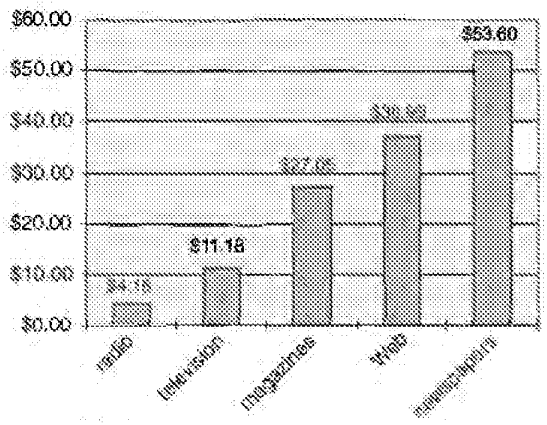


Figure 7.2 CPMs across media.
Source: AdKnowledge, Inc.

CPM Comparisons

How do these prices compare to the CPMs of traditional media? Take a look at Figure 7.2, and you'll see that Web CPMs tend to be higher. There are two justifications for this: not only is this is a new industry with no fixed market rate, but more importantly online advertising allows for a sophisticated level of targeting as demonstrated in Chapter 6. On the Web, advertisers can target by keyword, SIC code, company, geographic location, and time of day. Advertisers also have opportunities to target by preferences. Indeed, prices will continue to fluctuate as new targeting opportunities and new advertising units are developed.

Click-Through

The term *click-through* refers to the action of a viewer clicking on an ad. Paying per click-through is one of the pricing models derived from a direct marketing perspective.

Advertisers and publishers began debating the merits of the click-through model when, in late spring 1996, Procter & Gamble (P&G) negotiated with Yahoo! to purchase ads for five P&G brands on a click-through basis. This was a big move because it negated the value of Web advertising for brand building and put all the weight in Web advertising's direct marketing appeal. In effect, P&G was saying two things: first, that there was no value to simply having its banners seen—no brand-building value; they were putting a premium only on getting clicks. And second, that the burden of getting a banner clicked was on the publisher as much as on the creative. The online advertising world took notice because P&G is the nation's largest advertiser, and prior to its move, most pricing was CPM-based. Furthermore, the model held publishers accountable for response, but the publishers weren't producing the creative; and wasn't there some branding value in having a user see a banner?

According to John Nardone of Modem Media, Yahoo! did not simply agree to P&G's new model of paying per click-through without some concessions on P&G's part as well. P&G agreed to give Yahoo! input and veto power on the creative P&G submitted. Since Yahoo! was responsible for click-through, it wanted to be able to modify the banner ads if necessary to ensure their success. And by this time, Yahoo! knew a great deal about what worked creatively in a banner ad.

The ability to accurately measure click-through response to Web banner advertising has led many advertisers to push for pricing models based on these performance measures. The benefits of click-through are clear for the advertiser. The advertiser knows from the start what the Internet advertising campaigns will cost per click. "A saving grace for our small-site online advertising has been cost per click model," says Scott Sizemore of New Media Consulting. "Without arguing the merits and drawbacks between CPC [cost per click] and CPM, it has been our experience that many clients feel more comfortable beginning paid advertising campaigns in the Cost Per Click arena."

The benefit to publishers using click-through pricing is not always so clear. Some say that the advantage is that a pay-per-click model entices advertisers who otherwise wouldn't place ads on their site. However, the concept of paying per click-through has met with harsh criticism by many publishers for several reasons:

- A click-through pricing structure requires publishers to rely on the quality of advertisers' creatives to generate revenues. Traditionally, the role of a medium is to offer access to an audience, not to share in the responsibility for the quality of the advertisement itself. Since many advertisers are new to the medium, the publishers are being asked to gamble on the advertisers' learning curve, a burden Web publishers are reluctant to take on.
- This model encourages a publisher to run a banner until the advertisement receives the number of clicks an advertiser purchased. This could result in a publisher compromising its inventory and ultimately the user's experience on the site.
- This model completely dismisses a banner as having brand-building value.

The pricing for cost per click can run anywhere from $2 to five cents. As an advertiser, it's nice to buy on a cost-per-click basis because you have some level of security. However, advertisers need to remember that a click probably isn't the real goal, so they've got to come up with a metric to determine the quality of those clicks. For publishers it's important to maintain the right to veto creatives for cost-per-click campaigns. Publishers will know what works best on their site, and there is no reason to accept a creative that will have a low response rate for a cost-per-click campaign.

For those advertisers who are interested, there are some Web properties that generally accept advertising on a cost-per-click basis. One is the ad network ValueClick (www.valueclick.com). ValueClick has more than 10,000 member sites, and sites are generally paid fifteen cents per click by ValueClick (which of course then charges advertisers a higher rate). Advertisers can target based on affinities to better qualify their prospects (for example, targeting sports sites that are members of the network). ClickTrade (www.clicktrade.com) is a program in which advertisers can post the cost-per-

click advertising for which they would be willing to pay. Then sites that are part of the ClickTrade network can decide if they want to accept that deal. ClickTrade manages the process and charges a 30% fee on top of the price set by the advertiser. ClickTrade also offers cost-per-order and cost-per-lead advertising options. We discuss ClickTrade in more detail later in this chapter under "Affiliate Programs." Figure 7.3 is an offer as it would appear to a ClickTrade member site deciding what advertisements to accept.

GoTo.com (www.goto.com) is a pretty interesting cost-per-click search engine. Advertisers can buy keyword advertising on a cost-per-click basis. The advertiser who pays the most has her listing appear first. The interesting thing is that next to each search engine result, GoTo.com lists the amounts that the various advertisers have offered to pay per click to come up first. So, anyone using the search engine can see who is paying to come up early in the search results and how much they are paying. Figure 7.4 shows the search results on GoTo.com for the phrase "Star Trek."

There are many cost-per-click opportunities on the Web, but it's not always easy to find them. Mark Welch's Web Site Banner Advertising site (www.markwelch.com/bannerad/baf_click.htm) has a page where many cost-per-click opportunities are listed and evaluated (see Figure 7.5).

Flat Fee

Some publishers are charging a flat fee for advertising—a set amount charged either monthly or yearly for ad placement on a site. Thus, the fee is tied to a time period, not impressions or visitor activity. Flat-fee pricing is used in several circumstances.

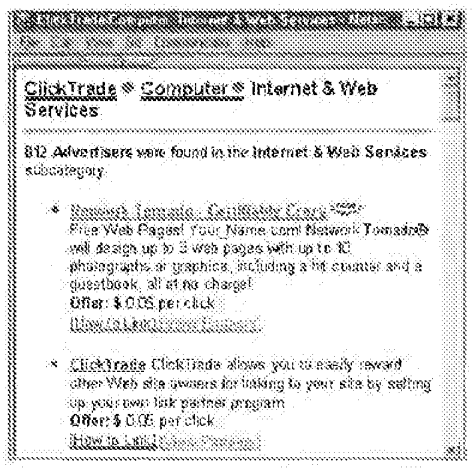


Figure 7.3 An offer as seen by a ClickTrade member site.

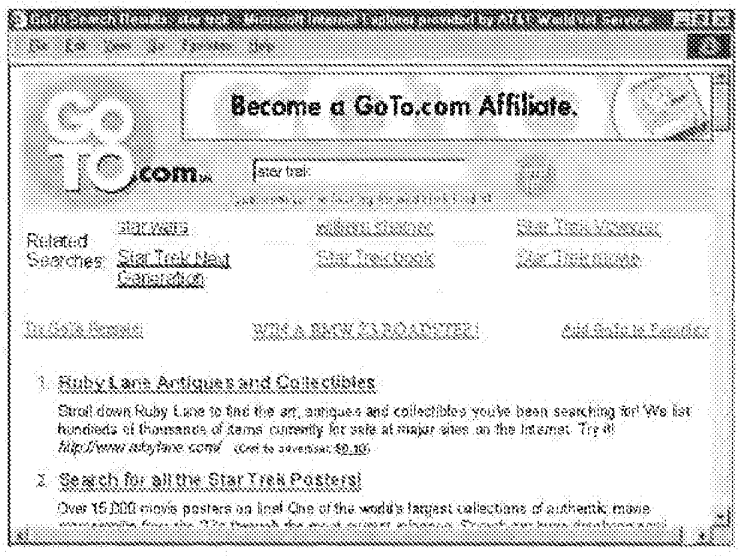


Figure 7.4 GoTo.com search results show users how much the advertiser pays for each click.

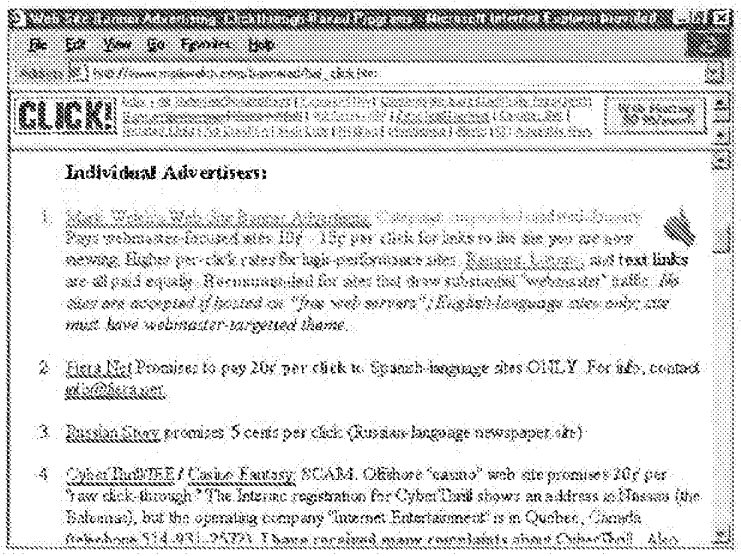


Figure 7.5 A page from Mark Welch's Web site on banner advertising.

Meta's Exhibit No. 1004
Page 00136

Seed money. Flat fees are often used when a publisher needs seed money, in which case the publisher often determines rates based on the amount of money needed. Clearly, this was the reason HotWired charged what it did for its first ads back in 1994. PriceSCAN (www.pricescan.com), a search engine that finds the lowest prices on goods, will often add new categories of products when an advertiser comes to them and offers sponsorship money in return for creating the content. Since PriceSCAN does not know the traffic of the category they are adding, they charge the sponsor an amount that will cover the cost and that they think will be reasonable given past traffic patterns. If the sponsor had not come forward, that specific category would not have been placed in the front of the development schedule.

Sites that are developing new content areas also charge a flat fee for seed money. These properties need to sell sponsorships before they know the amount of traffic the site will generate. Advertisers run the risk of paying a lot of money for a low volume of traffic, but there is also the potential to receive much more traffic than anticipated. Also, if the advertiser has a good relationship with the publisher that is creating the new content area, the advertiser can often get make-goods if the new property underdelivers on expectations.

Low pricing. The Year 2000 site (www.year2000.com), which covers the information systems problems that are predicted to arise when we reach the year 2000, is commercially supported by advertising that is bought on a flat fee. On the Year 2000 site, a large number of sites [illegible] the site is designed in frames, and each advertiser is given a text link on the home page in a sponsorship frame; the 100 advertisers are accessible on the home page, as shown in Figure 7.6. The Year 2000

REF. 85080 OED



Figure 7.6 The Year 2000 site uses flat rate pricing.

site has three packages, which range from $2,000 to $6,000 per year, depending largely on the type of link provided. The full package, which is the most popular, offers a link directly to the advertiser's home page. For clients without a home page or on a limited budget, there are less expensive options that offer links to mini-pages with an e-mail link. Charging a yearly fee reduces the time the site needs to administer the accounts. The effort to deliver traffic reports, renew, and bill advertisers on anything less than a yearly basis is not worthwhile when the advertising rate is so low (it would only be $166 to $500 per month). By implementing pricing policies in a manner that precludes significant administrative labor, the Year 2000 site can also charge less money. "We have been on solid financial ground since day one," reports Cliff Kurtzman, CEO of Tenagra, publisher of the site. The Year 2000 format generates a high volume of advertisers at a low fixed rate.

Premium placements. Flat-fee pricing is also used for premium placements on a particular Web site. For example, Netscape charges a search engine a flat fee to have a search listing on its popular search page. The amount, a dollar value over $5 million (both cash and barter), is now the norm for such high-profile premium placement opportunities.

Pay Per Viewer

Another model is to charge advertisers every time a user interacts with their brand. The new services that sell this type of advertising offer users money and other incentives, often similar to frequent flyer miles, to read advertisements and participate in interactive experiences with the advertiser's brand. Users can get paid for everything from visiting advertiser sites to filling out surveys to purchasing products. The philosophy is that the most valuable commodity on the Web is a person's attention and that the best way to acquire that is through direct incentives.

Pay Per Purchase

Taking the payment-per-response model one step further, some advertisers have negotiated to pay only when purchases result from their advertising. This places publishers at the mercy of both the ad creatives supplied and the consumers' interest in purchasing a particular product. This arrangement can be favorable to both the advertiser and the publisher in a few situations when the item is highly targeted, but in most cases, this model clearly benefits the advertiser. The publisher assumes all risk for the advertiser's campaign, and there is no value placed on the brand exposure that advertisers receive through ads that aren't clicked. But, if the advertiser has to choose between no ad revenue and pay-per-sale ad revenue, the pay-per sale model isn't all that bad.

In 1996, Orb Communications & Marketing, an interactive agency in New York, negotiated one of the first large-scale pay-per-purchase deals, for Nabisco Direct. The arrangement placed advertising for four Nabisco holiday gifts: Oreo and Mr. Peanut cookie jars, a holiday gift basket, and chocolate-covered Oreos. The advertising ran on such sites as Excite (www.excite.com), as well as on entertainment sites like MovieLink (www.movielink.com), Comedy Central (www.comcentral.com), and Entertainment

Meta's Exhibit No. 1004
Page 00138

Drive (www.edrive.com). Nabisco paid only a percentage of sales generated by each publisher's site.

Publishers are reluctant to accept this type of deal because they feel that it underestimates the brand-building capabilities of advertising on the Web, but some sites will accept a pay-per-purchase negotiated price.

Affiliate Programs

Another variation of this model is paying per sale through affiliate programs. Affiliate programs go by numerous names, including *partner* and *associates* programs, but they are most commonly referred to as affiliate programs. The typical affiliate program is run by an advertiser who would like publishers to run advertising on a cost per-action basis. Instead of a usual ad buy that is negotiated on a site-by-site basis, most affiliate programs are automated. A publisher will go to a merchant site, fill out a form to become an affiliate, and then have access to instructions for how to set up links, ads, or a virtual storefront that points toward the merchant. Figure 7.7 shows a sign-up page for an affiliate program. The merchant tracks the success of the advertising, and pays the participating publishers accordingly. The process is automated, so that the advertiser can have thousands of sites promoting his business, without having to invest the resources in negotiating with each site. To publicize affiliate programs, advertisers will often run an advertising campaign. They will also list their affiliate programs on the sites that have directories of these opportunities (see Table 7.1).

REF. 85080-OED

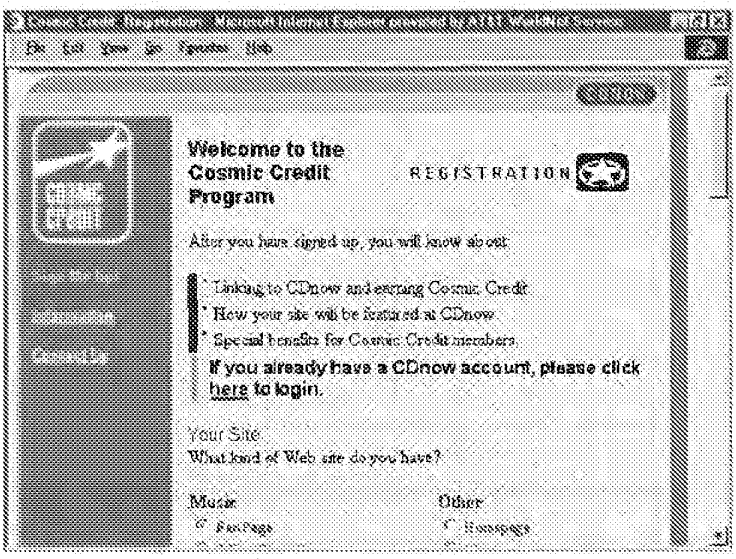


Figure 7.7 A sign-up page for the Cosmic Credit affiliate program.

Members of an affiliate program most often get paid a percentage of sales. This is a great way for advertisers to extend their reach with very little expense; advertisers only pay for orders that are generated. Also, the products can be sold within context. For example, a newspaper that belongs to Amazon.com's affiliate program can allow consumers to purchase a book right after they read the review. The newspaper will receive a percentage of the sale, and Amazon will probably successfully sell books through that type of contextual advertising.

"Affiliate marketing is simply direct marketing. In the cyberworld it allows marketers to be much more one-to-one, and products are marketed in context," says Gordon Hoffstein, COO of affiliate software maker Be Free. "There's no overhead. If it doesn't work, you don't pay."

However, advertisers must be aware that when publishing sites become affiliates, they can start to play a hand in promotion. The affiliate will try to promote the advertiser's store or products by running various ads. Instead of a logical link from a book review to purchasing the book, a site may place Amazon.com logos all over their site, and run Amazon.com banners that they create, to help generate more revenue. Since most merchants view sales as their primary objective, this isn't usually a problem, but the merchants must be prepared to occasionally have their brand advertised in a manner that they wouldn't choose themselves. As we mentioned in the newspaper example, affiliate programs usually go beyond a publisher simply running adver-



Figure 7.8 Amazon.com affiliate store within the Women.com site.

Table 7.1 Affiliate Program Resources

SITE	DESCRIPTION
ChannelSeven's Ad-Guide (www.ad-guide.com/Affiliate_Programs/)	Directory of advertising affiliate programs
Mark Welch's Web Site Banner Advertising (www.markwelch.com/bannerad/baf_commission.htm)	An annotated list of all affiliate programs, both domestic and international
ClickQuick (www.clickquick.com)	A listing of affiliate and pay-per-click programs
Associate Programs.com (www.associateprograms.com)	Directory of associate programs
Revenue Avenue (http://revenue.linkexchange.com)	Listing of affiliate programs

tisements. Publishers often create their own store within their site for the merchant. For example, Figure 7.8 shows the Women.com site, an affiliate of Amazon.com. Women.com built a store within their site. As you can see, the promotion is especially effective, because the advertisements are for books that are going to be of interest to Women.com viewers. Also, the site has the Women.com brand on it serving as a direct endorsement of the books.

The first granddaddy of all affiliate programs (and of the very first) is online bookseller Amazon.com. They offer 15 percent of sales on most books and 5 percent of sales on out-of-print books and other items. Like most affiliate programs, Amazon.com provides the graphics and technology for publishers to set up their own bookstores. Information on Amazon.com's affiliate program is available at www.amazon.com (click on the "join associates" link).

Some affiliate programs pay on a sliding scale. Bookseller Barnes and Noble (www.barnesandnoble.com) pays commissions between 3 and 7 percent depending on the amount of sales an affiliate site generates. Table 7.2 provides some examples of affiliate programs. A list of affiliate programs resource sites is available in Table 7.1.

We've listed some of the larger affiliate programs in Table 7.2, but there are literally thousands of affiliate programs with all sizes of sites and all types of product offerings.

There are hundreds of topic-specific affiliate programs: everything from Lobsternet (www.thelobsternet.com) (see Figure 7.9) to Viagra sales at KwikMed (http://kwikmed.com/viagra/101985/affiliate.html) (see Figure 7.10). For example, CarPrices (www.CarPrices.com/join.html) has an affiliate program for cars price quote requests, insurance quote requests, and other inquiries. They pay between $1 and $12.50 per information request or information purchase. A complete breakdown of CarPrices' pricing is given in Table 7.3.

Table 3.2 Sample of Affiliate Programs

COMPANY	DESCRIPTION	COMPENSATION PROGRAM (PERCENTAGES OF SALES UNLESS OTHERWISE SPECIFIED)
Amazon.com (www.amazon.com)	You can build a virtual bookstore based on each section, and recommend books, music, and magazines.	5%–15%
(www.barnesandnoble.com)	You can build a virtual bookstore based on each section, and recommend books, music, and magazines.	3%–7%
beyond.com (www.beyond.com)	An online software superstore. They have a large section of educational software to which you can link and recommend specific titles for specific curriculums, or even offer a "recommended tutorial" list for parents.	5%–10%
ClickTrade (division of LinkExchange) (clicktrade.linkexchange.com)	Offers a unique program where you can actually get paid for simply sending visitors to another site. ClickTrade will serve up the ads of the partners you choose and track how many visitors from your Web site clicked on the ad. They also pay on a per-lead and per-sale basis.	Varies depending on the company, ranging from $0.01 per click to $0.25 per sale or click.
Cyberian Outpost (www.outpost.com)	"The world's largest online computer store." They sell hardware, software, and associated technologies.	Terms are discussed after they review the site and are assured of a demographic fit.
Etoys (www.etoys.com)	Toys, toys, and more toys. This online toy store is growing in popularity.	They pay a huge 25% commission.
FileSafe (www.filesafe.net)	Online filing of your IRS returns for a low price. The site is known to be very safe and secure.	Variety of plans, $4.00 and up for referrals.

REF. 85080 OED

Meta's Exhibit No. 1004
Page 00142

Table 7.2 (Continued)

COMPANY	DESCRIPTION	COMPENSATION PROGRAM (PERCENTAGES OF SALES UNLESS OTHERWISE SPECIFIED)
iPrint (www.iprint.com)	Online ordering for stationery, business cards, and printed novelty items. Lower cost than many local quick print stores, and you can do some amazing fine-tuning of the images and text online.	4%–8%
LandscapeUSA (www.landscapeusa.com)	Grow your garden online? Not exactly, but everything you need to plan, and plant, and grow your garden can be purchased here.	8% of total sale (average order is $65, some over $500, according to the site)
MapQuest Map Store (www.mapstore.com)	The people who developed the most popular online map system also sell maps of all shapes and types, including USGS topographic maps, historic maps, and maps of planets and countries.	8% of sales during clicked session (program in transition)
Music Boulevard (www.musicblvd.com)	They sell . . . music. A large site dedicated to selling CDs online. Since Amazon also sells books, you may want to focus your efforts on the one-stop shopping of Amazon, but this could be a good alternate resource if there are not contract conflicts.	5%–15% depending on the method by which you link them, and the item itself
Reel.com (www.reel.com)	This site sells and rents movies.	8% of gross revenue (probably more likely on net revenue after they remove credit card fees, shipping, etc.)
Seattle's Finest Gourmet Coffee (www.seattlesfinest.com)	Buy Seattle's Finest coffee online.	20%

Continues

Meta's Exhibit No. 1004
Page 00143

Table 7.2 Sample of Affiliate Programs *(Continued)*

COMPANY	DESCRIPTION	COMPENSATION PROGRAM (PERCENTAGES OF SALES UNLESS OTHERWISE SPECIFIED)
SportsSite.com (www.sportsite.com)	An online superstore of sporting goods.	20% (program in transition)
Wall Street Journal Online (www.wsj.com)	Offers some special promotional programs where they will provide you with access to news elements to keep your site up to date daily, while also providing a revenue opportunity to you for referring customers to them for their subscription service.	$5.00 for every trial subscription you generate that results in a subscription
Ziff Davis University (ZDU) (www.zdu.com)	This online distance learning service focuses on high technology and other training programs. Some courses are independent training programs and others are instructor-led.	10% referral fee for those who sign up for one year of unlimited training

REF. 85080 OED

Source: 1998 InternetAdvocate.com (www.internetadvocate.com)

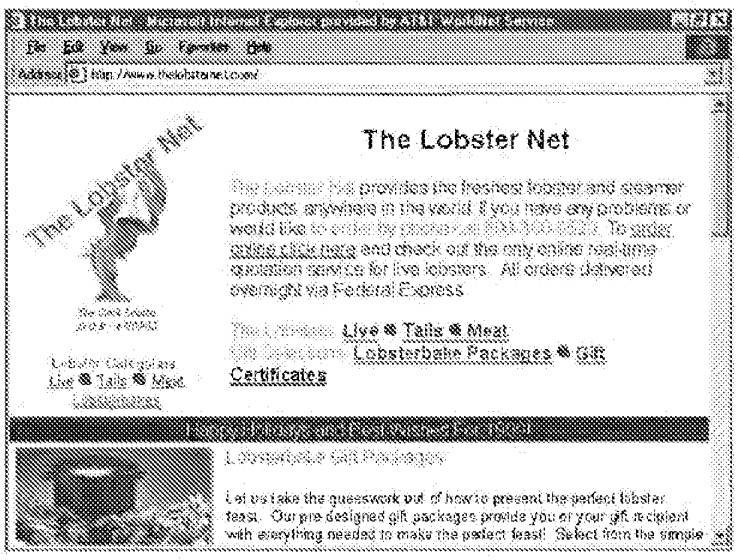


Figure 7.9 Lobster Net is one of the many affiliate programs available to publishers.

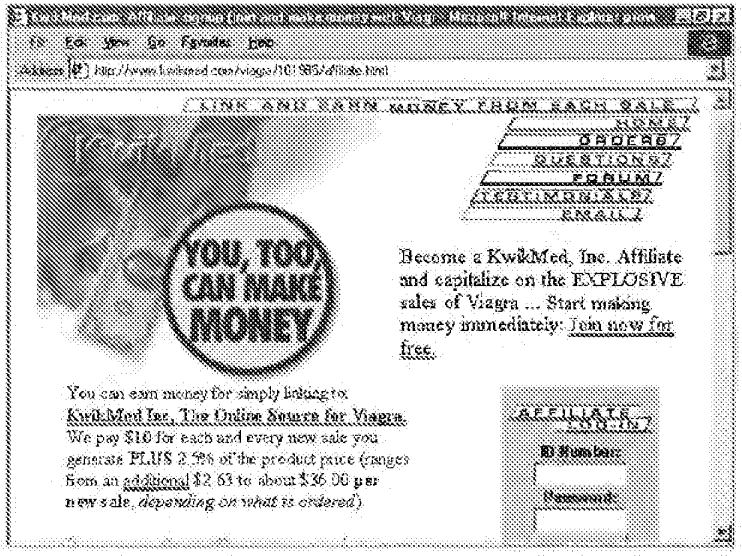


Figure 7.10 The screen to sign up for the KwikMed affiliate program for Viagra. The program pays $10 per order and 2.5% of the product price for a new sale.

Managing an Affiliate Program

To implement an affiliate program, publishers need software that (in addition to other capabilities) will manage the process of tracking orders, leads, or clicks back to a specific referring Web site. The ideal software package will also calculate how much each site is owed, handle the addition and subtraction of affiliate

Table 7.3 CarPrices Affiliate Payment Rates

ITEM	COST TO USERS	YOUR COMMISSION
Insurance Quote	Free	$1
Finance Quote	Free	$1
New Car Price Quote	Free	$3
Used Car Purchase Request	Free	$0.80
Personal Credit Check	$10.95	$1
Warranty Purchase	Varies	$12.50
Car Title Check	$12.50	$1
Custom Window Sticker	$12.95	$0.75
Car Secrets Book	$27.95	$4

sites, and other tasks associated with running an affiliate program. There are several solutions for managing an affiliate program outlined below. A comparison of features is provided in Table 7.4.

Some folks use home-brewed solutions. Unless you have a large technical staff or an ad agency that can handle this type of development, you're better off using packaged software. It's almost always cheaper, and it will allow you to focus on getting orders rather than on software development.

ClickTrade (www.clicktrade.com) (see Figure 7.11) is an affiliate manager that matches buyers and sellers of Web advertising (ClickTrade is part of the LinkExchange, now owned by Microsoft). Web advertisers can place cost-per-click, cost-per-lead, or cost-per-sale advertising through this affiliate program. An advertiser begins by signing up at the ClickTrade Web site and then entering information about the affiliate offer. For example, a merchant might list $10 per sale as the amount she would want to pay for the advertising. ClickTrade charges a 30 percent commission, so in addition to paying the sites a $10 fee for every sale, she'd pay a $3 fee to ClickTrade. ClickTrade lists the offer in their database of ad offers. Sites go through the database and choose affiliate programs in which they want to participate. If the advertiser chooses, the system can be set up so that after the site chooses a specific advertiser, the advertiser has to approve the site. ClickTrade tracks the clicks, orders, or leads. The advertiser pays ClickTrade in advance, and ClickTrade pays the sites. There is no sign-up fee for ClickTrade's affiliate program. They do ask that advertisers make a minimum $100 deposit for starting a pay-per-click program and $500 if they are starting a pay-per-lead/sale program. ClickTrade has affiliates generate revenue through banners, buttons, and text links that drive traffic to the merchant's site. The advertiser can send an e-mail to publishers (who choose to receive it) through ClickTrade's interface; the advertiser does not have direct contact with the publishers.

All ClickTrade advertisers automatically get featured in Revenue Avenue (http://revenue.linkexchange.com), a directory of revenue-generating offers from Web merchants. In addition, ClickTrade delivers information about the offer to over 30,000 sites who participate in their affiliate programs.

If a Web site wants to participate in an affiliate program run by ClickTrade, the site will sign up with ClickTrade and then be able to choose the affiliate program that is of interest. Participation entails placing merchant banners, buttons, and text links on their site.

LinkShare (www.linkshare.com) manages the entire process for sites and their affiliates. They have software that merchants install on their e-commerce engine. As with ClickTrade, merchants can benefit from having their affiliate program broadcast to LinkShare's network of affiliate sites. Merchants who use LinkShare pay a $5,000 setup fee and 2–3 percent of revenues from sales. The merchant sets up various stores from which the affiliate can choose. The affiliate can also choose a specific product, if there is one that the affiliate decides is best suited for his audience. This allows the merchant some control over how the affiliate's store will look. This differs from ClickTrade in that affiliates can set up complete stores, instead of using only banners, buttons, and text links.

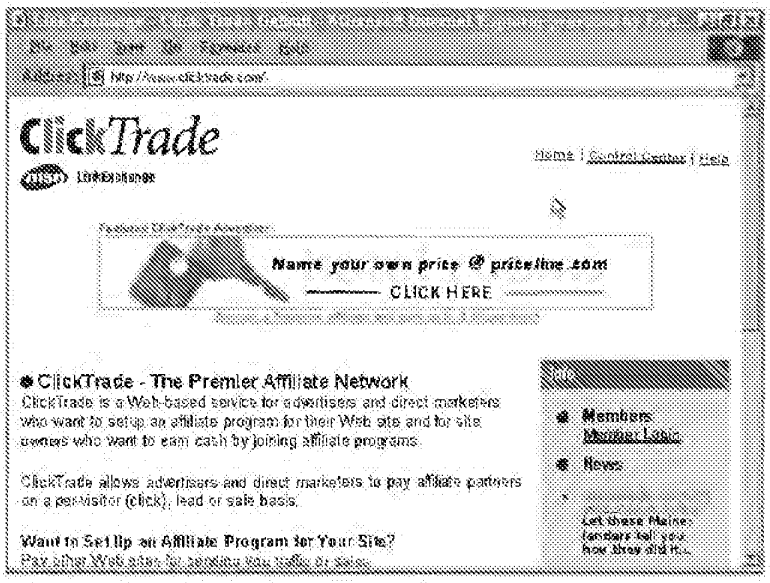


Figure 7.11 The ClickTrade affiliate program network

LinkShare's clients range from the largest e-commerce retailers like 1-800-FLOWERS (www.800flowers.com) to smaller retailers like Bulldog Store (www.bulldogstore.com), which sells products for bulldog owners. The advantage for merchants using LinkShare is that they are able to control how the Web sees their brand. The merchant, and not the affiliates, serves the banners and whichever store the affiliate chooses through LinkShare's service. This ensures that the merchant can remove old creatives and offers from affiliate sites. When a merchant signs up affiliates, the affiliates join the LinkShare network and then they can join the merchant's program. This means that as with ClickTrade, the merchant's affiliates will be exposed to many offers.

Be Free (www.befree.com) manages the enrolling, supporting, and tracking of affiliate networks. Be Free functions a little differently than LinkShare and ClickTrade: when an affiliate joins the network, that affiliate can either sign up through the merchant or through Be Free. If the affiliate signs up through the merchant, Be Free receives no information regarding that affiliate. Be Free generally charges 2 percent of affiliate sales revenue for these services, and focuses on large and medium-sized commerce sites. Clients include barnesandnoble.com, American Greetings, and GeoCities.

Be Free lists two components to its affiliate software: Be Free Affiliate Serving Technology (BFAST) and Be Free Intelligent Targeting (BFIT). BFAST manages the process of enrolling, supporting, tracking, and managing affiliate networks. BFIT is an engine that tracks sales from the advertisement to the point of purchase. If merchants want, they can also pay Be Free to manage the payment of affiliate sites.

Table 7.4 Comparison of Affiliate Vendor Features

	BE FREE	LINKSHARE
Key Clients	barnesandnoble.com, Travelocity, Reel.com	Cyberian Outpost, FAO Schwartz, Clinique
Number of Clients	30	100
Setup Fee	$5,000 to $7,500 for one-year contract; free of charge for three-year contract	$5,000
Percentage of Transaction	2–3 %	2–3%
Affiliate Outreach	Merchants receive 5,000 affiliates free of charge; fee of $14 for each additional affiliate Be Free generates for you.	6,326 affiliates; merchants have free access to all affiliates
Allows CPM Payment	Yes	Yes
Allows Click-through Payment	Yes	Yes
Allow Bounty Payment	Yes	Yes
Allow Loyalty Points (Click Rewards, My Points) Payment	No, but in development	Yes
Allow Hybrid Payment	Yes	Yes

Source: 1998 Jupiter Communications (www.jup.com)

With hundreds of affiliate programs out there, it is difficult to determine the right price to pay affiliates. Merchants generally give their affiliates a 5–15 percent revenue share. If there are a lot of merchants in your category, you will want to offer a higher revenue share to encourage sites to join your affiliate program instead of your competitions' affiliate programs. If you're a merchant of a very high-priced product (like a car), you might want to give a flat fee per car sold rather than a commission.

Auctions

Auctions are big on the Internet and are a good place to buy and sell excess ad inventory. Because the auctions are selling excess inventory, advertisers can usually get pretty good prices on well-recognized sites. For publishers it's a way to generate revenue on inventory that would otherwise go unsold.

The leader in this area is Adauction.com (www.adauction.com; see Figure 7.12). "Adauction.com is building a Web-based exchange for buying and selling media," says David

Tips for Running a Good Affiliate Program

Research pricing of other merchants in your market and be competitive. If you don't pay enough money, you won't have a high number of respondents.

Payments should be sent in a timely manner. When affiliates receive checks, they remember you and work harder to promote your offer. Remember, they can quadruple the sales they generate for you by giving you better positioning on their site. They often just need to know that they'll get paid.

Provide an area for affiliates to download advertisements that they can use to promote your merchandise. You know what works to help sell your products. And, the affiliates may not run their own advertisements to promote the special offer because they don't have time to design the ads. If you provide the ads, affiliates are usually happy to run them in an effort to generate more revenue.

Help your affiliates sell. In addition to providing graphics, work with them to create promotions and to teach them what helps sell your products. This can be done through an e-mail newsletter that reaches all of your affiliates, or a special Web site. Barnesandnoble.com has set up a Web site to help affiliates succeed, at www.affiliate.net.

Give affiliates a way to check the status of the revenue they've generated through online reporting. It is a nice bonus.

Wamsley, CEO and president. It's "an online marketplace where price is determined dynamically based on a real-time supply and demand model." Adauction.com:

- Creates an alternate channel for media companies to move additional inventory without conflicting with direct sales.
- Evolves from auction to exchange—rather than participating in periodic events, advertisers around the globe can always find quality inventory at the best price possible.
- Offers just-in-time buy opportunities based on selling media days before it expires.
- Attracts buyers—large and small—with a valuable price proposition.

Some of the sites that auction inventory through this service are Yahoo!, Netscape, E! Online, Jumbo, and TUCOWS. Advertisers who purchase inventory have a choice of serving ads through AdForce or DoubleClick's DART. As of the writing of this edition, Adauction.com was auctioning CPM buys and had plans to start offering cost-per-click, sponsorship, and package buys. Adauction.com gives sites a means to sell excess inventory; they sell ad campaigns that begin the month after the auction (if the auction is in January, the flight is in February).

"Adauction.com has earned high marks for uniquely offering much-needed help to Web publishers in optimizing banner inventory. We are committed to using their service

to generate revenue which we'd otherwise forfeit," said Ruth Chang, of Student Advantage Network (www.studentadvantage.com).

There are three ways to buy at Adauction.com:

Tune-In. Topic-specific auctions (finance, women's interests, travel, computers, etc.) two or three times each week. This is targeted inventory, featuring exclusive packages and special inventory for the target- and the brand-focused buyer.

MarketPlace. Features 15 categories and over 100 million impressions for sale, held the third week of each month. Great values on ROS inventory.

LastCall. For the buyer looking for plenty of eyeballs. A network buy at low CPMs. The biggest bang for the buck for broad reach and frequency.

The Internet ad auction has gone so well that Adauction.com is launching a similar trading platform for auctioning ad space for print, radio, cable, and infomercials.

Real-Time Ad Buying Through Networks

Auctions are not the only means for real-time buying. Flycast (www.flycast.com), an ad network with over 700 sites that reaches more than 12 million people each month, also sells ads in real time. Unlike ad networks like DoubleClick, Flycast does not focus on branded high-end CPM sites. Instead, they've amassed a large number of sites, and they're focused on selling this low-priced inventory to direct marketers. Buys on the Flycast network fall into four categories: Run of Network (all sites), Run of Category (specific affinity), Category Select (affinity and additional targeting like geographic area), or Site Select (choosing a specific site).

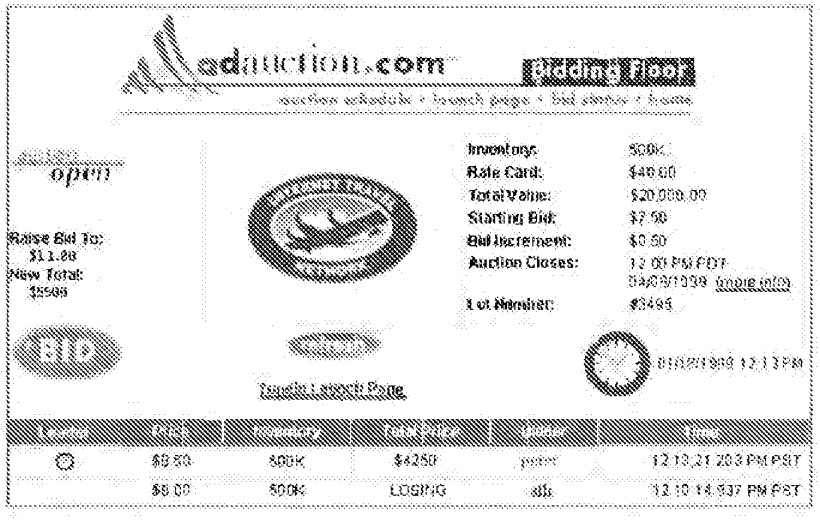


Figure 7.12 The Adauction.com bidding floor.

The Flycast philosophy, according to their Web site, is as follows:

> We begin with what we call zero-based media planning. This is where we use the actual results from a campaign, run without any targeting considerations, to create a response baseline. This baseline is used to determine where on the Flycast Network the ad campaign should be focused. The results from this first level of optimization form a behavioral targeting model. These real-time advertising responses provide our media consultants with the behavioral results to continue to optimize as needed to maximize the success of your campaign.

The run of network rate card price is a $6 CPM, and there are additional costs as advertisers choose more targeting options. Flycast's proprietary ValueTrack software enables advertisers to track the performance of a campaign from banner serve, to the "click," and then through to any type of back-end transaction, such as registrations, orders, or downloads.

Advertisers can negotiate with a Flycast representative and buy inventory weeks or months before a campaign begins, or buy inventory in Flycast's real-time bidding for space. In the real-time bidding, the site sets the minimum bid on the inventory, and the price increases if an advertiser needs to bid higher than someone else to secure the space. The real-time bidding usually occurs 48 hours before the ad buy will commence, so an advertiser has to be flexible; however, since this is a last-minute buy of inventory that would not otherwise be sold, the price is often reduced as much as 75 percent from rate card.

A Flycast network site can use Flycast's Site Registry to register the site and control what inventory is available on the Flycast Network, and at what price. The site can also control what types of ads it will accept (rich media, GIF, etc.). Plus, publishers can block any advertisers they do not want on their site.

Here is how Larry Braitman, vice president of business development and co-founder of Flycast, explains the system:

> At the top of every Web page is code that directs the server on how to fill the page. Every Web page that carries advertising is delivered with one or more holes in it. Code at the top of the page tells the site to fill the hole with a banner ad that has been purchased in the up-front market. If no banner ad has been sold up front, additional code tells the ad server to fill the hole with a banner ad that's sold in the real-time market, via an ad network like Flycast. If no one buys a banner to serve on that page in real time, a PSA (public service announcement) or other default banner can be served.

Like Adauction.com, Flycast is a good way for publishers to sell inventory that would otherwise go unused; for the sites who use these services, it's better to get something for these ad spaces rather than nothing.

Hybrid

Because online advertising can achieve various ad campaign objectives, many publishers are finding that a tiered pricing model makes the most sense. For example, a site could offer one price for impressions, a higher price for click-through, and an even higher price if an advertiser pays only when the click-through results in a sale. If publishers offer only one model, they are going to lose those advertisers who are determined

to use a different one. A hybrid of these models is one solution. Often a Web site will charge a low CPM as well as a cost per action. This is a compromise that will often make both sides happy.

Trends in Pricing

Since the first edition of this book was published, we have seen several significant pricing trends. Prime real estate is costing a lot more than it ever has, and less desirable ad inventory is cheaper than ever. As we mentioned earlier in this chapter, keywords are an example of prime inventory. The price for keywords has more than doubled on many sites over the past year. And, some search engines will only allow advertisers purchasing a certain level of advertising to purchase some of the most popular keywords. Hot, well-targeted content sites have also been able to command very high prices. iVillage is one example; their network of Web sites will rarely allow advertisers who won't commit to spending $100K or more on their site.

On the other hand, non-targeted sites that reach the same audience as a search engine have had real difficulty in maintaining their pricing. An untargeted buy of $50K on a search engine can get a CPM as low as $5. Search engines have enough excess inventory that they can offer these low CPM deals to entice advertisers. This makes it difficult for smaller sites to charge high rates. Unless they're very narrowly targeted, how can they justify to advertisers charging more than what the advertisers would pay to run on a search engine?

Problems with Comparing Value

Some of the same problems that impact Web measurement also stymie advertisers who try to compare the value of advertising on different Web sites.

One problem that must be solved is really one of semantics: There must be consistency in the definitions used by publishers' sites. For example, it is difficult to compare impressions when one site counts an impression only when the ad is fully downloaded before the user leaves the page, another site counts an impression when a user accesses the page (regardless of whether the ad is fully downloaded), and yet another counts an impression when an ad has been on the page for 8 seconds before another ad is rotated into its position.

There is also little comparative demographic data available, which makes it difficult to distinguish among the audiences of different Web sites. For example, what distinguishes users of the search engine Yahoo! from users of the search engine Excite? When buying traditional media, an advertiser knows the profile of the target audience—from age, gender, and zip code to such detailed information as single mothers in their thirties with average incomes over $50,000. At this point, Web publishers are only beginning to learn about their audiences. Moreover, these audiences are growing at such a quick rate that even accurate data seem outdated as soon as they are published.

Flexibility Is the Key

The slate of pricing models conforms to the specific goals of an advertiser's campaign. There is no one pricing model that is the best; instead, the model should reflect the campaign's objective and the site's revenue needs. The rule of thumb in ad pricing, as in everything Internet, is to be flexible.

How Much Do You Need to Spend on an Ad Campaign?

People are always trying to figure out how much they should budget for Internet advertising. Is $5,000 enough? Is $50,000 a month the magic number? The answer lies in your goals.

If you're trying to sell products or generate leads, there is no minimum budget. You can run a month-long test for as little as $1,000 to $5,000. Then you can use results to determine your cost per lead or cost per sale. If your cost per sale is lower than for your other media channels, you should start ramping up your advertising spending. However, keep in mind that as you increase spending, you need to watch your buys carefully to make sure your conversion rate doesn't decrease.

The determination of your budget will be based on the number of sales or leads you want to generate (assuming you have the budget to reach those goals). The cost per lead or sale will also determine if you should move some of your traditional advertising dollars to the Internet. This strategy also works if your goal is to build traffic.

If you are doing a serious direct marketing campaign to a very specific niche audience, you should expect to pay $10,000–$20,000 per month, buying on a cost-per-click and low CPM basis.

If you're doing a branding campaign to a small niche audience, you can do so with a budget of $150,000–$200,000 per year (we'd also recommend some print and direct mail to deliver your message).

If you're trying to reach a large consumer audience, that budget should be increased to $100,000 per month at the minimum.

If you want to dominate a specific category or appear under keyword placements, you'll have to spend more money. Many of the major search engines require minimum buys of $4,000 to $5,000 per month to purchase keywords. In general, branding campaigns take large amounts of money. A major branding campaign across the Web, buying premium CPM, can run from $250,000 to $600,000 a month.

If your budget is small—under $5,000 for the entire year—buying ads may not be the best use of your money. A better strategy would include banner exchanges and other more grassroots methods such as those discussed in Eric Ward's special section in this book, "How to Advertise for Free or Almost Free."

Challenges Dictating Our Long-Term Success

Liddy Manson (mansonl@washpost.com), VP of Advertising and Marketing
Washington Post Newsweek Interactive (www.washingtonpost.com)

In the online news world, we are watching advertising demand explode, as advertisers and agencies become more strategic and varied in their approach to media placement. As a medium we are discovering that the Web works extraordinarily well as a promotional and branding vehicle for a broad range of businesses from construction companies to diamond dealers to car dealers to employment agencies. We currently face several challenges that will dictate our long-term success in a highly competitive and innovative market:

1. Our single challenge as a medium is to perfect the art of matching buyers with sellers. The accountability and trackability of Internet advertising combined with the newness of the medium puts significant pressure on our medium to deliver measurable results in a quick time frame. One result of this trend is a rapid proliferation of advertising products, and I expect that this creativity will continue for quite some time.
2. Currently advertising supply outpaces demand by a factor of roughly 50 to 1. As a result, pricing is extremely volatile, buyers have a tremendous amount of power, and sites without renowned brands are struggling to survive. My prediction is that there will be a shakeout in 1999 of those less well capitalized sites, whose staying power will start to fade.
3. The medium is advancing so quickly, both creatively and technically, that media sales reach far beyond devising creative advertising solutions in a static world; sales reps must reach beyond what is possible today, and devise new products and services in order to stay on the forefront of the industry.

CHAPTER 9

Buying Online Ads

Many marketers jump into online advertising without a clear plan of action. Internet ad buying entails more than merely calling Yahoo! and placing a few banner ads. An online advertising campaign, like any media campaign, requires strategic planning to ensure that money is well spent and that objectives are achieved.

As in all marketing endeavors, it is best to start with the basics. In the case of an online ad campaign, this means being able to answer the following question in the affirmative: Is the targeted market online, and if so, where online? While there may be discrepancies regarding various statistics on the Internet, no one disputes that there are many millions of online users. However, if none of these surfers fall into your target market, then an online advertising campaign makes no sense; likewise, if the target audience's online whereabouts cannot be pinpointed, an online ad campaign is probably doomed to fail. Thus, the advertiser's challenge is to discover, first, whether there is an online audience for the product or service, and, second, where online this audience congregates.

This chapter is a guided tour through the ins and outs of running an online advertising campaign. It begins by discussing the steps you need to consider before you should spend a single cent on Internet advertising, and then spells out the step-by-step process of buying and running an online ad campaign.

Determining Campaign Goals

It is impossible to develop an effective marketing campaign if you don't have objectives. Every print, broadcast, or radio campaign has clearly defined goals, and interactive campaigns are no different. The key question you need to ask yourself is, "What do I want my audience to do?" Simply getting people to click on your ad is not enough.

Meta's Exhibit No. 1004
Page 00155

One of the great misconceptions about Internet advertising is that click-through rates are a measure of how well a campaign performed. A company can't gauge a campaign's success by click-throughs alone. As an advertiser, you are looking for something more. You may want to sell products, generate leads, or build brand. When we mention building brand, many marketers think that click-through is an effective measurement of brand impact. This is not true. What if 50 percent of the people who click on your ad leave your site after they get to the home page? You may have a great click-through rate, but we'd argue that 50 percent of the people who clicked may actually be getting a negative brand impact. They're clicking on an advertisement because they expect something. When they get to the jump page they don't see what they expected and may even feel deceived, so they leave. Now, we'll go through various campaign goals and various ways to measure success for each goal.

Traffic Building

Many campaigns are aimed at driving traffic. For example, Yahoo! wants people to use their Web site, so that they can sell more advertising and boast about their large number of visitors. Of course, their campaign isn't simply looking at the cost per visitor as a metric of success. Their real goal is to maximize the value of each visitor by reaching folks who will frequently return to their site. They try to accomplish these goals with advertisements that mention their customized options, free e-mail, auctions, and other features that tend to give consumers a reason to return. Yahoo! even runs advertising on its own site (called "housed ads") that entices users to try new features (like free e-mail) and to encourage visitors to return frequently (see Figure 9.1).

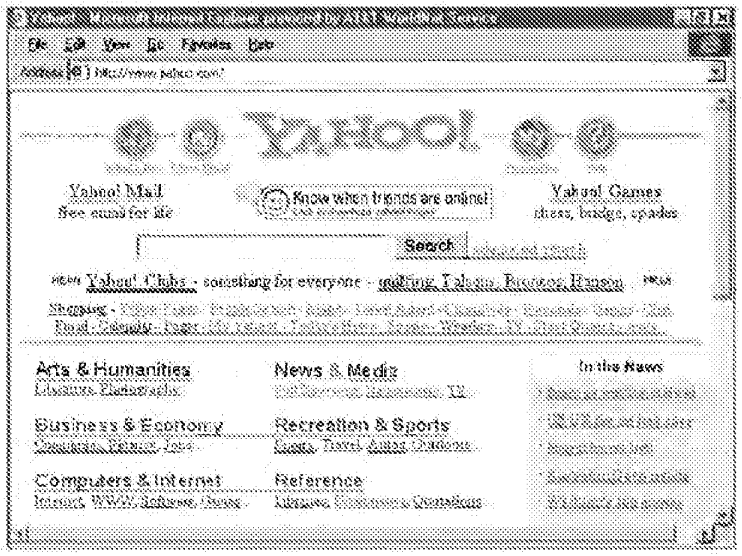


Figure 9.1 Yahoo! advertises its free e-mail service on the top left of its home page.

Preparing Your Web Site for Increased Traffic from Advertising

After you determine your goals, you still may not be ready to spend money on advertising. If you're planning to use a Web site to achieve your goals—by driving traffic to the site—prior to spending any advertising money, you need to ensure that your site is effective. Before launching a traditional campaign, advertisers test everything related to the campaign to ensure that they're using their budgets efficiently. The same process must be followed on the Web.

"I have spent many hours trying free banner exchanges, low budget promotional offers, link strategies, and yes, I even tried bulk e-mail, only to realize I hadn't spent the time to make my site appealing enough to hold anyone's interest," says Kevin Frazier, VP/ COO of Focus On U (www.focusonu.com). "This should be your primary goal. Develop your own business, cater to your customers, focus on improvements, so when you do get the results from your advertising, your business is there to support the demand. The irony is that many people believe that bigger is better, when the truth is, better will make you bigger."

Here are some ways to ensure your Web site is ready for success:

- **Analyze data in your log files.** As we mentioned in Chapter 4, all activity on your Web site is recorded in your log files. You can use this data to test the wording of different buttons. If you want people to register on your site, is it more effective if you say "register" or "register free"? Log data can also tell you what pages are the most frequent exit pages—the pages from which people leave your site—so that you can find out what can be done to retain more visitors. Since you can track how changes to your Web site affect user behavior and navigation, it is important to test and implement the most effective design prior to launching the campaign. Find the design that generates the most orders, leads, and other activity, and then invest in Internet advertising.
- **Use a focus group to see how consumers react to different implementations of your Web site.**
- **Use an expert panel to review your site.** Find people in your field who are experts on the Web and ask them to review your site and provide feedback.

Selling Products

E-commerce sites are growing quickly, and sales are one of the leading revenue sources for companies advertising on the Internet. A company selling products doesn't need as many people as possible visiting its Web site. That company needs the right traffic—consumers or businesses who will buy its products.

Businesses selling products also have the opportunity to achieve their goals without having to drive people to their Web sites. For example, Roxy Systems, a reseller of DIRECTV, doesn't always need to drive people to the Roxy Systems Web site. Sports fans are usu-

ally good prospects for DIRECTV, because DIRECTV allows them to watch any sports game that is televised. As you can see in Figure 9.2, Roxy Systems has placed its product in the CBS SportsLine store. This is the perfect way to reach sports consumers, who will probably also have a higher comfort level buying the product from a known brand like CBS SportsLine instead of lesser known technology brand like Roxy Systems.

CDnow is another example of a Web property with the goal of selling products. Figure 9.3 shows one of their banner advertisements containing an offer regarding a music sale. CDnow is focused on getting consumers interested in the sale (and purchasing music) and not just on click-through.

Generating Leads

Many direct marketers have found that it's easier to get prospects to request information online than it is to get them to order. Some consumers are fearful about giving their credit card information online or don't want to make a purchase at the time they see an advertisement or visit an e-commerce site. That has pushed many marketers to use their Web sites for lead generation. Prospects usually have the choice of requesting to receive a phone call, fax, or postal mailer with more information. Like sales, generating leads can be done through a marketer's Web site or through the site on which the marketer advertises. Often it is easier to generate leads through the advertiser's Web site, because you don't have to force consumers to leave the site if they want to request



Figure 9.2 Roxy Systems sells its satellite TV, DIRECTV, directly through CBS SportsLine by targeting the sports fanatic who wants to have access to all sports channels.



Figure 9.3 CDnow banner for "Music Lovers" promoting a discount through text right below banner.

more information. Figure 9.4 is an ad for Star Trek Universe that ran on fastWEB (www.fastweb.com). Consumers were able to request a free set of cards without leaving the fastWEB site. Figure 9.5 is an example of a lead generation banner that brings users to a form on the Star Trek Universe site.

Building Brand

Many companies are using Web advertising to build the brand recognition of their product or service, which may eventually lead to sales. To measure your brand-building efforts through banner advertising, you can conduct a focus group to determine how your banner ads affect brand recall. If your banner ads positively influence the impression customers have of your brand, you can assign a value to the number of impressions your advertisements receive. You may also decide that a good brand impact is made if someone looks at more than two pages of your Web site, and that a great brand impact is made if someone stays for more than 10 minutes on your Web site. Then you have two more measurements for success.

Building brand recognition does not require driving consumers to the advertiser's Web site. In fact, we believe that some of the most successful branding campaigns have involved advertising that integrates with the editorial and doesn't require a click. Our favorite examples of this type of advertisement are the now classic Dockers' sponsorship of HotWired's DreamJobs (see Figure 2.33), and the Maxwell House Coffee sponsorship of the crossword puzzle on NYTimes.com.

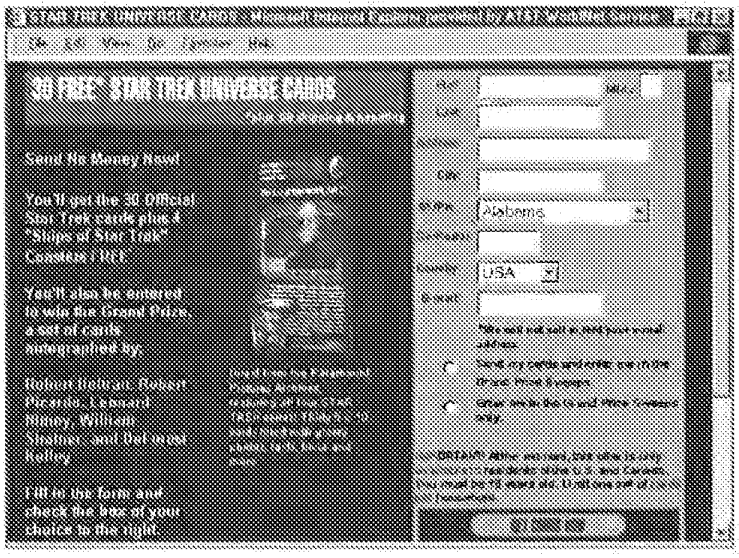


Figure 9.4 A lead generation form placed on fastWEB by Star Trek Universe.

Regardless of which goals you choose for your advertising campaign, every aspect of your campaign must be wrapped around those goals: your Web site or jump page design, measurement system, banner design, and site selection.

The Site Selection Process

There are thousands of Web sites and a variety of methods for paying for advertising on those Web sites. One of the most important factors in reaching your goals is choosing the best places to run your advertising. In this section we'll review some of the factors you should consider when choosing a Web site for your advertising.

Audience

The first and most important factor is whether or not your target audience can be reached through the Web site. The best advertising in the world won't be successful



Figure 9.5 Lead generation banner.

unless it reaches the right people. There are different tools for measuring who is using a Web site, which will be mentioned later in this chapter. We just wanted to point out that although the factors below are important in finding the best placement, they are all secondary to the audience on a site.

Reporting

A great benefit of the Internet is the flexibility to adjust campaigns to maximize performance. For example, if one banner isn't performing well, you want to be able to remove it from your campaign as soon as possible. To determine what actions will increase performance, you need frequent and timely access to advertising performance reports. We believe it is also critical to have your own method for measuring performance (determining sales and leads and other measures beyond what the clicks and impressions measured by the site). Moreover, you want to make sure that your reporting method works well with the site's system for delivering advertisements.

Ability to Quickly Rotate Creatives

The fastest ad reports in the world are ineffective if you can't make timely creative changes. Although it seems crazy, some sites require that you wait a certain number of days before they implement a creative change. Also, some sites claim they will make banner changes within one day, but it actually takes them at least a week to make the changes correctly. You'll learn which sites have these problems as you purchase advertising, and you will probably want to negotiate lower rates with these sites or avoid them altogether.

Banner Size and Placement on the Page

There are many different opportunities to buy advertising. One factor that contributes to the value of an ad placement is the size of the ad and its placement on the page. In each case, test what works best for your campaign. Here are some general guidelines to follow.

First, bigger ads tend to have a greater response rate. This may seem obvious, but sometimes in advertising stating the obvious is important.

Second, ads on the very top of the page, like in Figure 9.6, don't pull as well as those near the top. An ad at the very top can be mistaken for the top of the page instead of a clickable advertisement. Figure 9.7 is an example of an ideal placement: the ad is not the first element on the page but is near the top. Of course, all placements should be within the first fourteen inches of the screen, so that no one has to scroll to see your advertisement. Since many people don't scroll, ads below the fourteen-inch point are not usually worth purchasing. (You want your advertisement to be viewable at 640-by-480 resolution, because that is usually the lowest resolution at which someone will be viewing the Web.)

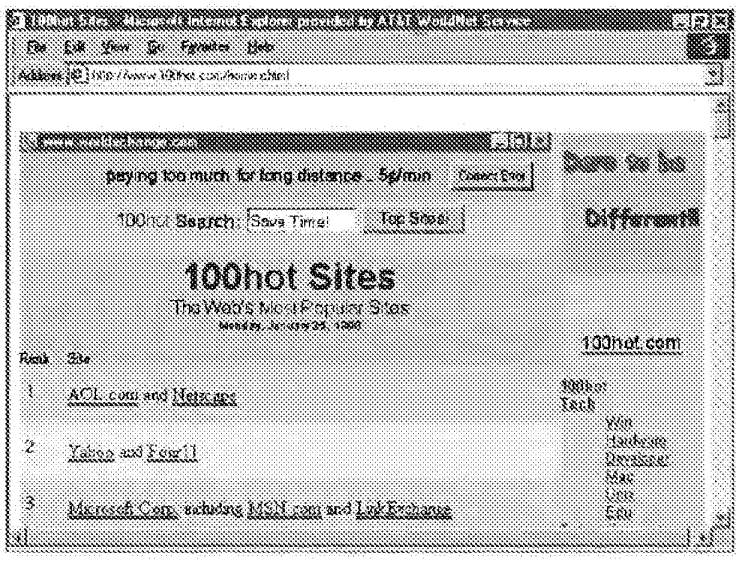


Figure 9.6 Ads placed too high on the screen that blend into the page often don't produce significant results.



Figure 9.7 This banner ad in the top portion of the Lycos page has good eye-catching placement.

In evaluating the placement on the page, you also want to know how visually busy the page is, because that will determine whether or not your ad gets noticed. Figure 9.8 shows Infoseek Spain's layout, an ideal page for placement; the page is very simple, so the ad has a high visibility. On the other hand, Figure 9.9 shows a strategy employed by ClickZ, which has many different advertisements on the page. The high number of ads for a single page makes it difficult for one advertiser to really stand out; but for those who want to reach the niche audience of online advertisers, ClickZ is important as one of the few sites to reach this market; so, the site can place a lot of ads on the page and even charge a premium for ad space.

Amount of Ad Inventory Available for Purchase

Depending on your goals and the specific site on which you're advertising, you may want to purchase anywhere from 10 to 100 percent of a site's inventory. Someone who is a category leader will want to purchase 100 percent of certain Web site inventory. For example, Claritin is the leading prescription allergy medication. They should purchase all of the allergy keyword inventory on search engines, because they will want anyone searching for allergies to always see an ad for Claritin. There are also allergy areas of medical sites, like the Mayo Clinic's allergy area, that Claritin should "own." These are targeted areas where they will reach their best prospects—people searching for information and solutions regarding their allergies.

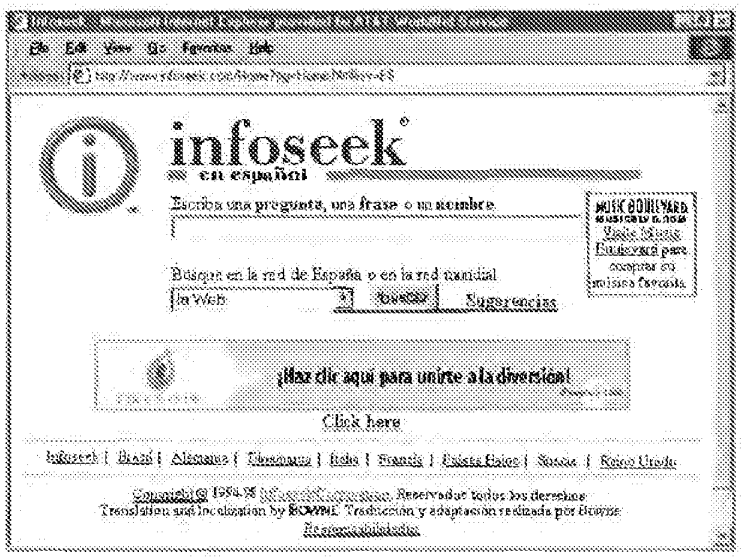


Figure 9.8 Infoseek Spain has a simple layout that helps bring attention to the ads on the page.

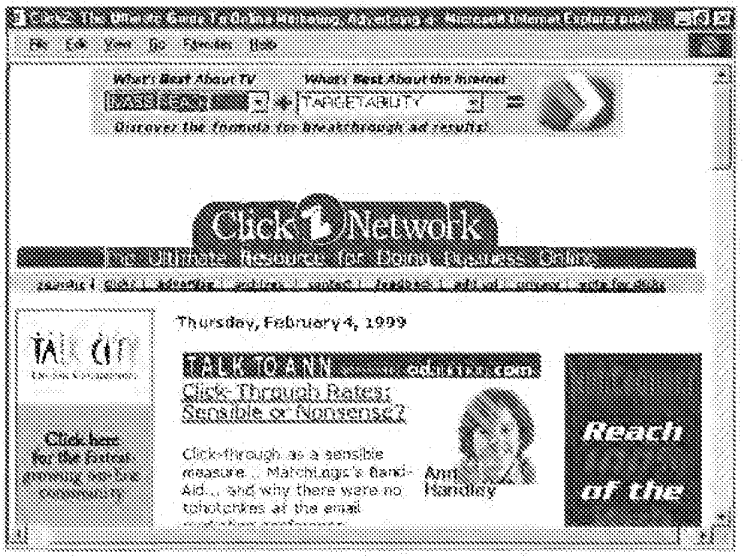


Figure 9.9 Click2 has ads to the right, left, and above the site's content.

On the other hand, there are situations where it is inadvisable to purchase more than 10 percent of the available inventory. If you want to sell a product, unless the site is a perfect fit, test the site with a purchase of no more than 10 percent of available inventory. There are a few reasons for this. First, good marketers will always test small amounts before committing more advertising dollars. Second, the response from certain types of sites will drop after more than 10 percent of the inventory has been purchased. This could happen on a site like CNN. CNN has high traffic, but it's reasonable to assume that much of that traffic is from people who check the news every day. So, if you buy a large portion of CNN's inventory, there is a good chance that the same people will see your ad many times. If you're trying to sell a product, your response goes down as your ad is seen multiple times by the same person rather than being seen by different people. Use the 10 percent figure for testing, then see what works—for some campaigns, a buy of 75 percent of a site's inventory will produce the desired results.

Performance Guarantees

When buying on a site for the first time (actually, anytime, but especially the first time) get some sort of guarantee regarding performance. For example, a site will often promise a certain number of orders or clicks. Not all sites will offer a performance guarantee, but it is worth asking for and certainly makes a site more attractive.

Technology Capabilities

Depending on your product and target audience, you may be looking to use certain technologies, like HTML banners or Java. Make sure the sites accept these creative formats before you sign the contract. Also keep your eyes open for any sites that try to charge a premium for running HTML and rich media ads. There are sites that charge advertisers an additional $1,000 fee to use HTML advertisements. We imagine that by the time this book has been printed, this practice will stop.

Targeting Capabilities

In Chapter 6 we mentioned many of the targeting capabilities of the Web. In evaluating a site you want to try to find the placement that helps you best target your audience.

Tips for Finding the Best Sites for Your Ad Campaign

Think like your audience. Put yourself in the shoes of your target audience. If you were them, where would you go on the Web?

Ask your audience. You can even survey your target audience to get an idea of some sites that they visit.

Look in your log files. As we mentioned in Chapter 4, the log files record activity on your Web site, including the site from which visitors came to your site. Many advertisers think that they shouldn't advertise on sites that already send them a lot of traffic. They don't think they should pay when they're getting free traffic from links already provided. As we already mentioned, we've found that in almost all cases the sites that are already sending traffic to an advertiser are the best sites on which to advertise.

Check online lists. There are Web publications that list sites that accept advertising. Some of those sites are SearchZ (www.searchz.com) and ChannelSeven (www.channelseven.com).

Check subscription databases. The Web publications that list ad sites are okay, but far more comprehensive information is available from subscription databases. The two best are Planner from AdKnowledge (www.adknowledge.com), and Interactive Advertising Source, which is put out by SRDS, the leader in providing ad rates for traditional ad vehicles (www.srds.com). Planner allows advertisers to include Simmons and MRI research as well as Media Metrix data in their site research (if the advertiser subscribes to those services).

Both online planning tools allow media buyers to search Web sites by criteria. Figures 9.10 and 9.11 both show Planner menus from which media buyers can begin researching Web sites that fit their media plans.

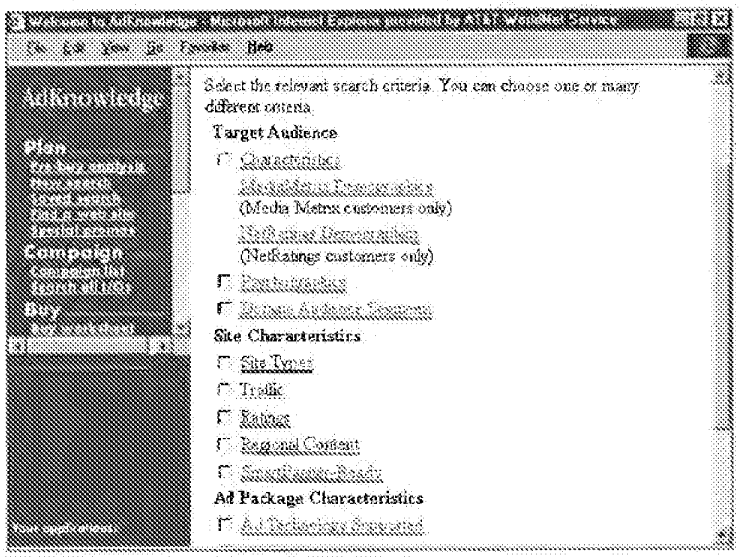


Figure 9.10 AdKnowledge planner menu helps you select your ad campaign buying profile.

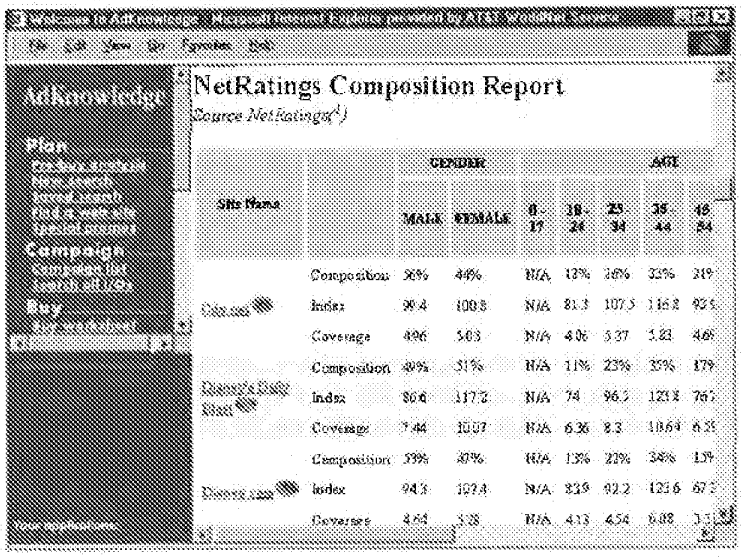


Figure 9.11 AdKnowledge recommends sites that meet your campaign profile needs.

Comparing Audience

There are several sources to help advertisers compare site audience. The leading providers of audience measurement information are Media Metrix (www.media-metrix.com) and Nielsen/NetRatings (www.netratings.com). Both of these companies track Web usage through panels. Panel participants have installed software from either of these companies on their computers, and the companies use that software to track the Web usage of these individuals. Advertisers who pay for these services receive reports on the most heavily trafficked Web sites as well as the most heavily trafficked Web sites in specific categories.

Since both of these services cost over $10,000, many small and medium-sized advertisers do not subscribe to them. One alternative is to use 100hot (www.100hot.com, see Figure 9.12), which is a Web site that offers free access to lists of the 100 top Web sites in specific categories—gaming, news, and so on. 100hot gathers its data through ISPs; the aggregate data of the sites users from some of the largest national ISPs visit is used to determine the highest traffic Web sites. We don't believe this is extremely accurate, but it is definitely worth investigating.

Types of Sites

As we've mentioned, and as you probably already know, there are hundreds of thousands of Web sites. Breaking those sites down into categories is difficult, but we think these categories can be used to classify most sites: (1) well-branded, highly recognizable

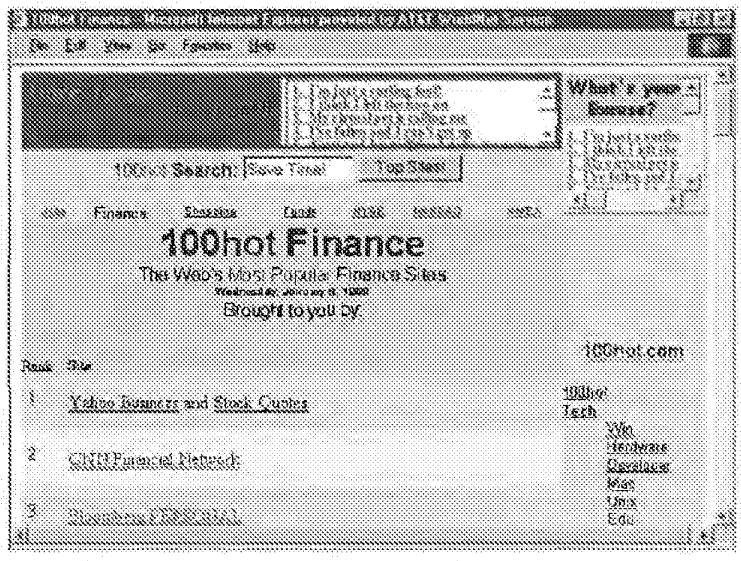


Figure 9.12 100hot offers free access to lists of the top 100 Web sites in specific categories.

sites, (2) high traffic sites with less- or non-recognizable brand names, (3) niche sites with low traffic, and (4) hobby and personal sites.

Well-Branded, Highly Recognizable Sites

These sites include the likes of Yahoo!, ESPN, and other similar sites that most people have heard of, and that receive a large number of visitors. Since these are the most recognizable sites, they are the ones where most advertisers first go to place their advertising. The prices are usually the least negotiable of any of the types of sites. For advertisers, these sites offer a good brand with which to associate their product or services and a sizable audience reach.

High-Traffic Sites with Less- or Non-Recognizable Brand Names

This category includes sites like Dogpile (www.dogpile.com, see Figure 9.13). If you've got a strong brand, you may not want to associate it with a site named Dogpile. However, the site gets a ton of traffic, so it is perfect for marketers selling products. Sites with low brand recognition and high traffic usually have a lot of unsold inventory that they will offer at a low price.

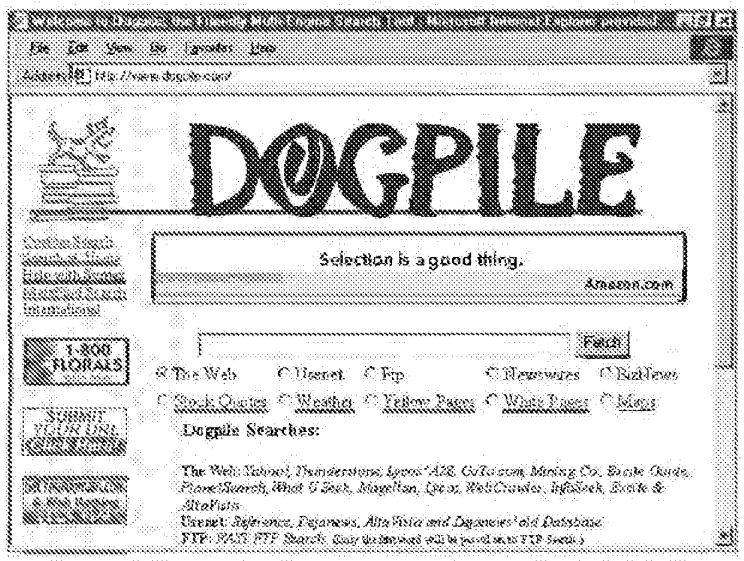


Figure 9.13 Sites with *names with attitude* like the search engine "Dogpile" can get tremendous traffic.

Niche Sites with Low Traffic

Sites in this category include ClickZ (www.clickz.com) and ChannelSeven (www.channelseven.com). Both sites are devoted to Internet advertising. They have a very high volume of traffic from Internet advertisers, but that doesn't make them high traffic sites in terms of the Internet. Since they reach such a targeted audience, there is a long line of advertisers willing to pay top dollar to advertise on their sites. ClickZ and ChannelSeven don't have to negotiate as much as other sites, because they offer a great means for reaching advertisers. Well-designed and marketed niche sites that have a good amount of traffic often have so much interest from advertisers that they can even clutter their pages with three or more ads and still be well worth the advertising expense.

Hobby and Personal Pages

These are often fan pages on a particular topic. They're not as well designed as the niche sites, and usually have much less traffic than them as well. This makes them a very inexpensive buy. Many of these sites are poorly designed, which makes advertisers unwilling to place their messages here. Also, the expense of managing an ad campaign on this type of site can be more trouble than it is worth. You have to decide the minimum amount of impressions that makes it worthwhile to negotiate with a site, deliver ads, check ads, and handle other management overhead.

Paying for Media Buys

Besides deciding on the type of site you are looking for, you should also look at the different types of media buys. Like any advertising medium, all prices are negotiable. Unless you're making a small buy on a large site, don't forget to negotiate. Below are some of the pricing models that sites are using. Chapter 7 provides a full discussion of pricing.

CPM

CPM, or cost per thousand, means that the advertiser will pay a set fee for every thousand times an ad banner is delivered. High traffic sites prefer to sell their advertising on a CPM basis. Lower traffic sites may want to price their advertising packages at a fixed sponsor price, since they don't have enough traffic to charge high ad rates based on CPM.

Cost Per Action

Cost per action means that advertisers pay for every time a specific action is taken. For example, you may pay on a per click or per order basis. Most sites don't like selling advertising on a per action basis. These types of buys compensate the sites for the performance of advertisements, but since the site has no control of the ad copy or the product, it also doesn't have control over how well the advertising works. Sites also

feel that if advertisers are only paying on a cost per action basis, there is no compensation for the branding exposure advertisers receive. However, this is a great buy for advertisers, if they can negotiate it.

Hybrid

A hybrid buy usually means that an advertiser pays a combination of CPM and cost per action. Since sites don't like to sell on a cost per action basis and advertisers clearly want to pay cost per action, a hybrid buy is a nice middle ground upon which both sides can usually agree. In a hybrid, the advertiser will usually pay a reduced CPM as well as some sort of a bonus for every click, sale, or order that is generated from the advertising campaign.

Pricing for Buys

There are a host of different types of buys with different pricing models. Here we're going to go through some of the basic pricing structures and rate card prices (see Table 9.1). The companion Web site includes sample rate cards to give you an idea of what various sites are charging for advertising space.

Search Engines and Directories

Sites, like Yahoo!, Lycos, and Excite, have a variety of buys. One is a general rotation buy. This means that the advertisement can appear anywhere in the search engine and at any time. The buy is the lowest priced of the search engine buys, because there is no targeting. The rate card price for these buys is $20 to $30 CPM. However, if an advertiser buys $25K or more for this type of a buy, there is usually substantial room for negotiation.

A more targeted search engine buy is in a topic area. For example, Figure 9.14 is the financial topic area on the GO Network (www.go.com). Marketers can purchase advertisements that rotate throughout topic areas like these. The rate card cost is typically around $40 CPM.

The most targeted search engine ad placement is by keyword. These placements allow for an advertisement to come up whenever a consumer performs a specific search. The price for these is usually between a $40 and a $70 CPM.

Content Sites

Content sites like Quote.com or CNN are usually more expensive than search engines. But content sites can be great, because you often have a very targeted audience (it's clear that ESPN is a great way to reach sports fans). However, unlike search engines, good content sites are visited to read articles or research information, which makes visitors less likely to click on an advertisement. Content site costs can be anywhere from a $25 to a $75 CPM.

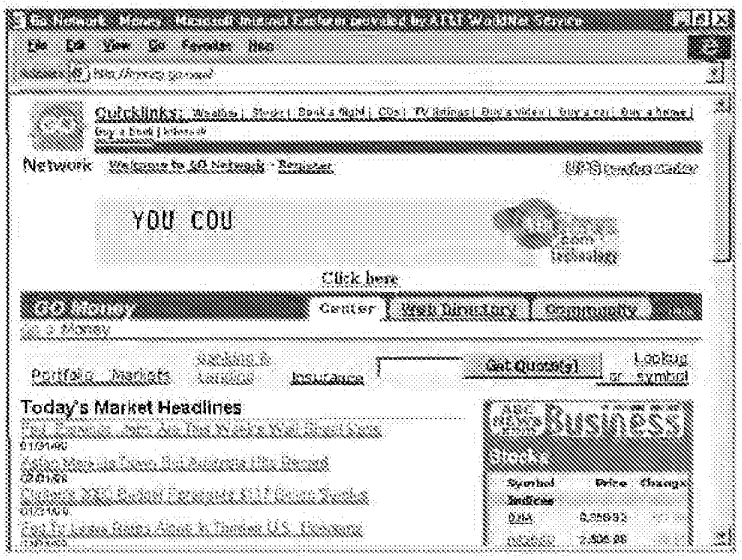


Figure 9.14 Financial topic area on the GO Network.

Network Buys

Networks sell advertising for a number of sites—see Table 10.2 in Chapter 10 for a list of ad networks. The network doesn't usually own the sites, but serves as the ad sales representative. This works well, because sites often don't have the experienced sales force (or the money to hire an experienced sales force) that will be able to sell advertisements to large advertisers and agencies. And, agencies can make large ad purchases on a variety of sites through one entity rather than negotiating with all of the network's sites individually. Another advantage to networks is that they offer great targeting capabilities. For example, if you want to target everyone using the Netscape 4 browser, you can quickly reach a large number of people by advertising across an entire network. Many networks represent enough sites that they actually offer more reach than the large search engine sites. As we mentioned earlier, network buys can range from a $30 CPM to a $35 CPM depending on targeting.

Free E-mail Services

Advertisers can purchase advertisements in many of the free e-mail services, which include Hotmail (www.hotmail.com), USA.net (www.usanet.net), Juno (www.juno.com, see Figure 9.15), and most of the major search engines. Many free e-mail services require registration from users; these services can then offer demographic targeting based on users' responses to the questions. For example, advertisers on Juno can target based on zip code, sex, and household income. The cost on these sites can range from a $10 CPM to a $35 CPM depending on how finely the advertising is targeted.

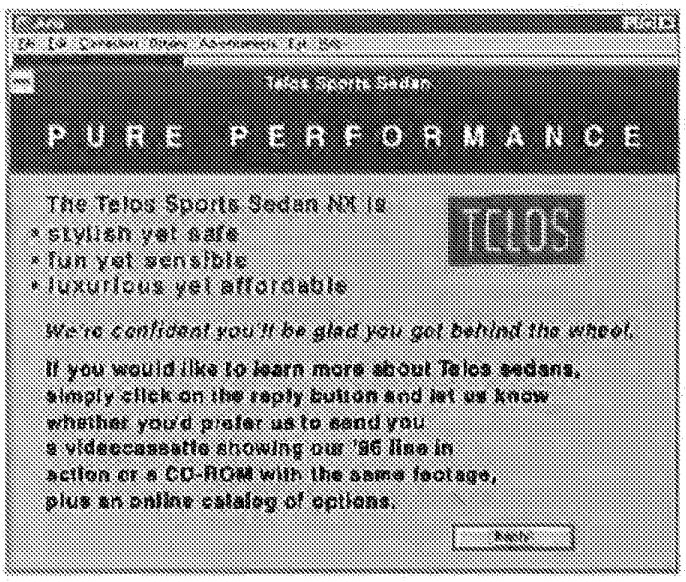


Figure 9.15 Juno offers its users free e-mail in return for ads on the e-mail reader.

Spot Buys

These are ad buys that allow the advertiser to buy advertising space in the open market. Usually, the entity selling the advertising sets a bottom price point, and then advertisers can submit bids. Flycast (www.flycast.com) and Adauction.com (www.adauction.com) both offer these services. Figure 9.16 shows a screen in which an advertiser is offering a price for advertising on Flycast.

So far we've talked about setting campaign goals and some of the differences you'll see between different media buys. Now we're going to cover budget allocation and testing.

Allocating Your Campaign Budget

We believe that as an advertiser you should commit a maximum of 50 percent of your budget to long-term buys. If you want to dominate a category, you may have to allocate more budget to long-term buys to lock competitors out of the best targeted inventory; otherwise, there is no reason to commit more than 50 percent of your budget up front. So, what should you do with your remaining budget? For starters, once the campaign begins, you will be able to see which sites are performing best. It's nice to have extra budget to allocate to those sites that are top performers. Also, there are constantly new opportunities. The Net is a rapidly changing environment, and you should always have funds

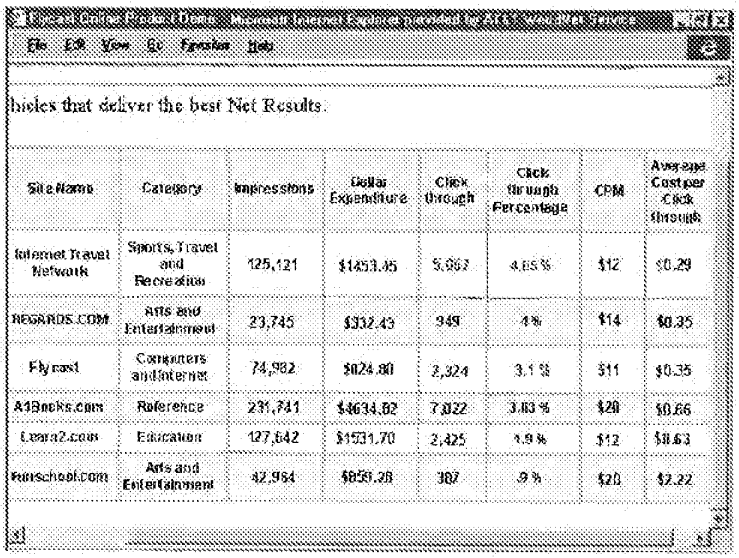


Figure 9.16 An advertiser is offered a price for advertising on Flycast.

remaining for new targeting opportunities and new ad venues. Last, there are always discounted inventory sales, and there will be nothing more frustrating than having no budget for when the inventory is at half price.

Cross-Media Promotion

There is a great benefit from advertising your Web property offline as well as online. From billboards to print ads to television commercials, ads for Web sites are appearing everywhere. In 1998 we saw many top Web sites go to TV to build brand awareness. From Amazon.com to HotBot, the irreverent humor of the Web met broadcast TV and it worked. TV advertising increased traffic and stock value.

Hollis Thomases from Web Ad.vantage (www.webadvantage.net) reminds us not to forget using targeted co-op print ads for the sole benefit of promoting Web site URLs. People who surf the Web also participate in other media, so let one help build on the other so that all roads do lead online.

For example, an American Airlines promotion placed by DDB–Dallas on the Broadcast.com site was part of a cross-media blitz that included national radio airtime, a golf "partnering" site, and a site to promote contest entries. It was the marketing mix that made the campaign work, with 93,000-plus visitors in the one-month campaign.

Continues

Cross-Media Promotion *(Continued)*

Maxim Communications International used a Web site as the destination for people responding to print, radio, and online ads about a real estate venture in Stowe, Vermont, www.stowehighlands.com. Even with a moderate budget, they were able to use the Internet to improve the effectiveness of their ads. "It also proves," says Keith Reynolds, president of Maxim Communications (www.maxim-com.com), "you don't need millions of hits to succeed if you have a narrow audience and a high-ticket product."

In addition to media costs, advertisers must also pay for creative. An agency will typically charge about $600 to $1,000 for an animated banner and $1,000 to $2,000 for an HTML banner (which includes the banner, script, and jump pages; resizes cost extra). We'd suggest that you run no fewer than three banners on a site. Having extra creative budget set aside is also very important. During the campaign, you'll learn a lot about

Table 9.1 Ad Rates and Offerings

Search Engines	$20–$50 CPM
Keyword Advertising	$40–$70 CPM
Local Advertising	$20–$80 CPM
City Guides	$20–$80 CPM
Advertising Networks	$10–$70 CPM
Auction Sites	$1–$25 CPM
Top 100 Web Sites	$25–$100 CPM
Opt-In Email Advertising	$.10–$.20 per address
Small Targeted Content Sites	$10–$80 CPM
Sponsored Content	$45–$85 CPM
Rich Media	$35–$100 CPM
Click-Through Pricing	$.20–$1.20
Range of Click-Throughs	.5%–15%
Average Click-Through	.5%–2.5%
Overall Average CPM	$27
Pay Per Lead	$1–$15
Pay Per Sale	5%–30% of purchase price

Compiled by Ad Resource (www.adresource.com) as of 6/9/98.

what creative works best, and you'll want to have the money to change your creative to maximize response.

Don't forget that you'll also have tracking and fulfillment costs. If you offer to mail something to people, you'll have to pay for shipping and handling, not to mention purchasing or manufacturing the product.

Testing

You've already heard from us about the Internet's great ability to efficiently test campaigns. But, of course, this only works if you actually take the time to test sites, ad placements, and the ads themselves. Two good ad-testing programs are TestIt! (www.doubleclick.net/advertisers/ad_rates/network/testit.htm) and FlyCast (www.flycast.com). A small test flight, a week before your campaign, is good for ensuring that you know as much as possible about what will work before you begin your actual campaign.

Dump Your Traditional "Partner" Agency Now

Todd Tweedy, Director of Client Services <todd@ webnet-marketing.com>
web-net marketing <www.webnet-marketing.com>

Interactive media shops have dominated the digital marketing arena because they understand that customers control the flow of information and digital agencies also realize how technology can enhance online media placements. These two facts are critical when assessing the value of interactive media pricing. That's why I believe you'll find more direct marketers in the digital marketing game. It's also the reason why the traditional advertising model of pushing messages is dead in the digital age.

Traditional shops have outsourced interactive media planning and placements for one reason: The big shops still don't get it even if say they are "integrated." And although it sounds wonderful for an advertiser and its "partner" agency to have a shared vision of where the advertiser's brand will go, if the partner agency doesn't know where to go online the advertiser's online placement will go nowhere. If you want an effective online media campaign, dump your "partner" agency now.

Why take such a radical step? Because most agencies are behind the curve. They don't have the resources in staff and intellectual capital to provide the service you need and deserve. Getting the most bang for your interactive investment dollar begins with raising the bar on the services your interactive media agency provides. Interactive media planning is more than making recommendations on specific media properties to include in your plan. Your interactive media placement agency should also lead contract negotiations; provide placement verification reports, daily performance reporting, and tracking; and test methodologies to sustain and optimize creative performance, trend reporting, reconciliation of media delivery, and analysis. It's amazing, but I still hear advertisers telling horror stories that the first data they received was after their campaign concluded.

The issue isn't whether the agency is old or new. The issue isn't whether the agency is economically priced. The issue is: Can the agency help you maximize your online investment? Whether the tactic is utilizing high-reach vehicles with run-of-site positioning or site-specific content targeting, the marketing value of online media is not solely driven by media costs. An agency that understands this is worth its weight in gold.

CHAPTER 10

Selling Online Advertising

You want to sell advertising on your Web site. Who doesn't? The American Association of Advertising Agencies professes that advertising is the engine of capitalism. And, with the Internet frontier wide open, does everyone want to jump on the bandwagon?

Before you throw all your cards on the table, recognize that conducting business on the Internet is much more than a matter of point and click. Planning, persistence, and hard work are essential. And, even with all that, there is no guarantee you will make a profit from selling ads. It's a buyer's market in cyberspace. There are thousands of sites selling advertising. To succeed, you not only have to be lucky, but cyber-savvy. This means taking steps to:

- Stand out from the multitude of publishers selling advertising
- Offer inventory that meets the needs and interests of your advertisers, whether that's banners, rich media, or sponsorships
- Work to build long-term relationships with advertisers
- Be flexible with pricing and creative

The selling philosophy presented in this chapter is simple: Add value through quality—quality in measurement, management, targeting, and delivery. Whether your site receives 100,000 or 1,000,000 visits per week, whether you cater to college students or car buyers, following this rule of thumb will have your site well on its way to attracting advertising dollars. The suggestions in this chapter will help you make your site advertising-friendly, smart, and successful.

Meta's Exhibit No. 1004
Page 00177

Identifying and Defining What You Have to Sell

The first step in developing an ad sales program is identifying and defining what you have to sell. As in any sales structure, you need to define your product. When you are selling online advertising, the ads are the product and need to be defined, classified, and inventoried. The entire site needs to be reviewed for ad format acceptance and ad placement possibilities so that site specifications can be determined, including:

Ad formats. The types of ads accepted: banners, buttons, logos, reciprocal linking, sponsorships.

Technical specifications. File size, dimensions.

Ad space locations. Run of site, run of category, section-specific.

Placement. Location of ads on each page (this could be literally anywhere on the page, although we recommend placing all advertising within the first fourteen inches of the page). Placement can also refer to how many ads are on a page.

Type. Static, animated, rich media.

Targeting. Demographic, collaborative filtering, behavioral, etc.

Special benefits. Exclusive agreements and other special opportunities.

Preparing Your Site's Infrastructure

Before you begin accepting advertising, you must set up your site's infrastructure. This means making certain you have the technical applications on your site to fulfill the needs of the professional advertising community, which include monitoring and measuring your site's traffic, deciding which ad models to use, and ad management. Some advertisers may also need an audit of your numbers.

Monitoring and Measuring Site Traffic

Part of knowing what you have to sell is knowing the traffic on your Web site. It is central that you employ effective measurement tools so that you can precisely calibrate the traffic to and from your site, as well as gauge the traffic patterns within your site. Being able to monitor and measure how many users are visiting your site, where they come from, and what they are doing once they get there will provide you with a blueprint for determining the best pages for placing paid advertising. For a full discussion of measurement software, revisit Chapter 4.

Ad Models

In designing your site, you must decide which ad models you'll accept: banners, buttons, product placement, interstitials, sponsorship, and so on. You can accept one or several ad types, as well as institute your own innovative ad models to intrigue cutting-edge advertisers. Chapter 2 provides a lengthy overview of all the ad models available for use.

Ad Management

Media buyers want to be confident that when they buy advertising on a site, the ad will go live at the designated time; that it will remain online for the proper period of time; and that they will receive performance reports as often as possible in an understandable format. Advertisers also like to be able to rotate multiple ads simultaneously. Fulfilling advertisers' ad management needs is a daunting task to perform manually—almost impossible to do well—even for the most organized individual. To ensure that the ads you sell are delivered professionally, it is wise to employ ad management tools; it's very difficult to sell advertising, and, the last thing you want to do is lose an advertiser because you couldn't manage the ad serving. For a full discussion of ad management tools and techniques, refer back to Chapter 5.

Auditing

If you hope to secure top-billing advertisers, you may have to provide them with the audit services they have come to expect in traditional advertising. Audit reports document the credibility of results in the form of an independent third-party analysis. Although currently, most online publishers do not provide audits, in the near future this will no doubt become standard practice. More about auditing is available in Chapter 4.

Know Your Audience

The true value of your site is its audience. If your site generates a good volume of traffic—100,000 to 1 million or more visits per month—then one of your site's asset is its traffic volume. Currently, sites with the most traffic receive the bulk of online advertising dollars. Even if getting to 100,000 visits is unrealistic for your site, you need to continually work to increase its traffic.

Increasing traffic sends two messages to advertisers. First, it says that the site is popular and servicing the needs and interests of its audience. Second, increasing traffic ensures media buyers that they will be able to continually expand their reach with the advertising they place on your Web site.

We've mentioned that it's the highest traffic sites that get the lion's share of the advertising dollars. However, even if your site is not in the top 1 percent (or even 10 percent) in traffic, you can still find buyers for your ad space if your site's visitors represent a target market. By catering to a specific audience, you can present yourself as the premier marketing opportunity for a specific consumer product or service. For example, if your site is designed for tennis enthusiasts, then it would be a prime location for advertising tennis products. If your site provides information about a particular disease, such as diabetes, or a medical condition, like high blood pressure, your site would be an excellent location for a pharmaceutical company to promote its new medications.

Advertisers want to know—indeed, they must know—who they will reach by advertising on your site. Consequently, you need to be able to provide a profile of your audience. This profile can be compiled through interactive strategies such as registration,

surveys, or even games and contests. You can also implement many of the new technologies available for greater targeting and personalization of ads. See Chapter 6 for more on targeting. The more detail you can include in your audience profile, the more valuable your advertising real estate. If your site can deliver an already targeted consumer market, your real estate has a higher value that generic "land" on another site.

The Credibility Issue

Credibility is transferable. A site with a loyal following is especially attractive to an advertiser. Advertisers expect that the goodwill and community you foster on your site will inspire response, resulting in sales, brand building, and leads.

Because audience loyalty is of prime importance, publishers must be certain that the advertising they accept won't upset their audience. Therefore, it is important to review ads, not only for spec adherence but also for contextual appropriateness. In most cases, this is a non-issue. But a site might want to establish a policy addressing unacceptable ads. Ads might be deemed unacceptable because they are too sexually provocative or contain other content that might be offensive to the site's particular audience. Or the format of the ad itself might be offensive. For example, there are a number of sites that don't allow advertisers to run banners that resemble computer-generated dialogue boxes with error messages (see Figure 10.1). When such ads ran, consumers thought the banners were actual error messages and clicked on them for that reason. When the users were transferred to the advertiser's site, they realized that they were deceived. Some sites have banned these ads because of the resulting user backlash. Publishers should also be careful not to add too many ads to a page before they are certain that it won't upset their audience.

What is appropriate for one audience may be inappropriate for another. This is especially true for children's sites. By serving an advertisement on your site, you are tacitly endorsing the company, product, or issue. You don't have to accept every ad just because someone is willing to pay for the impressions. Credibility is easier to lose than to regain. And, if you lose credibility with your audience, you may lose traffic and, subsequently, inventory available for sale.

Pricing Ads—Developing the Rate Structure

Once you know what you have to sell, then you need to decide what rate structure you are going to use. How are you going to charge for your ads: flat fee, CPM (cost per thousand), CPC (cost per click), PPS (pay per sale)? Chapter 7 provides a full discussion of pricing options. You can use one method, or offer a menu of rates for an à la carte approach to pricing. In fact, the rate menu approach is gaining popularity as ad campaigns become more customized and more representative of a partnership between the advertiser and the site property.



Figure 10.1 A banner that resembles an error dialogue box.

only outside sales force. They do this to limit confusion caused when an advertiser is approached by a site's internal sales force, by an ad network, and by a rep firm—each with different rates and different inventory.

Auctions

Another means to sell inventory is to auction impressions. In 1998 auctions came of age on the Internet. From consumer goods to impressions, the online community embraces the action, excitement, and potential savings involved in buying through an auction. The leader in the online auction space is Adauction.com (www.adauction.com).

Adauction.com only offers well-branded sites for auction. The auction process is very simple. Using an auction gives Web publishers an easy way to generate incremental revenue from their expiring ad inventory. Adauction.com provides exposure to the industry's top media buyers, who can use the service to save on quality Internet ad space. Adauction.com offers large quantities of banner ad impressions at prices significantly below rate card.

Words of Wisdom on Selling Online Ads from an Internet Advertising Director

Dave Coyne, Internet Advertising Director, Community Newspaper Company

- **Motivating existing staff to sell is a slow process due to need for intensive training, extensive handholding and other competing priorities for their time.**
- **Customers *hate* multiple sales reps from the same company, which is why we continue to try to leverage our existing sales force.**
- **Unlike even one year ago, we rarely find ourselves beginning an Internet sales call by giving an Internet lesson. What's more, many customers seem to be aware that having a site isn't sufficient and that advertising it is the most important part.**
- **Customers hardly ever bother to log in to check their banner stats.**
- **We have seen no cannibalization of our newspapers. Customer who have no more budget for print routinely find more money somewhere to spend on the Internet with us! We've also developed significant new business for the company in terms of new advertising accounts with no history advertising in our papers.**
- **Compensation issues are a bear!**
- **We use our standard newspaper advertising agreement with slight modification when we bother to use one at all. Although it has not been a significant issue for us so far, we are moving to tighten up our use of contracts to ensure it never becomes an issue.**

Adauction.com's Close-Out Auctions (which occur on the third Thursday of each month) offer single site buys or bundled content category packages, both of which run over a four-week flight period, starting the following month. To participate in an auction, you need to register as a bidder. Auctions are proving to be best way to unload remnant inventory.

Sales Staff

Selling is an art form, and nothing can replace raw talent, but with proper training, a young sales force can be groomed to become a dynamic team. A key component of this training is learning the medium itself.

Gina Carrubbo, executive vice president of sales at Women.com, is known for saying that her sales force is the hardest working crew at Women.com. Most online ad sales directors would make the same claim about their staff and their company. Selling online ads is tough. Really tough. Not only do you need to know your property, how to sell, and what to sell, but you also need to know the technology.

Sales Training

Whether you are looking for a job in online sales or hiring sales staff, whether you have experience in sales or are new to selling Web ads, learning the art and science of selling ads is an ongoing process. Online advertising and marketing is a new and burgeoning field, and business schools are only now adding it to the curriculum. A survey conducted by The Laredo Group found that 67 percent of sales reps have under two years of Internet experience, with 22 percent having less that two years of any sales experience at all. Clearly, training is key to ensuring you that your staff is schooled in the skills needed to sell ads efficiently and effectively. Most of the material available for training the sales staff of today and tomorrow exists either online, at conferences, or through publications.

Distance Learning

Distance learning is a great training tool for busy staff in different locations. Through online courses, staff can take a class at their convenience, whether that is 10 A.M. or 10 P.M. Ziff Davis University (ZDU) offers a course on Internet Advertising as well as online business strategies and Web site promotion (www.zdu.com). These classes are extremely inexpensive: under $10 per month.

Conferences

One of the best avenues for learning about online advertising is to attend one of the outstanding conferences and seminars held every year. These events are a mix of panel discussions, exhibit hall extravaganzas, and serious industry networking in the hall. Some of the key events can be found in Table 8.4. Another option is to attend the smaller and more focused advertising seminars. These events take more of a skills-building approach:

"Advertising and Marketing on the Internet." Inexpensive, one-day conferences held in cities across the country, providing a local and statewide perspective; produced by The Zeff Group (www.zeff.com).

"How to Buy and Sell Web Ads." A comprehensive training seminar on buying and selling Web ads for the novice and professional, taught by the online ad sales experts and produced by AdWeek (www.adweek.com) and the Laredo Group (www.laredogroup.com).

Online Publications

As presented in Tables 8.1–8.3, the online advertising discussion lists and newsletters are where all the cutting edge information in the industry gets discussed on a daily basis. The first step in learning about the industry is to become a member of the online advertising community. And there is no better way to immerse yourself in the community than reading and participating in the industry's key publications.

The Never-Ending Quest to Keep Good Sales Staff

There used to be an unwritten rule that longevity at a job was the norm. The Internet, with its own sense of time, has turned job security on its head; staff changes jobs at regular intervals. However, sales are built on relationships, and a high turnover rate hurts the bottom line. High turnover in sales positions is common right now. If you want to keep qualified salespeople, you should consider giving them a bigger piece of the equity pie, according to the Laredo Group 1998 survey, 95 percent of respondents said that more equity would decrease their chances of job hopping.

Compensation Plans

There are a variety of options available for compensating sales staff: salary, base salary plus commission, or straight commission. The most common compensation package is salary plus commission. A survey conducted by the Laredo Group in 1998 found that 55 percent of sales staff are following this model. The study also found that 25 percent are on straight salary, 12 percent on straight commission, and 8 percent receive a salary plus a commission but with a maximum they can earn.

Finding a Job in Online Ad Sales

Employment is a hot topic online, and no type of job is better marketed online than positions at online companies. You can pursue jobs by going through the front door or find other means to open doors. The front door method is to post your resume through online resume banks and dutifully surf the online job boards such as Monsterboard.com (www.monsterboard.com), under the category "Advertising/Marketing Media and Media Sales," or CareerPath.com (www.careerpath.com), which not only lists the classified ads from the nation's leading newspapers, but also lists ads taken from the Web sites of leading employers. Online sales job are also listed in the print and online versions of Advertising Age (www.adage.com) and AdWeek (www.adweek.com). Another option is to visit the Web sites of companies you find of interest and check out their job listings. If that doesn't work, then network your way in through the side door. And don't underestimate the value of selling yourself to the company by sending an unsolicited e-mail.

A full list of online advertising periodicals is given in the Resource Directory of the companion Web site.

Spread the Word about Your Advertising Opportunities

Just because you have ad space available, doesn't mean people will know to look at your site as a place to buy advertising. You need to do some advertising of your own. It is your responsibility to let media buyers know about your site. The following is a list of sites where you can post your rate cards.

You can list your site and rates the way you would in traditional media through SRDS. SRDS (www.srds.com) is the largest compiler of advertising placement information across all media, and has put together two extremely valuable resources for online advertising. The first is the SRDS Interactive Advertising Source, which is available in print and online (www.srds.com/cgi-bin/srds/iasonline/iasonline.cgi). Contact SRDS directly to be including in the source directory. They also have a free online database of 1,500 online media kits and Web sites for business publications and consumer magazines (www.srds.com/cgi-bin/srds/mediakitlink/mediakitlink.cgi).

You should also list your site with AdKnowledge (www.adknowledge.com), a database of Web sites that accept advertising, which is available to advertisers who pay AdKnowledge or use one of their services. And, finally, many online advertising resource sites—for example, SearchZ (www.searchz.com)—list media kits.

And, of course, don't overlook prominently promoting the advertising opportunities on your own site—on your home page and on every page. Make certain that a media buyer can easily find your rate card and contact information. If a prospective advertiser requests information, have a sales representative personally follow up with a phone call. (In spite of—or maybe because of—e-mail and voice mail, people appreciate when someone takes the time to establish personal contact; and it's especially valuable when you want to secure a sale.) In listing your advertising opportunities, itemize the creative opportunities you offer, from banners to full-section sponsorship; inform

Getting Your Ad Opportunities Noticed

With thousands of sites offering advertising opportunities, if you want to get noticed, you need to take extra steps to put your advertising opportunities on someone's radar screen. You can do this by:

- **Listing ad opportunities on your site and in appropriate rate card indexes**
- **Participating in ad networks**
- **Hiring a Web rep**
- **Putting together your own sales force**

potential advertisers of all that your site can do for them. And, most importantly, make sure you follow up in a timely manner.

Any veteran salesperson will tell you that sales are all based on relationships. Since the Web is a relationship technology, use all the resources at your disposal—both technological and personal—to give your site that "extra something" that will set it apart from the pack. And that "extra something" might just be that your advertising space is listed everywhere making it easy for the media buyer to find the information he needs to purchase ad space quickly on your site.

If you are just starting your ad program, you might want to put up your own "house ads" (ads that you use to advertise your products, services, or information available in specific sections of your site) until your list of paid advertisers grows.

One Last Word on Client Care and Feeding

With so much technology, how come it's still hard to connect with people? Selling online ads is like selling anything: the sale is in the relationship. The bottom line is take care of your clients. Provide multiple means of contact, not just an online form. Some advertisers still like to hear a human voice. Above all, answer e-mail requests promptly (see Figure 10.10).

This chapter presented the challenges of selling Web advertising in a way that we hope encourages responsible planning while clarifying the many opportunities available in this new medium. As we've pointed out, be flexible and ready to work with buyers; you could be—and should want to be—building the foundation for a profitable long-term advertising relationship.

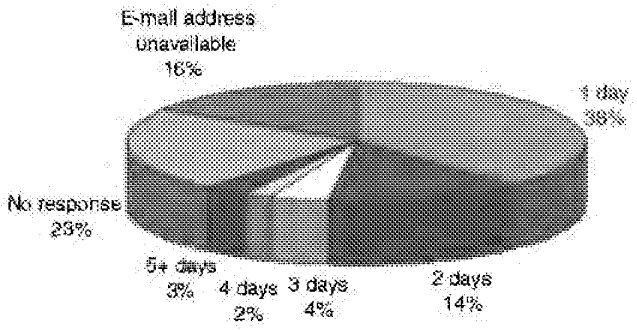


Figure 10.10 Response time to e-mailed question.

The Challenges and Future of Selling Internet Ads

Leslie Laredo, President (llaredo@laredogroup.com)
Laredo Group (www.laredogroup.com)

Internet Ad sales is one of the hardest media selling jobs. Sites and advertisers expect reps to know and understand not only their online property and media principles, but they must also be well versed in:

Technology: How Internet technology impacts their ability to create, target, and deliver advertising campaigns.

Media Measurement: Media measurement and research methodologies and why the numbers are often very different across the services and how caching and proxy servers impact the numbers.

Ad Delivery: The necessity and impact of ad operations and ad delivery options, including third-party ad serving by ad agencies and rich media technology enablers.

Business and Pricing Models: Various business models from CPM to performance-based pricing, plus hybrid options, and how to measure ROI with each.

Relationship Marketing: How to utilize advanced relationship marketing techniques within media campaigns as advertisers look to expand their online presence, improve customer relationships, increase customer involvement over a product life cycle, and increase up-sell opportunities for their brands.

Database Marketing: Selling the advantages and availability of database-derived targeting and ad delivery.

Targeting Technologies: The emerging ad targeting technologies and their impact on pricing and ROI.

Creative Technologies: Utilizing creative and rich media options for banners and evaluating the impact they have on ad effectiveness.

Ad Production: The production process for creating new ad units, micro-sites, and integrated ad programs. In traditional media selling, this would be equivalent to understanding the technologies of four-color separation, the economics of print production, the intricacies of broadcast transmission, and more. In addition it means expecting reps to be experts in all media (broadcast, print, direct marketing) and to play high level marketing consultant and technology experts.

Reps are also faced with a fast moving landscape that often means a rate card "du jour." Web sites add and change content and design frequently, and they employ new technologies. Reps must keep up with these changes and relearn the benefits of their media property.

Reps are also faced with a wide ranging and not well defined competitive environment. They are competing with like sites in their category as well as search engines, ad networks, rep firms, and traditional media.

The current shortage of seasoned Internet sales reps will continue and grow over the next five years. Most companies will not be able to hire an entire staff of seasoned Internet sales professionals so they must hire and train sales reps from other sales and media fields.

Many large, traditional media companies are watching their advertising base become more interested in their Internet properties. Advertisers are demanding integrated media proposals. Both the traditional and Internet sales organizations will have to work together and sell integrated programs. Such programs have to reach vertically across media brands (for example, a program that includes multiple Conde Nast and CondeNet properties) and horizontally to reach print and Internet audiences across one brand (for example, an integrated program that reaches Fortune magazine readers and visitors to Fortune.com).

The future of Internet selling means that traditional reps need to be conversant about their Internet properties. Internet reps need to work with their counterparts to jointly meet the needs of their common advertisers. Those companies with separate sales forces must combine their talents, experience, and sales goals—not a small feat.

From the Laredo Group's work (training the sales organizations of both Internet-only publishers and large media firms), we see a growing need to make traditional sales reps Internet literate. We also see the development of a "hybrid" or team selling approach. We believe the future success of large media companies requires training all sales organizations to sell the Internet to meet the needs of their advertisers and promote integrated media selling.

Solbright

URL: www.solbright.com

E-mail: info@solbright.com

Company: Solbright Inc.

Address: 420 Lexington Ave., Suite 300, New York, NY 10170

Phone: 212.673.6100 Fax: 212.673.6166

Description: Developer of software for managing online advertising buys, trafficking, and billing.

Straight Up!

URL: www.roianalysis.com

E-mail: sales@straightupsoftware.com

Company: Straight Up! Software

Address: 3550 General Atomics Court, Building 15, San Diego, CA 92121

Phone: 619.455.3124 Fax: 619.455.3595

Description: Internet tracking and reporting service that enables a Web marketer to fully measure and understand the impact and effectiveness of Web marketing options including banners, buttons, e-mails, media sites, interstitials, and hyperlinks.

trueVu

URL: www.matchlogic.com

E-mail: info@matchlogic.com

Company: MatchLogic

Address: 10333 Church Ranch Blvd., Westminster, CO 80021

Phone: 303.222.2000 Fax: 303.222.2001

Description: Solution for advertisers interested in delivering rich media ads. Provides site clearance, compatibility assurance, plug-in determination, targeting, efficient delivery, and centralized reporting on campaign-wide advertising traffic.

Ad Networks

@dVenture

URL: www.ad-venture.com

E-mail: sales@ad-venture.com

Company: Venture Communications Electronic Marketing Inc.

Address: 60 Madison Avenue, New York, NY 10010

Phone: 212.684.4800 Fax: 212.576.1129

Description: Places advertising on its network of business-to-business and consumer sites.

24/7 Media, Inc.

URL: www.247media.com

E-mail: info@247media.com

Company: 24/7 Media Inc.

Address: 1250 Broadway, 27th Floor, New York, NY 10001

Phone: 212.231.7100 Fax: 212.760.1774

Description: Reaching more than one out of three online users, 24/7 Media, Inc., provides Internet advertising and online direct marketing solutions for advertisers and Web publishers through its national network of sites.

Ad Club Network

URL: www.adclub.net

E-mail: members@adclub.net

Company: Ad Club Network

Address: 1100 Olive Way, Suite 1270, Seattle, WA 98101

Phone: 888.907.3224 Fax: 206.521.8808

Description: Ad network that offers real-time, online reporting, major advertiser banners, and a dedicated ad serving system, as well as a focused sales force.

Ad-Net

URL: www.ad-network.com

E-mail: admin@ad-network.com

Company: Victory Information Network

Address: P.O. Box 517, Cleveland, OH 44107

Phone: 216.961.8970 Fax: 216.732.9953

Description: An advertising network with hundreds of affiliated Web sites.

ADSmart

URL: www.adsmart.net

E-mail: info@adsmart.net

Company: ADSmart Corporation

Address: 100 Brickstone Sq., 1st Floor, Andover, MA 01810

Phone: 888.559.8222 Fax: 978.684.3618

Description: ADSmart.net automates the planning, executing, and managing of campaigns for Web advertisers, and facilitates the advertisers' selection of a highly targeted audience through its national ad network.

AdSoftware Network

URL: http://E216adnet.aureate.com/advertisers.html

E-mail: info@aureate.com

Company: Aureate Media

Address: 8777 Purdue Road, Suite 225, Indianapolis, IN 46268

Phone: 800.695.5369 Fax: 317.802.5100

Description: AdSoftware Network is a client-side software application. It allows advertisers to target their messages to users who are running the software.

Area IP

URL: www.areaip.com

E-mail: info@areaip.com

Company: Area Internet Publicidad S.A.

Address: Calle Travesera de Gracia 342-344, Barcelona, Spain 08025

Phone: 34.93.446.50.33 Fax: 34.93.446.50 34

Description: Spanish advertising network with some of the most important Spanish and South American Web sites.

ASIAD

URL: www.asiad.net/english/

E-mail: service@asiad.net

Company: Asiad Corp.

Address: 7F, 267 Hsin Yi Road, Sec. 4, Taipei, Taiwan, R.O.C.

Phone: 886.2.2755.1755 Fax: 886.2.2755.2795

Description: Taiwan's largest ad network, providing professional quality advertising and localization services to international clients.

Burst! Media

URL: www.burstmedia.com

E-mail: bhanna@burstmedia.com

Company: BURST! Media LLC

Address: 10 New England Executive Park, Burlington, MA 01803

Phone: 781.272.5544 Fax: 781.272.0897

Description: A banner brokering service that represents both small and large Web sites both national and local in focus.

Classified Warehouse

URL: www.classifiedwarehouse.com

E-mail: advertising@classifiedwarehouse.com

Company: AdOne Classified Network

Address: 361 Broadway, Suite 100, New York, NY 10013

Phone: 212.965.2900 Fax: 212.334.3307

Description: Classified Warehouse, from AdOne Classified Network, is one of the largest networks of classified advertising for employment, autos, transportation, and real estate for sale or rent.

ClickThrough Interactive

URL: www.clickthrough.ca

E-mail: sales@clickthrough.ca

Company: ClickThrough Interactive

Address: 86 Bloor Street West, Toronto, Ontario, Canada

Phone: 416.966.2542 Fax: 416.966.3097

Description: Canadian company with exclusive sales partnerships with many well-known and popular Web sites in Canada and around the world.

CLIQNOW!

URL: www.clicknow.com

E-mail: admin@clicknow.com

Company: 24/7

Address: P.O. Box 2614, Rome, GA 30164-2614

Phone: 888.213.5635 Fax: n/a

Description: Sells advertising on Web sites.

ContentZone

URL: www.contentzone.com

E-mail: csupport@contentzone.com

Company: 24/7 Media Inc.

Address: 1250 Broadway, 27th Floor, New York, NY 10001

Phone: 212.752.3928 Fax: 212.760.1774

Description: The ContentZone Network is a community of over 2,000 Web site owners and developers which targets sites by gender, age, income, and other criteria.

DMN Media

URL: www.dmnmedia.com

E-mail: comments@dmn.com

Company: Digital Music Network, Inc.

Address: 319 Main Street (Suite 1), Stoneham, MA 02180

Phone: 781.279.2895 Fax: 781.279.3507

Description: Online advertising network for music-oriented sites.

DoubleClick Direct

URL: www.doubleclick.net

E-mail: info@doubleclick.net

Company: DoubleClick

Address: 41 Madison Ave. 32nd Floor, New York, NY 10010

Phone: 212.271.2542 Fax: 212.889.0062

Description: A cost-per-action (per click, per lead, and per download) ad network.

DoubleClick Network

URL: www.doubleclick.net

E-mail: info@doubleclick.net

Company: Doubleclick

Address: 41 Madison Ave., 32nd Floor, New York, NY 10010

Phone: 212.271.2542 Fax: 212.889.0062

Description: An Internet advertising network that has local, national, and international audiences for its top brand name sites.

eAds

URL: www.eads.com

E-mail: info@eAds.com

Company: Electronic Advertising

Address: 5200 Keller Springs Road, Suite 1012, Dallas, TX 75248

Phone: 214.696.2252 Fax: 214.696.5726

Description: An ad network where the host site gets paid based on the number of responses the ad generates, not the number of impressions provided.

Flycast Network

URL: www.flycast.com

E-mail: info@flycast.com

Company: Flycast Communication Corporation

Address: 181 Fremont St., San Francisco, CA 94105

Phone: 415.977.1000 Fax: 415.977.1009

Description: Ad Network with over 700 member sites targeting national and local audiences.

HyperAsia

URL: www.hyperasia.com

E-mail: info@hyperasia.com

Company: HyperAsia

Address: 805 Third Avenue, 8th Floor, New York, NY 10022

Phone: 212.832.7000 Fax: 212.751.8888

Description: Develops targeted Chinese ads and places them strategically throughout its network of leading Chinese Web sites.

InterAd

URL: www.interadsales.net

E-mail: corporate@interadsales.net

Company: InterAd Holdings Ltd.

Address: 49-51 Carnaby Street, London, UK W1C 1PF

Phone: 44.171.287.7300 Fax: 44.171.437.0780

Description: Provides global representation for leading European and international Web sites, giving top quality services to advertisers and Web site clients alike.

Meta's Exhibit No. 1004
Page 00213

LinkExchange

URL: www.linkexchange.com

E-mail: admin@linkexchange.com

Company: Microsoft

Address: 217 Second St., San Francisco, CA 94105

Phone: 415.543.4435 Fax: 415.543.1483

Description: Largest ad network in the world.

Medius Network

URL: www.mediusnetwork.com

E-mail: media@medius.net

Company: Medius Interactive

Address: 830 4th Ave. S., Seattle, WA 98134

Phone: 206.442.9099 Fax: 206.442.9199

Description: Ad network of sports-related sites.

NarrowCast Media

URL: www.narrowcastmedia.com

E-mail: info@narrowcastmedia.com

Company: NarrowCast Media

Address: 12304 Santa Monica Blvd., Suite 101, Los Angeles, CA 90025

Phone: 310.979.4638 Fax: 310.979.4633

Description: NarrowCast Media is an advertising space brokerage and exchange service that sells banner advertising space on a click-through basis. Sites can earn cash or barter for their banner ad space inventory.

Net AD Net

URL: www.netadnet.com

E-mail: NetADNet@pobox.com.au

Company: Net Ad Net

Address: P.O. Box 521, Brentford Sq., Victoria, Australia 03131

Phone: 61.3.9803.3886 Fax: 61.3.9886.3990

Description: Australian ad banner network.

Meta's Exhibit No. 1004
Page 00214

PennyWeb

URL: www.pennyweb.com

E-mail: sales@pennyweb.com

Company: PennyWeb

Address: 468 N. Camden Dr., #200, Beverly Hills, CA 90210

Phone: 888.867.3669 Fax: 213.939.9591

Description: PennyWeb allows advertisers to have their own click-through banner campaigns.

Real Media

URL: www.realmedia.com

E-mail: rmsales@realmedia.com

Company: Real Media

Address: 32 E. 31st St., 9th Floor, New York, NY 10016

Phone: 212.725.4537 Fax: 212.725.4573

Description: Sells advertising on newspaper Web sites.

Safe-Audit

URL: www.safe-audit.com

E-mail: admin@safe-audit.com

Company: Global Market Ltd.

Address: 29 Fairholme Gardens, London, UK N3 3ED

Phone: 44.0181.346.8316 Fax: 44.0181.346.0770

Description: Banner ad network with audited statistics. Also offers banner design services.

TeknoSurf Adwave

URL: www.teknosurf.com

E-mail: john@teknosurf.com

Company: TeknoSurf Teknologies, Inc.

Address: 2315 Velvet Ridge Dr., Owings Mills, MD 21117

Phone: 410.825.6143 Fax: 410.825.6147

Description: Pay-per-click banner network.

ValueClick

URL: www.valueclick.com

E-mail: robyn@valueclick.com

Company: ValueClick

Address: 1333 De La Vina Street, Suite E, Santa Barbara, CA 93010

Phone: 805.965.0543 Fax: 805.564.7151

Description: Pay-per-click banner network with premium quality ad space on over 10,000 sites.

WebCounter

URL: www.digits.com

E-mail: questions@digits.com

Company: Computer Networking Services Inc.

Address: 322 Mall Blvd., Suite 306, Monroeville, PA 15146-2229

Phone: n/a Fax: n/a

Description: Provides ad banner opportunities on the sites of its 300,000-plus active subscribers.

The Women's Forum

URL: www.womensforum.com

E-mail: jodi@womensforum.com

Company: Women's Forum

Address: 253 E. Delaware Place, Suite 21 C, Chicago, IL 60611

Phone: 312.475.0100 Fax: 312.475.0116

Description: Ad network containing sites for women of all ages.

Alternative Ad Models

Alexa

URL: www.alexa.com

E-mail: info@alexa.com

Company: Alexa Internet

Address: The Presidio of San Francisco, P.O. Box 29141, San Francisco, CA 94129-0141

Phone: 415.561.6900 Fax: 415.561.6795

Description: A free advertising-support Web navigations service that provides Web site recommendations as the user surfs.

Glossary

The only thing harder than keeping up with what the Internet can do is keeping up with what things are called. Indeed, no one is shy in the online world when it comes to making up new words for tools, techniques, and technologies. The following glossary covers most of the terms and buzzwords presented in this book. However, by the time you flip through it, another new word will have entered the online vernacular. Recognizing the organic nature of Internet vocabulary terms, we encourage you to bookmark some online glossaries that will keep you up to date on the latest terms to hit the industry:

- The Marketing Manager's Plain English Internet Glossary (www.jaderiver.com/gloss1.htm)
- WebPromote Glossary (www.webpromote.com/tools/glossary.asp)

Ad request. When an ad is requested from the server. This happens when someone visits a Web page that has an advertisement and the surfer's browser asks the server to deliver the ad. For a variety of reasons the ad may not always be successfully served.

Ad view. Technically this is when an ad is delivered. However, it is often used to mean ad request.

ALT tag. Code that tells your browser to show specific text while a graphic is loading. When a banner is downloading, the ALT text can be reinforcing the ad's message with a simple line of text.

Auditor. A third party that verifies ad delivery and site traffic.

Bandwidth. How much information can be transferred through an Internet connection at a given time, usually measured in seconds.

Banner. A rectangular ad usually appearing at the top or the bottom of a Web page, commonly 468 x 60 pixels in size.

Meta's Exhibit No. 1004
Page 00217

Bookmark. A Web browser feature that serves as an address book for Web addresses.

Browser. The software that runs on your computer that translate the information transferred on the Internet into the Web pages that you see.

Button. A small, usually square or rectangular ad in the following dimensions: 125 × 125, 120 × 90, 120 × 60, 88 × 31, or 120 × 240 pixels.

Cache (pronounced *cash*). The temporary storing of a Web page on your computer following the page downloading. This speeds up downloading on repeated use, so that the next time you request that page, it is accessed from the cache on your computer instead of from the Web server. Proxy servers and ISPs may also cache pages to make access easier for their customers.

Click-through. The number of times an ad is clicked; when a user clicks on an ad that is linked to another page and that page downloads on to the user's browser.

Click-through rate (CTR). The percentage of times the ad is clicked on divided by the total number of times an ad is viewed.

Closed loop reporting. Measuring the achievement of goals beyond simple click-through—orders, leads, etc. that are generated from specific banners on specific sites.

Conversion. The rate at which a goal is achieved (for example, orders or leads).

Cookie. A text string that is placed on your computer (in your cookie file) by a Web site. The Web site can view the cookies it places when you move through the site, and use cookies as a means of tracking return visits and interests. (See Chapter 4.)

Cost per click (CPC). The price of placing an ad on a site based exclusively on how many times the ad is clicked on by users.

Cost per lead (CPL). The price of placing an ad on a site based on how many leads the ad generates.

Cost per sale (CPS). The price of placing an ad on a site based on how many sales result from the ad.

Counter. A program that a Web site uses to count the number of visits to a Web page.

CPM (cost per thousand). The cost of an ad for every 1,000 times the ad is shown on a Web page. (The "M" stands for the Latin mille, meaning "one thousand.")

Directory. A listing of Web sites organized into categories. The Yahoo! site is a directory.

Discussion list. An e-mail forum for people to discuss a particular topic. Discussion lists usually have a moderator who guides the discussion and ensures that it is pertinent to the discussion list's topic.

Domain name. A naming system that translates the numeric name given computers on the Internet into user-friendly addresses. Domain names are read from left to right and go from specific to general. For example, the Web address for Yahoo is www.yahoo.com and is dissected as follows. The Web addresses (URL) starts with the type of machine (www denotes a World Wide Web server), then the general domain name, which is yahoo, and then the top-level domain (.com for commercial). There are two types of top-level domains: categories (com, edu, org, etc.) and countries (uk—United Kingdom, au—Australia, etc.).

Dynamic rotation (or dynamic delivery). Ads delivered on a rotating or random basis, allowing different users to see different ads on a given page and for an ad to appear on more than one page. Ads can also be dynamically rotated through a particular section of a site or can be called up as part of a keyword search. (Contrast with *hardwired.*)

E-mail list. Similar to a mailing list, except that it contains e-mail addresses.

Exposures. The number of times that an advertisement is viewed. Commonly used interchangeably with *impressions.*

Extranet. A private network of networked computers that uses Internet protocols to share information between businesses. It is often part of a company's intranet that is shared with vendors and customers.

File. A computer document.

Firewall. A security barrier that separates a company's internal systems from outside systems.

Flame. A hostile note either posted in a discussion list or sent as an e-mail message.

Flame war. The exchange of hostile messages between two or more people.

Flight. An advertising campaign on a specific Web site for a specific time period. Also, to send out a creative to a Web site.

GIF or .gif (Graphic Interchange Format). A format for saving images for use on the Web. Mostly used for line art and simple logos and not meant for high quality art or photographic images, which require the more robust *JPEG* format.

Gross exposures. This is another term used to define an advertisement that is viewed.

Hardwired. Manner of placing an ad in a fixed position on a particular Web page so that it is delivered each time the page is downloaded (in contrast with *dynamic rotation*).

Hit. A line that is recorded in a log file when something is requested of a Web server. For example, if someone goes to a Web page that has five graphics, six hits will be recorded in the log file: one hit for the HTML page and one for each of the graphics. Hits are considered an inaccurate way to measure traffic and are no longer used by savvy advertisers.

HTML (HyperText Markup Language). The coding language used to make Web pages. HTML works by surrounding text with codes, called tags, that identify how the text should appear. For example this designates that the word **example** should be bold. HTML also lets you link one file to another.

Hypertext. Text that when clicked on causes the browser to jump to another file either on the same Web site or in another location on the Web. Browsers generally display hypertext by underlining, with the cursor changing into a hand when placed over hypertext or linked images.

Impressions. The number of times an ad is delivered. When an advertiser buys advertising on a CPM basis, the advertiser is paying for every 1,000 impressions that the site can deliver. Different sites measure an impression differently. Some sites count an impression when an ad is requested, others only when an ad is fully downloaded,

and there are other definitions that are used as well. A savvy advertiser will ask publishers how they define impressions.

Internet. The global network of networked computers.

Interstitial. An advertisement that interrupts the user. This can be a full page ad that pops up on the user's screen or a pop-up window.

Intranet. A network of networked computers that is contained within an enterprise. In other words, it is a private Web site usually servicing a single company. Its primary purpose is to share company information and computer resources with staff.

Java. The programming language developed by Sun Microsystems to create self-running applications that can be easily distributed through networks like the Web.

Java applets. Small Java programs used on Web pages to operate animation, calculators, and other tasks.

JPEG or .jpg (Joint Photographic Experts Group). A format for saving photographic and high quality images for use on the Web.

Keyword. A word or series of words that describe a page or site.

Link. A hypertext entry that lets a reader jump from one Web page to another page or file either on the same site or on another Web site.

Log file. A text file that records all activity on a server: the items that are requested (Web pages, graphics, etc.), the times they are requested, the browser used by the surfer, and other information (see Chapter 4). Traffic analysis software crunches log files to produce traffic reports that are easy to understand and use as a basis for measuring Web site success.

Meta tag. A tag that appears in the head portion of an HTML page after the title tag where you can list a series of keywords for your site.

Netiquette. Proper Internet behavior.

Newsgroup. A Usenet discussion group that talks about a specific topic

Page view. A common metric for measuring Web site traffic, usually representing the viewing of all of the elements that comprise a Web page (graphics, text, etc.). It is generally agreed that if someone looks at a Web page that has frames, a page view will be looking at all of the elements of the page. However, some advertisers will count each frame as a separate page view to boost their numbers for level of traffic.

Pay per click (PPC). An ad pricing structure by which the advertiser pays the publisher according to how many times an ad is clicked on by users (see *Cost per click*).

Pay per lead (PPL). An ad pricing structure by which the advertiser pays the publisher according to how many leads are generated by an ad, often determined by information submitted directly into the banner ad.

Pay per transaction (PPT). An ad pricing structure by which the advertiser pays the publisher according to the number of sales generated by an ad.

Plug-in. A small application file that is added on to a browser to allow it to play certain applications such as Shockwave and flash files. Microsoft Internet Explorer uses the term *ActiveX controls* instead of plug-in.

Reciprocal link. A link provided from one site to another in exchange for a link back.

Response rate. The percentage of impressions of an ad that result in the ad being clicked.

Rich-media ads. Ads incorporating video, audio, and other technology components beyond simple animation. Rich media is considered higher bandwidth advertising that delivers more of a brand impact than an animated banner advertisement.

Robot. A program that automatically surfs the Web. Search engines use robots to surf the Web and catalog different Web sites in their databases. This allows the Web pages to be found when someone performs a search. Robots are commonly referred to as bots and spiders.

Search engine. A tool for finding information on the Web. Users can type in exact text, or *keywords* representing their interests, to initiate the search. Popular search engines include Lycos, AltaVista, InfoSeek, and HotBot.

Shockwave. A program (see *Plug-in*) that allows Macromedia Director animate files to be played on a browser.

Signature. Text often appended to the end of someone's e-mail message to identify them. It can contain name, company name, and contact information. Many people use it as a form of advertising themselves when they post to an e-mail discussion list.

Site traffic. The amount of activity on a Web site. This is usually measured in page views or visitors.

Spam. On the Internet, not processed meat but unwanted and unsolicited messages, either on a discussion list or as an e-mail message.

Unique users. Individual visitors who have been to a Web site within a given amount of time.

URL (Uniform Resource Locator). A Web address.

Usenet. The Internet's bulletin board system of discussion groups on specific topics. There are more than 25,000 Usenet discussion groups today.

Visitor. A person who goes to a specific Web site. Sites often define their traffic levels in terms of the number of visitors they've had in a given time period.

Visits. Another way to compare and measure Web site traffic. A visit is activity on a Web site from a specific individual. That activity usually counts as a new visit if the individual has been away from the Web site for a period of fifteen minutes or longer.

Web page. A page of a Web site.

Web site. A collection of documents made available to users by a publisher on the Internet. These can include news and entertainment centers as well as corporate information sites.

Internet Business

$24.99 USA/$38.95 CAN

LIBRARY OF CONGRESS

How today's most successful and marketers maximize their online preser an too

0 006 206 247 2

Advertising on the Internet, Second Edition

This updated and expanded *Second Edition* of the bestselling guide to online advertising is must reading for everyone who wants to take advantage of the most important new advertising medium since television. A complete primer on online advertising for businesses of all sizes, it gets you up to speed on the crucial issues, hot new trends, and most effective new technologies in Internet advertising. Loaded with examples of some of today's most successful online advertising and marketing initiatives, it gives you the inside track on:

- Successful online ad models
- Market research online
- Direct marketing, including opt-in e-mail, promotions, and sweepstakes
- Targeting and personalization
- Internet advertising management tools
- Traffic measurement and gauging the effectiveness of your ads
- Buying and selling ads on your Web site
- Advertising locally
- Advertising to an international market
- Legal aspects of Internet advertising
- How to advertise for free, or almost free (contributed by Eric Ward, the "father" of grassroots advertising strategies)

BOBBIN ZEFF is President of Zeff Group, a firm specializing in strategic Internet advertising and marketing, research, training, and consulting.

BRAD ARONSON is President of i-frontier, a leading Internet advertising agency.



On the companion Web site at www.wiley.com/compbooks/zeff you'll find:

- Links to additional training and tool resources

Wiley Computer Publishing | Timely. Practical. Reliable.

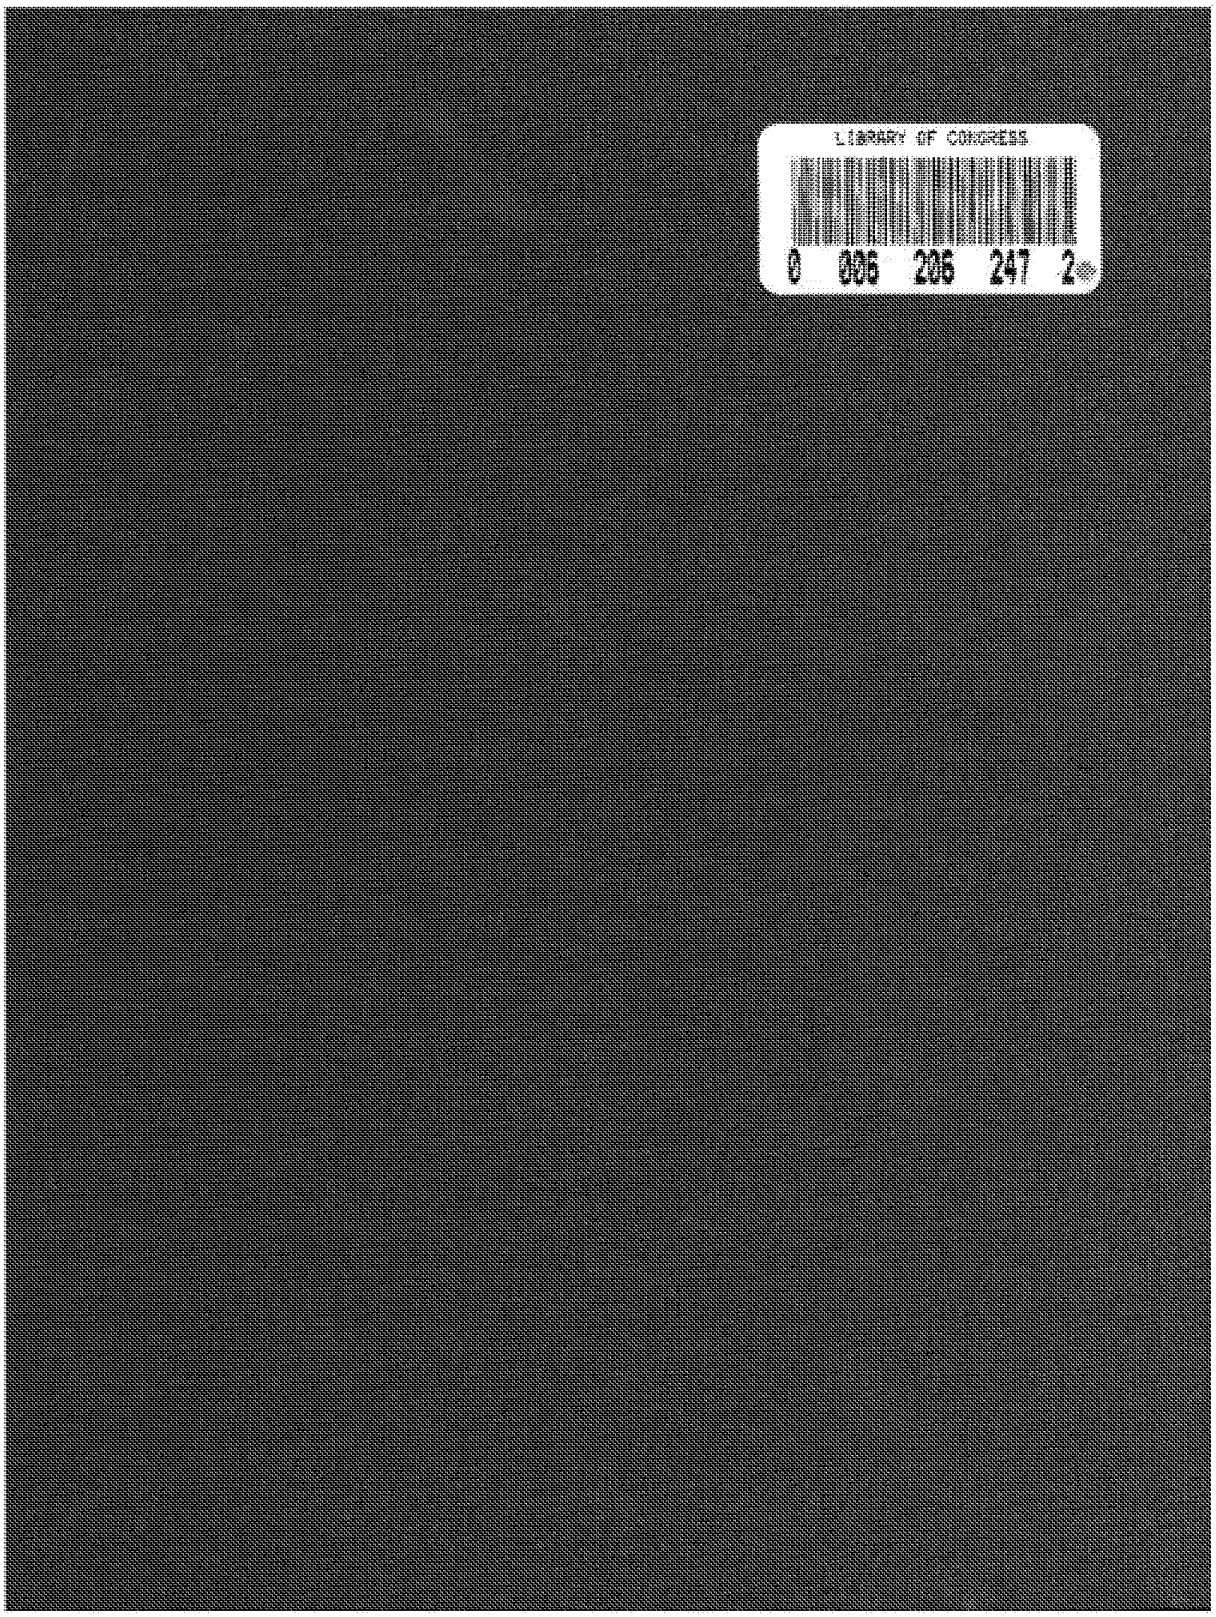
Visit our Web site at www.wiley.com/compbooks/

Cover Design: Tom Draper Design

John Wiley & Sons, Inc.
Professional/Trade Division
605 Third Avenue, New York, N.Y. 10158-0012
New York • Chichester • Weinheim
Brisbane • Singapore • Toronto

ISBN 0-471-34404-N

Meta's Exhibit No. 1004
Page 00222


LIBRARY OF CONGRESS
0 006 206 247 2

The primary considerations in determining your rate structure should be the size and diversity of your ad inventory, your revenue needs, and what the advertisers have come to expect and demand from the market.

Size and diversity of ad inventory. The size of your site's audience and how much ad space you have to offer are important factors in determining the best rate structure for your site. If your site only has 50,000 page views per month, then charging by CPM might not be in your best interest; if you offer one banner per page, the number of impressions you have on your site is 50,000, and if you charged a $32 CPM, the maximum revenue you could expect is $1,600. For a small site, a flat fee on a monthly, quarterly, or yearly basis offers the site more revenue potential. Also, multiple ad units may allow the publisher to sell more than one impression on a page. For example, the publisher could sell a banner impression on the top, a button on the right, and a text link somewhere else on the page. If the site is extremely well targeted, the publisher may even be able to get away with selling multiple banners on each page. A site's ad inventory (the total number of impressions available for sale) is rarely static. There are many factors that can influence the size of the inventory beyond the number of pages with ads and the number of ads per page. A site's inventory can increase or decrease by fluctuations in the number of visitors, the amount of time each visitor spends on the site, and the number of page views by a visitor. In short, if site traffic drops, ad inventory drops.

Revenue needs. This is always a loaded question. Advertising can generate revenue, but few sites rely solely on advertising. Advertising should be part of the site's overall revenue strategy. A diversified revenue approach seems to be the safest strategy on the Web. Putting that aside, keep your expectations realistic. Advertising is not the silver bullet, and only a handful of sites are making serious money. Start by trying to cover the costs of running the advertising program, and build from there.

Market structure. With so many sites accepting advertising, every site has competition. Your advertiser is probably reviewing your competition's advertising opportunities as well. Know the rates your competitors are charging and make your competitive advantage value, quality, and aggressive pricing.

As discussed in Chapter 7, there are several pricing models to choose from; CPM, flat-fee, click-through, and cost per action are among the most common. To decide this important issue, publishers need to take a realistic look at the value of the advertising they're offering. If you are building a site that will be in the top one percent of traffic generators with several million impressions per day, you can charge a lower CPM. For such high-traffic sites with broad appeal, the attraction to advertisers is simply volume.

But most sites do not fall into this category. Most cater to an identifiable group: teens with expendable income, golfers, model rocket hobbyists, doll collectors, and so on. And, this is what the Internet does best: provide access and exposure to niche markets. Such targeted sites can charge a higher CPM. They have a limited inventory to sell, but what they have is the best vehicle for reaching a particular audience.

If you are still in the process of building your site and need sponsors to provide working capital, then a flat-fee model would work best. Rob Frankel, marketing executive at Frankel & Anderson (www.frankel-anderson.com), sees a clear advantage for small sites to accept fixed pricing. "Fixed pricing gets you paid. Impression-based takes for-

ever to calculate and collect, and is almost never worth the same as fixed price. Take the money up front, not 'on the come.'" The flat-fee model also works best for small sites that don't generate enough traffic to charge reasonable CPMs. A flat fee of $250 will seem more reasonable to advertisers than a $125 CPM.

If your site generates a lot of response for advertisers, you may be able to charge a higher rate. You might be able to lure advertisers with an attractive test rate for first-time clients. If response on the site is good, the advertisers will be willing to pay a higher rate after the test. If your Web site is having trouble selling inventory, offering per-click-through advertising is usually a good way to get advertisers to test the site.

As noted earlier in the book, the market rate for impressions and click-throughs constantly changes. Therefore, the best way to find out the going rate is to compare your site with the rate sheet for similar sites.

If you are a small to medium-sized site, the best pricing model might be an individual arrangement with each advertiser. If you are a large site, this might be something you can provide only for your major advertisers. As with any sales deal, be flexible with your pricing and be prepared to customize on demand. However, if you offer someone a great deal to test your Web site, make sure he or she is aware that this is a one-time pricing break. You don't want to start a trend in which you lose out on revenues.

Since the Internet was built on barter, and it's a buyer's market, a rate sheet is often viewed as the opening bid, with the buyer and seller negotiating from there.

Once you determine the rate structure, then you are ready to calculate the rate for each ad product. The rate prices are determined by such factors as:

- size of site audience
- content of site
- ad's ability to reach and interact with the audience (targeting and ad format)
- ad placement (top or bottom of page)
- number of ads on a page
- dynamic delivery to targeted audience

Agency and Sales Commissions

In calculating your rates, don't forget agency commissions (customarily 15 percent). The agency gets their commission from the 15 percent discount that sites give agencies. The agency keeps the 15 percent so that there is no extra cost to the advertiser. This is how many agencies are paid for their services, so not offering this commission could make you a less attractive buy.

Discounts and Special Deals

As in retail, discounts and special deals work in selling advertising space. Remnant sales and end-of-quarter sales are always welcomed by media buyers.

It is not uncommon for a site to have remaining impressions at the end of month. According to Richard Hoy, moderator of the Online Advertising List, a "PR-centered" answer is to overdeliver. "The client feels like she got a little something for free, and you aren't really losing anything. I mean, what the heck are you going to do with those extra impressions anyway?" For example, if you have 100,000 impressions to sell, you may only sell 95,000 of them. Then, at the end of the month, you can give a client the 5,000 remaining impressions for "free." Because you keep track of the traffic to your site, you know you will have the impressions to deliver your minimum and more. Overdelivery builds goodwill and loyalty, or can be a safety net if something unexpected should occur. For example, if an ad campaign is pulling really poorly, the advertiser may ask for over-delivery to compensate for the poor performance.

Exclusivity is a standard in traditional advertising and a popular strategy in the arsenal of online advertising deals. Exclusivity means that for an increased price, an advertiser is the only advertiser in a specific category or is the only vendor of a type of product or service on a site. For example, a travel firm can be the exclusive travel agency for a site. Or a bookseller can be the exclusive bookstore of a site. The benefit to the advertiser is premium exposure. The benefit to the publisher is increased revenue.

Selling Day Parts

Day part advertising is starting to take hold online, and being able to offer it could give your site an advantage over competitors who don't. A commonly used tactic by television and radio advertising, careful timing of when an ad appears can result in better response. Auction sites are doing it: iDeal increases its use of banner ads in the hours before the midday bidding deadline for its auction site www.dealdeal.com. "If we get people to bid at the end of an auction," says Jamie Locke, iDeal co-founder, "that helps move prices into an acceptable range for merchants." Internet advertising veteran John Audette of MMG has found that lunch hour advertising can increase click-through rates by 25 percent. Other key times are right after work (5–6 P.M.) and after 10 P.M. In short, "Everything we know from traditional media about day parts must be translated to the Internet," says Laura Berland, of ORB Digital Direct. "There is not just a right place but also a right time to convey your message."

Competitive Analysis

It is important that you know your competition so that you can anticipate and answer advertisers' questions. Review other sites in your category, and know their audience and advertising program. Your advertisers will. Even if you see a world of difference between your site and the other sites in your category, your advertisers may not. It's best to take the philosophy that those sites are great, but your site is different, and it is that difference that will help advertisers reach their campaign objectives.

The Web offers a veritable cornucopia of tools and services for conducting competitive research and learning from other advertising campaigns. If you want to learn about banner ad and its relation to CTR (click-through rate), check out the CTR Comparison site at www.whitepalm.com/fourcorners/ctrcomparison.shtml and see how banners that ran on Yahoo!, LinkExchange, and other sites performed. If you're interested in

following an ad's placement, you can use NetRating's Online Observer (www.netratings.com), which offers ad banner tracking. This service shows you the ad, where it ran on any particular week, how many unique impressions were delivered, the demographics of those impressions, and the click-through rates achieved. You can even drill down to the individual pages on which the ads were posted.

Media Kit

Whether you are selling advertising for a Web site or an e-mail newsletter, you need a media kit that spells out your advertising opportunities (see the companion Web site for sample media kits). The online version should be a link accessible on every page and titled "media kit," "advertising opportunities," "how to advertise with us," or even just "advertising" or "sponsorships."

What makes a good media kit? One that works. Barbara Feldman, syndicated columnist and publisher of the site "Surfing the Net with Kids" (www.surfnetkids.com), has found that her simple and straightforward online media kit is producing results. "I'm selling at least one ad campaign (of various sizes, of course) per month with *only* in-bound leads. In fact, I can't even use the word 'sell.' I'm just taking orders from those that have already read my media kit."

There are nine key elements that need to be included in an online advertising program media kit. These elements can be presented as separate sections or consolidated into sections of a single page:

1. Site overview and features
2. Contact information
3. Advertising and sponsorship programs
4. Advertising rates
5. Site traffic
6. Audience demographics
7. Production specifications
8. Delivery specifications
9. Reporting

Site Overview and Features

The media kit needs an introductory page, which serves as a springboard to all of its elements. The page does not need to be lengthy. Its purpose is to set the tone and offer a text overview of the pieces in the media kit as well as a list of links to each of these elements.

Most importantly, this is the page where you put your site's best features forward. In short, this is where the hard sell occurs, with all the other pages serving as supporting documents. For a sample media kit introductory page that uses a road map format, see Figure 10.2 from Women.com. For a sample of a page that spells out the value of advertising on the site, see Figure 10.3 from PriceSCAN.

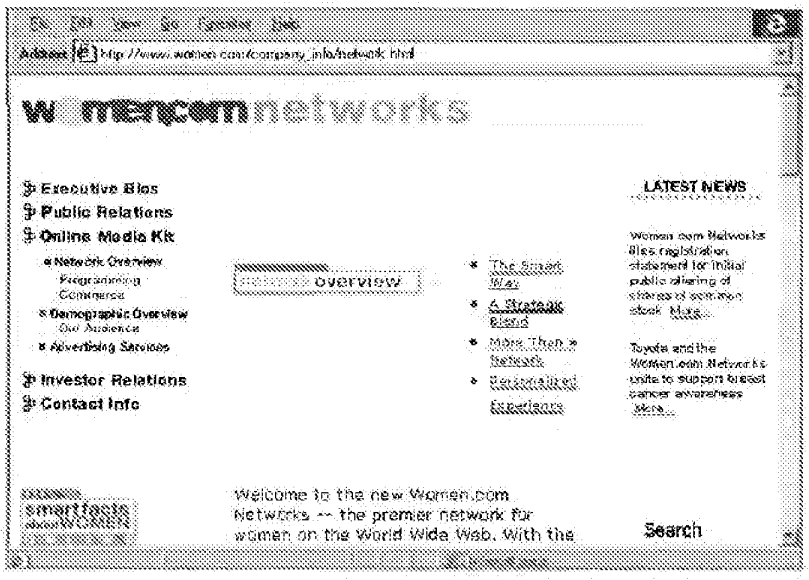


Figure 10.2 Women.com media kit uses a road map format listing the sections of the media kit in the left column.

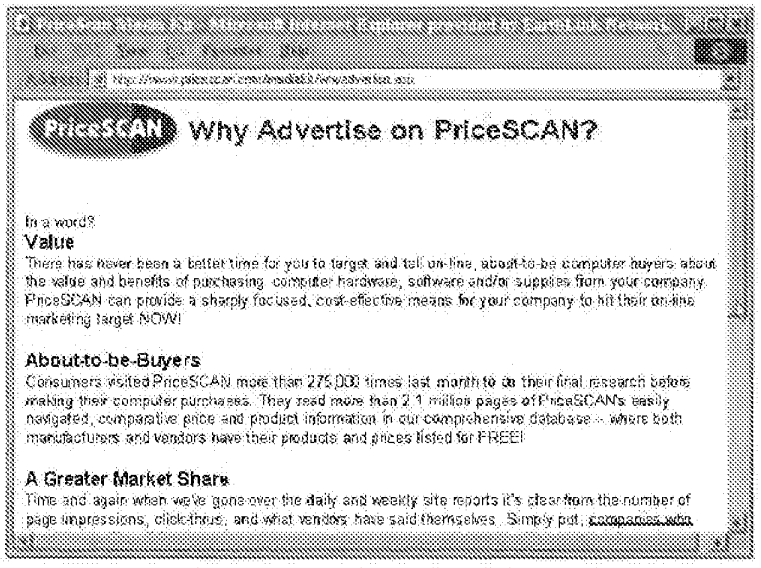


Figure 10.3 PriceSCAN media kit spells out the value of advertising on the site on the introductory page.

Contact Information

This can either be a separate page or incorporated in a page with other information. Wherever this information is placed, it should be easy to find and list full contact information for the person or persons handling the advertising and sponsorship program. Full contact information includes name, title, address, phone, fax, and e-mail (on the Web page the e-mail should be a mailto: link). Keep in mind that people buying ad space often like to deal directly with a person; a generic contact fill-in form may be efficient for you, but sales is about building relationships. There is nothing warm, fuzzy, or personal about a form. A person at the other end of the telephone is the best approach.

Advertising and Sponsorship Programs

This is where the advertising and sponsorship options are spelled out for potential advertisers. It can be in any design format—tables or text—but keep it simple and provide answers to all question you can anticipate. For example, in the Inc. Online media kit, the banner advertising program is separated out from the sponsorship program, with each receiving its own page (see Figure 10.4). A site offering advertising on many different areas of the site, such as Talk City with its multiple communities, will often spend time going into the value of the offerings in each location (see Figure 10.5).

You might also want to list some of your current advertisers, especially if the list is long and impressive. Nothing motivates an advertiser like competition—not to mention that it indicates that your site is a proven entity. However, it is important to realize that if you list



Figure 10.4 Inc. Online media kit has a separate section for banner advertising.

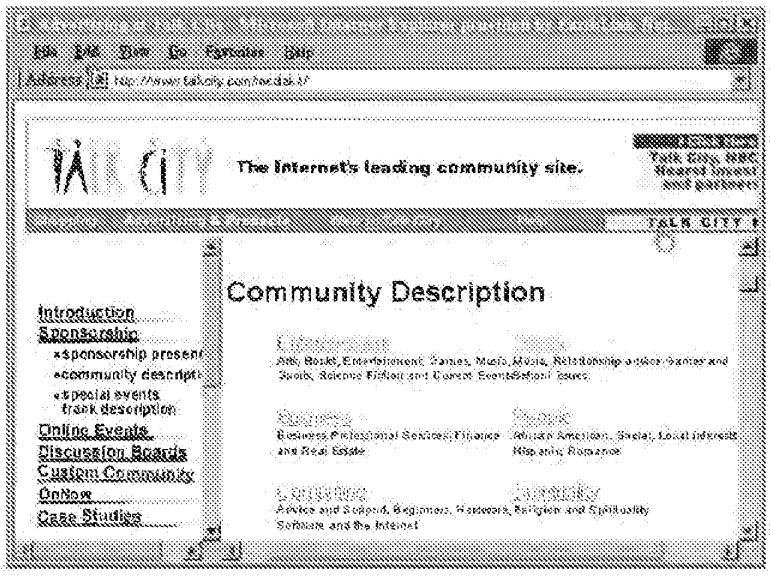


Figure 10.5 Talk City media kit has a separate section for each of its community areas.

your advertisers, that information will be accessed by your current and future competitors, who will begin to call on your clients (if they don't do so already).

Advertising Rates

This is the page where the rates are listed, and it is referred to as the "rate card." Place your rates up front and center; don't make the buyer dig for this information. There should be no mystery as to your fees. Also, date your rate card so the media buyer knows that the rates are for the current quarter or time period. The LATINOLink rate card is simple—one page, no graphics. It lists the CPM, the cost per month, insertion due dates, commission policy, and right of refusal policy, with complete contact information on the bottom of the page (see Figure 10.6). Sample rate cards are provided in this book's companion Web site.

Site Traffic

This is the area where information from the Web site measurement software is displayed. Charts and graphs work particularly well in this section. Information of particular interest to advertisers includes impressions throughout the site and in all categories where you sell advertising, and traffic growth over the past three to six months. As the saying goes, a picture is worth a thousand words, and nowhere does a chart work better than in reflecting site traffic. In the HomeArts media kit, the one page that is shortest is the one that documents the site traffic; a simple chart and table say it all (see Figure 10.7).



Figure 10.6 LATINOLink chose a simple and straightforward rate card format.

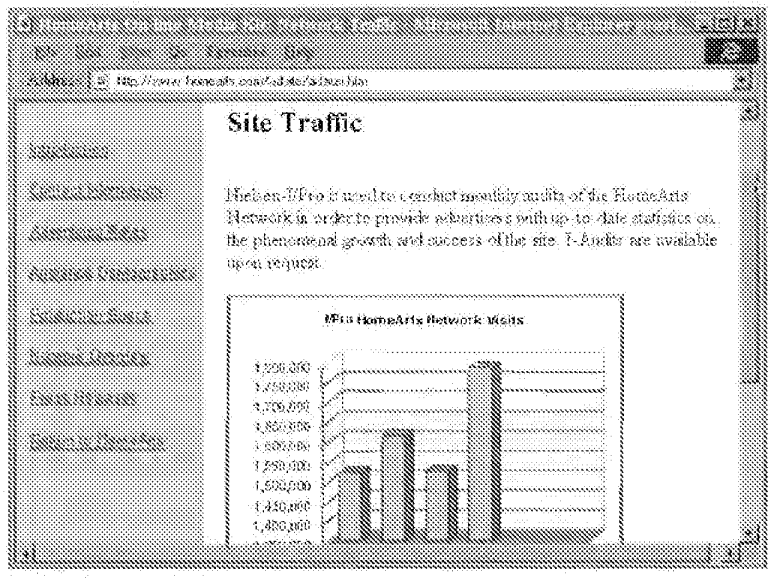


Figure 10.7 Site traffic presented as a simple chart and table in the HomeArts media kit.

Audience Demographics

This is where information on the site's audience is presented, such as age, gender, income, and geographic distribution.

Production Specifications

This is the area where the design specifications are provided for each ad product. The information needed includes:

Ad dimensions. Whether it is a button (120 x 60 pixels), a banner (468 x 60), or other dimensions.

Ad due date. How soon before a flight the creative must be delivered.

File format. Most common is GIF, with many sites accepting JPEG. Indicate type of looping allowed and rich media rules.

File size. Most sites require 12 KB or less.

ALT text. The specs for the text that will appear in the ad space before the ad downloads.

Delivery Specifications

This section delineates how ads need to be sent to the site. Options include e-mail, Web, or ftp.

E-mail. Name of the person to whom the e-mail should be sent, and how the ads should be named.

Web. Instructions for sending an e-mail with URL for pickup location of banners.

FTP. FTP address and file location (name and password as well, if necessary).

Reporting

This section denotes how you are going to provide the advertiser with reporting on the performance of their ad or sponsorship; whether this information is accessible in real time or on a daily or weekly basis; and whether this information will be e-mailed or is accessible online. Also provide information if your site is professionally audited by one of the industry firms (as discussed in Chapter 4). If your site is audited, publish this information as well. Inc. Online does this on one screen and makes the information short, sweet, and to the point (see Figure 10.8).

Deciding the Format of Your Media Kit

There is no clear-cut rule as to whether you need both an online and print version of the media kit. Many find an online-only version sufficient. Others find a value in having a print version as well. "We've found that hard copy media kits work better in terms of sales than online only," says Eric Ng of Student.com (www.student.com). Still others don't want to make their rates available online, because of a fear that their competition will see.

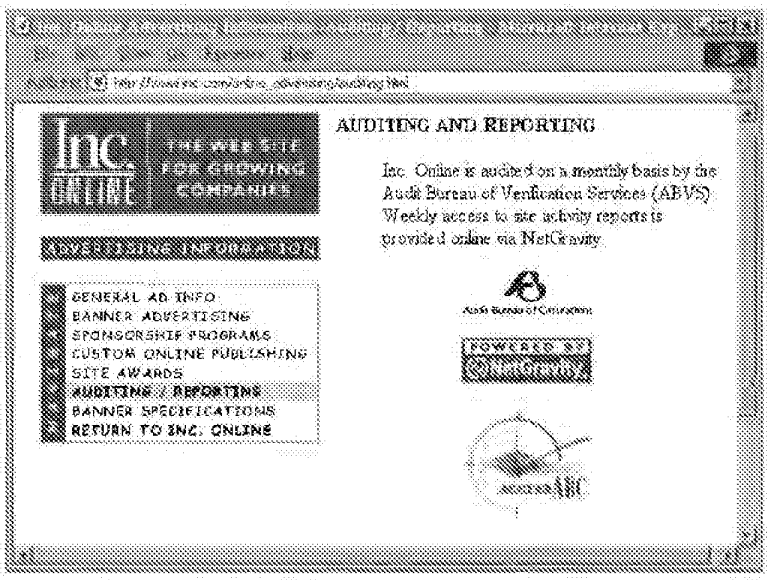


Figure 10.8 Inc. Online's audit information: short, sweet, and to the point.

Be sure that the media kit is full of good information, and doesn't simply look good on the outside. A fancy design won't answer a media buyer's questions. Dana Blankenhorn, editor of a-clue.com and veteran journalist, puts it this way: "The money spent in packaging a press announcement or media kit is often inversely proportional to its importance. An expensively wrapped announcement of little real news value is a sure sign of a company out of ideas." Whether you go paper or plastic, online-only or online and print versions, it's what's inside that counts. What's a solid media kit? One that works.

A good media kit is essential, but it doesn't substitute for strong sales. "It's hard to say exactly how much of our success had to do with the media kit," says Ann LeRoy, vice president of Marketing and Business Development at PriceSCAN, "and how much can be attributed to my kick-ass determination to make us profitable, by working, calling, writing, and pitching some more." As LeRoy states, it takes a lot more than a media kit to generate ad sales.

Sign on the Dotted Line

Consult your legal counsel to draw up an insertion order for your advertising space. The insertion order should include the specifications of the ad placement, name of advertiser, duration of the flight, cost, start and stop dates, and any other special orders. The contract should also include a payment schedule, as well as any other terms and conditions of particular importance to you and your business. A sample insertion order from Townonline (www.townonline.com) is given in Figure 10.9, which is an excellent example, but not meant to be a boilerplate. You should consult your legal counsel to have an insertion order developed for your particular needs.

Information Media Buyers Need to Know

- Size of audience
- Demographics of audience
- Targeting capabilities for ads and for content
- Flight schedules (daily, weekly, monthly)
- Guarantees and make-goods
- Reporting and measurement capabilities
- How the site defines terms like "impression" and "page view"
- What gets measured
- Availability of reports and access to data (real time)
- Rates
- Inventory available
- Ad specs
- Joint marketing opportunities (online and off)
- Sponsorship (co-branding) opportunities
- Content integration (product placements)

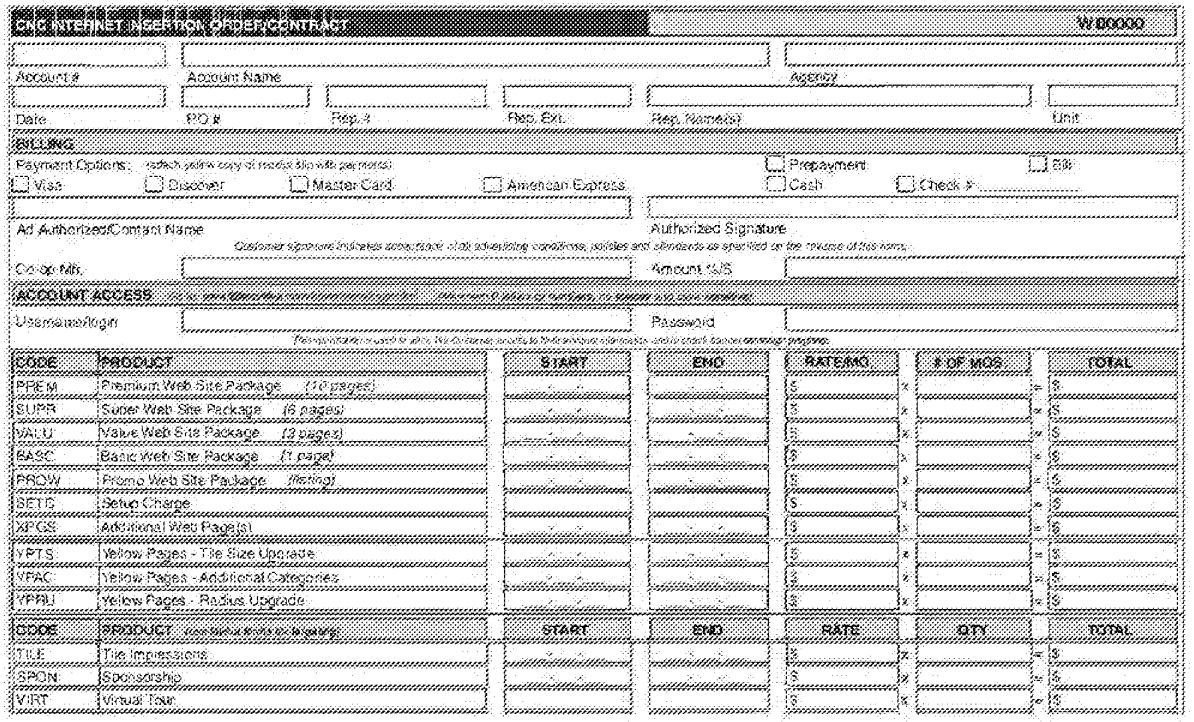
Selling Strategies

Selling can be the most exciting part of the advertising process, but it is also the most difficult to do right. In addition to readying your site to accept advertising, you need to get the right people to sell the ads. Those people may be an in-house sales force, a representation firm, an ad network, or an online ad auction service. It may even be all of them. To determine what will work best for you and your company, examine the benefits of all the different sales strategies, and remain flexible so you can implement an alternative approach if one proves to be ineffective.

Site Representation Firm

A site representation firm—informally known as a rep firm—specializes in selling advertising and sponsorship packages for a select slate of clients (see Table 10.1). These firms usually have an exclusive agreement with a site to act as its sales representative. Representation firms charge a commission for their services, with fees ranging from 35 to 40 percent. They carefully select properties they represent and often have a monthly impression minimum for their sites (averaging one million impressions per month). It should be noted that the monthly minimum does vary, and firms do occasionally represent sites with smaller monthly impressions if those sites are market leaders of a clearly defined and lucrative target audience. Rep firms also generally require third-party research to validate, identify, and document audience size and demographics.

Representation firms can have different clientele focuses. For example, HerAgency.com (www.heragency.com) focuses exclusively on the female demographic online by partnering with high-quality, female-focused Web sites to assist advertisers in reaching women



CNC INTERNET INSERTION PREPAY CONTRACT W/00000

Account # | Account Name | Agency

Date | P.O.# | Rep. # | Rep. Ext. | Rep. Name(s) | Unit

BILLING

Payment Options: (attach yellow copy of media slip with payments) ☐ Prepayment ☐ Bill

☐ Visa ☐ Discover ☐ MasterCard ☐ American Express ☐ Cash ☐ Check #

Ad Authorized/Contact Name | Authorized Signature

Customer signature indicates acceptance of all advertising conditions, policies and standards as specified on the reverse of this form.

Co-op Mfr. | Amount %/$

ACCOUNT ACCESS

Username/login | Password

CODE	PRODUCT	START	END	RATE/MO		# OF MOS		TOTAL
PREM	Premium Web Site Package (10 pages)			$	x		=	$
SUPR	Super Web Site Package (6 pages)			$	x		=	$
VALU	Value Web Site Package (3 pages)			$	x		=	$
BASC	Basic Web Site Package (1 page)			$	x		=	$
PROW	Promo Web Site Package (listing)			$	x		=	$
SETC	Setup Charge			$	x		=	$
XPGS	Additional Web Page(s)			$	x		=	$
YPTS	Yellow Pages - Tile Size Upgrade			$	x		=	$
YPAC	Yellow Pages - Additional Categories			$	x		=	$
YPRU	Yellow Pages - Radius Upgrade			$	x		=	$

CODE	PRODUCT (see reverse for listing)	START	END	RATE		QTY		TOTAL
TILE	Tile Impressions			$	x		=	$
SPON	Sponsorship			$	x		=	$
VIRT	Virtual Tour			$	x		=	$

Figure 10.9 Insertion order form from Townonline.

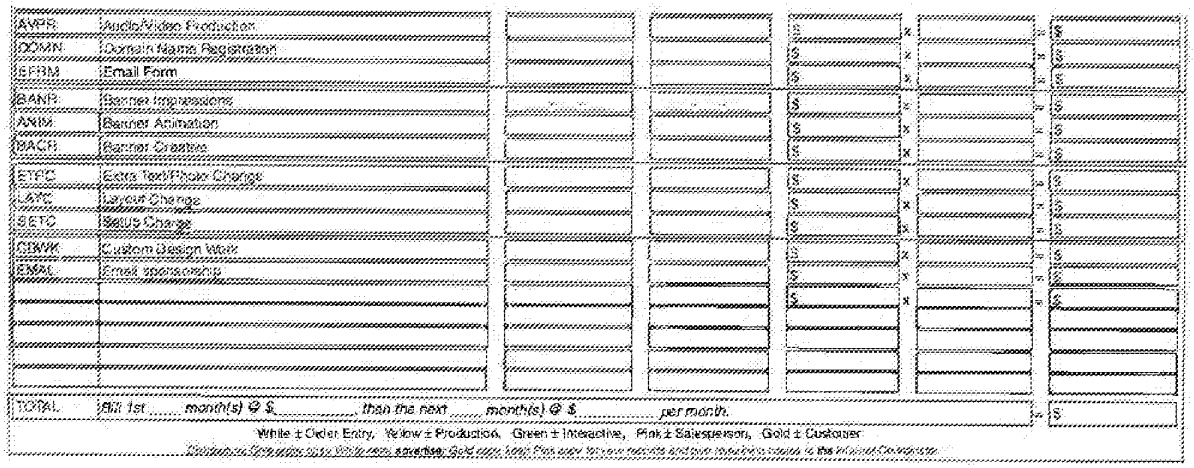


Figure 10.9 *(Continued)*

Page 666 of 1032

Meta's Exhibit No. 1004
Page 00193

Terms and Conditions

1. All Products are subject to and governed by the rates, conditions, standards, terms, guidelines and policies contained in this CNC Internet Insertion Order/Contract ("Contract") and the applicable Rate Card.

2. CNC and Advertiser will cooperatively develop the "Product" specified in this Contract. CNC will accept text, graphics or other information from Advertisers to create "Content" for use with the Product. No Content will be posted unless first approved by Advertiser.

3. CNC may provide Advertiser with a "Design Document." A Design Document approved by Advertiser will become a part of this Contract.

4. Advertiser will use the minimum time period or quantity agreed upon and, if such time period or quantity is not used as specified, Set Up Fee charges will be due and the cost of the Product shall revert to the non-discounted rate retroactive to the Start Date of this Contract. **This Contract in no way binds CNC to specific rates; rather, it applies discounts on the prevailing Rate Card rates. In the event of a rate revision, the Advertiser may, at its option, terminate this Agreement without penalty upon written notice to CNC prior to the effective date of the scheduled rate revision. Advertisers will receive thirty (30) days advance notice of any rate increases. Products canceled after ordering will be charged at the rate for the minimum time period available, or for a cancellation fee as noted on the Rate Card.** Invoices for Products are due and payable upon receipt, and will be considered overdue thirty (30) days from the date of billing. Advertiser, and any third party or agency used in the order of Products, shall be jointly and severally liable for such monies as are due and payable to CNC for Products ordered by such agency or third party. CNC reserves the right to suspend any Product if invoices remain overdue for more than thirty (30) days.

5. The Advertiser shall retain all proprietary rights and interests in any Content submitted for use in the Product. Advertiser hereby grants to CNC a non-exclusive, non-restricted, worldwide license to use, or create derivative works from, all Content it submits or develops for use in providing the Product.

6. All Content is accepted and posted entirely on the representation by Advertiser that it is authorized to publish the Content and subject matter thereof, that all Content submitted is in strict conformance with the Content Guidelines and that the Content is neither unlawful nor violates the rights of any person. In consideration for providing the Product, the Advertiser (and its agent) agrees to defend, indemnify and hold CNC, its affiliates and personnel, harmless from and against any loss or expense (including, without limitation, fees and expenses of counsel), arising out of any claim, actual or threatened, resulting from the posting of Content, including, without limitation, those resulting from claims or suits for defamation, violation of privacy, trade secrets or publicity rights, plagiarism, or copyright, patent or trademark infringement. Advertiser acknowledges that CNC is relying on such indemnity and would be unwilling to enter into this Contract in its absence. CNC reserves the right to suspend any Product for any reason, and may suspend a Product if CNC becomes aware of any potential violation of third party legal rights.

7. Neither CNC, its affiliates, nor any of their personnel, will be liable to the Advertiser for: any failure to include Content, or any delay in providing the Product; any typographical, linking, formatting or other error occurring in providing the Product; any claims or threatened claims relating to Content; the existence, non-existence, functioning or non-functioning of any hypertext links provided through the Product; any material available through any hypertext links requested by Advertiser; and any interruption or unavailability of the Product as a result of technical difficulties, Acts of God or otherwise. In the event of a continuous interruption or unavailability of the Product in excess of twenty-four (24) hours, Advertiser must notify CNC. Advertiser will be entitled to an extension of the amount of time specified in this Contract for a period of time equivalent to the time from CNC' receipt of such notice until resumption of normal services. This shall be Advertiser's sole remedy in such cases. Advertiser acknowledges that CNC is relying on this limitation of liability and would be unwilling to enter into this Contract in the absence of such limitation.

8. Under no circumstances shall CNC, its affiliates or their personnel be liable to the Advertiser under any theory for any indirect, incidental, consequential, punitive, special or exemplary damages.

Figure 10.9 Insertion order form from Townonline. *(continues)*

Meta's Exhibit No. 1004
Page 00194

9. This Contract shall remain in effect for the time period specified herein (the "Term"). Either party may renew this Contract, at the then current Rate Card rates, by giving the other party notice thirty (30) days before the expiration of the then current Term.

10. It is expressly agreed that CNC and Advertiser are acting hereunder as independent contractors and neither shall be deemed to be the agent of the other. Under no circumstances shall any of the employees of one party be deemed to be the employees of the other for any purpose. Neither party has any right or authority to assume or create any obligations of any kind or to make any representation or warranty on behalf of the other party, whether express or implied, or to bind the other party in any respect whatsoever. Advertiser and its agency agree not to make promotional, merchandising or advertising references to CNC, its newspapers or Town Online without prior express, written approval.

11. All notices required or permitted under this Contract shall be deemed delivered when in writing and delivered by a reputable mail carrier to the persons on the face of this Contract at the billing address. This Contract represents the final and complete agreement of the parties. It shall be governed by Massachusetts law and shall not be assigned by Advertiser.

12. All advertising is accepted subject to the Advertiser's strict adherence to CNC's credit terms. Advertisers will be required to make application for credit and may be required to pay in advance until such time as credit has been established. CNC reserves the right to cancel advertising at any time without notice for noncompliance with the credit terms. A fifteen dollar ($15.00) charge will be applied for all returned checks.

Content Guidelines

a. Town Online reserves the right to refuse to post, or to remove after posting, any Content which in its sole judgment is not in conformance with these Content Guidelines. However, Town Online has no duty to review the Content, and Advertiser remains solely responsible for all Content at all times.

b. No Content will be posted without the express permission of Advertiser. Any self-published Content will be presumed to have been reviewed and approved. All Content is subject to federal, state and local laws.

c. The Content may not include any material that: is grossly offensive, including bigotry, racism, discrimination, hatred or profanity; is pornographic or obscene; is defamatory or libelous; results in an invasion of privacy; violates another's rights of publicity; promotes or provides instructional information about illegal activities or physical harm or injury to any group, individual, institution or property; or infringes on a proprietary interest of another, including but not limited to, copyrights, trademarks, servicemarks or patents.

d. Advertiser may not use the Product for any purpose not expressly permitted in this Contract without first obtaining the express, written approval of Town Online.

Figure 10.9 *(Continued)*

throughout the Web. TMC (www.dottmc.com) has mostly travel clients. These specialized focuses help the firms tone their expertise in a specific area, benefiting the advertisers and publishers alike.

Pros and Cons

There are both pros and cons to using a representation firm. The number one benefit is that a representation firm can serve as a site's sales force, eliminating the need to hire an internal sales force. Another benefit is for sites with cash flow issues, since payment is based on commission, requiring less cash outlay for staff. On the negative side, representation firms have an exclusivity requirement, such that no internal sales people or ad network can be used in conjunction with them. Also, representation firms only want to represent the top tier sites, making their services inaccessible to smaller sites with more limited budgets.

Meta's Exhibit No. 1004
Page 00195

Table 10.1 Representation Firms

REPRESENTATIVE FIRM	CLIENT ROSTER
2CAN Media (www.2canmedia.com)	123 Greetings, 32bit.com, Accountingjobs.com, Accountingnet.com, AFI 100 Top Movies, America's Most Wanted Ameritrade, Arthur Frommer's Encyclopedia of Travel, Bed & Breakfast Online, Body and Mind, Business Financial Network, BusinessWire, Car Wizard, Christian Science e-Monitor, Cinescape Online, City Search, College Club USA, Cool Site of the Day Cyberswap, Developer Shed Discjockey.com, Entrepreneur Magazine Online, FinanceCenter, Frommer's, Fun School, Galaxy, GoRacing, Headbone.com, Historic Wings, IEE Network, Infohiway, Internet Travel Network, IPPA, Kiplinger Online, Lease Source, Live Online, Mamma.com, MCP, Metroville.com, Motor Trend Online, Mysterynet.com, National Enquirer Online, National Geographic Online Net London, NewsBlues, On Money, Oscar Knows, PCGame, Peoplelink, Petersen Events, Pure Baseball.com, Pure Sports games.com, Radcity.com, Realnames, Rick Dees, S&P Personal Wealth, School Sucks, Smart Money, Surflink, Teen Online, The Case.com, The Wedding Channel, TheCase.com/Kids, Total News, Truck Trends, UBL, Ultimate TV, Variety, WebFlyer
Cybereps (www.cybereps.com)	Andover Network, CAD InfoNet, Cambridge Information Network (CIN), ClickZ Network, Cool Tool, CYBER Reports, Free Realtime, Hit The Beach!, Inside Dynamic HTML, Nettaxi.com, Stewart Cheifet Productions, Teen.com, The Internet Movie Database, Weather24, Win95 Mag & Go Windows NT, YNot Creations
Ken Margolis Associates (www.kenmargolis.com)	Channel A, CLIQCOLLEGE, CLIQFINANCIAL, CLIQGOLF, CLIQHOME, CLIQKIDS, CLIQSPORTS, CLIQTECH, CLIQTRAVEL, CollegeNET, Concentric Network, CyberKids, CyberTeens, Digital Music Network, GamePen, Gavin, MysteryNet, Save The Earth Artrock Auction, Sweet!, Unified Gamers Online, Women's Forum, Yack!
Online Media Partners (www.onlinemp.com)	Adletts, Childbirth.org, ClickTV, Disney Nation, Learn2, MDG, New Media Communications, San Francisco Giants, WWWomen
.tmc (www.dottmc.com)	BizTravel.com, Consumers Digest Online, LeisurePlanet.com, NetStations, ShopNow.com, TheTrip.com, Travelbase, TravelNow, TravelWeb,

Meta's Exhibit No. 1004
Page 00196

Table 10.1 *(Continued)*

REPRESENTATIVE FIRM	CLIENT ROSTER
.tmc *(continued)* (www.dottmc.com)	WebOffers.com; Metropolitan-Networks: Hotelguide, EventGuide, DiningGuide, RetailGuide, and NightGuide
Unique Media Services (www.uniquemedia.net)	Childbirth.org, ClickTV, Learn2, San Francisco Giants, WWWomen
WebMedia (www.webmedia-sales.com)	Acclaim.net, Country Music Association, Information Please, Marvel, PGA.com, TalkCity, Traveller
WinStar Interactive (www.winstar.com/ims/index.html)	1travel.com, AllRecipes.com, Biztravel.com, Bloomberg Web Networks, Christmas98.com, Classical Insites, drkoop.com, Film Scouts, Gist TV Listing, HomeShark, Intellicast, Jazz Central Station, Moms Online, Music Boulevard, PGA-TOUR.com, Rocktropolis, SportsFan Online, Student.com, Sundance Channel, Yuckiest Site on the Internet

Ad Networks

The ad network universe just keeps growing, both in number of networks and size of networks (Table 10.2). Many of the larger networks' criteria for site selection is short and simple: at least 1 million page views per month, quality content, and a major brand name on the Web. We've touched on ad networks briefly in several other places throughout this book. To recap, an ad network brings together a group of sites with advertising inventory to sell in order to offer greater access and ease to publishers and advertisers. For publishers, this means access to more advertisers through the network's sales force;

Table 10.2 Top Ad Networks

	MEMBER SITES
LinkExchange (www.linkexchange.com)	400,000
ContentZone (www.contentzone.com)	3000
Flycast (www.flycast.com)	700
Real Media (www.realmedia.com)	560
DoubleClick (www.doubleclick.com)	300
24/7 Media (www.247media.com)	200
CLIQNOW! (www.clicknow.com)	85
@dVenture Network (www.adventure.com)	70
ADSmart (www.adsmart.com)	33

for the advertisers, this means that with one buy they can purchase impressions on a host of sites. As the cliché promises, there is power in numbers, and ad networks bring the power of a selling block to individual sites.

Benefits of Ad Networks

As already described, participating in an ad network has benefits for the advertiser and publisher alike. Media buyers like working with ad networks for many reasons, which is certainly reason enough to place inventory on a network.

For a Web site publisher, using an ad network provides three primary benefits. First, an ad network can substitute or augment in-house sales departments with the network's dedicated, aggressive, and online ad management sales team. Second, having an ad network carry all of a site's ad inventory can eliminate the need for a site to invest in sophisticated and expensive ad management software. And finally, an ad network can move remnant inventory. To summarise:

BENEFITS TO THE SITES

- Substitute or augment in-house sales departments with the network's dedicated, aggressive, and online ad management savvy sales team
- Eliminate the need for a site to invest in sophisticated and expensive ad management software by running the ads off of the network
- Sell leftover or remnant inventory

BENEFITS TO ADVERTISERS

- Consistency in ad delivery reporting
- Can work with one sales rep for multiple site buys
- Competitive pricing because of the network's economy of scale
- Targeting options over a slate of sites, reducing the number of wasted impressions

Simplifying the Buying Process

Media buyers need to purchase enough ads to reach the target audience enough times to get the desired response and be able to measure that response. The typical ad agency represents dozens of clients with dozens of products and dozens of target audiences. An ad network provides a means to buy in large enough volume while controlling the segmentation of the audience, and, in some cases, controlling the frequency and reach for maximum penetration. With an ad network, the media buyer can call in the morning and have an ad running on a slate of sites that same day. For the media buyer, networks help navigate through the high volume of sites offering advertising.

Experienced Sales Force

Buying ads separately on the ever-growing number of sites can be a daunting task; you have to make your way through separate sales forces, individual rate sheets, and a myriad of audience demographics. When you work with a network, you are in contact

with just one experienced sales force made up of people who know the online advertising business. And the ad space offered on the network usually represents prime online real estate. Ad networks function to simplify the ad buying experience for advertisers, who otherwise would face many daunting ad placement challenges on the Web such as overwhelming site choices, lack of legitimate, verifiable information, and inconsistent ad packages, rates, terms, and ad creative specifications. Ad networks streamline the entire process, benefiting the buyer and seller alike. This is not to say that ad networks are the only means through which ad agencies purchase online media. It is just often a part of their media mix.

The benefits to the publisher are equally as great. By participating in a reputable network, a site is in the hands of an experiences sales force. Thus, the publisher is free to devote time and resources to producing the best content possible. In response, some sites choose not to retain an in-house sales force in order to curb costs of hiring, training, and maintaining a sales team. Naturally, other costs are incurred, because networks command a hefty commission on all advertising space that they sell.

Inventory

Instead of buying piecemeal, an advertiser or an agency can buy large quantities of ad space over several sites. Studies show that results are better for ad placement over a wide variety of sites, in contrast to placing advertising within many layers of the same site. Networks provide a convenient way to implement these far-reaching yet targeted buys. For the publisher, working with an ad network can help move remnant or left-over inventory. It can be a great way to augment the sale of inventory.

Membership Criteria

What are networks looking for in choosing members sites? First, they want sites that service an identifiable audience, from financial to sports to parenting. Second, they seek sites with large monthly impressions levels. Most networks want a traffic level of a half million to a million-plus monthly impressions. This traffic volume signifies that the site has the draw and can deliver an audience. Third, networks want to know whether the site's audience will respond to ads. And finally, networks are looking for major online brands, because many advertisers want to place their ads on brand name sites.

Here are some issues for site publishers to address when considering working with a network. First, if your site is part of a network, you will have little control over the specific ads shown on the site. You can specify categories of ad types, but you will not receive each ad separately to review. Equally important, an ad network sales force focuses on the selling of categories of sites and not specific sites, so your site will not be receiving individual attention. And, if you are a small site, remember that the major networks are unavailable to sites with less than one million impressions per month.

Servicing Small to Medium-Sized Sites

If your monthly audience is under 500,000 impressions, you can still join an ad network. There are networks where the criteria for inclusion are not based on audience

size. These networks cater to the small and medium-sized sites and service hundreds of sites just like yours. For example, Burst Media (www.burstmedia.com) will represent sites that have as few as 1,000 requests per month.

Targeting Buyers on a Budget

Networks that cater to the small site and the buyers on a budget make it possible to get a good buy whether your Web advertising budget is $25 or $1,000 per month. What, you might ask, can someone get with just $25? At LinkExchange, 2,000 impressions. According to Tony Hsieh, co-founder of LinkExchange, "If you compare that with a $5,000 buy on Yahoo!, the CPM is a better buy on us."

Barter Networks

One of the most popular types of ad networks is one that offers a barter program. Under a barter program, sites exchange ad placements. The equation works as follows: when a site places a certain number of the network's impressions on its pages, the site receives a certain number of its ads on other network properties. One of the first barter networks, and the largest ad network in existence (over 400,000 member sites), is LinkExchange (www.linkexchange.com). Members of LinkExchange can either purchase ad space on the network or participate in the banner exchange program. LinkExchange members receive an ad placement on the network for every two ads displayed on their site (making it a 2:1 ratio).

Barter networks differ from each other in many ways. Some networks are distinguished by their exchange ratio: 1:1 or 2:1 ratio. Others differ by the size of ads they place: buttons (120 × 90) or banners (486 × 60). Some banner exchanges revolve around a topic or theme, such as sports (Basketball Banner Exchange, www.verticaljump.com/banners/, and SoccerSwap, www.soccerswap.com), gender (The Women's Hyberbanner, women.hyperbanner.net), or age (Kids Hyberbanner, kids.hyperbanner.net). There are banner exchanges for minority-owned sites such as La Camara Hispana (http://camarahispana.com), and networks focusing on European countries like EuroBanner (www.eurobanner.com). Lists of banner exchange programs can be found on SearchZ (www.searchz.com) and Adbility (http://markwelch.com/bannerad/bat_exchange.htm).

Differences Between Ad Networks and Representation Firms

An ad network and a representation firm both broker ad space, but differ in the focus of the sales process. Both serve as supplemental or replacement sales forces, but a rep firm represents the site itself. An ad network, on the other hand, is usually selling its network categories to an advertiser and presents all of its member sites as a group. In other words, the network may represent four sports-related sites, so they would sell advertisers space on the network's sports category sites, not on one individual site.

A network most often sells its inventory based on the strength of the consolidation of its member sites into categories, such as financial sites or sports sites. A repping firm, on the other hand, represents the individual sites themselves. It is the sales force for a selected group of Web sites. Whereas ad networks are looking for volume of sites, rep firms only represent a small list of sites because of the individual nature of the work. In addition, reps often request exclusive arrangements with sites whereby they are a site's

Network Working Group J. Mogul (Stanford)
Request for Comments: 950 J. Postel (ISI)
August 1985

Internet Standard Subnetting Procedure

Status Of This Memo

This RFC specifies a protocol for the ARPA-Internet community. If subnetting is implemented it is strongly recommended that these procedures be followed. Distribution of this memo is unlimited.

Overview

This memo discusses the utility of "subnets" of Internet networks, which are logically visible sub-sections of a single Internet network. For administrative or technical reasons, many organizations have chosen to divide one Internet network into several subnets, instead of acquiring a set of Internet network numbers. This memo specifies procedures for the use of subnets. These procedures are for hosts (e.g., workstations). The procedures used in and between subnet gateways are not fully described. Important motivation and background information for a subnetting standard is provided in RFC-940 [7].

Acknowledgment

This memo is based on RFC-917 [1]. Many people contributed to the development of the concepts described here. J. Noel Chiappa, Chris Kent, and Tim Mann, in particular, provided important suggestions. Additional contributions in shaping this memo were made by Zaw-Sing Su, Mike Karels, and the Gateway Algorithms and Data Structures Task Force (GADS).

Roku Exhibit 1038
Page 1 of 18

1. Motivation

The original view of the Internet universe was a two-level hierarchy: the top level the Internet as a whole, and the level below it individual networks, each with its own network number. The Internet does not have a hierarchical topology, rather the interpretation of addresses is hierarchical. In this two-level model, each host sees its network as a single entity; that is, the network may be treated as a "black box" to which a set of hosts is connected.

While this view has proved simple and powerful, a number of organizations have found it inadequate, and have added a third level to the interpretation of Internet addresses. In this view, a given Internet network is divided into a collection of subnets.

The three-level model is useful in networks belonging to moderately large organizations (e.g., Universities or companies with more than one building), where it is often necessary to use more than one LAN cable to cover a "local area". Each LAN may then be treated as a subnet.

There are several reasons why an organization might use more than one cable to cover a campus:

- Different technologies: Especially in a research environment, there may be more than one kind of LAN in use; e.g., an organization may have some equipment that supports Ethernet, and some that supports a ring network.

- Limits of technologies: Most LAN technologies impose limits, based on electrical parameters, on the number of hosts connected, and on the total length of the cable. It is easy to exceed these limits, especially those on cable length.

- Network congestion: It is possible for a small subset of the hosts on a LAN to monopolize most of the bandwidth. A common solution to this problem is to divide the hosts into cliques of high mutual communication, and put these cliques on separate cables.

- Point-to-Point links: Sometimes a "local area", such as a university campus, is split into two locations too far apart to connect using the preferred LAN technology. In this case, high-speed point-to-point links might connect several LANs.

An organization that has been forced to use more than one LAN has three choices for assigning Internet addresses:

Roku Exhibit 1038
Page 2 of 18

1. Acquire a distinct Internet network number for each cable; subnets are not used at all.

2. Use a single network number for the entire organization, but assign host numbers without regard to which LAN a host is on ("transparent subnets").

3. Use a single network number, and partition the host address space by assigning subnet numbers to the LANs ("explicit subnets").

Each of these approaches has disadvantages. The first, although not requiring any new or modified protocols, results in an explosion in the size of Internet routing tables. Information about the internal details of local connectivity is propagated everywhere, although it is of little or no use outside the local organization. Especially as some current gateway implementations do not have much space for routing tables, it would be good to avoid this problem.

The second approach requires some convention or protocol that makes the collection of LANs appear to be a single Internet network. For example, this can be done on LANs where each Internet address is translated to a hardware address using an Address Resolution Protocol (ARP), by having the bridges between the LANs intercept ARP requests for non-local targets, see RFC-925 [2]. However, it is not possible to do this for all LAN technologies, especially those where ARP protocols are not currently used, or if the LAN does not support broadcasts. A more fundamental problem is that bridges must discover which LAN a host is on, perhaps by using a broadcast algorithm. As the number of LANs grows, the cost of broadcasting grows as well; also, the size of translation caches required in the bridges grows with the total number of hosts in the network.

The third approach is to explicitly support subnets. This does have a disadvantage, in that it is a modification of the Internet Protocol, and thus requires changes to IP implementations already in use (if these implementations are to be used on a subnetted network). However, these changes are relatively minor, and once made, yield a simple and efficient solution to the problem. Also, the approach avoids any changes that would be incompatible with existing hosts on non-subnetted networks.

Further, when appropriate design choices are made, it is possible for hosts which believe they are on a non-subnetted network to be used on a subnetted one, as explained in RFC-917 [1]. This is useful when it is not possible to modify some of the hosts to support subnets explicitly, or when a gradual transition is preferred.

Roku Exhibit 1038
Page 3 of 18

2. Standards for Subnet Addressing

This section first describes a proposal for interpretation of Internet addresses to support subnets. Next it discusses changes to host software to support subnets. Finally, it presents a procedures for discovering what address interpretation is in use on a given network (i.e., what address mask is in use).

2.1. Interpretation of Internet Addresses

Suppose that an organization has been assigned an Internet network number, has further divided that network into a set of subnets, and wants to assign host addresses: how should this be done? Since there are minimal restrictions on the assignment of the "local address" part of the Internet address, several approaches have been proposed for representing the subnet number:

1. Variable-width field: Any number of the bits of the local address part are used for the subnet number; the size of this field, although constant for a given network, varies from network to network. If the field width is zero, then subnets are not in use.

2. Fixed-width field: A specific number of bits (e.g., eight) is used for the subnet number, if subnets are in use.

3. Self-encoding variable-width field: Just as the width (i.e., class) of the network number field is encoded by its high-order bits, the width of the subnet field is similarly encoded.

4. Self-encoding fixed-width field: A specific number of bits is used for the subnet number.

5. Masked bits: Use a bit mask ("address mask") to identify which bits of the local address field indicate the subnet number.

What criteria can be used to choose one of these five schemes? First, should we use a self-encoding scheme? And, should it be possible to tell from examining an Internet address if it refers to a subnetted network, without reference to any other information?

An interesting feature of self-encoding is that it allows the

Roku Exhibit 1038
Page 4 of 18

address space of a network to be divided into subnets of different sizes, typically one subnet of half the address space and a set of small subnets.

For example, consider a class C network that uses a self-encoding scheme with one bit to indicate if it is the large subnet or not and an additional three bits to identify the small subnet. If the first bit is zero then this is the large subnet, if the first bit is one then the following bits (3 in this example) give the subnet number. There is one subnet with 128 host addresses, and eight subnets with 16 hosts each.

To establish a subnetting standard the parameters and interpretation of the self-encoding scheme must be fixed and consistent throughout the Internet.

It could be assumed that all networks are subnetted. This would allow addresses to be interpreted without reference to any other information.

This is a significant advantage, that given the Internet address no additional information is needed for an implementation to determine if two addresses are on the same subnet. However, this can also be viewed as a disadvantage: it may cause problems for networks which have existing host numbers that use arbitrary bits in the local address part. In other words, it is useful to be able to control whether a network is subnetted independently from the assignment of host addresses.

The alternative is to have the fact that a network is subnetted kept separate from the address. If one finds, somehow, that the network is subnetted then the standard self-encoded subnetted network address rules are followed, otherwise the non-subnetted network addressing rules are followed.

If a self-encoding scheme is not used, there is no reason to use a fixed-width field scheme: since there must in any case be some per-network "flag" to indicate if subnets are in use, the additional cost of using an integer (a subnet field width or address mask) instead of a boolean is negligible. The advantage of using the address mask scheme is that it allows each organization to choose the best way to allocate relatively scarce bits of local address to subnet and host numbers. Therefore, we choose the address-mask scheme: it is the most flexible scheme, yet costs no more to implement than any other.

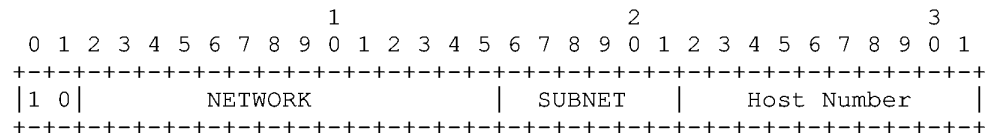
Roku Exhibit 1038
Page 5 of 18

For example, the Internet address might be interpreted as:

<network-number><subnet-number><host-number>

where the <network-number> field is as defined by IP [3], the <host-number> field is at least 1-bit wide, and the width of the <subnet-number> field is constant for a given network. No further structure is required for the <subnet-number> or <host-number> fields. If the width of the <subnet-number> field is zero, then the network is not subnetted (i.e., the interpretation of [3] is used).

For example, on a Class B network with a 6-bit wide subnet field, an address would be broken down like this:



Since the bits that identify the subnet are specified by a bitmask, they need not be adjacent in the address. However, we recommend that the subnet bits be contiguous and located as the most significant bits of the local address.

Special Addresses:

From the Assigned Numbers memo [9]:

> "In certain contexts, it is useful to have fixed addresses with functional significance rather than as identifiers of specific hosts. When such usage is called for, the address zero is to be interpreted as meaning "this", as in "this network". The address of all ones are to be interpreted as meaning "all", as in "all hosts". For example, the address 128.9.255.255 could be interpreted as meaning all hosts on the network 128.9. Or, the address 0.0.0.37 could be interpreted as meaning host 37 on this network."

It is useful to preserve and extend the interpretation of these special addresses in subnetted networks. This means the values of all zeros and all ones in the subnet field should not be assigned to actual (physical) subnets.

> In the example above, the 6-bit wide subnet field may have any value except 0 and 63.

Roku Exhibit 1038
Page 6 of 18

Please note that there is no effect or new restriction on the addresses of hosts on non-subnetted networks.

2.2. Changes to Host Software to Support Subnets

In most implementations of IP, there is code in the module that handles outgoing datagrams to decide if a datagram can be sent directly to the destination on the local network or if it must be sent to a gateway.

Generally the code is something like this:

```
IF ip_net_number(dg.ip_dest) = ip_net_number(my_ip_addr)
    THEN
        send_dg_locally(dg, dg.ip_dest)
    ELSE
        send_dg_locally(dg,
                        gateway_to(ip_net_number(dg.ip_dest)))
```

(If the code supports multiply-connected networks, it will be more complicated, but this is irrelevant to the current discussion.)

To support subnets, it is necessary to store one more 32-bit quantity, called my_ip_mask. This is a bit-mask with bits set in the fields corresponding to the IP network number, and additional bits set corresponding to the subnet number field.

The code then becomes:

```
IF bitwise_and(dg.ip_dest, my_ip_mask)
                          = bitwise_and(my_ip_addr, my_ip_mask)
    THEN
        send_dg_locally(dg, dg.ip_dest)
    ELSE
        send_dg_locally(dg,
               gateway_to(bitwise_and(dg.ip_dest, my_ip_mask)))
```

Of course, part of the expression in the conditional can be pre-computed.

It may or may not be necessary to modify the "gateway_to" function, so that it too takes the subnet field bits into account when performing comparisons.

To support multiply-connected hosts, the code can be changed to

Roku Exhibit 1038
Page 7 of 18

keep the "my_ip_addr" and "my_ip_mask" quantities on a per-interface basis; the expression in the conditional must then be evaluated for each interface.

2.3. Finding the Address Mask

How can a host determine what address mask is in use on a subnet to which it is connected? The problem is analogous to several other "bootstrapping" problems for Internet hosts: how a host determines its own address, and how it locates a gateway on its local network. In all three cases, there are two basic solutions: "hardwired" information, and broadcast-based protocols.

Hardwired information is that available to a host in isolation from a network. It may be compiled-in, or (preferably) stored in a disk file. However, for the increasingly common case of a diskless workstation that is bootloaded over a LAN, neither hardwired solution is satisfactory.

Instead, since most LAN technology supports broadcasting, a better method is for the newly-booted host to broadcast a request for the necessary information. For example, for the purpose of determining its Internet address, a host may use the "Reverse Address Resolution Protocol" (RARP) [4].

However, since a newly-booted host usually needs to gather several facts (e.g., its IP address, the hardware address of a gateway, the IP address of a domain name server, the subnet address mask), it would be better to acquire all this information in one request if possible, rather than doing numerous broadcasts on the network. The mechanisms designed to boot diskless workstations can also load per-host specific configuration files that contain the required information (e.g., see RFC-951 [8]). It is possible, and desirable, to obtain all the facts necessary to operate a host from a boot server using only one broadcast message.

In the case where it is necessary for a host to find the address mask as a separate operation the following mechanism is provided:

> To provide the address mask information the ICMP protocol [5] is extended by adding a new pair of ICMP message types, "Address Mask Request" and "Address Mask Reply", analogous to the "Information Request" and "Information Reply" ICMP messages. These are described in detail in Appendix I.
>
> The intended use of these new ICMP messages is that a host, when booting, broadcast an "Address Mask Request" message. A

Roku Exhibit 1038
Page 8 of 18

gateway (or a host acting in lieu of a gateway) that receives this message responds with an "Address Mask Reply". If there is no indication in the request which host sent it (i.e., the IP Source Address is zero), the reply is broadcast as well. The requesting host will hear the response, and from it determine the address mask.

Since there is only one possible value that can be sent in an "Address Mask Reply" on any given LAN, there is no need for the requesting host to match the responses it hears against the request it sent; similarly, there is no problem if more than one gateway responds. We assume that hosts reboot infrequently, so the broadcast load on a network from use of this protocol should be small.

If a host is connected to more than one LAN, it might have to find the address mask for each.

One potential problem is what a host should do if it can not find out the address mask, even after a reasonable number of tries. Three interpretations can be placed on the situation:

1. The local net exists in (permanent) isolation from all other nets.

2. Subnets are not in use, and no host can supply the address mask.

3. All gateways on the local net are (temporarily) down.

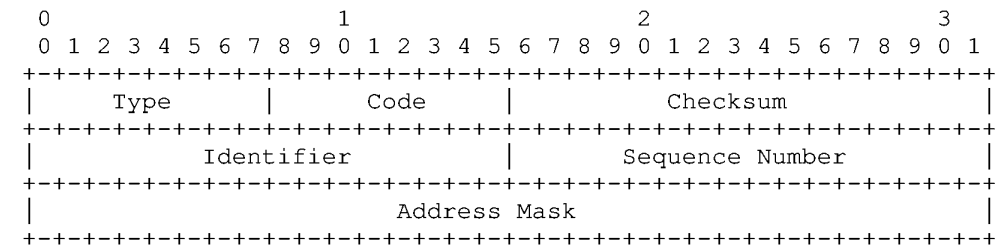
The first and second situations imply that the address mask is identical with the Internet network number mask. In the third situation, there is no way to determine what the proper value is; the safest choice is thus a mask identical with the Internet network number mask. Although this might later turn out to be wrong, it will not prevent transmissions that would otherwise succeed. It is possible for a host to recover from a wrong choice: when a gateway comes up, it should broadcast an "Address Mask Reply"; when a host receives such a message that disagrees with its guess, it should change its mask to conform to the received value. No host or gateway should send an "Address Mask Reply" based on a "guessed" value.

Finally, note that no host is required to use this ICMP protocol to discover the address mask; it is perfectly reasonable for a host with non-volatile storage to use stored information (including a configuration file from a boot server).

Roku Exhibit 1038
Page 9 of 18

Appendix I. Address Mask ICMP

Address Mask Request or Address Mask Reply



IP Fields:

Addresses

The address of the source in an address mask request message will be the destination of the address mask reply message. To form an address mask reply message, the source address of the request becomes the destination address of the reply, the source address of the reply is set to the replier's address, the type code changed to AM2, the address mask value inserted into the Address Mask field, and the checksum recomputed. However, if the source address in the request message is zero, then the destination address for the reply message should denote a broadcast.

ICMP Fields:

Type

AM1 for address mask request message

AM2 for address mask reply message

Code

0 for address mask request message

0 for address mask reply message

Checksum

The checksum is the 16-bit one's complement of the one's

Roku Exhibit 1038
Page 10 of 18

complement sum of the ICMP message starting with the ICMP Type. For computing the checksum, the checksum field should be zero. This checksum may be replaced in the future.

Identifier

An identifier to aid in matching requests and replies, may be zero.

Sequence Number

A sequence number to aid in matching requests and replies, may be zero.

Address Mask

A 32-bit mask.

Description

A gateway receiving an address mask request should return it with the address mask field set to the 32-bit mask of the bits identifying the subnet and network, for the subnet on which the request was received.

If the requesting host does not know its own IP address, it may leave the source field zero; the reply should then be broadcast. However, this approach should be avoided if at all possible, since it increases the superfluous broadcast load on the network. Even when the replies are broadcast, since there is only one possible address mask for a subnet, there is no need to match requests with replies. The "Identifier" and "Sequence Number" fields can be ignored.

Type AM1 may be received from a gateway or a host.

Type AM2 may be received from a gateway, or a host acting in lieu of a gateway.

Roku Exhibit 1038
Page 11 of 18

Appendix II. Examples

These examples show how a host can find out the address mask using the ICMP Address Mask Request and Address Mask Reply messages. For the following examples, assume that address 255.255.255.255 denotes "broadcast to this physical medium" [6].

1. A Class A Network Case

For this case, assume that the requesting host is on class A network 36.0.0.0, has address 36.40.0.123, that there is a gateway at 36.40.0.62, and that a 8-bit wide subnet field is in use, that is, the address mask is 255.255.0.0.

The most efficient method, and the one we recommend, is for a host to first discover its own address (perhaps using "RARP" [4]), and then to send the ICMP request to 255.255.255.255:

Source address:	36.40.0.123
Destination address:	255.255.255.255
Protocol:	ICMP = 1
Type:	Address Mask Request = AM1
Code:	0
Mask:	0

The gateway can then respond directly to the requesting host.

Source address:	36.40.0.62
Destination address:	36.40.0.123
Protocol:	ICMP = 1
Type:	Address Mask Reply = AM2
Code:	0
Mask:	255.255.0.0

Suppose that 36.40.0.123 is a diskless workstation, and does not know even its own host number. It could send the following datagram:

Source address:	0.0.0.0
Destination address:	255.255.255.255
Protocol:	ICMP = 1
Type:	Address Mask Request = AM1
Code:	0
Mask:	0

36.40.0.62 will hear the datagram, and should respond with this datagram:

Roku Exhibit 1038
Page 12 of 18

```
   Source address:          36.40.0.62
   Destination address:     255.255.255.255
   Protocol:                ICMP = 1
   Type:                    Address Mask Reply = AM2
   Code:                    0
   Mask:                    255.255.0.0
```

Note that the gateway uses the narrowest possible broadcast to reply. Even so, the over use of broadcasts presents an unnecessary load to all hosts on the subnet, and so the use of the "anonymous" (0.0.0.0) source address must be kept to a minimum.

If broadcasting is not allowed, we assume that hosts have wired-in information about neighbor gateways; thus, 36.40.0.123 might send this datagram:

```
   Source address:          36.40.0.123
   Destination address:     36.40.0.62
   Protocol:                ICMP = 1
   Type:                    Address Mask Request = AM1
   Code:                    0
   Mask:                    0
```

36.40.0.62 should respond exactly as in the previous case.

```
   Source address:          36.40.0.62
   Destination address:     36.40.0.123
   Protocol:                ICMP = 1
   Type:                    Address Mask Reply = AM2
   Code:                    0
   Mask:                    255.255.0.0
```

2. A Class B Network Case

For this case, assume that the requesting host is on class B network 128.99.0.0, has address 128.99.4.123, that there is a gateway at 128.99.4.62, and that a 6-bit wide subnet field is in use, that is, the address mask is 255.255.252.0.

The host sends the ICMP request to 255.255.255.255:

```
   Source address:          128.99.4.123
   Destination address:     255.255.255.255
   Protocol:                ICMP = 1
   Type:                    Address Mask Request = AM1
   Code:                    0
   Mask:                    0
```


Roku Exhibit 1038
Page 13 of 18

The gateway can then respond directly to the requesting host.

Source address:	128.99.4.62
Destination address:	128.99.4.123
Protocol:	ICMP = 1
Type:	Address Mask Reply = AM2
Code:	0
Mask:	255.255.252.0

In the diskless workstation case the host sends:

Source address:	0.0.0.0
Destination address:	255.255.255.255
Protocol:	ICMP = 1
Type:	Address Mask Request = AM1
Code:	0
Mask:	0

128.99.4.62 will hear the datagram, and should respond with this datagram:

Source address:	128.99.4.62
Destination address:	255.255.255.255
Protocol:	ICMP = 1
Type:	Address Mask Reply = AM2
Code:	0
Mask:	255.255.252.0

If broadcasting is not allowed 128.99.4.123 sends:

Source address:	128.99.4.123
Destination address:	128.99.4.62
Protocol:	ICMP = 1
Type:	Address Mask Request = AM1
Code:	0
Mask:	0

128.99.4.62 should respond exactly as in the previous case.

Source address:	128.99.4.62
Destination address:	128.99.4.123
Protocol:	ICMP = 1
Type:	Address Mask Reply = AM2
Code:	0
Mask:	255.255.252.0

Roku Exhibit 1038
Page 14 of 18

3. A Class C Network Case (illustrating non-contiguous subnet bits)

For this case, assume that the requesting host is on class C network 192.1.127.0, has address 192.1.127.19, that there is a gateway at 192.1.127.50, and that on network an 3-bit subnet field is in use (01011000), that is, the address mask is 255.255.255.88.

The host sends the ICMP request to 255.255.255.255:

Source address:	192.1.127.19
Destination address:	255.255.255.255
Protocol:	ICMP = 1
Type:	Address Mask Request = AM1
Code:	0
Mask:	0

The gateway can then respond directly to the requesting host.

Source address:	192.1.127.50
Destination address:	192.1.127.19
Protocol:	ICMP = 1
Type:	Address Mask Reply = AM2
Code:	0
Mask:	255.255.255.88.

In the diskless workstation case the host sends:

Source address:	0.0.0.0
Destination address:	255.255.255.255
Protocol:	ICMP = 1
Type:	Address Mask Request = AM1
Code:	0
Mask:	0

192.1.127.50 will hear the datagram, and should respond with this datagram:

Source address:	192.1.127.50
Destination address:	255.255.255.255
Protocol:	ICMP = 1
Type:	Address Mask Reply = AM2
Code:	0
Mask:	255.255.255.88.

If broadcasting is not allowed 192.1.127.19 sends:

Roku Exhibit 1038
Page 15 of 18

```
         Source address:          192.1.127.19
         Destination address:     192.1.127.50
         Protocol:                ICMP = 1
         Type:                    Address Mask Request = AM1
         Code:                    0
         Mask:                    0
```

192.1.127.50 should respond exactly as in the previous case.

```
         Source address:          192.1.127.50
         Destination address:     192.1.127.19
         Protocol:                ICMP = 1
         Type:                    Address Mask Reply = AM2
         Code:                    0
         Mask:                    255.255.255.88
```

Appendix III. Glossary

Bridge

A node connected to two or more administratively indistinguishable but physically distinct subnets, that automatically forwards datagrams when necessary, but whose existence is not known to other hosts. Also called a "software repeater".

Gateway

A node connected to two or more administratively distinct networks and/or subnets, to which hosts send datagrams to be forwarded.

Host Field

The bit field in an Internet address used for denoting a specific host.

Internet

The collection of connected networks using the IP protocol.

Local Address

The rest field of the Internet address (as defined in [3]).

Network

A single Internet network (which may or may not be divided into subnets).

Roku Exhibit 1038
Page 16 of 18

Network Number

The network field of the Internet address.

Subnet

One or more physical networks forming a subset of an Internet network. A subnet is explicitly identified in the Internet address.

Subnet Field

The bit field in an Internet address denoting the subnet number. The bits making up this field are not necessarily contiguous in the address.

Subnet Number

A number identifying a subnet within a network.

Appendix IV. Assigned Numbers

The following assignments are made for protocol parameters used in the support of subnets. The only assignments needed are for the Internet Control Message Protocol (ICMP) [5].

ICMP Message Types

AM1 = 17

AM2 = 18

Roku Exhibit 1038
Page 17 of 18

References

[1] Mogul, J., "Internet Subnets", RFC-917, Stanford University, October 1984.

[2] Postel, J., "Multi-LAN Address Resolution", RFC-925, USC/Information Sciences Institute, October 1984.

[3] Postel, J., "Internet Protocol", RFC-791, USC/Information Sciences Institute, September 1981.

[4] Finlayson, R., T. Mann, J. Mogul, M. Theimer, "A Reverse Address Resolution Protocol", RFC-903, Stanford University, June 1984.

[5] Postel, J., "Internet Control Message Protocol", RFC-792, USC/Information Sciences Institute, September 1981.

[6] Mogul, J., "Broadcasting Internet Datagrams", RFC-919, Stanford University, October 1984.

[7] GADS, "Towards an Internet Standard Scheme for Subnetting", RFC-940, Network Information Center, SRI International, April 1985.

[8] Croft, B., and J. Gilmore, "BOOTP -- UDP Bootstrap Protocol", RFC-951, Stanford University, August 1985.

[9] Reynolds, J., and J. Postel, "Assigned Numbers", RFC-943, USC/Information Sciences Institute, April 1985.

Roku Exhibit 1038
Page 18 of 18

Network Working Group
Request for Comments: 1531
Category: Standards Track

R. Droms
Bucknell University

October 1993

Dynamic Host Configuration Protocol

Status of this memo

This RFC specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

The Dynamic Host Configuration Protocol (DHCP) provides a framework for passing configuration information to hosts on a TCP/IP network. DHCP is based on the Bootstrap Protocol (BOOTP) [7], adding the capability of automatic allocation of reusable network addresses and additional configuration options [19]. DHCP captures the behavior of BOOTP relay agents [7, 23], and DHCP participants can interoperate with BOOTP participants [9].

Table of Contents

1. Introduction. 2
1.1 Related Work. 4
1.2 Problem definition and issues 4
1.3 Requirements. 5
1.4 Terminology 6
1.5 Design goals. 6
2. Protocol Summary 8
2.1 Configuration parameters repository 10
2.2 Dynamic allocation of network addresses 11
3. The Client-Server Protocol 11
3.1 Client-server interaction - allocating a network address. . . 12
3.2 Client-server interaction - reusing a previously allocated network address 17
3.3 Interpretation and representation of time values. 19
3.4 Host parameters in DHCP 19
3.5 Use of DHCP in clients with multiple interfaces 20
3.6 When clients should use DHCP. 20
4. Specification of the DHCP client-server protocol 21
4.1 Constructing and sending DHCP messages. 21
4.2 DHCP server administrative controls 23
4.3 DHCP server behavior. 24

Roku Exhibit 1041
Page 1 of 39

4.3.1 DHCPDISCOVER message. 24
4.3.2 DHCPREQUEST message 27
4.3.3 DHCPDECLINE message 29
4.3.4 DHCPRELEASE message 29
4.4 DHCP client behavior. 29
4.4.1 Initialization and allocation of network address. 29
4.4.2 Initialization with known network address 33
4.4.3 Initialization with a known DHCP server address 34
4.4.4 Reacquisition and expiration. 34
4.4.5 DHCPRELEASE 35
5. Acknowledgments. 35
6. References 36
7. Security Considerations. 37
8. Author's Address 38
A. Host Configuration Parameters 39

List of Figures

1. Format of a DHCP message 9
2. Format of the 'flags' field. 10
3. Timeline diagram of messages exchanged between DHCP client and servers when allocating a new network address. 15
4. Timeline diagram of messages exchanged between DHCP client and servers when reusing a previously allocated network address. . 18
5. State-transition diagram for DHCP clients. 31

List of Tables

1. Description of fields in a DHCP message. 14
2. DHCP messages. 16
3. Fields and options used by DHCP servers. 25
4. Fields and options used by DHCP clients. 32

1. Introduction

The Dynamic Host Configuration Protocol (DHCP) provides configuration parameters to Internet hosts. DHCP consists of two components: a protocol for delivering host-specific configuration parameters from a DHCP server to a host and a mechanism for allocation of network addresses to hosts.

DHCP is built on a client-server model, where designated DHCP server hosts allocate network addresses and deliver configuration parameters to dynamically configured hosts. Throughout the remainder of this document, the term "server" refers to a host providing initialization parameters through DHCP, and the term "client" refers to a host requesting initialization parameters from a DHCP server.

Roku Exhibit 1041
Page 2 of 39

A host should not act as a DHCP server unless explicitly configured to do so by a system administrator. The diversity of hardware and protocol implementations in the Internet would preclude reliable operation if random hosts were allowed to respond to DHCP requests. For example, IP requires the setting of many parameters within the protocol implementation software. Because IP can be used on many dissimilar kinds of network hardware, values for those parameters cannot be guessed or assumed to have correct defaults. Also, distributed address allocation schemes depend on a polling/defense mechanism for discovery of addresses that are already in use. IP hosts may not always be able to defend their network addresses, so that such a distributed address allocation scheme cannot be guaranteed to avoid allocation of duplicate network addresses.

DHCP supports three mechanisms for IP address allocation. In "automatic allocation", DHCP assigns a permanent IP address to a host. In "dynamic allocation", DHCP assigns an IP address to a host for a limited period of time (or until the host explicitly relinquishes the address). In "manual allocation", a host's IP address is assigned by the network administrator, and DHCP is used simply to convey the assigned address to the host. A particular network will use one or more of these mechanisms, depending on the policies of the network administrator.

Dynamic allocation is the only one of the three mechanisms that allows automatic reuse of an address that is no longer needed by the host to which it was assigned. Thus, dynamic allocation is particularly useful for assigning an address to a host that will be connected to the network only temporarily or for sharing a limited pool of IP addresses among a group of hosts that do not need permanent IP addresses. Dynamic allocation may also be a good choice for assigning an IP address to a new host being permanently connected to a network where IP addresses are sufficiently scarce that it is important to reclaim them when old hosts are retired. Manual allocation allows DHCP to be used to eliminate the error-prone process of manually configuring hosts with IP addresses in environments where (for whatever reasons) it is desirable to manage IP address assignment outside of the DHCP mechanisms.

The format of DHCP messages is based on the format of BOOTP messages, to capture the BOOTP relay agent behavior described as part of the BOOTP specification [7, 23] and to allow interoperability of existing BOOTP clients with DHCP servers. Using BOOTP relaying agents eliminates the necessity of having a DHCP server on each physical network segment.

Roku Exhibit 1041
Page 3 of 39

1.1 Related Work

There are several Internet protocols and related mechanisms that address some parts of the dynamic host configuration problem. The Reverse Address Resolution Protocol (RARP) [10] (through the extensions defined in the Dynamic RARP (DRARP) [5]) explicitly addresses the problem of network address discovery, and includes an automatic IP address assignment mechanism. The Trivial File Transfer Protocol (TFTP) [20] provides for transport of a boot image from a boot server. The Internet Control Message Protocol (ICMP) [16] provides for informing hosts of additional routers via "ICMP redirect" messages. ICMP also can provide subnet mask information through the "ICMP mask request" message and other information through the (obsolete) "ICMP information request" message. Hosts can locate routers through the ICMP router discovery mechanism [8].

BOOTP is a transport mechanism for a collection of configuration information. BOOTP is also extensible, and official extensions [17] have been defined for several configuration parameters. Morgan has proposed extensions to BOOTP for dynamic IP address assignment [15]. The Network Information Protocol (NIP), used by the Athena project at MIT, is a distributed mechanism for dynamic IP address assignment [19]. The Resource Location Protocol RLP [1] provides for location of higher level services. Sun Microsystems diskless workstations use a boot procedure that employs RARP, TFTP and an RPC mechanism called "bootparams" to deliver configuration information and operating system code to diskless hosts. (Sun Microsystems, Sun Workstation and SunOS are trademarks of Sun Microsystems, Inc.) Some Sun networks also use DRARP and an auto-installation mechanism to automate the configuration of new hosts in an existing network.

In other related work, the path minimum transmission unit (MTU) discovery algorithm can determine the MTU of an arbitrary internet path [14]. Comer and Droms have proposed the use of the Address Resolution Protocol (ARP) as a transport protocol for resource location and selection [6]. Finally, the Host Requirements RFCs [3, 4] mention specific requirements for host reconfiguration and suggest a scenario for initial configuration of diskless hosts.

1.2 Problem definition and issues

DHCP is designed to supply hosts with the configuration parameters defined in the Host Requirements RFCs. After obtaining parameters via DHCP, a host should be able to exchange packets with any other host in the Internet. The parameters supplied by DHCP are listed in Appendix A.

Roku Exhibit 1041
Page 4 of 39

Not all of these parameters are required for a newly initialized host. A client and server may negotiate for the transmission of only those parameters required by the client or specific to a particular subnet.

DHCP allows but does not require the configuration of host parameters not directly related to the IP protocol. DHCP also does not address registration of newly configured hosts with the Domain Name System (DNS) [12, 13].

DHCP is not intended for use in configuring routers.

1.3 Requirements

Throughout this document, the words that are used to define the significance of particular requirements are capitalized. These words are:

- "MUST"

  This word or the adjective "REQUIRED" means that the item is an absolute requirement of this specification.

- "MUST NOT"

  This phrase means that the item is an absolute prohibition of this specification.

- "SHOULD"

  This word or the adjective "RECOMMENDED" means that there may exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood and the case carefully weighed before choosing a different course.

- "SHOULD NOT"

  This phrase means that there may exist valid reasons in particular circumstances when the listed behavior is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behavior described with this label.

Roku Exhibit 1041
Page 5 of 39

- "MAY"

  This word or the adjective "OPTIONAL" means that this item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because it enhances the product, for example; another vendor may omit the same item.

1.4 Terminology

This document uses the following terms:

- "DHCP client"

  A DHCP client is an Internet host using DHCP to obtain configuration parameters such as a network address.

- "DHCP server"

  A DHCP server is an Internet host that returns configuration parameters to DHCP clients.

- "BOOTP relay agent"

  A BOOTP relay agent is an Internet host or router that passes DHCP messages between DHCP clients and DHCP servers. DHCP is designed to use the same relay agent behavior as specified in the BOOTP protocol specification.

- "binding"

  A binding is a collection of configuration parameters, including at least an IP address, associated with or "bound to" a DHCP client. Bindings are managed by DHCP servers.

1.5 Design goals

The following list gives general design goals for DHCP.

- DHCP should be a mechanism rather than a policy. DHCP must allow local system administrators control over configuration parameters where desired; e.g., local system administrators should be able to enforce local policies concerning allocation and access to local resources where desired.

Roku Exhibit 1041
Page 6 of 39

- Hosts should require no manual configuration. Each host should be able to discover appropriate local configuration parameters without user intervention and incorporate those parameters into its own configuration.

- Networks should require no hand configuration for individual hosts. Under normal circumstances, the network manager should not have to enter any per-host configuration parameters.

- DHCP should not require a server on each subnet. To allow for scale and economy, DHCP must work across routers or through the intervention of BOOTP/DHCP relay agents.

- A DHCP host must be prepared to receive multiple responses to a request for configuration parameters. Some installations may include multiple, overlapping DHCP servers to enhance reliability and increase performance.

- DHCP must coexist with statically configured, non-participating hosts and with existing network protocol implementations.

- DHCP must interoperate with the BOOTP relay agent behavior as described by RFC 951 and by Wimer [21].

- DHCP must provide service to existing BOOTP clients.

The following list gives design goals specific to the transmission of the network layer parameters. DHCP must:

- Guarantee that any specific network address will not be in use by more than one host at a time,

- Retain host configuration across host reboot. A host should, whenever possible, be assigned the same configuration parameters (e.g., network address) in response to each request,

- Retain host configuration across server reboots, and, whenever possible, a host should be assigned the same configuration parameters despite restarts of the DHCP mechanism,

- Allow automatic assignment of configuration parameters to new hosts to avoid hand configuration for new hosts,

- Support fixed or permanent allocation of configuration parameters to specific hosts.

Roku Exhibit 1041
Page 7 of 39

2. Protocol Summary

From the client's point of view, DHCP is an extension of the BOOTP mechanism. This behavior allows existing BOOTP clients to interoperate with DHCP servers without requiring any change to the clients' initialization software. A separate document details the interactions between BOOTP and DHCP clients and servers [9]. There are some new, optional transactions that optimize the interaction between DHCP clients and servers that are described in sections 3 and 4.

Figure 1 gives the format of a DHCP message and table 1 describes each of the fields in the DHCP message. The numbers in parentheses indicate the size of each field in octets. The names for the fields given in the figure will be used throughout this document to refer to the fields in DHCP messages.

There are two primary differences between DHCP and BOOTP. First, DHCP defines mechanisms through which clients can be assigned a network address for a fixed lease, allowing for serial reassignment of network addresses to different clients. Second, DHCP provides the mechanism for a client to acquire all of the IP configuration parameters that it needs in order to operate.

DHCP introduces a small change in terminology intended to clarify the meaning of one of the fields. What was the "vendor extensions" field in BOOTP has been re-named the "options" field in DHCP. Similarly, the tagged data items that were used inside the BOOTP "vendor extensions" field, which were formerly referred to as "vendor extensions," are now termed simply "options."

DHCP defines a new 'client identifier' option that is used to pass an explicit client identifier to a DHCP server. This change eliminates the overloading of the 'chaddr' field in BOOTP messages, where reply messages and as a client identifier. The 'client identifier' option may contain a hardware address, identical to the contents of the 'chaddr' field, or it may contain another type of identifier, such as a DNS name. Other client identifier types may be defined as needed for use with DHCP. New client identifier types will be registered with the IANA [18] and will be included in new revisions of the Assigned Numbers document, as well as described in detail in future revisions of the DHCP Options [2].

Roku Exhibit 1041
Page 8 of 39

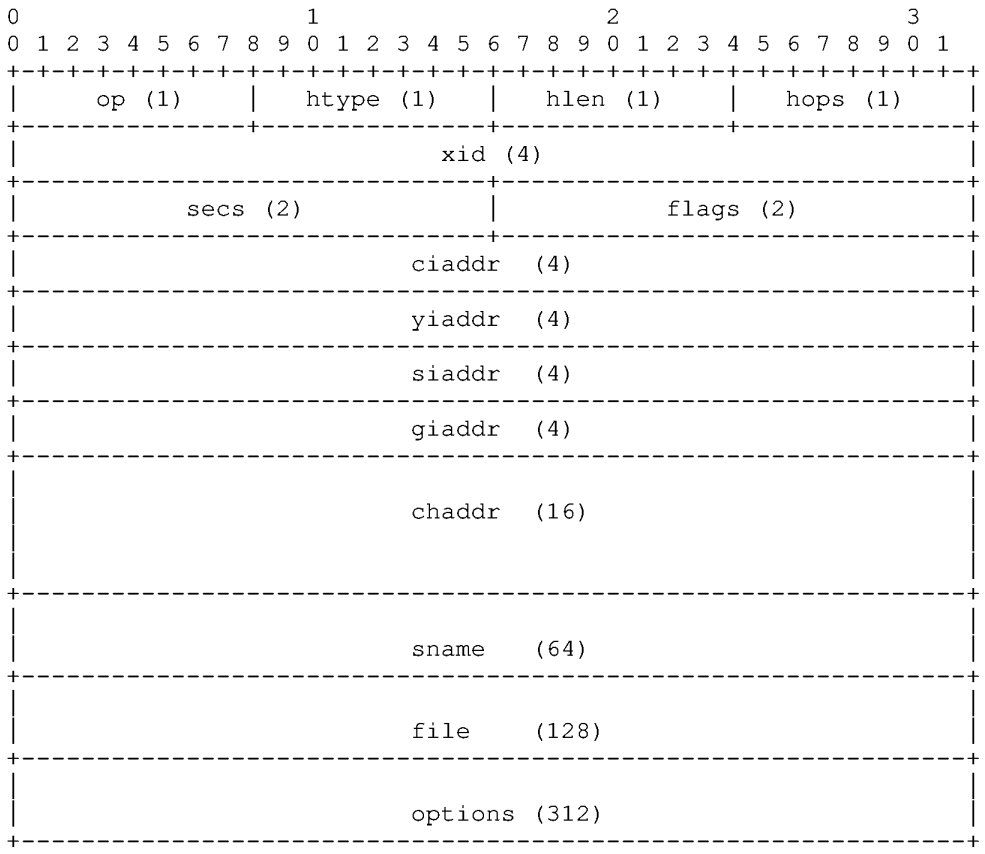


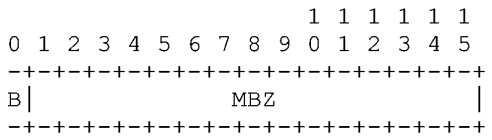
Figure 1: Format of a DHCP message

DHCP clarifies the interpretation of the 'siaddr' field as the address of the server to use in the next step of the client's bootstrap process. A DHCP server may return its own address in the 'siaddr' field, if the server is prepared to supply the next bootstrap service (e.g., delivery of an operating system executable image). A DHCP server always returns its own address in the 'server identifier' option.

The options field is now variable length, with the minimum extended to 312 octets. This brings the minimum size of a DHCP message up to 576 octets, the minimum IP datagram size a host must be prepared to accept [3]. DHCP clients may negotiate the use of larger DHCP messages through the 'Maximum DHCP message size' option. The options field may be further extended into the 'file' and 'sname' fields.

Roku Exhibit 1041
Page 9 of 39

A new option, called 'vendor specific information', has been added to allow for expansion of the number of options that can be supported [2]. Options encapsulated as 'vendor specific information' must be carefully defined and documented so as to allow for interoperability between clients and servers from diferent vendors. In particular, vendors defining 'vendor specific information' MUST document those options in the form of the DHCP Options document, MUST choose to represent those options either in data types already defined for DHCP options or in other well-defined data types, and MUST choose options that can be readily encoded in configuration files for exchange with servers provided by other vendors. Options included as 'vendor specific options' MUST be readily supportable by all servers.



B: BROADCAST flag

MBZ: MUST BE ZERO (reserved for future use)

Figure 2: Format of the 'flags' field

DHCP uses the 'flags' field [21]. The leftmost bit is defined as the BROADCAST (B) flag. The semantics of this flag are discussed in section 4.1 of this document. The remaining bits of the flags field are reserved for future use. They MUST be set to zero by clients and ignored by servers and relay agents. Figure 2 gives the format of the

2.1 Configuration parameters repository

The first service provided by DHCP is to provide persistent storage of network parameters for network clients. The model of DHCP persistent storage is that the DHCP service stores a key-value entry for each client, where the key is some unique identifier (for example, an IP subnet number and a unique identifier within the subnet) and the value contains the configuration parameters for the client.

For example, the key might be the pair (IP-subnet-number, hardware-address), allowing for serial or concurrent reuse of a hardware address on different subnets, and for hardware addresses that may not be globally unique. Alternately, the key might be the pair (IP-subnet-number, hostname), allowing the server to assign parameters intelligently to a host that has been moved to a different subnet or

Roku Exhibit 1041
Page 10 of 39

has changed hardware addresses (perhaps because the network interface failed and was replaced).

A client can query the DHCP service to retrieve its configuration parameters. The client interface to the configuration parameters repository consists of protocol messages to request configuration parameters and responses from the server carrying the configuration parameters.

2.2 Dynamic allocation of network addresses

The second service provided by DHCP is the allocation of temporary or permanent network (IP) addresses to hosts. The basic mechanism for the dynamic allocation of network addresses is simple: a client requests the use of an address for some period of time. The allocation mechanism (the collection of DHCP servers) guarantees not to reallocate that address within the requested time and attempts to return the same network address each time the client requests an address. In this document, the period over which a network address is allocated to a client is referred to as a "lease" [11]. The client may extend its lease with subsequent requests. The client may issue a message to release the address back to the server when the client no longer needs the address. The client may ask for a permanent assignment by asking for an infinite lease. Even when assigning "permanent" addresses, a server may choose to give out lengthy but non-infinite leases to allow detection of the fact that the host has been retired.

In some environments it will be necessary to reassign network addresses due to exhaustion of available addresses. In such environments, the allocation mechanism will reuse addresses whose lease has expired. The server should use whatever information is available in the configuration information repository to choose an address to reuse. For example, the server may choose the least recently assigned address. As a consistency check, the allocation mechanism may probe the reused address, e.g., with an ICMP echo request, before allocating the address, and the client will probe the newly received address, e.g., with ARP.

3. The Client-Server Protocol

DHCP uses the BOOTP message format defined in RFC 951 and given in table 1 and figure 1. The 'op' field of each DHCP message sent from a client to a server contains BOOTREQUEST. BOOTREPLY is used in the 'op' field of each DHCP message sent from a server to a client.

The first four octets of the 'options' field of the DHCP message contain the (decimal) values 99, 130, 83 and 99, respectively (this

Roku Exhibit 1041
Page 11 of 39

is the same magic cookie as is defined in RFC 1395). The remainder of the 'options' field consists a list of tagged parameters that are called "options". All of the "vendor extensions" listed in RFC 1395 are also DHCP options. A separate document gives the complete set of options defined for use with DHCP [2].

Several options have been defined so far. One particular option - the "DHCP message type" option - must be included in every DHCP message. This option defines the "type" of the DHCP message. Additional options may be allowed, required, or not allowed, depending on the DHCP message type.

Throughout this document, DHCP messages that include a 'DHCP message type' option will be referred to by the type of the message; e.g., a DHCP message with 'DHCP message type' option type 1 will be referred to as a "DHCPDISCOVER" message.

3.1 Client-server interaction - allocating a network address

The following summary of the protocol exchanges between clients and servers refers to the DHCP messages described in table 2. The timeline diagram in figure 3 shows the timing relationships in a typical client-server interaction. If the client already knows its address, some steps may be omitted; this abbreviated interaction is described in section 3.2.

1. The client broadcasts a DHCPDISCOVER message on its local physical subnet. The DHCPDISCOVER message may include options that suggest values for the network address and lease duration. BOOTP relay agents may pass the message on to DHCP servers not on the same physical subnet.

2. Each server may respond with a DHCPOFFER message that includes an available network address in the 'yiaddr' field (and other configuration parameters in DHCP options). Servers need not reserve the offered network address, although the protocol will work more efficiently if the server avoids allocating the offered network address to another client. The server unicasts the DHCPOFFER message to the client (using the DHCP/BOOTP relay agent if necessary) if possible, or may broadcast the message to a broadcast address (preferably 255.255.255.255) on the client's subnet.

3. The client receives one or more DHCPOFFER messages from one or more servers. The client may choose to wait for multiple responses. The client chooses one server from which to request configuration parameters, based on the configuration parameters offered in the DHCPOFFER messages. The client broadcasts a

Roku Exhibit 1041
Page 12 of 39

DHCPREQUEST message that MUST include the 'server identifier' option to indicate which server it has selected, and may include other options specifying desired configuration values. This DHCPREQUEST message is broadcast and relayed through DHCP/BOOTP relay agents. To help ensure that any DHCP/BOOTP relay agents forward the DHCPREQUEST message to the same set of DHCP servers that received the original DHCPDISCOVER message, the DHCPREQUEST message must use the same value in the DHCP message header's 'secs' field and be sent to the same IP broadcast address as the original DHCPDISCOVER message. The client times out and retransmits the DHCPDISCOVER message if the client receives no DHCPOFFER messages.

4. The servers receive the DHCPREQUEST broadcast from the client. Those servers not selected by the DHCPREQUEST message use the message as notification that the client has declined that server's offer. The server selected in the DHCPREQUEST message commits the binding for the client to persistent storage and responds with a DHCPACK message containing the configuration parameters for the requesting client. The combination of 'chaddr' and assigned network address constitute an unique identifier for the client's lease and are used by both the client and server to identify a lease referred to in any DHCP messages. The 'yiaddr' field in the DHCPACK messages is filled in with the selected network address.

   If the selected server is unable to satisfy the DHCPREQUEST message (e.g., the requested network address has been allocated), the server SHOULD respond with a DHCPNAK message.

   A server may choose to mark addresses offered to clients in DHCPOFFER messages as unavailable. The server should mark an address offered to a client in a DHCPOFFER message as available if the server receives no DHCPREQUEST message from that client.

Roku Exhibit 1041
Page 13 of 39

FIELD	OCTETS	DESCRIPTION
op	1	Message op code / message type. 1 = BOOTREQUEST, 2 = BOOTREPLY
htype	1	Hardware address type, see ARP section in "Assigned Numbers" RFC; e.g., '1' = 10mb ethernet.
hlen	1	Hardware address length (e.g. '6' for 10mb ethernet).
hops	1	Client sets to zero, optionally used by relay-agents when booting via a relay-agent.
xid	4	Transaction ID, a random number chosen by the client, used by the client and server to associate messages and responses between a client and a server.
secs	2	Filled in by client, seconds elapsed since client started trying to boot.
flags	2	Flags (see figure 2).
ciaddr	4	Client IP address; filled in by client in DHCPREQUEST if verifying previously allocated configuration parameters.
yiaddr	4	'your' (client) IP address.
siaddr	4	IP address of next server to use in bootstrap; returned in DHCPOFFER, DHCPACK and DHCPNAK by server.
giaddr	4	Relay agent IP address, used in booting via a relay-agent.
chaddr	16	Client hardware address.
sname	64	Optional server host name, null terminated string.
file	128	Boot file name, null terminated string; "generic" name or null in DHCPDISCOVER, fully qualified directory-path name in DHCPOFFER.
options	312	Optional parameters field. See the options documents for a list of defined options.

Table 1: Description of fields in a DHCP message

Roku Exhibit 1041
Page 14 of 39

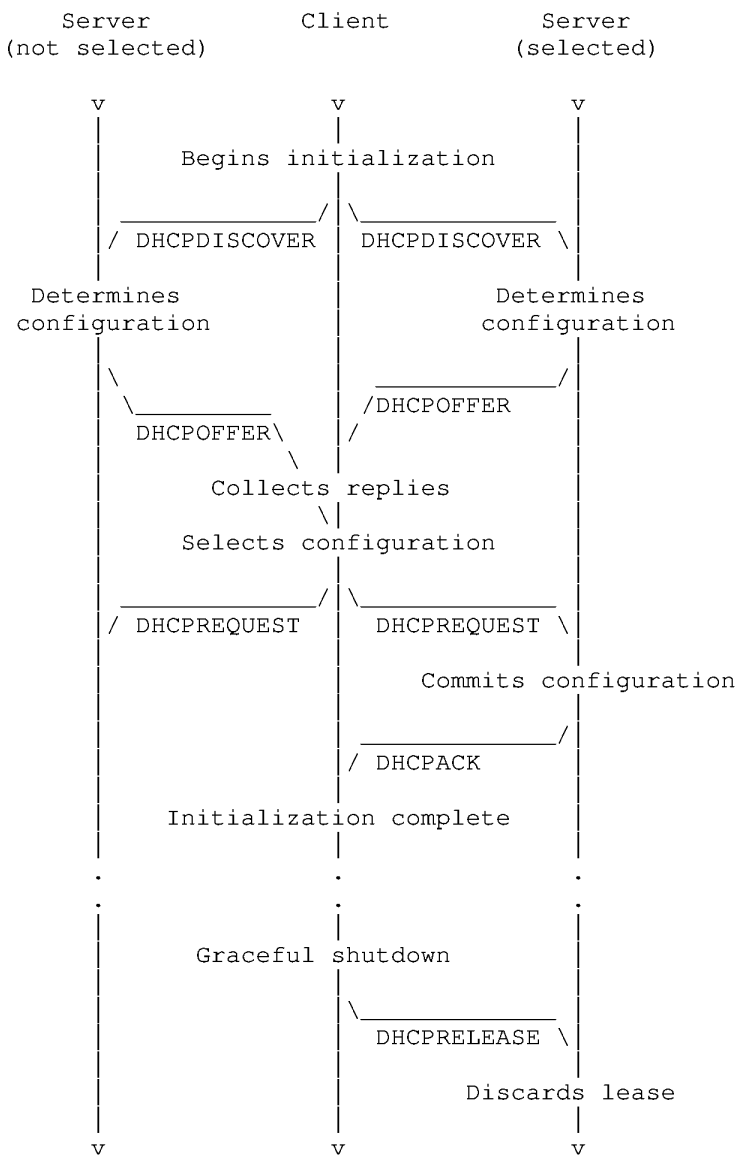


Figure 3: Timeline diagram of messages exchanged between DHCP client and servers when allocating a new network address

Roku Exhibit 1041
Page 15 of 39

Message	Use
DHCPDISCOVER	Client broadcast to locate available servers.
DHCPOFFER	Server to client in response to DHCPDISCOVER with offer of configuration parameters.
DHCPREQUEST	Client broadcast to servers requesting offered parameters from one server and implicitly declining offers from all others.
DHCPACK	Server to client with configuration parameters, including committed network address.
DHCPNAK	Server to client refusing request for configuration parameters (e.g., requested network address already allocated).
DHCPDECLINE	Client to server indicating configuration parameters (e.g., network address) invalid.
DHCPRELEASE	Client to server relinquishing network address and cancelling remaining lease.

Table 2: DHCP messages

5. The client receives the DHCPACK message with configuration parameters. The client performs a final check on the parameters (e.g., ARP for allocated network address), and notes the duration of the lease and the lease identification cookie specified in the DHCPACK message. At this point, the client is configured. If the client detects a problem with the parameters in the DHCPACK message, the client sends a DHCPDECLINE message to the server and restarts the configuration process. The client should wait a minimum of ten seconds before restarting the configuration process to avoid excessive network traffic in case of looping.

   If the client receives a DHCPNAK message, the client restarts the configuration process.

   The client times out and retransmits the DHCPREQUEST message if the client receives neither a DHCPACK or a DHCPNAK message. The client retransmits the DHCPREQUEST according to the retransmission algorithm in section 4.1. If the client receives neither a DHCPACK or a DHCPNAK message after ten retransmissions of the DHCPREQUEST message, the client reverts to INIT state and restarts the initialization process. The client SHOULD notify the user that the

Roku Exhibit 1041
Page 16 of 39

initialization process has failed and is restarting.

6. The client may choose to relinquish its lease on a network address by sending a DHCPRELEASE message to the server. The client identifies the lease to be released by including its network address in the 'ciaddr' field and its hardware address in the 'chaddr' field.

3.2 Client-server interaction - reusing a previously allocated network address

If a client remembers and wishes to reuse a previously allocated network address (allocated either by DHCP or some means outside the protocol), a client may choose to omit some of the steps described in the previous section. The timeline diagram in figure 4 shows the timing relationships in a typical client-server interaction for a client reusing a previously allocated network address.

1. The client broadcasts a DHCPREQUEST message on its local subnet. The DHCPREQUEST message includes the client's network address in the 'ciaddr' field. DHCP/BOOTP relay agents pass the message on to DHCP servers not on the same subnet.

2. Servers with knowledge of the client's configuration parameters respond with a DHCPACK message to the client.

   If the client's request is invalid (e.g., the client has moved to a new subnet), servers may respond with a DHCPNAK message to the client.

3. The client receives the DHCPACK message with configuration prameters. The client performs a final check on the parameters (as in section 3.1), and notes the duration of the lease and the lease identification cookie specified in the DHCPACK message. At this point, the client is configured.

   If the client detects a problem with the parameters in the DHCPACK message, the client sends a DHCPDECLINE message to the server and restarts the configuration process by requesting a new network address. This action corresponds to the client moving to the INIT state in the DHCP state diagram, which is described in section 4.4.

Roku Exhibit 1041
Page 17 of 39

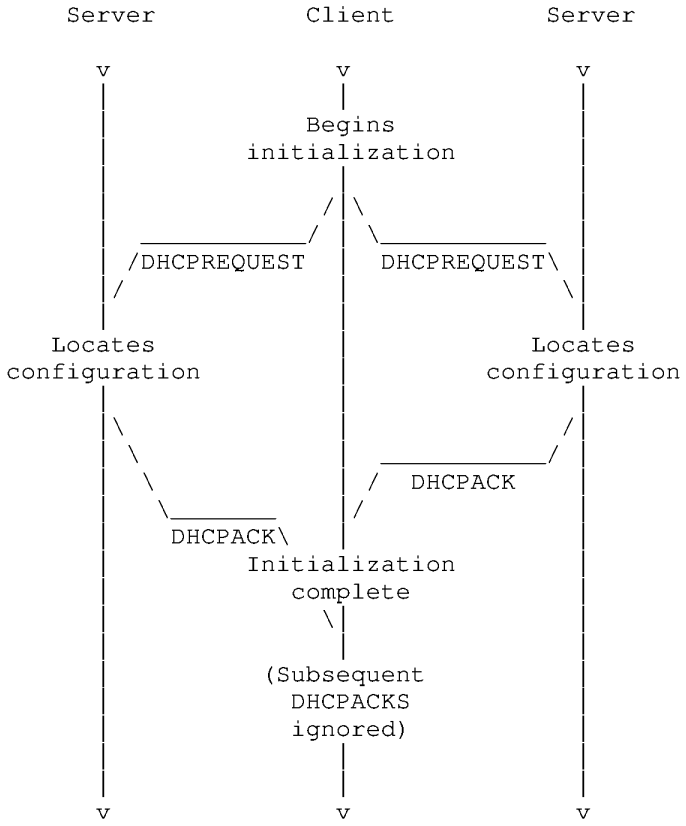


Figure 4: Timeline diagram of messages exchanged between DHCP client and servers when reusing a previously allocated network address

If the client receives a DHCPNAK message, it cannot reuse its remembered network address. It must instead request a new address by restarting the configuration process, this time using the (non-abbreviated) procedure described in section 3.1. This action also corresponds to the client moving to the INIT state in the DHCP state diagram.

The client times out and retransmits the DHCPREQUEST message if the client receives neither a DHCPACK nor a DHCPNAK message. The time between retransmission MUST be chosen according to the algorithm given in section 4.1. If the client receives no answer after transmitting 4 DHCPREQUEST messages, the client MAY choose to use the previously allocated network address and

Roku Exhibit 1041
Page 18 of 39

configuration parameters for the remainder of the unexpired lease. This corresponds to moving to BOUND state in the client state transition diagram shown in figure 5.

4. The client may choose to relinquish its lease on a network address by sending a DHCPRELEASE message to the server. The client identifies the lease to be released with the lease identification cookie.

   Note that in this case, where the client retains its network address locally, the client will not normally relinquish its lease during a graceful shutdown. Only in the case where the client explicitly needs to relinquish its lease, e.g., the client is about to be moved to a different subnet, will the client send a DHCPRELEASE message.

3.3 Interpretation and representation of time values

A client acquires a lease for a network address for a fixed period of time (which may be infinite). Throughout the protocol, times are to be represented in units of seconds. The time value of 0xffffffff is reserved to represent "infinity". The minimum lease duration is one hour.

As clients and servers may not have synchronized clocks, times are represented in DHCP messages as relative times, to be interpreted with respect to the client's local clock. Representing relative times in units of seconds in an unsigned 32 bit word gives a range of relative times from 0 to approximately 100 years, which is sufficient for the relative times to be measured using DHCP.

The algorithm for lease duration interpretation given in the previous paragraph assumes that client and server clocks are stable relative to each other. If there is drift between the two clocks, the server may consider the lease expired before the client does. To compensate, the server may return a shorter lease duration to the client than the server commits to its local database of client information.

3.4 Host parameters in DHCP

Not all clients require initialization of all parameters listed in Appendix A. Two techniques are used to reduce the number of parameters transmitted from the server to the client. First, most of the parameters have defaults defined in the Host Requirements RFCs; if the client receives no parameters from the server that override the defaults, a client uses those default values. Second, in its initial DHCPDISCOVER or DHCPREQUEST message, a client may provide the

Roku Exhibit 1041
Page 19 of 39

server with a list of specific parameters the client is interested in.

The client SHOULD include the 'maximum DHCP message size' option to let the server know how large the server may make its DHCP messages. The parameters returned to a client may still exceed the space allocated to options in a DHCP message. In this case, two additional options flags (which must appear in the 'options' field of the message) indicate that the 'file' and 'sname' fields are to be used for options.

The client can inform the server which configuration parameters the client is interested in by including the 'parameter request list' option. The data portion of this option explicitly lists the options requested by tag number.

In addition, the client may suggest values for the network address and lease time in the DHCPDISCOVER message. The client may include the be assigned, and may include the 'IP address lease time' option to suggest the lease time it would like. No other options representing "hints" at configuration parameters are allowed in a DHCPDISCOVER or DHCPREQUEST message. The 'ciaddr' field is to be filled in only in a DHCPREQUEST message when the client is requesting use of a previously allocated IP address.

If a server receives a DHCPREQUEST message with an invalid 'ciaddr', the server SHOULD respond to the client with a DHCPNAK message and may choose to report the problem to the system administrator. The server may include an error message in the 'message' option.

3.5 Use of DHCP in clients with multiple interfaces

A host with multiple network interfaces must use DHCP through each interface independently to obtain configuration information parameters for those separate interfaces.

3.6 When clients should use DHCP

A host should use DHCP to reacquire or verify its IP address and network parameters whenever the local network parameters may have changed; e.g., at system boot time or after a disconnection from the local network, as the local network configuration may change without the host's or user's knowledge.

If a host has knowledge of a previous network address and is unable to contact a local DHCP server, the host may continue to use the previous network address until the lease for that address expires. If the lease expires before the host can contact a DHCP server, the

Roku Exhibit 1041
Page 20 of 39

host must immediately discontinue use of the previous network address and may inform local users of the problem.

4. Specification of the DHCP client-server protocol

In this section, we assume that a DHCP server has a block of network addresses from which it can satisfy requests for new addresses. Each server also maintains a database of allocated addresses and leases in local permanent storage.

4.1 Constructing and sending DHCP messages

DHCP clients and servers both construct DHCP messages by filling in fields in the fixed format section of the message and appending tagged data items in the variable length option area. The options area includes first a four-octet 'magic cookie' (which was described in section 3), followed by the options. The last option must always be the 'end' option.

DHCP uses UDP as its transport protocol. DHCP messages from a client to a server are sent to the 'DHCP server' port (67), and DHCP messages from a server to a client are sent to the 'DHCP client' port (68).

DHCP messages broadcast by a client prior to that client obtaining its IP address must have the source address field in the IP header set to 0.

If the 'giaddr' field in a DHCP message from a client is non-zero, the server sends any return messages to the 'DHCP server' port on the DHCP relaying agent whose address appears in 'giaddr'. If the 'giaddr' field is zero, the client is on the same subnet, and the server sends any return messages to either the client's network address, if that address was supplied in the 'ciaddr' field, or to the client's hardware address or to the local subnet broadcast address.

If the options in a DHCP message extend into the 'sname' and 'file' fields, the 'option overload' option MUST appear in the 'options' field, with value 1, 2 or 3, as specified in the DHCP options document [2]. If the 'option overload' option is present in the 'options' field, the options in the 'options' field MUST be terminated by an options field. The options in the 'sname' and 'file' fields (if in use as indicated by the 'options overload' option) MUST begin with the first octet of the field, MUST be terminated by an 'end' option, and MUST be followed by 'pad' options to fill the remainder of the field. Any individual option in the 'options', 'sname' and 'file' fields MUST be entirely contained in

Roku Exhibit 1041
Page 21 of 39

that field. The options in the 'options' field MUST be interpreted first, so that any 'option overload' options may be interpreted. The 'file' field MUST be interpreted next (if the options), followed by the 'sname' field.

DHCP clients are responsible for all message retransmission. The client MUST adopt a retransmission strategy that incorporates a randomized exponential backoff algorithm to determine the delay between retransmissions. The delay before the first retransmission MUST be 4 seconds randomized by the value of a uniform random number chosen from the range -1 to +1. Clients with clocks that provide resolution granularity of less than one second may choose a non-integer randomization value. The delay before the next retransmission MUST be 8 seconds randomized by the value of a uniform number chosen from the range -1 to +1. The retransmission delay MUST be doubled with subsequent retransmissions up to a maximum of 64 seconds. The client MAY provide an indication of retransmission attempts to the user as an indication of the progress of the configuration process. The protocol specification in the remainder of this section will describe, for each DHCP message, when it is appropriate for the client to retransmit that message forever, and when it is appropriate for a client to abandon that message and attempt to use a different DHCP message.

Normally, DHCP servers and BOOTP relay agents attempt to deliver DHCPOFFER, DHCPACK and DHCPNAK messages directly to the client using unicast delivery. The IP destination address (in the IP header) is set to the DHCP 'yiaddr' address and the link-layer destination address is set to the DHCP 'chaddr' address. Unfortunately, some client implementations are unable to receive such unicast IP datagrams until the implementation has been configured with a valid IP address (leading to a deadlock in which the client's IP address cannot be delivered until the client has been configured with an IP address).

A client that cannot receive unicast IP datagrams until its protocol software has been configured with an IP address SHOULD set the BROADCAST bit in the 'flags' field to 1 in any DHCPDISCOVER or DHCPREQUEST messages that client sends. The BROADCAST bit will provide a hint to the DHCP server and BOOTP relay agent to broadcast any messages to the client on the client's subnet. A client that can receive unicast IP datagrams before its protocol software has been configured SHOULD clear the BROADCAST bit to 0. The BOOTP clarifications document discusses the ramifications of the use of the BROADCAST bit [21].

A server or relay agent sending or relaying a DHCP message directly to a DHCP client (i.e., not to a relay agent specified in the

Roku Exhibit 1041
Page 22 of 39

'giaddr' field) SHOULD examine the BROADCAST bit in the 'flags' field. If this bit is set to 1, the DHCP message SHOULD be sent as an IP broadcast using an IP broadcast address (preferably 255.255.255.255) as the IP destination address and the link-layer broadcast address as the link-layer destination address. If the BROADCAST bit is cleared to 0, the message SHOULD be sent as an IP unicast to the IP address specified in the 'yiaddr' field and the link-layer address specified in the 'chaddr' field. If unicasting is not possible, the message MAY be sent as an IP broadcast using an IP broadcast address (preferably 255.255.255.255) as the IP destination address and the link-layer broadcast address as the link-layer destination address.

4.2 DHCP server administrative controls

DHCP servers are not required to respond to every DHCPDISCOVER and DHCPREQUEST message they receive. For example, a network administrator, to retain stringent control over the hosts attached to the network, may choose to configure DHCP servers to respond only to hosts that have been previously registered through some external mechanism. The DHCP specification describes only the interactions between clients and servers when the clients and servers choose to interact; it is beyond the scope of the DHCP specification to describe all of the administrative controls that system administrators might want to use. Specific DHCP server implementations may incorporate any controls or policies desired by a network administrator.

In some environments, a DHCP server will have to consider the values of the 'chaddr' field and/or the 'class-identifier' option included in the DHCPDISCOVER or DHCPREQUEST messages when determining the correct parameters for a particular client. For example, an organization might have a separate bootstrap server for each type of client it uses, requiring the DHCP server to examine the 'class-identifier' to determine which bootstrap server address to return in the 'siaddr' field of a DHCPOFFER or DHCPACK message.

A DHCP server must use some unique identifier to associate a client with its lease. The client may choose to explicitly provide the identifier through the 'client identifier' option. If the client does not provide a 'client identifier' option, the server MSUT use the contents of the 'chaddr' field to identify the client.

DHCP clients are free to use any strategy in selecting a DHCP server among those from which the client receives a DHCPOFFER message. The client implementation of DHCP should provide a mechanism for the user to select directly the 'class-identifier' value.

Roku Exhibit 1041
Page 23 of 39

4.3 DHCP server behavior

A DHCP server processes incoming DHCP messages from a client based on the current state of the binding for that client. A DHCP server can receive the following messages from a client:

- o DHCPDISCOVER
- o DHCPREQUEST
- o DHCPDECLINE
- o DHCPRELEASE

Table 3 gives the use of the fields and options in a DHCP message by a server. The remainder of this section describes the action of the DHCP server for each possible incoming message.

4.3.1 DHCPDISCOVER message

When a server receives a DHCPDISCOVER message from a client, the server chooses a network address for the requesting client. If no address is available, the server may choose to report the problem to the system administrator and may choose to reply to the client with a DHCPNAK message. If the server chooses to respond to the client, it may include an error message in the 'message' option. If an address is available, the new address should be chosen as follows:

- o The client's previous address as recorded in the client's binding, if that address is in the server's pool of available addresses and not already allocated, else
- o The address requested in the 'Requested IP Address' option, if that address is valid and not already allocated, else
- o A new address allocated from the server's pool of available addresses.

Roku Exhibit 1041
Page 24 of 39

Field	DHCPOFFER	DHCPACK	DHCPNAK
'op'	BOOTREPLY	BOOTREPLY	BOOTREPLY
'htype'	(From "Assigned Numbers" RFC)		
'hlen'	(Hardware address length in octets)		
'hops'	0	0	0
'xid'	'xid' from client DHCPDISCOVER message	'xid' from client DHCPREQUEST message	'xid' from client DHCPREQUEST message
'secs'	0	0	0
'ciaddr'	0	'ciaddr' from DHCPREQUEST or 0	'ciaddr' from DHCPREQUEST or 0
'yiaddr'	IP address offered to client	IP address assigned to client	0
'siaddr'	IP address of next bootstrap server	IP address of next bootstrap server	0
'flags'	if 'giaddr' is not 0	then 'flags' from client message else 0	
'giaddr'	0	0	0
'chaddr'	'chaddr' from client DHCPDISCOVER message	'chaddr' from client DHCPREQUEST message	'chaddr' from client DHCPREQUEST message
'sname'	Server host name or options	Server host name or options	(unused)
'file'	Client boot file name or options	Client boot file name or options	(unused)
'options'	options	options	

Option	DHCPOFFER	DHCPACK	DHCPNAK
Requested IP address	MUST NOT	MUST NOT	MUST NOT
IP address lease time	MUST	MUST	MUST NOT
Use 'file'/'sname' fields	MAY	MAY	MUST NOT
DHCP message type	DHCPOFFER	DHCPACK	DHCPNAK
Parameter request list	MUST NOT	MUST NOT	MUST NOT
Message	SHOULD	SHOULD	SHOULD
Client identifier	MUST NOT	MUST NOT	MUST NOT
Class identifier	MUST NOT	MUST NOT	MUST NOT
Server identifier	MUST	MAY	MAY
Maximum message size	MUST NOT	MUST NOT	MUST NOT
All others	MAY	MAY	MUST NOT

Table 3: Fields and options used by DHCP servers

Roku Exhibit 1041
Page 25 of 39

As described in section 4.2, a server MAY, for administrative reasons, assign an address other than the one requested, or may refuse to allocate an address to a particular client even though free addresses are available.

While not required for correct operation of DHCP, the server should not reuse the selected network address before the client responds to the server's DHCPOFFER message. The server may choose to record the address as offered to the client.

The server must also choose an expiration time for the lease, as follows:

- IF the client has not requested a specific lease in the DHCPDISCOVER message and the client already has an assigned network address, the server returns the lease expiration time previously assigned to that address (note that the client must explicitly request a specific lease to extend the expiration time on a previously assigned address), ELSE

- IF the client has not requested a specific lease in the DHCPDISCOVER message and the client does not have an assigned network address, the server assigns a locally configured default lease time, ELSE

- IF the client has requested a specific lease in the DHCPDISCOVER message (regardless of whether the client has an assigned network address), the server may choose either to return the requested lease (if the lease is acceptable to local policy) or select another lease.

Once the network address and lease have been determined, the server constructs a DHCPOFFER message with the offered configuration parameters. It is important for all DHCP servers to return the same parameters (with the possible exception of a newly allocated network address) to ensure predictable host behavior regardless of the which server the client selects. The configuration parameters MUST be selected by applying the following rules in the order given below. The network administrator is responsible for configuring multiple DHCP servers to ensure uniform responses from those servers. The server MUST return to the client:

Roku Exhibit 1041
Page 26 of 39

- The client's network address, as determined by the rules given earlier in this section, and the subnet mask for the network to which the client is connected,

- The expiration time for the client's lease, as determined by the rules given earlier in this section,

- Parameters requested by the client, according to the following rules:

  -- IF the server has been explicitly configured with a default value for the parameter, the server MUST include that value in an appropriate option in the 'option' field, ELSE

  -- IF the server recognizes the parameter as a parameter defined in the Host Requirements Document, the server MUST include the default value for that parameter as given in the Host Requirements Document in an appropriate option in the 'option' field, ELSE

  -- The server MUST NOT return a value for that parameter,

- Any parameters from the existing binding that differ from the Host Requirements documents defaults,

- Any parameters specific to this client (as identified by the contents of 'chaddr' in the DHCPDISCOVER or DHCPREQUEST message), e.g., as configured by the network administrator,

- Any parameters specific to this client's class (as identified by the contents of the 'class identifier' option in the DHCPDISCOVER or DHCPREQUEST message), e.g., as configured by the network administrator; the parameters MUST be identified by an exact match between the client's 'client class' and the client class identified in the server,

- Parameters with non-default values on the client's subnet.

The server inserts the 'xid' field from the DHCPDISCOVER message into the 'xid' field of the DHCPOFFER message and sends the DHCPOFFER message to the requesting client.

4.3.2 DHCPREQUEST message

A DHCPREQUEST message may come from a client responding to a DHCPOFFER message from a server, or from a client verifying a previously allocated IP address. If the DHCPREQUEST message contains a 'server identifier' option, the message is in response to a

Roku Exhibit 1041
Page 27 of 39

DHCPOFFER message. Otherwise, the message is a request to renew or extend an existing lease.

Consider first the case of a DHCPREQUEST message in response to a DHCPOFFER message. If the server is identified in the 'server identifier' option in the DHCPREQUEST message, the server checks to confirm that the requested parameters are acceptable. Usually, the requested parameters will match those returned to the client in the DHCPOFFER message; however, the client may choose to request a different lease duration. Also, there is no requirement that the server cache the parameters from the DHCPOFFER message. The server must simply check that the parameters requested in the DHCPREQUEST are acceptable. If the parameters are acceptable, the server records the new client binding and returns a DHCPACK message to the client.

If the requested parameters are unacceptable, e.g., the requested lease time is unacceptable to local policy, the server sends a DHCPNAK message to the client. The server may choose to return an error message in the 'message' option.

If a different server is identified in the 'server identifier' field, the client has selected a different server from which to obtain configuration parameters. The server may discard any information it may have cached about the client's request, and may free the network address that it had offered to the client.

Note that the client may choose to collect several DHCPOFFER messages and select the "best" offer. The client indicates its selection by identifying the offering server in the DHCPREQUEST message. If the client receives no acceptable offers, the client may choose to try another DHCPDISCOVER message. Therefore, the servers may not receive a specific DHCPREQUEST from which they can decide whether or not the client has accepted the offer. Because the servers have not committed any network address assignments on the basis of a DHCPOFFER, servers are free to reuse offered network addresses in response to subsequent requests. As an implementation detail, servers should not reuse offered addresses and may use an implementation-specific timeout mechanism to decide when to reuse an offered address.

In the second case, when there is no 'server identifier' option, the client is renewing or extending a previously allocated IP address. The server checks to confirm that the requested parameters are acceptable. If the parameters specified in the DHCPREQUEST message match the previous parameters, or if the request for an extension of the lease (indicated by an extended 'IP address lease time' option) is acceptable, the server returns a DHCPACK message to the requesting client. Otherwise, the server returns a DHCPNAK message to the

Roku Exhibit 1041
Page 28 of 39

client. In particular, if the previously allocated network address in the 'ciaddr' field from the client does not match the network address recorded by the server for that client, the server sends a DHCPNAK to the client.

A DHCP server chooses the parameters to return in a DHCPACK message according to the same rules as used in constructing a DHCPOFFER message, as given in section 4.3.1.

4.3.3 DHCPDECLINE message

If the server receives a DHCPDECLINE message, the client has discovered through some other means that the suggested network address is already in use. The server MUST mark the network address as not allocated and SHOULD notify the local system administrator of a possible configuration problem.

4.3.4 DHCPRELEASE message

Upon receipt of a DHCPRELEASE message, the server marks the network address as not allocated. The server should retain a record of the client's initialization parameters for possible reuse in response to subsequent requests from the client.

4.4 DHCP client behavior

Figure 5 gives a state-transition diagram for a DHCP client. A client can receive the following messages from a server:

- DHCPOFFER
- DHCPACK
- DHCPNAK

Table 4 gives the use of the fields and options in a DHCP message by a client. The remainder of this section describes the action of the DHCP client for each possible incoming message. The description in the following section corresponds to the full configuration procedure previously described in section 3.1, and the text in the subsequent section corresponds to the abbreviated configuration procedure described in section 3.2.

4.4.1 Initialization and allocation of network address

The client begins in INIT state and forms a DHCPDISCOVER message. The client should wait a random time between one and ten seconds to desynchronize the use of DHCP at startup. The client sets 'ciaddr'

Roku Exhibit 1041
Page 29 of 39

to all 0x00000000. The client MAY request specific parameters by including the 'parameter request list' option. The client MAY suggest a network address and/or lease time by including the 'requested IP address' and 'IP address lease time' options. The client MUST include its hardware address in the 'chaddr' field for use in delivery of DHCP reply messages. The client MAY include a different unique identifier in the 'client identifier' option. If the client does not include the

The client generates and records a random transaction identifier and inserts that identifier into the 'xid' field. The client records its own local time for later use in computing the lease expiration. The client then broadcasts the DHCPDISCOVER on the local hardware broadcast address to the all-ones IP broadcast address and 'DHCP server' UDP port.

If the 'xid' of an arriving DHCPOFFER message does not match the 'xid' of the most recent DHCPDISCOVER message, the DHCPOFFER message must be silently discarded. Any arriving DHCPACK messages must be silently discarded.

The client collects DHCPOFFER messages over a period of time, selects one DHCPOFFER message from the (possibly many) incoming DHCPOFFER messages (e.g., the first DHCPOFFER message or the DHCPOFFER message from the previously used server) and extracts the server address from the 'server identifier' option in the DHCPOFFER message. The time over which the client collects messages and the mechanism used to select one DHCPOFFER are implementation dependent. The client may perform a check on the suggested address to ensure that the address is not already in use. For example, if the client is on a network that supports ARP, the client may issue an ARP request for the suggested request. When broadcasting an ARP request for the suggested address, the client must fill in its own hardware address as the sender's hardware address, and 0 as the sender's IP address, to avoid confusing ARP caches in other hosts on the same subnet. If the network address appears to be in use, the client sends a DHCPDECLINE message to the server and waits for another DHCPOFFER. As the client does not have a valid network address, the client must broadcast the DHCPDECLINE message.

Roku Exhibit 1041
Page 30 of 39

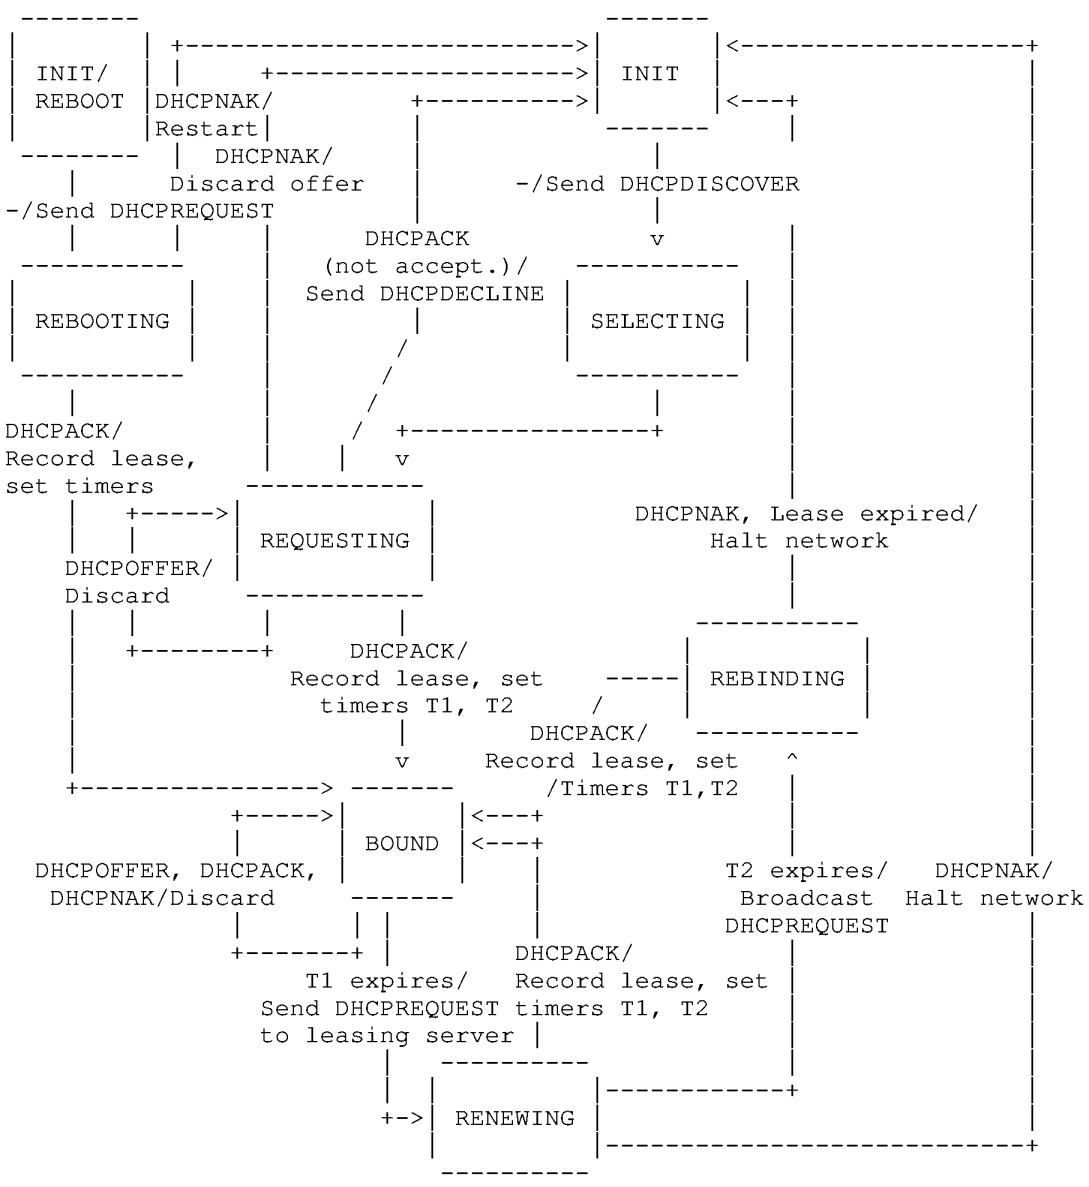


Figure 5: State-transition diagram for DHCP clients

Roku Exhibit 1041
Page 31 of 39

Field	DHCPDISCOVER	DHCPREQUEST	DHCPDECLINE, DHCPRELEASE
'op'	BOOTREQUEST	BOOTREQUEST	BOOTREQUEST
'htype'	(From "Assigned Numbers" RFC)		
'hlen'	(Hardware address length in octets)		
'hops'	0	0	0
'xid'	selected by client	selected by client	selected by client
'secs'	(opt.)	(opt.)	0
'flags'	Set 'BROADCAST' flag if client requires broadcast reply	Set 'BROADCAST' flag if client requires broadcast reply	
	0		
'ciaddr'	0	previously allocated newtork address	ciaddr
'yiaddr'	0	0	0
'siaddr'	0	0	0
'giaddr'	0	0	0
'chaddr'	client's hardware address	client's hardware address	client's hardware address
'sname'	options, if indicated in 'sname/file' option; otherwise unused	options, if indicated in 'sname/file' option; otherwise unused	(unused)
'file'	options, if indicated in 'sname/file' option; otherwise 'generic' name or null	options, if indicated in 'sname/file' option; otherwise 'generic' name or null	(unused)
'options'	options	options	(unused)

Roku Exhibit 1041
Page 32 of 39

Option	DHCPDISCOVER	DHCPREQUEST	DHCPDECLINE, DHCPRELEASE
Requested IP address	MAY	MUST NOT	MUST NOT
IP address lease time	MAY	MAY	MUST NOT
Use 'file'/'sname' fields	MAY	MAY	MAY
DHCP message type	DHCPDISCOVER	DHCPREQUEST	DHCPDECLINE/ DHCPRELEASE
Client identifier	MAY	MAY	MAY
Class identifier	SHOULD	SHOULD	MUST NOT
Server identifier	MUST NOT	MUST (after DHCPDISCOVER), MUST NOT (when renewing)	MUST
Parameter request list	MAY	MAY	MUST NOT
Maximum message size	MAY	MAY	MUST NOT
Message	SHOULD NOT	SHOULD NOT	SHOULD
Site-specific	MAY	MAY	MUST NOT
All others	MUST NOT	MUST NOT	MUST NOT

Table 4: Fields and options used by DHCP clients

If the parameters are acceptable, the client records the address of the server that supplied the parameters from the 'server identifier' field and sends that address in the 'server identifier' field of a DHCPREQUEST broadcast message. Once the DHCPACK message from the server arrives, the client is initialized and moves to BOUND state. The DHCPREQUEST message contains the same 'xid' as the DHCPOFFER message. The client records the lease expiration time as the sum of the time at which the original request was sent and the duration of the lease from the DHCPOFFER message. The client SHOULD broadcast an ARP reply to announce the client's new IP address and clear any outdated ARP cache entries in hosts on the client's subnet.

4.4.2 Initialization with known network address

The client begins in INIT-REBOOT state and sends a DHCPREQUEST message with the 'ciaddr' field set to the client's network address. The client may request specific configuration parameters by including the random transaction identifier and inserts that identifier into the computing the lease expiration. The client MUST NOT incldue a 'server identifier' in the DHCPREQUEST message. The client then broadcasts the DHCPREQUEST on the local hardware broadcast address to the 'DHCP server' UDP port.

Once a DHCPACK message with an 'xid' field matching that in the client's DHCPREQUEST message arrives from any server, the client is

Roku Exhibit 1041
Page 33 of 39

initialized and moves to BOUND state. The client records the lease expiration time as the sum of the time at which the DHCPREQUEST message was sent and the duration of the lease from the DHCPACK message.

4.4.3 Initialization with a known DHCP server address

When the DHCP client knows the address of a DHCP server, in either INIT or REBOOTING state, the client may use that address in the DHCPDISCOVER or DHCPREQUEST rather than the IP broadcast address. If the client receives no response to DHCP messages sent to the IP address of a known DHCP server, the DHCP client reverts to using the IP broadcast address.

4.4.4 Reacquisition and expiration

The client maintains two times, T1 and T2, that specify the times at which the client tries to extend its lease on its network address. T1 is the time at which the client enters the RENEWING state and attempts to contact the server that originally issued the client's network address. T2 is the time at which the client enters the REBINDING state and attempts to contact any server.

At time T1 after the client accepts the lease on its network address, the client moves to RENEWING state and sends (via unicast) a DHCPREQUEST message to the server to extend its lease. The client generates a random transaction identifier and inserts that identifier into the 'xid' field in the DHCPREQUEST. The client records the local time at which the DHCPREQUEST message is sent for computation of the lease expiration time. The client MUST NOT include a 'server identifier' in the DHCPREQUEST message.

Any DHCPACK messages that arrive with an 'xid' that does not match the When the client receives a DHCPACK from the server, the client computes the lease expiration time as the sum of the time at which the client sent the DHCPREQUEST message and the duration of the lease in the DHCPACK message. The client has successfully reacquired its network address, returns to BOUND state and may continue network processing.

If no DHCPACK arrives before time T2 (T2 > T1) before the expiration of the client's lease on its network address, the client moves to REBINDING state and sends (via broadcast) a DHCPREQUEST message to extend its lease. The client sets the 'ciaddr' field in the DHCPREQUEST to its current network address. The client MUST NOT include a 'server identifier' in the DHCPREQUEST message.

Times T1 and T2 are configurable by the server through options. T1

Roku Exhibit 1041
Page 34 of 39

defaults to (0.5 * duration_of_lease). T2 defaults to (0.875 * duration_of_lease). Times T1 and T2 should be chosen with some random "fuzz" around a fixed value, to avoid synchronization of client reacquisition.

In both RENEWING and REBINDING state, if the client receives no response to its DHCPREQUEST message, the client should wait one-half the remaining time until the expiration of T1 (in RENEWING state) and T2 (in REBINDING state) down to a minimum of 60 seconds, before retransmitting the DHCPREQUEST message.

If the lease expires before the client receives a DHCPACK, the client moves to INIT state, MUST immediately stop any other network processing and requests network initialization parameters as if the client were uninitialized. If the client then receives a DHCPACK allocating that client its previous network address, the client SHOULD continue network processing. If the client is given a new network address, it MUST NOT continue using the previous network address and SHOULD notify the local users of the problem.

4.4.5 DHCPRELEASE

If the client no longer requires use of its assigned network address (e.g., the client is gracefully shut down), the client sends a DHCPRELEASE message to the server. Note that the correct operation of DHCP does not depend on the transmission of DHCPRELEASE messages.

5. Acknowledgments

Greg Minshall, Leo McLaughlin and John Veizades have patiently contributed to the the design of DHCP through innumerable discussions, meetings and mail conversations. Jeff Mogul first proposed the client-server based model for DHCP. Steve Deering searched the various IP RFCs to put together the list of network parameters supplied by DHCP. Walt Wimer contributed a wealth of practical experience with BOOTP and wrote a document clarifying the behavior of BOOTP/DHCP relay agents. Jesse Walker analyzed DHCP in detail, pointing out several inconsistencies in earlier specifications of the protocol. Steve Alexander reviewed Walker's analysis and the fixes to the protocol based on Walker's work. And, of course, all the members of the Dynamic Host Configuration Working Group of the IETF have contributed to the design of the protocol through discussion and review of the protocol design.

Roku Exhibit 1041
Page 35 of 39

6. References

[1] Acetta, M., "Resource Location Protocol", RFC 887, CMU, December 1983.

[2] Alexander, S., and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 1533, Lachman Technology, Inc., Bucknell University, October 1993.

[3] Braden, R., Editor, "Requirements for Internet Hosts -- Communication Layers", STD 3, RFC 1122, USC/Information Sciences Institute, October 1989.

[4] Braden, R., Editor, "Requirements for Internet Hosts -- Application and Support, STD 3, RFC 1123, USC/Information Sciences Institute, October 1989.

[5] Brownell, D, "Dynamic Reverse Address Resolution Protocol (DRARP)", Work in Progress.

[6] Comer, D., and R. Droms, "Uniform Access to Internet Directory Services", Proc. of ACM SIGCOMM '90 (Special issue of Computer Communications Review), 20(4):50--59, 1990.

[7] Croft, B., and J. Gilmore, "Bootstrap Protocol (BOOTP)", RFC 951, Stanford and SUN Microsystems, September 1985.

[8] Deering, S., "ICMP Router Discovery Messages", RFC 1256, Xerox PARC, September 1991.

[9] Droms, D., "Interoperation between DHCP an BOOTP" RFC 1534, Bucknell University, October 1993.

[10] Finlayson, R., Mann, T., Mogul, J., and M. Theimer, "A Reverse Address Resolution Protocol", RFC 903, Stanford, June 1984.

[11] Gray C., and D. Cheriton, "Leases: An Efficient Fault-Tolerant Mechanism for Distributed File Cache Consistency", In Proc. of the Twelfth ACM Symposium on Operating Systems Design, 1989.

[12] Mockapetris, P., "Domain Names -- Concepts and Facilities", STD 13, RFC 1034, USC/Information Sciences Institute, November 1987.

[13] Mockapetris, P., "Domain Names -- Implementation and Specification", STD 13, RFC 1035, USC/Information Sciences Institute, November 1987.

Roku Exhibit 1041
Page 36 of 39

[14] Mogul J., and S. Deering, "Path MTU Discovery", RFC 1191, November 1990.

[15] Morgan, R., "Dynamic IP Address Assignment for Ethernet Attached Hosts", Work in Progress.

[16] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, USC/Information Sciences Institute, September 1981.

[17] Reynolds, J., "BOOTP Vendor Information Extensions", RFC 1497, USC/Information Sciences Institute, August 1993.

[18] Reynolds, J., and J. Postel, "Assigned Numbers", STD 2, RFC 1340, USC/Information Sciences Institute, July 1992.

[19] Jeffrey Schiller and Mark Rosenstein. A Protocol for the Dynamic Assignment of IP Addresses for use on an Ethernet. (Available from the Athena Project, MIT), 1989.

[20] Sollins, K., "The TFTP Protocol (Revision 2)", RFC 783, NIC, June 1981.

[21] Wimer, W., "Clarifications and Extensions for the Bootstrap Protocol", RFC 1532, Carnegie Mellon University, October 1993.

7. Security Considerations

DHCP is built directly on UDP and IP which are as yet inherently insecure. Furthermore, DHCP is generally intended to make maintenance of remote and/or diskless hosts easier. While perhaps not impossible, configuring such hosts with passwords or keys may be difficult and inconvenient. Therefore, DHCP in its current form is quite insecure.

Unauthorized DHCP servers may be easily set up. Such servers can then send false and potentially disruptive information to clients such as incorrect or duplicate IP addresses, incorrect routing information (including spoof routers, etc.), incorrect domain nameserver addresses (such as spoof nameservers), and so on. Clearly, once this seed information is in place, an attacker can further compromise affected systems.

Malicious DHCP clients could masquerade as legitimate clients and retrieve information intended for those legitimate clients. Where dynamic allocation of resources is used, a malicious client could claim all resources for itself, thereby denying resources to legitimate clients.

Roku Exhibit 1041
Page 37 of 39

8. Author's Address

Ralph Droms
Computer Science Department
323 Dana Engineering
Bucknell University
Lewisburg, PA 17837

Phone: (717) 524-1145
EMail: droms@bucknell.edu

Roku Exhibit 1041
Page 38 of 39

A. Host Configuration Parameters

IP-layer_parameters,_per_host:_		
Be a router	on/off	HRC 3.1
Non-local source routing	on/off	HRC 3.3.5
Policy filters for non-local source routing	(list)	HRC 3.3.5
Maximum reassembly size	integer	HRC 3.3.2
Default TTL	integer	HRC 3.2.1.7
PMTU aging timeout	integer	MTU 6.6
MTU plateau table	(list)	MTU 7
IP-layer_parameters,_per_interface:_		
IP address	(address)	HRC 3.3.1.6
Subnet mask	(address mask)	HRC 3.3.1.6
MTU	integer	HRC 3.3.3
All-subnets-MTU	on/off	HRC 3.3.3
Broadcast address flavor	0x00000000/0xffffffff	HRC 3.3.6
Perform mask discovery	on/off	HRC 3.2.2.9
Be a mask supplier	on/off	HRC 3.2.2.9
Perform router discovery	on/off	RD 5.1
Router solicitation address	(address)	RD 5.1
Default routers, list of:		
router address	(address)	HRC 3.3.1.6
preference level	integer	HRC 3.3.1.6
Static routes, list of:		
destination	(host/subnet/net)	HRC 3.3.1.2
destination mask	(address mask)	HRC 3.3.1.2
type-of-service	integer	HRC 3.3.1.2
first-hop router	(address)	HRC 3.3.1.2
ignore redirects	on/off	HRC 3.3.1.2
PMTU	integer	MTU 6.6
perform PMTU discovery	on/off	MTU 6.6
Link-layer_parameters,_per_interface:_		
Trailers	on/off	HRC 2.3.1
ARP cache timeout	integer	HRC 2.3.2.1
Ethernet encapsulation	(RFC 894/RFC 1042)	HRC 2.3.3
TCP_parameters,_per_host:_		
TTL	integer	HRC 4.2.2.19
Keep-alive interval	integer	HRC 4.2.3.6
Keep-alive data size	0/1	HRC 4.2.3.6

Key:

MTU = Path MTU Discovery (RFC 1191, Proposed Standard)
RD = Router Discovery (RFC 1256, Proposed Standard)

Roku Exhibit 1041
Page 39 of 39

#1 Best Seller
Over 400,000 Sold

Over 20,000
terms defined

NEWTON'S TELECOM DICTIONARY

The Official Dictionary of Telecommunications Networking and the Internet

16th EXPANDED & UPDATED EDITION

BY HARRY NEWTON

Roku Exhibit 1070
Page 1 of 4

NEWTON'S TELECOM DICTIONARY
copyright © 2000 Harry Newton
email: Harry_Newton@TechnologyInvestor.com
personal web site: www.HarryNewton.com

All rights reserved under International and Pan-American Copyright conventions, including the right to reproduce this book or portions thereof in any form whatsoever.

Published by CMP Books
An Imprint of CMP Media Inc.
12 West 21 Street
New York, NY 10010

ISBN 1-57820-053-9

July, 2000

Sixteenth and a Half Edition, Expanded and Updated

For individual orders, and for information on special discounts for quantity orders, please contact:

CMP Books
6600 Silacci Way
Gilroy, CA 95020
Tel: 800-LIBRARY or 408-848-3854
Fax: 408-848-5784
Email: telecom@rushorder.com

Distributed to the book trade in the U.S. and Canada by
Publishers Group West
1700 Fourth St., Berkeley, CA 94710

Manufactured in the United States of America

NEW
TEL
DICT

The Official Dic
Net

16th and a Ha

NEWTON's TELECOM DICTIONARY

The Official Dictionary of Telecommunications
Networking and Internet

16th and a Half Updated, Expanded and Much Improved Edition

Roku Exhibit 1070
Page 3 of 4

ices (generally GPS) into automobiles. ple is GM's OnStar system, which auto- distance if the vehicle is in an accident, also perform such functions as remote tracking stolen vehicles, provide roadside ssistance.com

ce Exchange MTX. The Northern SO (Mobile Telephone Switching Office).

ted The term used to describe a call of the call is a mobile (i.e. cellular) tele-

cellular telephone equipment installed in of a transceiver, control head, handset

8" cellular phone service covering 83% opulation. It is a consortium of six Bell s some Canadian cellular operators. The ne calling a subscriber of one of these es his call automatically routed to the where in Mobilink that subscriber was. ded to know where the person was you a bunch of complex codes to get to him s a paging system. The technical details ally referred to as the Mobilex interface can be obtained from any operator of a ek or from Ericsson directly. Mobilex is etary system but the specification is made available under a royalty-free sis in the public domain. All Mobilex net- at three frequency families, 80 MHz, 400 frequencies and channels are generally f government authority; each local oper- s specific channel assignments for its Mobilex networks around the world, its continents. Some of the networks are ed is, you may buy network service from hile others are owned and operated by ser use and benefit.

ding Call. A cellular phone term. MOC efers where the outbound leg of a call s.

on Also known as Differential Mode Dispersion can be thought of as the blur- d in a fiber by the several modes (paths) ke as it propagates through (transverse- l pulse in a digital optical fiber system. al pulse created by a light source, such erate (LED) or a Laser Diode, with later the high-speed, long-haul systems. As s through the thin and very pure exper- sturally spreads out. The inner core is ed of "cladding," glass of a slightly differ- ich serves to bend the light signal back ons of the light pulse travel more or less t the core while other portions of the s, they travel different distances from ions, some portions of the light pulse ers. The result of Modal Dispersion is ension," which is a blurring or smearing ses. In other words, the pulses overrun er individual identity. The thicker the s the effect. The longer the cable, the aster the transmission speed, the short- se pulses, and the worse the effect. As Single Mode Fiber (SMF) has a thinner inner core than MultiMode Fiber (MMF), the effect of modal dispersion is less with SMF. See also Collimation, MMF and SMF.

Modal dispersion can be thought of as the blurring of the input signal in a fiber by the several modes that may propagate down a fiber. Each mode may take a separate path down the fiber, and thus that signal may arrive at slightly different times. The dispersion depends on the fiber's internal characteristics and its length. In short, modal dispersion is pulse rounding in lightwave communications that takes place because of the slightly different paths followed by the laser light rays as they arrive at the detector slightly out of phase.

Modal Distribution In an optical fiber operating at a single wavelength, the number of modes supported by the fiber, and their propagation time differences. In an optical fiber operating at multiple wavelengths simultaneously, the separation in wavelengths among the modes being supported by the fiber.

Modal Loss In an open waveguide, such as an optical fiber, a loss of energy on the part of an electromagnetic wave due to obstacles outside the waveguide, abrupt changes in direction of the waveguide, or other anomalies, that cause changes in the propagation mode of the wave in the waveguide.

Mode 1. A mode is a stable propagation of state of an electromagnetic wave in an optical fiber. In rough terms a mode can be thought of bundles of light rays of the same wavelength, that enter a fiber at the same angle. See Modal Dispersion.

2. Mode is essentially a switch inside a computer that makes it run like another computer, usually an older one.

Mode S Discrete addressable secondary radar system with data link.

Modem 1. Acronym for MOdulator/DEModulator. Conventional modems comprise equipment which converts digital signals to analog signals and vice versa. Modems are used to send digital data signals over the analog PSTN (Public Switched Telephone Network. Although the carrier switches (e.g., central offices and tandem offices) are typically digital, as is the backbone transmission network (e.g., T-carrier), the local loop always is analog unless the user orders a more costly digital loop (e.g., ISDN or T-1). Therefore, the PSTN is analog as far as most people are concerned.

Conventional modems work like this. Your PC outputs data in the form of "1's" and "0's" which are represented by varying levels of voltage. The modem converts the digital data signal into variations of the analog sine wave so the data can be transmitted over the PSTN. A matching modem on the other end reverses the process in order to present the target device with a digital bit stream. The modulation techniques include some combination of Amplitude Modulation (AM), Frequency Modulation (FM) and Phase Modulation (PM), also known as Phase Shift Keying (PSK). Used in combination, these techniques allow multiple bits to be represented with a single (or single set) of sine waves. In this fashion, compression is accomplished, which allows more data to be transmitted in the same period of time and which therefore reduces the connect time and the associated cost of the data transfer. Contemporary, conventional modems are standardized by the ITU-T as part of the "V" series of standards. Such modems are characterized by error detection and correction mechanisms, adaptive equalization, internal dialing, and numerous other sophisticated capabilities. 56 Kbps modems are the latest development in the world of conventional modems; they remain to be standardized. The term "modem" also is applied (and correctly so, in the purely technical sense) to ISDN TAs (Terminal Adapters), ADSL TUs (Terminating Units), line drivers and short-haul modems. The last two, in fact, are voltage converters.

See also Line Driver, Modem Eliminator, Modem Pool, Modem Standards, Modulation Protocols, Serial Port, Short-Haul Modem and 56 Kbps Modem.

2. According to the Vermonter's Guide to Computer Lingo, modem is what landscapers do to dem lawns. (This is a joke.)

Modem Bonding A term which describes the bonding, or linking, of two 56 Kbps modems over two phone lines to double the performance. This process is accomplished through matching devices, one on each end of the connection; each modem operates at its maximum achievable rate, with the aggregate rate being roughly double that of each individual modem. Theoretically, modem bonding can yield speeds of as much as 128 Kbps downstream and 97.2 Kbps upstream, although the FCC currently limits maximum downstream performance through each modem to 53 Kbps, for a total of 106 Kbps. Modem bonding technologies are proprietary. See 56 Kbps Modem.

Modem Cowboy A slang term for someone who typically lives in the mountain states of Western America and does most of his work via modem.

Modem Eliminator A wiring device designed to replace two modems; it connects equipment over a distance of up to several hundred feet. In asynchronous systems, this is a simple cable. Here is a specific application using a Modem Eliminator: You can connect a PC to a printer, or a PC to another printer using a cable. But you can only go a certain distance — maybe 100 feet. After that, the traditional solution has been to use a modem and go over traditional phone lines. Instead, you can connect the two devices directly by wire using a Modem Eliminator. There are two advantages of a Modem Eliminator over a normal modem. The eliminator is cheaper and it can often transmit faster. According to Glacgal Communications, there are many cases where it is either unnecessary, cumbersome or too expensive to interconnect terminals using modems or line drivers in an experimental or a very short-haul environment. A modem eliminator functionally resembles two modems back-to-back on a leased line and therefore saves the cost of two modems and a line in many situations.

Modem Fallback When the telephone line quality is not good enough to accommodate the top rated speed of a modem — for example 14,400 bps — the modem drops down to lower speeds — initially to 9,600 bps, then if necessary, to 4,800 bps, or even down to 2,400 bps.

Modem Pair Delay The one way delay added by a pair of modems to the propagation delay on a communications line.

Modem Pool A collection of modems which a user can dial up from his terminal, access one and use that one to make a data call over the switched telephone network. Modem pools are obviously designed to allow many users to share few modems, thus saving on modems. Now that modems have become less expensive, the advantages of modem pooling are no longer as great. There are also some advantages in having a modem right next to your terminal or computer — namely you can see how it is functioning. And modems have lights to indicate what they're doing. One of the most useful lights on most modems indicates whether the line is "off hook" or not. It is possible for your computer to instruct your modem to hang up the line and for your modem to forget to do it, leaving you with a huge phone bill. One problem with giving people their own modems is they (the modems) have a tendency

Roku Exhibit 1070
Page 4 of 4

Electronic Acknowledgement Receipt

EFS ID:	46328672
Application Number:	17714965
International Application Number:	
Confirmation Number:	7004
Title of Invention:	Targeted television advertisements based on online behavior
First Named Inventor/Applicant Name:	Roy Shkedi
Customer Number:	26362
Filer:	Louis J. Hoffman/Michelle Lisa
Filer Authorized By:	Louis J. Hoffman
Attorney Docket Number:	14-TV1-CON7
Receipt Date:	03-AUG-2022
Filing Date:	06-APR-2022
Time Stamp:	00:23:56
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Non Patent Literature	Zwicky_NPL_6-2000.pdf	6792196 e2ac0ab489d0d283bc4fb17c47f5f57b6fca4902	no	46

Warnings:

Information:					
2	Non Patent Literature	Hahn_NPL_3-28-1996.pdf	13808886 9a695b4b0432ace8675d009b6d60bd16d468eb37	no	74
Warnings:					
Information:					
3	Non Patent Literature	Srisuresh_NPL_RFC2663_8-1999.pdf	774540 a9fd25242a66f031eb5e78be96b89ffbd56397b3	no	30
Warnings:					
Information:					
4	Non Patent Literature	RFC793_9-1981.pdf	908374 42b12c4f98a51e106f52606f9170ca9599f02bae	no	91
Warnings:					
Information:					
5	Non Patent Literature	Zeff-Aronson-Part1.pdf	10486833 f444aaa5f44d72524538 4c011e6bc7f07d9aa225	no	75
Warnings:					
Information:					
6	Non Patent Literature	Zeff-Aronson-Part2-reduced.pdf	8152657 45333269752a72445b8aa8224fc82c19d29c79b0	no	105
Warnings:					
Information:					
7	Non Patent Literature	Zeff-Aronson-Part3-reduced-part2.pdf	1578413 e42990775e445789e17637e37d76496382fb3372	no	23
Warnings:					
Information:					
8	Non Patent Literature	Zeff-Aronson-Part3-reduced-part1-pdfA.pdf	11080403 8fd55eaa1a5af65734b8fe51547c49ac000ae80a	no	20
Warnings:					
Information:					

9	Non Patent Literature	RFC_950.pdf	610172 268505e51da94711f3674610a3f3669bdf79463c	no	18
Warnings:					
Information:					
10	Non Patent Literature	RFC_1531.pdf	652470 0c689f8e59af8a47455b9b6066616d392ec94fdc	no	39
Warnings:					
Information:					
11	Non Patent Literature	Newtons_Telecom_Dictionary.pdf	1645054 fefee3b66301ffe2651fed04b9421d5e2330dcef	no	4
Warnings:					
Information:					
		Total Files Size (in bytes):	56489998		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/08a (01-22)
Approved for use through 05/31/2024. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (**Not for submission under 37 CFR 1.99)**		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Ekpo, Nnenna
	Attorney Docket Number	TV1-CON7

U.S.PATENTS

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	9110903		08/18/2005	Martinez et al.	

If you wish to add additional U.S. Patent citation information please click the Add button.

U.S.PATENT APPLICATION PUBLICATIONS

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20050278731		12/15/2005	Cameron et al.	
	2	20080120178		05/22/2008	Martinez et al.	
	3	20050160002		07/21/2005	Roetter et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] ¡	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (**Not for submission under 37 CFR 1.99)**		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Ekpo, Nnenna
	Attorney Docket Number	TV1-CON7

If you wish to add additional Foreign Patent Document citation information please click the Add button

NON-PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Ekpo, Nnenna
	Attorney Docket Number	TV1-CON7

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Louis J. Hoffman/	Date (YYYY-MM-DD)	2022-09-16
Name/Print	Louis J. Hoffman	Registration Number	38918

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	46637954
Application Number:	17714965
International Application Number:	
Confirmation Number:	7004
Title of Invention:	Targeted television advertisements based on online behavior
First Named Inventor/Applicant Name:	Roy Shkedi
Customer Number:	26362
Filer:	Louis J. Hoffman/Michelle Lisa
Filer Authorized By:	Louis J. Hoffman
Attorney Docket Number:	14-TV1-CON7
Receipt Date:	16-SEP-2022
Filing Date:	06-APR-2022
Time Stamp:	19:16:12
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Transmittal Letter	22-09-16-TV1-CON7-IDS_Cover.pdf	492635 d3d303ee8e651f92d56b207cb730fb4d162f8703	no	2
Warnings:					

Information:					
2	Information Disclosure Statement (IDS) Form (SB08)	22-09-16-TV1-CON7-Supp-IDS.pdf	466658 d89e5a279a2efbc4234fbe63c8685bf85e38ec9a	no	4
Warnings:					
Information:					
This is not an USPTO supplied IDS fillable form					
			Total Files Size (in bytes):	959293	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor	:	Shkedi, Roy	Art Unit	:	3693 (in parent)
Serial No.	:	17/714,965	Examiner	:	Ekpo, Nnenna (in parent)
Filing Date	:	04/06/2022	Conf. No.	:	7004
Title	:	Targeted television advertisements based on online behavior			

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Filed via EFS-Web on September 16, 2022

INFORMATION DISCLOSURE STATEMENT

Dear Sir:

Assignee discloses the references listed on attached form PTO/SB/08.

Also, assignee wishes to direct the examiner to the co-pending proceedings in the table below, as well as the references cited in those proceedings. All listed applications and the listed patent share the same first inventor:

Serial No.	Filing Date	Patent Challenged
IPR2022-01507	9/7/2022	Re: USP 10,321,198 (same family as SN 17/401,078, previously cited)
IPR2022-01506	9/7/2022	Re: USP 8,566,164 (same family as SN 17/401,078, previously cited)
IPR2022-01420	8/12/2022	Re: parent USP 8,677,398
IPR2022-01505	9/7/2022	Re: parent USP 10,715,878

If the Office has any questions, please feel free to contact assignee's undersigned attorney of record.

Respectfully submitted,

INTENT IQ, LLC.

by its attorney

Dated: September 16, 2022

/Louis J. Hoffman/
Louis J. Hoffman
Reg. No. 38,918

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, Arizona 85260
(480) 948-3295

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
17/714,965	04/06/2022	Roy Shkedi	14-TV1-CON7

CONFIRMATION NO. 7004

26362
LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, AZ 85260

PUBLICATION NOTICE

OC000000136815226

Title: Targeted television advertisements based on online behavior

Publication No.US-2022-0360862-A1
Publication Date:11/10/2022

NOTICE OF PUBLICATION OF APPLICATION

The above-identified application will be electronically published as a patent application publication pursuant to 37 CFR 1.211, et seq. The patent application publication number and publication date are set forth above.

The publication may be accessed through the USPTO's publically available Searchable Databases via the Internet at www.uspto.gov. The direct link to access the publication is currently http://www.uspto.gov/patft/.

The publication process established by the Office does not provide for mailing a copy of the publication to applicant. A copy of the publication may be obtained from the Office upon payment of the appropriate fee set forth in 37 CFR 1.19(a)(1). Orders for copies of patent application publications are handled by the USPTO's Public Records Division. The Public Records Division can be reached by telephone at (571) 272-3150 or (800) 972-6382, by facsimile at (571) 273-3250, by mail addressed to the United States Patent and Trademark Office, Public Records Division, Alexandria, VA 22313-1450 or via the Internet.

In addition, information on the status of the application, including the mailing date of Office actions and the dates of receipt of correspondence filed in the Office, may also be accessed via the Internet through the Patent Electronic Business Center at www.uspto.gov using the public side of the Patent Application Information and Retrieval (PAIR) system. The direct link to access this status information is currently https://portal.uspto.gov/pair/PublicPair. Prior to publication, such status information is confidential and may only be obtained by applicant using the private side of PAIR.

Further assistance in electronically accessing the publication, or about PAIR, is available by calling the Patent Electronic Business Center at 1-866-217-9197.

Office of Data Managment, Application Assistance Unit (571) 272-4000, or (571) 272-4200, or 1-888-786-0101

To: Louis@valuablepatents.com,donald@valuablepatents.com,dsalavi@northwestpatent.com
From: PAIR_eOfficeAction@uspto.gov
Cc: PAIR_eOfficeAction@uspto.gov
Subject: Private PAIR Correspondence Notification for Customer Number 26362

Nov 11, 2022 03:27:04 AM

Dear PAIR Customer:

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, AZ 85260
UNITED STATES

The following USPTO patent application(s) associated with your Customer Number, 26362 , have new outgoing correspondence. This correspondence is now available for viewing in Private PAIR.

The official date of notification of the outgoing correspondence will be indicated on the form PTOL-90 accompanying the correspondence.

Disclaimer:
The list of documents shown below is provided as a courtesy and is not part of the official file wrapper. The content of the images shown in PAIR is the official record.

Application	Document	Mailroom Date	Attorney Docket No.
17714965	NTC.PUB	11/10/2022	14-TV1-CON7

To view your correspondence online or update your email addresses, please visit us anytime at https://sportal.uspto.gov/secure/myportal/privatepair.

If you have any questions, please email the Electronic Business Center (EBC) at EBC@uspto.gov with 'e-Office Action' on the subject line or call 1-866-217-9197 during the following hours:

Monday - Friday 6:00 a.m. to 12:00 a.m.

Thank you for prompt attention to this notice,

UNITED STATES PATENT AND TRADEMARK OFFICE

PATENT APPLICATION INFORMATION RETRIEVAL SYSTEM

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

U.S. PATENT DOCUMENTS

Examiner Initials	Cite No.	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
_____	1	5105184	04/14/1992	Pirani et al.
_____	2	5220655	06/15/1993	Tsutsui
_____	3	5359508	10/25/1994	Rossides
_____	4	5442633	08/15/1995	Perkins et al.
_____	5	5572643	11/05/1996	Judson
_____	6	5710884	01/20/1998	Dedrick
_____	7	5724424	03/03/1998	Gifford
_____	8	5724521	03/03/1998	Dedrick
_____	9	5751956	05/12/1998	Kirsch
_____	10	5774670	06/30/1998	Montulli
_____	11	5794210	08/11/1998	Goldhaber et al.
_____	12	5809242	09/15/1998	Shaw et al.
_____	13	5819092	10/06/1998	Ferguson et al.
_____	14	5819285	10/06/1998	Damico et al.
_____	15	5835718	11/10/1998	Blewett
_____	16	5838314	11/17/1998	Neel et al.
_____	17	5848396	12/08/1998	Gerace
_____	18	5855008	12/29/1998	Goldhaber et al.
_____	19	5870546	02/09/1999	Kirsch
_____	20	5870550	02/09/1999	Wesinger Jr. et al.
_____	21	5901287	05/04/1999	Bull et al.
_____	22	5918014	06/29/1999	Robinson
_____	23	5933811	08/03/1999	Angles et al.
_____	24	5937390	08/10/1999	Hyodo
_____	25	5937392	08/10/1999	Alberts
_____	26	5948061	09/07/1999	Merriman et al.
_____	27	5959623	09/28/1999	van Hoff et al.
_____	28	5963915	10/05/1999	Kirsch
_____	29	5978013	11/02/1999	Jones et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

_____	30	5995943	11/30/1999	Bull et al.
_____	31	5999912	12/07/1999	Wodarz et al.
_____	32	6002394	12/14/1999	Schein et al.
_____	33	6014638	01/11/2000	Burge et al.
_____	34	6016504	01/18/2000	Arnold et al.
_____	35	6018768	01/25/2000	Ullman et al.
_____	36	6026368	02/15/2000	Brown et al.
_____	37	6026369	02/15/2000	Capek
_____	38	6026374	02/15/2000	Chess
_____	39	6047310	04/04/2000	Kamakura et al.
_____	40	6049835	04/11/2000	Gagnon
_____	41	6055573	04/25/2000	Gardenswartz et al.
_____	42	6073241	06/06/2000	Rosenberg et al.
_____	43	6128663	10/03/2000	Thomas
_____	44	6133912	10/17/2000	Montero
_____	45	6134532	10/17/2000	Lazarus et al.
_____	46	6157946	12/05/2000	Itakura et al.
_____	47	6160989	12/12/2000	Hendricks et al.
_____	48	6177931	01/23/2001	Alexander et al.
_____	49	6188398	02/13/2001	Collins-Rector et al.
_____	50	6189008	02/13/2001	Easty et al.
_____	51	6205432	03/20/2001	Gabbard et al.
_____	52	6216112	04/10/2001	Fuller et al.
_____	53	6216129	04/10/2001	Eldering
_____	54	6285987	09/04/2001	Roth et al.
_____	55	6289318	09/11/2001	Barber
_____	56	6308202	10/23/2001	Cohn et al.
_____	57	6327574	12/04/2001	Kramer et al.
_____	58	6366298	04/02/2002	Haitsuka et al.
_____	59	6377986	04/23/2002	Philyaw
_____	60	6385592	05/07/2002	Angles et al.
_____	61	6401075	06/04/2002	Mason et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

_____	62	6408278	06/18/2002	Carney et al.
_____	63	6415270	07/02/2002	Rackson et al.
_____	64	6446261	09/03/2002	Rosser
_____	65	6449657	09/10/2002	Stanbach Jr. et al.
_____	66	6463533	10/08/2002	Calamera et al.
_____	67	6463585	10/08/2002	Hendricks et al.
_____	68	6490587	12/03/2002	Easty et al.
_____	69	6536041	03/18/2003	Knudson et al.
_____	70	6633850	10/14/2003	Gabbard et al.
_____	71	6654725	11/25/2003	Langheinrich et al.
_____	72	6684194	01/27/2004	Eldering et al.
_____	73	6718551	04/06/2004	Swix et al.
_____	74	6771290	08/03/2004	Hoyle
_____	75	6810527	10/26/2004	Conrad et al.
_____	76	6837436	01/04/2005	Swartz et al.
_____	77	6845396	01/18/2005	Kanojia et al.
_____	78	6847992	01/25/2005	Haitsuka et al.
_____	79	6868392	03/15/2005	Ogasawara
_____	80	6907566	06/14/2005	McElfresh et al.
_____	81	6985882	01/10/2006	Del Sesto
_____	82	7013290	03/14/2006	Ananian
_____	83	7038637	05/02/2006	Eller et al.
_____	84	7051351	05/23/2006	Goldman et al.
_____	85	7092926	08/15/2006	Cerrato
_____	86	7100183	08/29/2006	Kunkel et al.
_____	87	7146329	12/05/2006	Conkwright et al.
_____	88	7152237	12/19/2006	Flickinger et al.
_____	89	7158943	01/02/2007	Van der Riet
_____	90	7188085	03/06/2007	Pelletier
_____	91	7260823	08/21/2007	Schlack et al.
_____	92	7272853	09/18/2007	Goodman et al.
_____	93	7299195	11/20/2007	Tawakol et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

_____	94	7302419	11/27/2007	Conkwright et al.
_____	95	7139723	01/13/2000	Conkwright et al.
_____	96	7305459	12/04/2007	Klemba et al.
_____	97	7580855	08/25/2009	Fernandes
_____	98	7640214	12/29/2009	Black et al.
_____	99	7756744	07/13/2010	Schiff et al.
_____	100	7805332	09/28/2010	Wilson
_____	101	7806329	10/05/2010	Dmitriev et al.
_____	102	7856372	12/21/2010	Ullah
_____	103	7856373	12/21/2010	Ullah
_____	104	7913275	03/22/2011	Hendricks et al.
_____	105	7930252	04/19/2011	Bender et al.
_____	106	7937383	05/03/2011	Hintze et al.
_____	107	7966647	06/21/2011	Igoe et al.
_____	108	7971260	06/28/2011	Spalink et al.
_____	109	8015184	09/06/2011	Zito et al.
_____	110	8015286	09/06/2011	Jenkins
_____	111	8019881	09/13/2011	Sandhu et al.
_____	112	8024765	09/20/2011	Ramanathan et al.
_____	113	8079048	12/13/2011	Howcroft
_____	114	8086491	12/27/2011	Matz et al.
_____	115	8180674	05/15/2012	Madhavan
_____	116	8700464	04/15/2014	McAllister et al.
_____	117	8725571	05/13/2014	McAllister et al.
_____	118	8762201	06/24/2014	Noonan
_____	119	8880677	11/04/2014	Julia et al.
_____	120	20010034646	10/25/2001	Hoyt et al.
_____	121	20010042002	11/15/2001	Koopersmitch
_____	122	20010049620	12/06/2001	Blasko
_____	123	20020010653	01/24/2002	Lee
_____	124	20020010928	01/24/2002	Sahota
_____	125	20020013943	01/31/2002	Haberman et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

_____	126	20020019769	02/14/2002	Barrits et al.
_____	127	20020032603	03/14/2002	Yeiser
_____	128	20020046105	04/18/2002	Gardenswartz et al.
_____	129	20020049968	04/25/2002	Wilson et al.
_____	130	20020059094	05/16/2002	Hosea et al.
_____	131	20020059602	05/16/2002	Macrae et al.
_____	132	20020078444	06/20/2002	Krewin et al.
_____	133	20020082910	06/27/2002	Kontogouris
_____	134	20020082923	06/27/2002	Merriman et al.
_____	135	20020083043	06/27/2002	Hoshi et al.
_____	136	20020103850	08/01/2002	Moyer et al.
_____	137	20020104086	08/01/2002	Tomsen et al.
_____	138	20020112013	08/15/2002	Walsh
_____	139	20020120933	08/29/2002	Knudson et al.
_____	140	20020120935	08/29/2002	Huber et al.
_____	141	20020123928	09/05/2002	Eldering et al.
_____	142	20020124249	09/05/2002	Shintani
_____	143	20020124253	09/05/2002	Eyer et al.
_____	144	20020129362	09/12/2002	Chang et al.
_____	145	20020129368	09/12/2002	Schlack et al.
_____	146	20020147645	10/02/2002	Alao et al.
_____	147	20020194058	12/19/2002	Eldering
_____	148	20020194604	12/19/2002	Sanchez et al.
_____	149	20030023481	01/30/2003	Calvert et al.
_____	150	20030041156	02/27/2003	Pickover et al.
_____	151	20030050863	03/13/2003	Radwin
_____	152	20030051242	03/13/2003	Donnelly
_____	153	20030066078	04/03/2003	Bjorgan et al.
_____	154	20030093311	05/15/2003	Knowlson
_____	155	20030126597	07/03/2003	Darby et al.
_____	156	20030131355	07/10/2003	Berenson et al.
_____	157	20030145323	07/31/2003	Hendricks et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

_____	158	20030149975	08/07/2003	Eldering et al.
_____	159	20030187726	10/02/2003	Bull et al.
_____	160	20030187949	10/02/2003	Bhatt et al.
_____	161	20030208756	11/06/2003	Macrae et al.
_____	162	20030208758	11/06/2003	Schein et al.
_____	163	20030226141	12/04/2003	Krasnow et al.
_____	164	20030229900	12/11/2003	Reisman
_____	165	20030231207	12/18/2003	Huang
_____	166	20040002938	01/01/2004	Deguchi
_____	167	20040003406	01/01/2004	Billmaier
_____	168	20040010546	01/15/2004	Klug et al.
_____	169	20040031058	02/12/2004	Reisman
_____	170	20040044566	03/04/2004	Bostelmann et al.
_____	171	20040073915	04/15/2004	Dureau
_____	172	20040078809	04/22/2004	Drazin
_____	173	20040107437	06/03/2004	Reichardt et al.
_____	174	20040117827	06/17/2004	Karaoguz et al.
_____	175	20040123314	06/24/2004	Bova
_____	176	20040163101	08/19/2004	Swix et al.
_____	177	20040172650	09/02/2004	Hawkins et al.
_____	178	20050021387	01/27/2005	Gottfurcht
_____	179	20050049957	03/03/2005	Vakili et al.
_____	180	20050076393	04/07/2005	Sussman
_____	181	20050125289	06/06/2005	Beyda et al.
_____	182	20050125290	06/09/2005	Beyda et al.
_____	183	20050132405	06/16/2005	AbiEzzi et al.
_____	184	20050165638	07/28/2005	Piller
_____	185	20050165643	07/28/2005	Wilson et al.
_____	186	20050165644	07/28/2005	Beyda et al.
_____	187	20050177420	08/11/2005	Tanahashi
_____	188	20050187823	08/25/2005	Howes
_____	189	20060018466	01/26/2006	Adelstein et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

_____	190	20060116924	06/01/2006	Angles et al.
_____	191	20060137015	06/22/2006	Fahrny et al.
_____	192	20060184989	08/17/2006	Slothouber
_____	193	20060229930	10/12/2006	Gottfurcht
_____	194	20060253323	11/09/2006	Phan et al.
_____	195	20060259357	11/16/2006	Chiu
_____	196	20060271690	11/30/2006	Banga et al.
_____	197	20060271953	11/30/2006	Jacoby et al.
_____	198	20060271959	11/30/2006	Jacoby et al.
_____	199	20060287916	12/21/2006	Starr et al.
_____	200	20070027901	02/01/2007	Chan et al.
_____	201	20070038516	02/15/2007	Apple et al.
_____	202	20070039020	02/15/2007	Cansler et al.
_____	203	20070073585	03/29/2007	Apple et al.
_____	204	20070078774	04/05/2007	Brown
_____	205	20070078849	04/05/2007	Slothouber
_____	206	20070088603	04/19/2007	Jouppi et al.
_____	207	20070089127	04/19/2007	Flickinger et al.
_____	208	20070100690	05/03/2007	Hopkins
_____	209	20070130015	06/07/2007	Starr et al.
_____	210	20070157228	07/05/2007	Bayer et al.
_____	211	20070198327	08/23/2007	Yazdani et al.
_____	212	20070208619	09/06/2007	Branam et al.
_____	213	20070208789	09/06/2007	Reichman
_____	214	20070214048	09/13/2007	Chan et al.
_____	215	20070220553	09/20/2007	Branam
_____	216	20070233571	10/04/2007	Eldering et al.
_____	217	20070244750	10/18/2007	Grannan et al.
_____	218	20070266403	11/15/2007	Ou et al.
_____	219	20070277220	11/29/2007	Shikuma et al.
_____	220	20070283384	12/06/2007	Haeuser et al.
_____	221	20070294721	12/20/2007	Haeuser et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

_____	222	20080016540	01/17/2008	Savoor et al.
_____	223	20080040742	02/14/2008	Howcroft et al.
_____	224	20080046942	02/21/2008	Merlin
_____	225	20080052392	02/28/2008	Webster et al.
_____	226	20080082413	04/03/2008	Madhavan
_____	227	20080109306	05/08/2008	Maigret et al.
_____	228	20080109376	05/08/2008	Walsh et al.
_____	229	20080109843	05/08/2008	Ullah
_____	230	20080155613	06/26/2008	Benya et al
_____	231	20080181225	07/31/2008	Zampiello
_____	232	20080183573	07/31/2008	Muschetto
_____	233	20080189360	08/07/2008	Kiley et al.
_____	234	20080195664	08/14/2008	Maharajh et al.
_____	235	20080201772	08/21/2008	Mondaeev et al.
_____	236	20080222283	09/11/2008	Ertugrul et al.
_____	237	20080243822	10/02/2008	Campbell et al.
_____	238	20080244076	10/02/2008	Shah et al.
_____	239	20080255944	10/16/2008	Shah et al.
_____	240	20080262901	10/23/2008	Banga et al.
_____	241	20080263633	10/23/2008	Banga et al.
_____	242	20080275785	11/06/2008	Altberg et al.
_____	243	20080307112	12/11/2008	Andreasson et al.
_____	244	20090037949	02/05/2009	Birch
_____	245	20090049469	02/19/2009	Small et al.
_____	246	20090106100	04/23/2009	Mashinsky
_____	247	20090171780	07/02/2009	Aldrey et al.
_____	248	20090204706	08/13/2009	Ertugrul et al.
_____	249	20100076848	03/25/2010	Stefanik et al.
_____	250	20110125592	05/26/2011	Mesaros
_____	251	20110167043	07/07/2011	Hintze et al.
_____	252	20110258049	10/20/2011	Ramer et al.
_____	253	20110307548	12/15/2011	Fisk et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

_____	254	20130254035	09/26/2013	Ramer et al.
_____	255	6832207	12/14/2004	Shkedi
_____	256	6925440	08/02/2005	Shkedi
_____	257	6973436	12/06/2005	Shkedi
_____	258	7072853	07/04/2006	Shkedi
_____	259	7428493	09/23/2008	Shkedi
_____	260	7454364	11/18/2008	Shkedi
_____	261	7747745	06/29/2010	Shkedi
_____	262	7822637	10/26/2010	Shkedi
_____	263	7822639	10/26/2010	Shkedi
_____	264	7861260	12/28/2010	Shkedi
_____	265	7890609	02/15/2011	Shkedi
_____	266	7979307	07/12/2011	Shkedi
_____	267	8051444	11/01/2011	Shkedi
_____	268	8200822	06/12/2012	Shkedi
_____	269	8204783	06/19/2012	Shkedi
_____	270	8204965	06/19/2012	Shkedi
_____	271	8239264	08/07/2012	Shkedi
_____	272	8244574	08/14/2012	Shkedi
_____	273	8244582	08/14/2012	Shkedi
_____	274	8244583	08/14/2012	Shkedi
_____	275	8244586	08/14/2012	Shkedi
_____	276	8280758	10/02/2012	Shkedi
_____	277	8494904	07/23/2013	Shkedi
_____	278	8566164	10/22/2013	Shkedi et al.
_____	279	8589210	11/19/2013	Shkedi
_____	280	8595069	11/26/2013	Shkedi et al.
_____	281	8600815	12/03/2013	Shkedi
_____	282	8607267	12/10/2013	Shkedi
_____	283	8683502	03/25/2014	Shkedi et al.
_____	284	8713600	04/29/2014	Shkedi
_____	285	8775249	07/08/2014	Shkedi

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

_____	286	8959146	02/17/2015	Shkedi
_____	287	8997138	03/31/2015	Shkedi
_____	288	9071886	06/30/2015	Shkedi
_____	289	9078035	07/07/2015	Shkedi et al.
_____	290	9083853	07/14/2015	Shkedi
_____	291	9131282	09/08/2015	Shkedi
_____	292	9208514	12/08/2015	Shkedi
_____	293	9226019	12/29/2015	Shkedi
_____	294	9271024	02/23/2016	Shkedi et al.
_____	295	9351053	05/24/2016	Shkedi
_____	296	9369779	06/14/2016	Shkedi
_____	297	9508089	11/29/2016	Shkedi
_____	298	9571865	02/14/2017	Shkedi
_____	299	9591380	03/07/2017	Shkedi et al.
_____	300	9756372	09/05/2017	Shkedi
_____	301	9800917	10/24/2017	Shkedi
_____	302	9813778	11/07/2017	Shkedi
_____	303	9830615	11/28/2017	Shkedi
_____	304	10009636	06/26/2018	Shkedi
_____	305	10026100	07/17/2018	Shkedi
_____	306	10134054	11/20/2018	Shkedi
_____	307	10178442	01/08/2019	Shkedi et al.
_____	308	10306282	05/28/2019	Shkedi
_____	309	10321198	6/11/2019	Shkedi et al.
_____	310	10405058	09/03/2019	Shkedi
_____	311	10475073	11/12/2019	Shkedi
_____	312	10628857	04/21/2020	Shkedi
_____	313	10645438	05/05/2020	Shkedi
_____	314	10715878	07/14/2020	Shkedi
_____	315	10771860	09/08/2020	Shkedi et al.
_____	316	10839423	11/17/2020	Shkedi
_____	317	20050086112	04/21/2005	Shkedi

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

_____	318	20110035256	02/10/2011	Shkedi et al.
_____	319	20110040607	02/17/2011	Shkedi
_____	320	20090172723	07/02/2009	Shkedi et al.

FOREIGN PATENT DOCUMENTS

Examiner Initials		Document Number	Publication Date MM/DD/YYYY	Name of Patentee	Translation Yes / No
_____	321	WO 2005/046237	05/19/2005	Gutta et al.	No
_____	322	KR-10-0793513	01/03/2008	ICU Research & Industrial Cooperation Group (ETRI)	Yes
_____	323	KR 10-2009-0113456	08/14/2012	SK Planet Co., Ltd.	Yes
_____	324	JP2007-274246	10/18/2007	Casio Computer Co. Ltd.	Yes
_____	325	JP 2006-004441	01/05/2006	Yokogawa Electric Corp.	Yes
_____	326	WO0184440	11/08/2001	PSM Global, Inc.	No
_____	327	WO2004086680	10/07/2004	Thomson Licensing S.A.	Yes
_____	328	JP2003-331183	11/21/2003	NTT Communications KK	Yes
_____	329	WO1997021183A1	06/12/1997	Bell Communications Research Inc.	No
_____	330	WO1999057660A1	11/11/1999	IWEB LTD.	No
_____	331	WO00/004709	01/27/2000	Ellis	No
_____	332	WO 01/054034 A1	07/26/2001	Phillips et al.	No

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	333	Engage Technologies, Inc.; Form S-1A; filed with the SEC; pp. 36-50; July 19, 1999.
_____	334	Chen, Andy; "Behavioral Matchmaking, Part 3: GPS and Behavioral Targeting"; from website http://www.clickz.com; November 24, 2004.
_____	335	Battelle, John; "The Search - How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture"; pp. 167-171; 2005.
_____	336	"BSkyB in Google Link-up"; Financial Times (FT.com); December 6, 2006.
_____	337	"Google moves into TV ads with BSkyB Deal"; Guardian Unlimited; December 7, 2006.
_____	338	Chen, Andy; "The New Behavioral Wunder"; from webpage http://www. clickz.com; April 11, 2007.
_____	339	"NetZero Launches Ad-Backed Access Service"; from website www.clickz. com; October 20, 1998.
_____	340	"Nissan, BellSouth, Music Blvd. Among Latest Advertisers to Sign Up With NetZero"; press release from www.irconnect.com; November 16, 1998.
_____	341	"LookSmart and NetZero Extend and Enhance Winning Partnership"; press release found on www.shareholder.com; February 9, 2000.
_____	342	"NetZero Guarantees Increase in Web Site Traffic for New Adverstisers"; press release from www.irconnect.com; circa 2000.
_____	343	Saunders, Christopher; "NetZero to Launch Mobile Ad Platform"; from website www.clickz.com; March 29, 2001.
_____	344	"180Solutions Launches All-In-One Online Ad Solution - Metricsdirect Outperforms Conventional Cost-Per-Click Models"; from www.zango.com; March 30, 2004.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	345	Keegan, Paul; "The Man Who Can Save Advertising"; from http://money.cnn.com; November 1, 2004.
_____	346	"Visible World Targets Advertising Industry"; from http://informitv.com; February 20, 2005.
_____	347	"Venture Market Summary"; VentureWire Alert; March 29, 2005.
_____	348	Olsen, S.; "TiVo is in talks with Internet search giants Google and Yahoo over possible deal ... television and the Web..."; www.news.com; April 18, 2005.
_____	349	"This Ad's for You -- Just You"; BusinessWeek online; June 28, 2005.
_____	350	"Cable's Big Bet On Hyper-Targeting. Time Warner will test new software that sends different ads to different viewers."; BusinessWire online; July 4,2005.
_____	351	"Comcast, Cox Buy Interactive TV Company"; from Yahoo.com Financial News; July 13, 2005.
_____	352	"TiVo Launches New Interactive Advertising Technology"; press release from www.tivo.com; July 18, 2005.
_____	353	Grant, Peter; "To Ward Off New Competitors, Comcast Builds a Mini Internet"; Wall Street Journal online; October 13, 2005.
_____	354	"TiVo Announces First Advertising Search Product for Television"; press release from www.tivo.com; November 28, 2005.
_____	355	"TiVo Begins Rollout of Online Services Now Accessible Directly on the TV"; press release from www.tivo.com; December 1, 2005.
_____	356	"Comcast, Time Warner Cable hopeful on network DVRs"; from www.reuters.com; March 30, 2006.
_____	357	"TiVo Launches Television's New Advertising Search Product"; press release from www.tivo.com; May 8, 2006.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	358	"TiVo and Brightcove Partner to Deliver Broadband Content to TiVo Subscribers; press release from www.tivo.com; May 10, 2006.
_____	359	"TiVo Debuts New Lexus Car Configuration Application"; press release from www.tivo.com; May 11, 2006.
_____	360	Newcomb, Kevin; "BMW Uses TiVo Technology to Improve Product Placement"; from www.clickz.com; June 21, 2006.
_____	361	Wingfield, Nick; "New TiVo Feature Lets Families Share Videos On TV Via the Web"; The Wall Street Journal Online; November 14, 2006.
_____	362	Kaye, Kate; "TiVo Broadens CBS Offerings, Adds Web-to-TV Services"; from www.clickz.com; November 15, 2006.
_____	363	Grant, Peter and Marr, Merissa; "Comcast Nears Pact With Disney To Offer ABC Shows On Demand"; The Wall Street Journal Online; November 21, 2006.
_____	364	"TiVo to Insert Ads At End of Probrams"; The Wall Street Journal Online; November 28, 2006.
_____	365	FAQ from Zango Ad Services; http://adservices.zango.com; 2007.
_____	366	"NetZero and Juno's sophisticated targeting capabilities enable advertisers to hit an exact audience of qualified consumers"; from www.untd.com; 2007.
_____	367	Shukla, Anuradha; "Visible World Offers 'Any Screen' Solution for Customized Video Advertising"; from www.tmcnet.com; February 6, 2007.
_____	368	Tam, Pui-Wing; "TiVo, Amazon to Deliver Web Video to TV Sets"; The Wall Stree Journal Online; February 7, 2007.
_____	369	Story, Louise; "The Web drives an advertising boom without the need for agencies"; International Herald Tribune (iht.com); February 8, 2007.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	370	Todd, Daniel; "The New World of Desktop Advertising"; DMNews.com; March 28, 2007.
_____	371	"Medio Systems to Launch Mobile Performance Ad Network - Search and recommendation technologies propel advancements in mobile advertising", article from Mobilestreams website, March 19, 2007.
_____	372	HUNT, TIMOTHY J., "Moving Target", Financial Post, October 1, 2000.
_____	373	Luening E.; "Free ISP NetZero beefs up ad services"; web page at http://www.news.com/Free-ISP-NetZero-beefs-up-ad-services/2100-1023_3-241067.html; January 2, 2002.
_____	374	BERGSTEIN B.; "Ad-targeted System Monitors Your Interests with ISP's Help"; from website www.siliconvalley.com; December 10, 2007.
_____	375	WHITE B.; "Watching What You See on the Web"; Wall Street Journal Online; December 6, 2007.
_____	376	"What's New: The Latest On Technology Deals"; from Dow Jones VentureWire; as published on Wall Street Journal Online; October 22, 2007.
_____	377	"DoubleClick Launches New Marketing Solutions"; from website www.clickz.com; October 6, 1998.
_____	378	Poikselktä, Miikka, et al.; "IMS Concepts"; The IMS: IP 1-15 Multimedia Concepts and Services, Second Edition; John Wiley & Sons, Ltd.; 2006; pp. 43-116.
_____	379	Baseview.com website (ClassManagerPro) from www.archive.org, archived January 16, 1997
_____	380	S-1/A SEC Filing, filed by FlyCast Communications Corp. on April 30, 1999
_____	381	24/7 RealMedia, The Science of Digital Marketing, Company Overview (2007)
_____	382	aaddzz.com webpages from www.archive.org, archived January 30, 1998

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	383	AdForce, Inc., User Guide, A Complete Guide to AdForce, Version 2.6 (1998)
_____	384	Fawcett, Adrienne W., Publishers thankful but nervous about boom, Magazinedotcom, Hot spending boosts …, adage.com, November 1, 1999
_____	385	Ad-Star.com website archive from www.archive.org (archived on April 12, 1997 and February 1, 1997)
_____	386	Chang et al., "Goodies" in exchange for Consumer Information on the Internet, The Economics and Issues, Proceedings of the Thirty-First Annual Hawaii International Conference on System Sciences, 1998
_____	387	avenuea.com webpage, FAQ, www.avenuea.com/corpinfo/faq.html, archived October 2, 1999
_____	388	Baek et al., Implementation of Smart Home Control using LabVIEW and PDA," IEEE (2004)
_____	389	Bardram et al., Experiences from Real-World Deployment of Context-Aware Technologies in a Hospital Environment, Center for Pervasive Healthcare (2006)
_____	390	Advertising-A Unified Approach to Adverstising Management, Baseview Products, Inc. brochure, October 14, 1999
_____	391	Ribeiro-Neto et al., Extracting Semi-Structured Data Through Examples, Proceedings of the Eighth international Conference on Information and Knowledge Management (1999)
_____	392	Raskutti and Beitz, Acquiring Customer Preferences for Information Filtering, a Heuristic-Statistical Approach, Fifth Annual Conference on User Modeling (1996)
_____	393	AdKnowledge, Campaign Manager-Quick Reference Card (1996)
_____	394	Emert, 2 Web Ad Firms Decide to Marry, Palo Alto's Focalink, ClickOver Merging, San Francisco Chronicle, November 17, 1997

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	395	Castro et al., A Probabilistic Room Location Service for Wireless Networked Environments, Ubicomp 2001, LNCS 2201, pp. 18-34 (2001)
_____	396	Chiu et al., LiteMinutes, An Internet-Based System for Multimedia Meeting Minutes," WWW10, May 2-5, 2001
_____	397	Cox et al., IntelliBadge, Towards Providing Location-Aware Value-Added Services at Academic Conferences," UbiComp 2003, LNCS 2864, pp. 264-280 (2003)
_____	398	Cuzzocrea et al., Enriching Conceptual Modeling of XML-Aware Adaptive Web Systems with Object-Oriented Constructs and UML Diagrams," EC-WEB 2004, LNCS 3182, pp. 248-258 (2004)
_____	399	Elderbrock & Borwankar, Building Successful Internet Businesses, IDG Worldwide, 1996
_____	400	David W. Schumann & Ester Thorson, Advertising and the World Wide Web, Lawrence Erlbaum Associates, 1999
_____	401	Williamson, Debra Aho, Web ads mark 2nd birthday with decisive issues ahead, adage.com, October 21, 1996
_____	402	Williamson, Debra Aho, Targeting distinguishes AdForce from the pack, adage.com, March 1, 1999
_____	403	Doubleclick.com webpages from www.archived.org, archived February 5, 1998
_____	404	Tittel et al., Foundations of World Wide Web Programming with HTML & CGI, IDG Books Worldwide (1995)
_____	405	Genco et al., An Augmented Campus Design for Context-aware Service Provision, SIGUCCS'05 (November 6-9, 2005
_____	406	Genco, et al. An Agent-Based Service Network for Personal Mobile Devices, IEEE (2006)

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	407	Hess et al., A Context-Aware Data Management System for Ubiquitous Computing Applications, Proceedings of the 23rd Intl. Conf. on Distributed Computing Systems (2003)
_____	408	Graham, Ian S., HTML Sourcebook, A Complete Guide to HTML, John Wiley & Sons (1995)
_____	409	Aaddzz Brochure, Information Access Technologies, Inc. (1997)
_____	410	Iwasaki et al., Azim, Direction Based Service using Azimuth Based Position Estimation, ICDCS (2004)
_____	411	Gilbert, SPECIAL REPORT, Agencies centralize Web ad serving, adage.com, Mar 1, 1999
_____	412	Palmer & Eriksen, Digital Newspapers Explore Marketing on the Internet, Communications of the ACM, Vol. 42, No. 9, September 1999
_____	413	Angwin, Julia, Focalink Will Monitor World of Online Ads, San Francisco Chronicle, February 7, 1996
_____	414	Gallagher & Parsons, A Framework for Targeting Banner Advertising on the Internet, Proceedings of the Thirtieth Hawaii International Conference on System Sciences, January 1997
_____	415	Cleland, Kim, Accipiter latest entrant in race to manage ads; company sets sights on helping content providers target audience, Adage.com, September 9, 1996
_____	416	Cleland, Kim, New Player joins fray of Web ad targeters, adage.com, February 24, 1997
_____	417	Bayne, Kim M., AdKnowledge rolls out Web ad evaluation tool, adage.com, July 8, 1999
_____	418	Langheinrich et al., Unintrusive Customization Techniques for Web Advertising, COMPUTER NETWORKS 31, 1259–1272, May 1999

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	419	Maekawa et al., A Java-Based Information Browsing System in a Remote Display Environment, IEEE (2004)
_____	420	Meeker, Mary, The Internet Advertising Report, Harper Business, 1997
_____	421	McCarthy et al., UniCast, OutCast & GroupCast, Three Steps Toward Ubiquitous, Peripheral Displays, Ubicomp 2001, LNCS 2201, pp. 332-345 (2001)
_____	422	Nakajima et al., Universal Interaction with Networked Home Appliances, ICDCS'02, IEEE (2002)
_____	423	Nakazawa et al., A Bridging Framework for Universal Interoperability in Pervasive Systems, ICDCS'06 IEEE (2006)
_____	424	Yeager & McGrath, Web Server Technology, The Advance Guide for World Wide Web Information Providers, Morgan Kaufmann Publishers (1996)
_____	425	Nawyn et al., Embedding Behavior Modification Strategies into a Consumer Electronic Device, A Case Study, Ubicomp 2006, LNCS 4206, pp. 297-314 (2006)
_____	426	net.Genesis & Devra Hall, Build a Website, The Programmer's Guide to Creating, Building, and Maintaining a Web Presence, Prima Publishing (1995)
_____	427	NetGravity Launches AdServer 3.5 with Beefed Up Targeting, clickz.com, September 10, 1998
_____	428	Novak & Hoffman, New Metrics for New Media, Toward the Development of Web Measurement Standards, Project 2000, elab.vanderbilt.edu, September 26, 1996
_____	429	Ponnekanti et al., ICrafter, A Service Framework for Ubiquitous Computing Environments," Ubicomp 2001, LNCS 2201, pp. 56-75 (2001)
_____	430	RealMedia.com webpages, www.archive.org, archived on May 15, 19977 and June 7, 1997

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	431	Bruner, Rick E., Interactive, NetGravity to launch new ad service bureau, adage.com, December 1, 1997
_____	432	Bruner, Rick E., Small networks chase per-click ad business, Aaddzz, ValueClick compete, AdvertisingAge, September 8, 1997
_____	433	Zeff & Aronson, Advertising on the Internet, 1st Edition, John Wiley & Sons (1997)
_____	434	Zeff & Aronson, Advertising on the Internet, 2nd Edition, John Wiley & Sons (1999)
_____	435	Ruta et al., Ontology driven resource discovery in a Bluetooth based m-marketplace, CEC/EEE'06, IEEE (2006)
_____	436	Smetters et al., Instant Matchmaking, Simple and Secure Integrated Ubiquitous Computing Environments, Ubicomp 2006, LNCS 4206, pp. 477-794 (2006)
_____	437	Sogo et al., LightHolder, A Platform for Real-World Mobile Commerce, CEC/EEE'06, IEEE (2006)
_____	438	Sumi et al., Digital Assistant for Supporting Conference Participants, An Attempt to Combine Mobile, Ubiquitous and Web Computing, Ubicomp 2001, LNCS 2201, pp. 156-175 (2001)
_____	439	Singh & Dalal, Web Home Pages as Advertisements, Communications of the ACM, August 1999
_____	440	The ClickZ Network, Netscape Selects Real Media's Open AdStream, January 13, 1998
_____	441	The ClickZ Network, Real Media Signs More Media Companies, December 30, 1998
_____	442	Trevor et al., Issues in Personalizing Shared Ubiquitous Devices, UbiComp 2002, LNCS 2498, pp. 56-72 (2002)

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	443	www.Spinbox.com webpages, www.archive.org, archived beginning Nov. 11, 1998
_____	444	Bakos, The Emerging Role of Electronic Marketplaces on the Internet, Communications of the ACM, Vol. 41, No. 8, August 1998
_____	445	Youiji Kohda & Susumu Endo, Ubiquitous Advertising on the WWW, Merging Advertisement on the Browser, Fifth International World Wide Web Conference (May 1996)
_____	446	Zhang et al., Smart Alarm Clock – A Networked Home Appliance with Bluetooth Connection, IEEE (2004)
_____	447	Baseview.com webpages from www.archive.org, archived January 16, 1997
_____	448	AdForce, Inc., Quick Reference Guide, Version 3.0 (1999)
_____	449	AdManagerPro; Ad Entry Manual; User's Manual V.1.05; 1997
_____	450	AdManagerPro; Ad Entry Manual; User's Manual V.2.0; 1997
_____	451	AdManagerPro; Administration and Receivables Manual; Administrator's Manual Version 1.0.5; November 1997
_____	452	AdManagerPro; Administration and Receivables Manual; Administrator's Manual Version 2.0; June 1998
_____	453	Commercialization of the World Wide Web, The Role of Cookies, vanderbilt.edu web pages, archive.org, February 1, 1997
_____	454	Crabtree et al., Adaptive Personal Agents, Personal Technologies (1998)
_____	455	Dalton, Gregory, OPS, Answer to Cookies?, InformationWeek, October 13, 1997
_____	456	Dunn et al.; Privacy and Profiling on the Web; Submitted to W3C on June 2, 1997

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	457	firefly.com web pages, archive.org, archived 1998-2000
_____	458	Gilbert, Jennifer, CMGI's Web dealmaking hits big-time, Advertising Age, October 4, 1999
_____	459	Hersley et al., Implementation of OPS Over HTTP, Submitted to W3C on June 2, 1997
_____	460	Hersley et al., Proposal for an Open Profiling Standard, Submitted to W3C on June 2, 1997
_____	461	Hersley et al., Standard Practices of OPS Systems, Submitted to W3C on June 2, 1997
_____	462	junkbusters.com web pages; archive.org; archived April 1997 through November 2000
_____	463	Karpinski, Richard; ICE application to revolutionize e-commerce; Business Marketing, March 1998
_____	464	matchlogic.com web pages, archive.org, April 13, 2000
_____	465	Roscheisen, R. Martin; A Network-Centric Design for Relationship-Based Rights Management (Dissertation); December 1997
_____	466	BARRY M. LEINER ET AL.; Brief History of the Internet; Internet Society; October 15, 2012
_____	467	KRISTOL, D. & MONTULLI, L.; HTTP State Management Mechanism; Internet Engineering Task Force; available at https://tools.ietf.org/html/rfc2109; February 1997
_____	468	TUNGATE, MARK; Adland: A Global History of Advertising, 2d Ed. (2013)

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	469	JOHNSON, BRADLEY; Interactive: Planet Oasis Gives Web Sites Promotion Clout; Advertising Age; available at http://adage.com/article/news/interactive-planet-oasis-web-sites-promotion-clout/78467; July 8, 1996
_____	470	JONES, DUNCAN; The History of Pay Per Click (PPC) Marketing; available at https://www.duncanjonesnz.com/history-of-ppc-marketing; September 23, 2014
_____	471	D.A.R.P.A., Internet Protocol; DARPA Internet Program Protocol Specification; Internet Engineering Task Force; available at https://tools.ietf.org/html/rfc791; September 1981
_____	472	IEEE Standard 802.3; IEEE Standard for Information technology-Telecommunications and information exchange between systems-Local and metropolitan area networks-Specific requirements; Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications; 2002 Edition
_____	473	IEEE Standard 802.11; IEEE Standard for Information technology-Telecommunications and information exchange between systems-Local and metropolitan area networks-Specific requirements; Part 11: Wireless Medium Access Control (MAC) and Physical Layer Specifications; 1999 Edition
_____	474	LO, Shih-Hsiang et al.; Controlling Digital TV Set-Top Box with Mobile Devices via an IP Network; National Taiwan University; pp. 935-942; May 20, 2006
_____	475	FORD, Bryan et al.; Persistent Personal Names for Globally Connected Mobile Devices; OSDI '06 Paper; pp. 233-248; Nov. 2006
_____	476	FALLAHKHAIR, S. et al.; Development of a cross-platform ubiquitous language learning service via mobile phone and interactive television; Journal of Computer Assisted Learning, 23, pp. 312-325; 2007
_____	477	LYU, Jihye et al.; Design of Open APIs for Personalized IPTV Service; ICACT 2007; February 12-14, 2007

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	478	GEIER, Eric; The Business Side of Hosting a Hotspot; http://www.ciscopress.com/articles/printerfriendly/680821; January 12, 2007
_____	479	MOON, Aekyung et al.; Context-Aware Active Services in Ubiquitous Computing Environments; ETRI Journal, Volume 29, Number 2; April 2007
_____	480	The Gresham Hotel, http://ambassadordublin.com/the-gresham-hotel; February 24, 2007
_____	481	Time Warner Cable, Using BIAP Systems' Interactive Technology, Expands Availability of 'Fantasy Football Tracker'; http://www.timewarner.com/newsroom/press-releases/2006/09/19/time-warner-cable-using-biap-systems-interactive-technology; September 19, 2006
_____	482	YANG, Herong; 'SendMail' Mail Transfer Agent Developed by Eric Allman, http://www.herongyang.com/Computer-History/Sendmail-Mail-Transfer-Agent-by-Eric-Allman-1983.html; 1983
_____	483	DEDRICK, Rick; Interactive Electronic Advertising; IEEE; 1994
_____	484	DEDRICK, Rick; A Consumption Model for Targeted Electronic Advertising; IEEE Multimedia; 1995
_____	485	ROBERTSON, Scott et al.; Dual Device User Interface Design: PDAs and Interactive Television; Applied Research, US WEST Advanced Technologies; 1996
_____	486	GORDON, Mary Ellen et al.; Consumer attitudes towards Internet; International Marketing Review, Vol. 14, Issue 5; pp. 362-375; 1997
_____	487	BEIGL, Michael et al.; The UbicompBrowser; Telecooperation Office, University of Karlsruhe; 1998
_____	488	RISDEN, Kirsten et al; Interactive Advertising: Patterns of Use and Effectiveness; CHI '98, 18-23; pp. 219-224; April 1998

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	489	DREZE, Xavier and ZUFRYDEN, Fred; Is Internet Advertising ready For Prime Time?; April 5, 1998
_____	490	McCANDLESS, Michael; Web Advertising; IEEE Intelligent Systems; May/June 1998
_____	491	AGGARWAL, Charu C. et al.; A Framework for the Optimizing of WWW Advertising; June 3, 1998
_____	492	CHAN, M.C. and WOO, T.Y.; Next-generation wireless data services: Architecture and experience; Personal Commuincations, IEEE, 6(1); 1999
_____	493	SAMAR, Vipin; Single Sign-on Using Cookies for Web Applications; IEEE 8th International Workshops on Enabling Technologies: Infrastucture for Collaborative Enterprises (WET ICE '99); June 16-18, 1999
_____	494	SHEN, Andrew; Online Profiling Project – Comment, P994809; November 8, 1999
_____	495	SIEW, David and YI, Xun; Agent-mediated Internet Advertising; IEEE; 2000
_____	496	PHAM, Thai-Lai et al; A Situated Computing Framework for Mobile and Ubiquitous Multimedia Access using Small Screen and Composite Devices; January 2000
_____	497	TOMLIN, John A.; An Entropy Approach To Unintrusive Targeted Advertising On The Web; IBM Research Divison; June 2000
_____	498	SMYTH, Barry and COTTER, Paul; A Personalized Television Listings Service; Communications of the ACM; Vol. 43, No. 8, pp. 107-111; August 2000
_____	499	GOLDEN, G. Richard III; Service Advertisement and Discovery: Enabling Universal Device Cooperation; IEEE Internet Computing; September-October 2000
_____	500	BOLCER, Gregory Alan et al.; Peer-to-Peer Architectures and the Magi™ Open-Source Infrastructure; Endeavors Technology, Inc.; December 6, 2000

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	501	LEKAKOS, George et al.; Information Systems In The Living Room: A Case Study Of Personalized Interactive TV Design; Global Co-Operation in the New Millenium; The 9th European Conference on Information Systems; Bled, Slovenia, June 27-29, 2001
_____	502	JUELS, Ari; Targeted Advertising… And Privacy Too; RSA Laboratories; 2001
_____	503	LEKAKOS, Georgios et al.; An Integrated Approach to Interactive and Personalized TV Advertising; eL TRUN- The eBusiness Center, Department of Management Science & Technology, Athens University of Economics and Business; 2001
_____	504	PRAMATARIS, K. et al.; Personalized Interactive TV Advertising: The iMEDIA Business Model; Electronic Markets; Vol. 11 (1): 1-9; 2001
_____	505	BOZIOS, Theodoros et al.; Advanced Techniques for Personalized Advertising in a Digital TV Environment: The iMEDIA System; INTRACOM S.A., R&D Division, Development Programmes Department; 2001
_____	506	KOUTSORODI, Artemis et al.; A Management Architecture For Enabling Wireless System Operation In The B3g Context And Exploitation Of The Abc Concept; National Technical University of Athens, School of Electrical and Computer Engineering; 2001
_____	507	NGAI, E.W.T.; Selection of web sites for online advertising using the AHP; Information & Management 40; pp. 233-242; 2001
_____	508	BERNSTEIN, D.J.; E-mail re: Internet host SMTP server survey; http://cr.yp.to/surveys/smtpsoftware.txt; November 29, 1996
_____	509	GERSHMAN, Anatole; Ubiquitous Commerce – Always On, Always Aware, Always Pro-active; The 2002 International Symposium on Applications and the Internet (SAINT 2002); January 30, 2002

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	510	WANG, Helen J. and KATZ, Randy H.; Mobility Support in Unified Communication Networks; Computer Science Department, University of California, Berkeley; January 31, 2001
_____	511	NG, Vincent and HO, Mok Kwan; An Intelligent Agent For Web Advertisements; International Journal of Foundations of Computer Science; Vol. 13, No. 4; pp. 531-554; 2002
_____	512	CHORIANOPOULOS, Kostas et al.; Metrics For Advertisement Effectiveness Measurement In The Interactive TV Environment: The iMEDIA Case; 14th Bled Electronic Commerce Conference; June 25-26, 2001
_____	513	MISHRA, Pinki; Coming to a TV Set Near You; Catalog Age, 18, 8; ProQuest Central, pg. 61; July 2001
_____	514	BONCELLA, Robert; Internet Privacy – At Home and At Work; Communications of the Association for Information Systems; Volume 7, Article 14; September 2001
_____	515	KRAVETS; Robin et al.; A Cooperative Approach to User Mobility; Department of Computer Science, University of Illinois; October 2001
_____	516	ADYA, Atul et al.; Analyzing the Browse Patterns of Mobile Clients; Microsoft Research; November 1, 2001
_____	517	HASHMAN, Steve and KNUDSEN, Steven; The Application of Jini™ Technology to Enhance the Delivery of Mobile Services; PsiNaptic; December 2001
_____	518	Equal Access Networks for real broadband access; Cisco Systems white paper; 2002
_____	519	HOFER, Thomas et al.; Context-Awareness on Mobile Devices - the Hydrogen Approach; Proceedings of the 36th Hawaii International Conference on System Sciences; 2002

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	520	CARTER, Casey and KRAVETS, Robin; User Devices Cooperating to Support Resource Aggregation; Department of Computer Science, University of Illinois; 2002
_____	521	LEKAKOS, Georgios and GIAGLIS, George M.; Delivering Personalized Advertisements in Digital Television: A Methodology and Empirical Evaluation; January 2002
_____	522	PRASAD, Ashutosh et al.; Advertising versus pay-per-view in electronic media; International Journal of Research in Marketing 20, 13-30; 2003
_____	523	PODNAR, Ivana et al.; Mobile Push: Delivering Content to Mobile Users; July 2, 2002
_____	524	Coactive TV: CoTVTM Enables Powerful Interactive TV Now; press release at http://www.teleshuttle.com/cotv/CoTVPR902.htm; September 19, 2002
_____	525	Coactive TV, A new way to enable and profit from Interactive TV now; archived at https://web.archive.org/web/20020920101402/http://www.teleshuttle.com/cotv/default.htm; September 20, 2002
_____	526	REISMAN, Richard R.; Coactive Television Interactive TV, the Web, and 'Co-vergence': A User and Content-Centric Re-visioning; white paper at: http://www.teleshuttle.com/cotv/CoTVWP2.htm; September 2002
_____	527	REKIMOTO; Jun et al.; Proximal Interactions: A Direct Manipulation Technique for Wireless Networking; Interaction Laboratory, Sony Computer Sciences Laboratories, Inc.; 2003
_____	528	BAE, Sung Min et al.; Fuzzy Web Ad Selector Based on Web Usage Mining; IEEE Intelligent Systems; 2003
_____	529	RATSIMOR, Olga et al.; eNcentive: A Framework for Intelligent Marketing in Mobile Peer-To-Peer Environments; Department of Computer Science and Electrical Engineering, University of Maryland; 2003

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	530	LEE, Jeongkeun et al.; Hybrid gateway advertisement scheme for connecting mobile ad hoc networks to the Internet; 2003
_____	531	YUAN, Soe-Tsyr and TSAO, Y.W.; A recommendation mechanism for contextualized mobile advertising; Expert Systems with Applications 24, pp. 399-414; 2003
_____	532	BILCHEV, G. and MARSTON, D.; Personalised Advertising – exploiting the distributed user profile; BT Technology Journal, Vol. 21, No. 1; January 2003
_____	533	HANSEN, Lars M.; Blocking spam with Sendmail; http://www.hansenonline.net/networking/sendmail-spam.html; February 2003
_____	534	VANEM, Erik et al.; Managing Heterogeneous Services And Devices With The Device Unifying Service; Telenor R&D; March 24, 2003
_____	535	ITO, Masaki et al.; Smart Furniture: Improvising Ubiquitous Hot- spot Environment; Graduate School of Media and Governance, Keio University; May 19, 2003
_____	536	MAHLER, Alwin and STEINFIELD, Charles; The Evolving Hot Spot Market for Broadband Access; October 2003
_____	537	SCHULZRINNE, Henning et al.; Ubiquitous Computing in Home Networks; IEEE Communications Magazine; November 2003
_____	538	LEE, Choonhwa et al.; Konark: A system and protocols for device independent, peer-to-peer discovery and delivery of mobile services; IEEE Xplore Document; November 2003
_____	539	SPANGLER, William E. et al.; Using Data Mining To Profile Tv Viewers; Communications of the ACH; Vol. 46, No. 12; pp. 67-72; December 2003
_____	540	PRATA, Alcina et al.; iTV Enhanced System for Generating Multi- Device Personalized Online Learning Environments; 2004

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	541	THAWANI, Amit et al.; Viewing Characteristics based Personalized Ad Streaming in an Interactive TV Environment; Applied Research Group, Satyam Computer Services, Ltd.; 2004
_____	542	HRISTOVA, Nataliya and O'HARE, G.M.P.; Ad-me: Wireless Advertising Adapted to the User Location, Device and Emotions; Proceedings of the 37th Hawaii International Conference on System Sciences; 2004
_____	543	BRAUN, Elmar et al.; Accessing Web Applications With Multiple Context-Aware Devices; Department of Computer Science, Darmstadt University of Technology; 2004
_____	544	THAWANI, Amit et al.; Context Aware Personalized Ad Insertion in an Interactive TV Environment; Applied Research Group, Satyam Computer Services Ltd.; January 2004
_____	545	BAKER, Loren; Behavioral Targeting and Contextual Advertizing; Search Engine Journal; September 1, 2004
_____	546	PHELPS, Joseph E. et al.; Viral Marketing or Electronic Word-of-Mouth Advertising: Examining Consumer Responses and Motivations to Pass Along Email; Journal of Advertising Research; Vol. 44, No. 4; December 2004
_____	547	FALLAHKHAIR, Sanaz; Dual Device User Interface Design for Ubiquitous Language Learning: Mobile Phone and Interactive Television (iTV); Proceedings of the 2005 IEEE International Workshop on Wireless and Mobile Technologies in Education; 2005
_____	548	BELLMAN, Steven et al.; Duane; Interactive Television Advertising: A Research Agenda; ANZMAC 2005 Conference: Advertising/Marketing Communications Issues; 2005
_____	549	YU, Zhiwen et al.; User Preference Learning for Multimedia Personalization in Pervasive Computing Environment; KES 2005, LNAI 3682; pp. 236-242; 2005

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	550	KIM, Munjo et al.; Target Advertisement Service Using TV Viewers' Profile Inference; ; PCM 2005, Part T, LNCS 3767, pp. 202-211; 2005
_____	551	HARRIS, Amy and IRELAND, Greg; Enabling IPTV: What Carriers Need to Know to Succeed; white paper from IDC Analyze the Future; May 2005
_____	552	ATIS IPTV Exploratory Group Report and Recommendation to the TOPS Council; Alliance for Telecommunications Industry Solutions; July 2005
_____	553	HARTE, Lawrence; 10 Things You Need to Know About IP Television; iptvmagazine.com; July 2005
_____	554	RIBEIRO-NETO, Berthier et al.; Impedance Coupling in Content-targeted Advertising; SIGIR '05; August 15-19, 2005
_____	555	OpenTV® IPTV Solutions; white paper; OpenTV, Inc.; September 2005
_____	556	WAGNER, Ben; Driving IPTV Growth: The Challenges and Perspectives; Texas Instruments; September 2005
_____	557	JENSEN, Jens F.; Interactive Television: New Genres, New Format, New Content; Proceedings of the Second Australasian Conference on Interactive Entertainment; pp. 89-96; November 23, 2005
_____	558	EPIC – Privacy and Human Rights Report 2006; http://www.worldlii.org/int/journals/EPICPrivHR/2006/PHR2006-Behavior.html; 2006
_____	559	DRAKE, Eddie C.; Next-Generation Architecture for IPTV Middleware; Myrio Corporation; 2006
_____	560	MATE; Sujeet et al.; Movable-Multimedia: Session Mobility in Ubiquitous Computing Ecosystem; MUM '06; December 4-6, 2006
_____	561	GAL-OR, Esther et al.; Targeted Advertising Strategies on Television; Management Science 52(5); pp. 713-725; 2006

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	562	YOON, Jeong Hyun et al.; TV-Anytime based personalized bi-directional metadata service system; 2006
_____	563	KIM, Jong Woo and DU, Stephen; Design for an Interactive Television Advertising System; Proceedings of the 39th Hawaii International Conference on System Sciences; 2006
_____	564	DE CASTRO, Josephine E. and SHIMAKAWA, Hiromitsu; Mobile Advertisement System Utilizing User's Contextual Information; Proceedings of the 7th International Conference on Mobile Data Management; 2006
_____	565	CAGENIUS, Torbjörn et al.; Evolving the TV experience: Anytime, anywhere, any device; Ericsson Review No. 3; pp. 107-111; 2006
_____	566	TASKIRAN, Cuneyt et al.; Seamless Content Consumption Across Media Delivery Channels and Devices Using Media Bundles; SCTE Conference on Emerging Technologies; January 23-25, 2007
_____	567	MERISAVO, Marko and VESANEN, Jari; The effectiveness of targeted mobile advertising in selling mobile services: an empirical study; Int. J. Mobile Communications, Vol. 4, No. 2, pp. 119-127; 2006
_____	568	FERNÁNDEZ, Yolanda Blanco et al.; AVATAR: An Improved Solution for Personalized TV based on Semantic Inference; Department of Telematics Engineering, University of Vigo; January 11, 2006
_____	569	Front Porch WiFi Messaging and Advertising Solution; Front Porch; February 24, 2006
_____	570	YANG, Wan-Shiou et al.; Mining Social Networks for Targeted Advertising; Proceedings of the 39th Hawaii International Conference on System Sciences; February 24, 2006
_____	571	HARWOOD, Susie; The Revolution Masterclass on cross-media planning; Revolution; London; pp. 62-64; May 2006

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	572	SPANGLER, William E. et al.; Exploring The Privacy Implications Of Addressable Advertising And Viewer Profiling; Communications of the ACM, Vol. 49, No. 5; May 2006
_____	573	Yahoo's Invalidity Contentions (pages related to parent patents in this family); AlmondNet, Inc. et al. v. Yahoo! Inc., Civil Action No. 16-cv-1557-ILG-SMG, U.S. District Court for the Eastern District of New York; June 30, 2017
_____	574	Exhibits E-1 through E-14 (as revised) of Yahoo's Invalidity Contentions cited above; June 30, 2017 and August 22, 2017

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Electronic Acknowledgement Receipt

EFS ID:	47120810
Application Number:	17714965
International Application Number:	
Confirmation Number:	7004
Title of Invention:	Targeted television advertisements based on online behavior
First Named Inventor/Applicant Name:	Roy Shkedi
Customer Number:	26362
Filer:	Louis J. Hoffman/Michelle Lisa
Filer Authorized By:	Louis J. Hoffman
Attorney Docket Number:	14-TV1-CON7
Receipt Date:	30-NOV-2022
Filing Date:	06-APR-2022
Time Stamp:	20:14:01
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Applicant summary of interview with examiner	22-11-30-IIQ-TV1_CON7-Interview_Report.pdf	522274 8a2009c888b2f9228dd80c039e3b96790ac9535f	no	3

Warnings:

Information:					
2	Information Disclosure Statement (IDS) Form (SB08)	22-11-30-IIQ-TV1_CON7-FORM_1449-REPLACEMENT.pdf	662830 f68100c5fb9138f27443505a617238d9d57193e4	no	33
Warnings:					
Information:					
This is not an USPTO supplied IDS fillable form					
			Total Files Size (in bytes):	1185104	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor	:	Shkedi, Roy	Art Unit	:	2425
Serial No.	:	17/714,965	Examiner	:	Ekpo, Nnenna
Filing Date	:	04/06/2022	Conf. No.	:	7004
Title	:	Targeted television advertisements based on online behavior			

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Filed via EFS-Web on November 30, 2022

INTERVIEW REPORT

Dear Sir:

Beginning with a voicemail left by the examiner on October 18, 2022, the undersigned and the examiner spoke by phone on November 2 and 7, 2022, regarding co-pending Application No. 17/877,870. The examiner told the undersigned that the Publications Branch had identified a formatting problem in the forms PTO-1449 (modified) filed on August 5, 2022. Upon investigation, the undersigned determined, and brought to the examiner's attention, that the problem was greater than mere formatting: On account of a clerical error, applicant's representative had accidentally filed the wrong list entirely in both co-pending Application No. 17/877,870 and in this application. Specifically, in this application, the PTO-1449 (modified) filed on April 8, 2022, had been intended to list all references that the examiner had reviewed in the parent applications, but the list actually filed related to references cited in parents to a *co-pending* application of which applicant has an interest rather than the references cited in the parents to *this* application.

The undersigned therefore asked the examiner for permission to replace the 21-page form PTO-1449 (modified) filed on April 8, 2022, with the correct version, as well as to fix the formatting.

Applicant requests the examiner to confirm her review of the 33-page form PTO-1449 (modified) being filed concurrently so that the proper reference list appears on the face of this patent. Per 37 C.F.R. § 1.97(b)(3), an IDS "shall be considered by the Office if filed … [b]efore the mailing of a first Office action on the merits."

Applicant also repeats the request that the examiner disregard the 21-page form PTO-1449 (modified) filed on April 8, 2022, so that those patents are not printed unnecessarily on the face of this patent. However, the examiner should not disregard the cover sheet submitted on April 8, 2022, and this submission does not supersede those citations or those on any of the supplemental IDS filings filed after April 8, 2022, all of which information should be considered in the RCE per MPEP 609.02(II)(A)(3).

The undersigned thanked the examiner for the careful examination and consideration of this application, including her review of the parent applications and her review of the copending proceedings identified in the various IDS filings. Applicant and the undersigned also thank the examiner for her prompt action.

Please feel free to telephone the undersigned if it would in any way advance completion of this application.

Respectfully submitted,

INTENT IQ, LLC

by its attorney

Dated: November 30, 2022

/Louis J. Hoffman/
Louis J. Hoffman
Reg. No. 38,918

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, Arizona 85260
(480) 948-3295

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	17714965
Filing Date	2022-04-06
First Named Inventor	Shkedi, Roy
Art Unit	2425
Examiner Name	Ekpo, Nnenna
Attorney Docket Number	TV1-CON7

U.S.PATENTS

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button.

U.S.PATENT APPLICATION PUBLICATIONS

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20070061838		2007-03-15	Grubbs et al	
	2	20050108091		2005-05-19	Sotak et al	

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1	WO 2007/117613			2007-10-18	Ferguson		☐

If you wish to add additional Foreign Patent Document citation information please click the Add button

NON-PATENT LITERATURE DOCUMENTS

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (**Not for submission under 37 CFR 1.99)**		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Ekpo, Nnenna
	Attorney Docket Number	TV1-CON7

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Ekpo, Nnenna
	Attorney Docket Number	TV1-CON7

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Louis J. Hoffman/	Date (YYYY-MM-DD)	2022-11-30
Name/Print	Louis J. Hoffman	Registration Number	38918

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
18 October 2007 (18.10.2007)

PCT

(10) International Publication Number
WO 2007/117613 A2

(51) International Patent Classification:
H04N 7/10 (2006.01) *A63F 1/00* (2006.01)
H04N 7/025 (2006.01)

(21) International Application Number:
PCT/US2007/008601

(22) International Filing Date: 6 April 2007 (06.04.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/789,590 6 April 2006 (06.04.2006) US

(71) Applicant and
(72) Inventor: **FERGUSON, Kenneth, H.** [US/US]; 10618 Beach Mill Road, Great Falls, VA 22066 (US).

(74) Agent: **BLACKMON, Robert, N.**; Merek, Blackmon & Voorhees, LLC, 673 S. Washington St., Alexandria, VA 22314 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declaration under Rule 4.17:
— *of inventorship (Rule 4.17(iv))*

Published:
— *without international search report and to be republished upon receipt of that report*
— *with information concerning one or more priority claims considered void*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: MEDIA CONTENT PROGRAMMING CONTROL METHOD AND APPARATUS

(57) Abstract: A method and system for the control, aggregation, and management of television programming and Internet content (both traditional and video sources) and more specifically to the customization of media choices and content based in part on the dynamic editing of content according to user preferences. Affiliate Groups can be used to mark, recommend or provide selective editing of video and other media that can be received by users. Users may join Affiliate Groups and selectively view or have content automatically filtered from the data received in their premises such that the display of the media is less than all of the media received at the premises based on Affiliate Group recommendations. Menus and displays may be created that show preferred media content, additional content generated by the Affiliate Groups. Channel schedules may be pre-filtered or provide customized warnings about objectionable material.

WO 2007/117613 A2

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Media Content Programming Control Method and Apparatus

CROSS REFERENCE TO RELATED APPLICATIONS

[001] This application claims the benefit of U.S. Provisional Application 60/789,590, filed 4/6/2006, entitled Media Content Programming Control Method and Apparatus, which is incorporated herein by reference.

BACKGROUND OF THE INVENTION

1. FIELD OF THE INVENTION

[002]The present invention relates to the control, aggregation, and management of television programming and Internet content (both traditional and video sources) and more specifically to the customization of media choices and content based in part on the dynamic editing of content according to user preferences. The functionality derives from an innovative meta-data generation and delivery system providing content management at sub-program granularity.

2. DESCRIPTION OF THE PRIOR ART

[003]With the profusion of television entertainment programming sources now widely available through broadcast, cable, satellite, and Internet distribution systems, television viewers have an overwhelming variety of entertainment options. As a result of the tremendous variety, viewers find that many of the entertainment options are impossible to find or to navigate without editorial assistance. For example, one viewer may only have an interest in current event programming such as news or talk shows or sports and has no desire to watch children's shows or music videos. Families have the additional requirement of determining and displaying only age-appropriate content depending on time of day, family members present, and other

personal and usage patterns and requirements. It would be advantageous for that viewer to be able to have a service that automatically locates, stores and recommends desired programming while at the same time hiding, blocking, filtering or screening unwanted programming.

[004] Some methods of channel selection elimination are known. For example, satellite television uses setup guides and channel categorization as methods for eliminating the display of channel selections by a crude process of eliminating the appearance of channels from the on-screen program guide. A drawback to such a system is that some channels have a variety of programming that may include some desirable and some undesirable programming. Eliminating the entire channel from the onscreen guide also eliminates the opportunity to choose desirable programming from that channel when it is offered.

[005] Thus it can be seen that there is a need for a service and apparatus that can customize program selection options to present to the viewer all desired available programming options at any particular time from all sources while eliminating unwanted or inappropriate channels, programs, and content. This need will increase dramatically as Internet deliver of video content converges with traditional delivery mechanisms (broadcast, cable, and satellite). Without content regulation (broadcasters are regulated by the Federal Communications Commission) and technical and/or economic regulation (broadcasters, cable TV companies, and satellite companies are all regulated at the federal, state, or local

level either technically or economically as public utilities) and without the need for economic concentration and resulting large barriers to market entry, there will be a massive increase in the amount of available video programming, and a huge increase in diversity and range of quality of available video content.

[006] Another important video management feature is the modification or elimination of undesirable content that occurs within otherwise desirable programs. Many viewers find certain portions of otherwise entertaining television programs objectionable. For example, elimination of foul language, sexual content, nudity, and violence is necessary to protect children from exposure to potentially harmful content. Moreover, many adults find such scenes unnecessary and detracting from their enjoyment of the program. Sometimes the undesirable content is not limited to the editorial portion of program itself, but can include the commercial advertisements previously inserted and displayed within the programs as well. Prior devices have attempted to accomplish editing using data such as meta-tags to provide a personalized edition of the media program and identify objectionable scenes, but such a system only works when the incoming media stream is encoded by the original content provider. Further, the metadata itself may be compromised by the lack of independence in its creation (such as the MPAA television rating system, which is generated by the producers of the programs), rendering any possible

editing system based on such metadata untrustworthy to viewers and therefore of little or no use.

[007]A more effective approach is to separate entirely the creation of editorial metadata from the technical infrastructure and tools needed to deliver and utilize the judgments embodied in the metadata. This allows users of a single technology infrastructure to rely on the editorial judgment of one or more independent affinity or community groups such as parents' organizations, church groups, community and social groups, business organizations, sports clubs, and so forth. The proposed invention provides a robust way of differentiating and assigning and later re-associating user choices with regard to independently-authored editorial content. This involves using a combination of industry-specific identifiers for video/film content (such as those specified in SMTPE standards), URIs/URLs for Internet-based video, and UUIDs and/or URIs (universally unique identifiers and/or uniform resource identifiers) within the system to compactly and uniquely represent editorial content creators and metadata sets and map those reliably to user choices.

[008]Under the proposed invention independently created metadata can use a time index or bookmark or digital "fingerprint" to indicate a point in time from the beginning of the media program or a sub-part of the program. Alternatively, the metadata may reference a frame index, offset, chapter reference, scene reference, or other positional indicator within the media program, including any

media sub-stream. The metadata can arrive at the video display device via a completely independent communications channel (such as over a TCP/IP network protocol connection) or a physically-embodied digital medium such as a CD or DVD disc for use after the digital stream embodying the video/audio content has arrived separately at the customer premise equipment (CPE). Or the metadata can be independently embedded upstream of the CPE in a non-intrusively and non-destructive manner within a digital video stream by means of the extensible sub-stream/sub-channel model included in all modern video transmission formats such as MPEG-2 (Motion Picture Expert Group standard 2) transport streams and program streams. In either case, the metadata must interoperate correctly with traditional "open" media streams as well as encrypted and protected media content distributed using "digital rights management" (DRM) tools and supporting infrastructure.

[009] Beyond creation of editorial metadata and its delivery to the CPE either outside or inside the video stream to use with pre-recorded content, another highly useful application of the technology would be real-time or semi-real-time filtering and blocking of video content. In such a system, media editing would be conducted by humans monitoring the video feed at, for example, a central control station whose editorial decisions would result in a control signal or editing command signal being sent by a network to activate or deactivate video display/recording CPE located at the viewer's premises. The

human monitors view the same programs and transmit the control signals to all of the subscribers' homes simultaneously. If the video stream is being played with a slight delay, the command signals would reach the CPE prior to the time the undesired content was actually displayed to users. Various strategies are proposed to create and maintain (and, if content is skipped altogether, rebuild) a buffer of streaming semi-real-time content at the CPE. Over time, more powerful computational and "artificial intelligence" approaches can be used to automate or semi-automate (automate with human oversight) the real-time editorial decision-making process.

[010]To obtain such service, the viewer would need to subscribe and permit the service provider to place a control device (whether embodied in hardware or software) on the viewer's video display/recording equipment which would operate in conjunction with the control signal received from the central control station. In this system, the control signal is applied during the broadcast of the station as the program is being transmitted, either live or preferably with a slight delay.

[011]A system capable of combining pre-determined user preferences with editorial metadata in order to locate, aggregate, highlight, and filter or block video content could also be greatly enhanced if the actual behavior of the viewer were fed back into the system via a system-monitored usage-based feedback loop. The proposed invention does just that, maintaining (with the user's permission and with the necessary

security and privacy controls in place) a complete history of user media choices and behavior (channel selection, length of time viewing, and another automatically acquired information, as well as optional ratings and rankings of particular viewing choices). This user-generated metadata is fed back into the system, which applies artificial intelligence rules to infer user viewing preferences, allowing the system automatically to provide ever more targeted and customized aggregations, recommendations, and filtering behavior as the system is used. This user data can also be aggregated and jointly analyzed on the basis of affinity groups (AGs) with which the user self-identifies in order to further increase the power of the system. The user can also be notified (based on common elements in their usage of media) of previously unknown social groups with which they may have an affinity.

[012] The system also provides an ideal platform for custom advertising. The system's intimate knowledge of the user's interest and activities—known by information collected directly from the user, from affinity group information, and from data gathered directly from the usage of both video and Internet systems—provides a powerful platform for customized and highly targeted advertising. The system's content editing/substitution technologies—whether utilized during playback of recorded programming or on a "live" content—can be used to substitute advertising dynamically. Finally, the logging of user behavior by software on the CPE will provide

an unprecedented the level of detailed feedback available to advertisers.

[013] The notion of "customer premises equipment" is undergoing rapid change with respect to displaying video content. Video content can be displayed not only on television sets directly or via "set-top box" television-centric computers, but digital video recorders/personal video recorders (DVR/PVRs), personal computers and laptops, portable DVD players, portable media players (whether music/audio only or audio/video), mobile phones with multimedia capabilities, and so forth in an explosion of media-capable and increasingly interconnected digital devices. The proposed video aggregation, recommendation, and filtering technologies will be applied to other non-traditional devices as the system is enhanced to include them within its purview.

[014] Although individuals and families have special needs with respect to the finding, managing, and filtering/blocking video content from all sources, most users will also be using non-video Internet content at the same time and often on the same device or closely related devices. "Internet content" includes, but is not limited to, web pages made up of text, pictures, audio, embedded video, etc.; email consisting of similar content types; instant messages consisting of similar content types, and so forth. There is a major overlap between users' interests and concerns regarding video content and their interests and concerns regarding Internet content. Thus the proposed invention will manage Internet content using the

same knowledge base about the user's interests and concerns as it uses for video services. As with video usage, the user's Internet usage patterns will be aggregated into the same knowledge base and analyzed in order to improve the delivery of desired content and the filtering or blocking of undesired content, as well as improved recommendations and filtering in the video service (for example, if the user is using the Internet to do browsing on automobile racing, the video system will automatically record and propose to the video user programs on that same topic). In addition, the Internet usage behavior of large numbers of people who self-identify with one or more AGs can be used to further customize aggregation, recommendation, and filtering, since it is likely that users within self-chosen AGs will have similar interests, preferences, and values with respect to Internet content.

[015]With regard to the technological approach to Internet content management, the proposed invention will install no software (or minimal software) on client computer systems. Instead, it will run all of the aggregation, filtering, and usage-capture logic on network-based servers that logically sit between the client computer and the Internet. This approach has many theoretical and practical advantages, and is fundamentally enabled by the fact that the users' video service provider (VSP) is also likely to be their Internet service provider (ISP), and thus it is straightforward for the proposed invention to place the "business logic" of Internet content aggregation and filtering in the edge network operated

by the ISP. However, even in cases where users choose to use the Internet from other locations outside the home, a very small layer of software installed on their computers could be used to enable network-based content management and filtering.

[016] In view of the foregoing, it will be seen that there is a need for a system that provides highly customized video content management, aggregation, recommendation, and filtering for use by subscribers to broadcast, cable, satellite, and Internet-delivered video, as well as associated non-video Internet content. The power of the system is dramatically multiplied because users will be encouraged to associate with one or more AGs, whose aggregate behavior can be used to further customize the system for each individual user or group of users (typically, a single household). Users of the system set the parameters of the system and can be allowed to choose to by-pass it, so rather than in any way restricting the use of media the system enhances user choice and provides more viewing and media management options than are currently available.

[017] None of the prior art, taken either singly or in combination, is seen to describe the instant invention as claimed.

SUMMARY OF THE INVENTION

[018] An object of the invention is to achieve a means of pooling all available content delivery systems and screening and customizing the content available according to the subscriber's tastes and viewing conveniences to empower the consumer to have more personalized control over their media selections.

[019] An object of the invention is to provide to viewers an ability to receive an entertainment signal with an editorial application service capable of selectively screening out unwanted content.

[020] Another object of the invention is to provide a live editorial application service that can promote programming content, which better serves the individual subscriber.

[021] Yet another object of the invention is to provide broadband delivery systems for content delivery.

[022] Still another object is to provide electronic circuitry to receive signal input from both a television programming provider such as cable or satellite and a broadband connection for receiving content editing services.

[023] It is a further object of the invention to provide a system capable of editing and selecting media content provided by broadcast systems such as television, radio and computer networks such as the Internet.

[024] It is another object of the invention to provide various levels of selective screening of content, which selectively allow full viewing of modified viewing of the same content.

[025] Yet another object of the invention is to provide to subscribers a customized on screen programming guide which displays only desired programming as well as offering additional recommendation information to further personalize to the subscribers unique needs.

[026] Still another object of the invention is to provide to subscribers edited playback of programming, which deletes undesired scenes and language from television programs after broadcast of the original programming.

[027] Yet another object of the invention is to provide a personal video recorder that receives and records an entire program so that a subscriber can receive and store programs that match their unique needs and interests.

[028] Another object of the invention is to provide to subscribers Internet based programming content in addition to television programming content on the same television screen.

[029] These and other objects of the present invention will be readily apparent upon review of the following detailed description of the invention and the accompanying drawings. These objects of the present invention are not exhaustive and are not to be construed as limiting the scope of the claimed invention. Further, it must be understood that no one embodiment of the present invention need include all of the aforementioned objects of the present invention. Rather, a

given embodiment may include one or none of the aforementioned objects. Accordingly, these objects are not to be used to limit the scope of the claims of the present invention.

[030] In summary, the invention is directed to a custom television programming service with human editors evaluating and editing the playback of television programming content. Preferably, the editors evaluate each program individually for content based upon viewer expectations. Programming may be modified during playback to exclude objectionable subject matter, including scenes or language. An onscreen programming guide is displayed that preferably indicates programming meeting predetermined content preferences as well as automatically recording and storing those programs that meet predetermined criteria and evaluation by the live editorial desk specializing in the content genre(not unlike a newspaper service organizing its editorial desks by genre of interest). The onscreen programming guide includes a listing of programs that have been previewed and predetermined to meet the subscriber's expectations of programming content. Programming such as weekly television shows that normally meet content expectations can be displayed with a highlighted color when such shows contain objectionable subject matter. Editing functions can be overridden using an administrator access code.

[031] It is an object of the invention to provide improved elements and arrangements thereof in an apparatus for the purposes

described which is inexpensive, dependable and fully effective in accomplishing its intended purposes.

[032] These objects of the present invention are not exhaustive and are not to be construed as limiting the scope of the claimed invention. Further, it must be understood that no one embodiment of the present invention need include all of the aforementioned objects of the present invention. Rather, a given embodiment may include one or none of the aforementioned objects. Accordingly, these objects are not to be used to limit the scope of the claims of the present invention.

[033] These and other objects of the present invention will become readily apparent upon further review of the following specification and drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[034]Figure 1 is a diagrammatic view of the overall system according to a preferred embodiment of the invention.

[035]Figure 2 is a diagrammatic view of the overall system according to a second embodiment of the invention.

[036]Figure 3 is a diagrammatic view of the set-top box/digital video recorder (STB/DVR) client software architecture and functionality.

[037]Figure 4 is a diagrammatic view of the editorial meta-data creation, distribution, and usage process.

[038]Figure 5 is a diagrammatic view of the Internet content management and filtering subsystem.

[039]Figures 6A&B are to a table showing meta-data types and media types and their interrelation.

[040]Similar reference characters denote corresponding features consistently throughout the attached drawings.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT(s)

1. System Overview

[041]As depicted in Figure 1, the preferred embodiment of the invention includes a number of components and subsystems. At the heart of the system is a centralized data center (02) at which a large fault-tolerate server cluster (04) manages and orchestrates three databases: rich supplementary information about video/audio/multimedia sources that describes its contents in a structured fashion ("meta-data") (08), user-supplied information about himself/herself and family as well as their affinity groups ("user information") (10), and records of the user's multimedia and Internet usage behavior ("usage information") (12). Alternatively, to achieve massive system scale, the data center (02) can be regionalized and remain "close" to a geographically restricted set of users, with a "single image" multi-master replicated database used for meta-data (08) and only certain crucial but limited system-wide parameters stored across datacenters in a replicated multi-master fashion with respect to user information (10) and usage information (12).

[042]To make use of the invention, the user would contact either his/her video service provider or the system provider to initiate the service. The initial contact to acquire the system could be made by telephone or using a computer system. But even if initial contact is by telephone, preferably the user will utilize some kind of computer system, whether

personal computer (PC) or set-top box (STB) video computer, to configure and deploy the system, for example to upload the user's preferences to the system.

[043] In its preferred embodiment the user uses the software and services that are parts of the invention to acquire and activate the service provided by the invention. This approach is preferred because the system needs to acquire a significant amount of information about the user. Some of that information need is ameliorated by the user's self-identification with an AG (as discussed below). But a better approach is to use the system itself to create a user account and begin a series of automated or semi-automated processes that result in the full deployment and usage of the system.

[044] The preferred "boot-strapping" model for using the system to configure and deploy the system is based on the ubiquity of the Internet and world-wide web. The user would connect to the Internet and use a web browser on use his/her PC (56) to attempt to create a new system account via a system-provided web server (06). The system would ask the user to identify his/her video service provider (VSP). If the system has the appropriate partnership with the user's VSP, the account creation would be permitted. The user would then have an account created within the main system database (10). If the user's VSP is not a partner, a suitable VSP may be recommended, or the refused subscriber's (non-personalized) data could be used as evidence to the VSP to subscribe to the invention's concepts and systems.

[045]Another "boot-strapping" model for deploying the system would be to use the interactive TV features provided by some (but not all) VSPs. In some systems, users are able to browse, purchase, and download new applications to their STB (52). The system of the present invention would preferably be listed as a downloadable application from the VSP application server (22), the user would purchase ("license") the service, software or application, and the system would download (26) and initialize the client application. When started, the client application running on the STB (52) may then initiate an interaction with the user in order to acquire the needed information, upload the acquired information to the system server (04) (preferably via the tier 2 server (24), which would in turn upload the information to the main system server (04)), which would then be able fully to configure and deploy the service.

[046]Whether signing up for the service via a web browser or the more limited user interface provided by a STB application, the user would preferably initially enter his/her affinity group (AG) (if any) or groups (church/mosque, parents' organization, community groups, sports groups, etc.), allowing the system to create a number of reasonable default settings based on information provided by the leadership of the AG. Alternatively, non-AG, non-group default selections could be provided as options or a questionnaire could be provided to select, change or accept system defaults.

[047] AGs will have pre-created their accounts in the system so that as subscribers join, they may select one or more AGs of interest. An AG can choose to be open (anyone can join) or closed (the user must provide credentials to join), and those policies would be applied as appropriate to the initial user registration. AGs have a strong incentive to invest in the system because they are provided their own "virtual channels" for content distribution to their target audiences. "Virtual channels" are simply mechanisms within the system invention by which the digital video content produced by or otherwise of interest to the AG are delivered to the user's set-top box/digital video recording (STB/DVR) device (52) and presented to users in a manner almost identical to broadcast TV channels delivered by standard means. As described in detail in Section (6), AGs that also create meta-data for their members are given the functional name "meta-data creation organizations" (MCOs) in the system overview diagram (70, 80, 82).

[048] Beyond the already partially customized settings provided by his/her initial choice of AGs, the user would also provide the system with as much additional information or customization data as he/she is willing to provide, such as family members (each of whom can in turn have customized system settings), their ages, hobbies and avocations, media and entertainment interests, and level of content filtering desired (high, medium, or low) and in which appropriate categories (for example, separate settings for nudity, sexual

content, violence, and portrayal of alcohol/drug abuse), and user's ability to opt out of or change system settings and functions (such as filtering). The user would also configure settings appropriate to Internet content management and filtering, and optionally create managed and filtered email accounts, home pages, photo and video-sharing sub-sites for his/her family members, and so forth. The shared data portions (which may or may not be filtered as well) of the account such as, for example, uploaded photo and video content, can be automatically shared with members of an AG (this would be typical of a closed group), or the system can require that individuals receive permission from the user even if they are members of the same AG (typical of an open group). The user may also create the appropriate accounts and passwords that would allow the systems various content filtering features to be by-passed or altered as needed and appropriate.

[049]Once the account is created using the preferred method (web-based from PC (56) to system web server (06)), the system would optionally download to the PC a small piece of software ("client shim") that subsequently would be utilized to enhance the overall performance of the system in several respects (detailed below). The system server (04) would then notify the VSP that the user has signed up for the service. This could be done in an entirely software-driven, automated fashion if the IT infrastructure of the VSP is sufficiently sophisticated and network-connected. The VSP datacenter (20) would then push down the client portion of the invention to the user's STB/DVR

(52) which would automatically install it (automated client software push and installation is a reasonably ubiquitous feature of modern digital TV systems).

[050] At the same time, the main system server (04) would push down an appropriate subset of the acquired user information (10) to a respective tier 2 server, if any (24) located within the VSP datacenter and private network (20) where it is fully accessible to the STB/DVR (52). The interaction between the main server and the tier 2 server would take place over a private (or encrypted virtual private) network connection (14). As a result of this downstream replication, the bulk of interaction between the system and the STB/DVR would take place in a relatively "local" way between VSP datacenter (20) and the home (50), which is of course only one of many thousands of homes (40, 42) potentially using the invention within a single VSP environment. The VSP is, in turn, only one of dozens or hundreds of additional VSPs (44) that also provide the service that is the preferred embodiment of the invention.

[051] Once installed, the STB/DVR client portion of the invention ("STB client") would immediately undertake a number of steps to explore its environment and begin to customize the user experience. This is especially helpful to determine alternate or better communication methods available between the system and the user, such as broadband Internet connections.

[052]First, the STB client would contact the system tier 2 server (24) located at a pre-configured network address within the private IP network operated by the VSP. The client would register itself and pass up to the tier 2 server relevant information such as its own hardware and software capacities and levels, its IP address and media access control (MAC) address, relevant installed software and hardware, etc.

[053]Next, the STB client would examine local network connections and determine if it had any additional connectivity, such as a connection to the home local area network (LAN) (62). If this connectivity were available, it would then try to connect to the system datacenter (02) through the default gateway on that subnet. If successful, the STB client has auto-discovered an alternative path (36, 34, 18) to the main system server, as well as to other Internet resources. And the system server will have discovered and stored the client IP address, which is most likely the public address of the home gateway (60) due to almost-universal use of network address translation (NAT) in home environments. This data can be used in a number of useful ways. For example, if a browser client connects to the system servers from the same public IP address as an STB client, the system will know that the STB client and the web client are co-located behind one or more NAT devices (60), most likely on the same LAN (62). Conversely, if account modifications take place from a different IP address than the public IP of the STB client, the servers can note the discrepancy and possibly log it, allowing

(for example) a parent to be aware of modifications to the account were made or attempted from a non-home-based computer.

[054]Whether or not Internet connectivity exists, the STB client would then preferably broadcast to the subnet and capture the MAC addresses of all devices within the broadcast domain (typically, the home LAN). If the "client shim" system software is resident on one or more PCs, it would respond to a special broadcast and identify itself to the STB client, thus enabling the STB client to find an associated PC in a totally automated way, even if there is no Internet connectivity via the LAN. That PC can later serve as the source of multimedia content for display and interaction from the STB client.

[055]Next, the STB client would upload its VSP-provided, as well as home LAN, MAC and IP addresses, the MAC address and account identifier returned by any discovered "client shim," and (assuming the user has given permission for maximal network discovery when setting up his/her account) all other MAC addresses discovered on the LAN through the tier 2 server (24) to the main server (04). Since MAC addresses are (in general) globally unique identifiers, this stored data can be utilized to auto-discover and associate in-home (50) resources during the course of the usage of the entire system, including not only the video service but the customized Internet content management and filtering sub-service as well.

[056]While this auto-discovery and auto-configuration work is going on, the tier 2 server (24) will be downloading to the STB client (52): (a) customized "home page" content

appropriate to the user and his/her interests, (b) the meta-data provided by an AG editorial team (70, 80, 82) as needed to modify the electronic program guide (EPG) that is normally provided by the user's VSP, and (c) a welcome video from the system itself and/or a welcome video from the user's AG(s). These operations and their results will be examined in greater detail in Sections (4) and (5) below ("User premises functionality"). But before examining the preferred embodiment of the invention at the user site, an alternative architecture for the overall system will be described.

2. Alternative system architecture

[057] So far the system invention has been described in terms of a preferred embodiment, in which a three-tiered model is employed and the VSP is a necessary partner in the development and deployment of the system. As an alternative, the system can be deployed in a more "free-standing" way in which the cooperation of the VSP is not necessary and the system works completely independently of the video delivery system and mechanism. Figure 2 shows a diagram of this alternative architecture. The alternative in turn has two sub-alternatives, one based on relatively high-speed and more or less continuous Internet connectivity ("broadband"), and the other based on slow and intermittent network connectivity ("dial-up").

[058] In the alternative architecture, the video receiving and playing device on which the invention client software is running is not a STB provided by a VSP but some kind of

alternative TV device generically called a STB/DVR/PC (52), either (a) a dedicated piece of video display and playback hardware with local storage provided by the system provider with the requisite client software pre-installed; or (b) a special-purpose PC-like computer running a standard operating system like Linux, Apple Computer OS.X, or Microsoft Windows and dedicated to the STB/DVR function; or (c) a more general purpose PC running one of those operating systems and powering the television monitor but capable of more full-function usage as well; or (d) some other plausible alternative computer capable of running the system client software.

[059] In the case of (a), the client software is bundled with the hardware. The user would purchase both the hardware and the service either at a retail store or website, or directly from the system provider by logging in to the systems user interaction server (06) from a PC (56) via the Internet (18). In either case, the user's system account would be pre-created and associated with the hardware ID (MAC address or other UUID) of the device. After the hardware arrives at the home, it would be unpackaged and connected to the video system (20) via coaxial cable (or other television system medium) and also connected either to a phone jack for dial-up access, or to an existing LAN (62, wired or wireless) in the home (50) for access to the Internet. The client application would connect to the user interaction server (06) and automatically associate itself with the user account by hardware ID. All

preconfigured options for hardware/software device would then be downloaded and installed.

[060] In the cases of (b) or (c), the user would install the system client application on the hardware/OS using OS-specific tools and techniques. These might include inserting a CD/DVD or USB drive containing the client software and following the auto-run/auto-load model of the client operating system to install and start the system client application. In this case the system client application would provide the user with a hardware ID to manually associate with his/her user account. After that process was completed via web browser (56) and system web/database server (06, 10), then the client application could auto-configure as above.

[061] In the case of (d), techniques similar to the above would be used to minimize user configuration and automate the process of setting up the system as much as possible.

[062] Once the client software is installed and operating on the STB/DVR/PC (52), no further differences in the four kinds of devices will be discussed, as they will not be material from the perspective of the system invention and its preferred and other embodiments.

3. Alternative architecture: broadband versus dial-up connectivity

[063] There are two different ways the system would operate in the case of having broadband (60, 62, 34) versus dial-up (58, 48, 49); access to the Internet. The broadband case is a superset of the dial-up case all system features are available

in the broadband case, but only some will work in the dial-up case due to the dramatically reduced bandwidth.

[064]Broadband sub-architecture: In the broadband case, most or all features available in the preferred embodiment (three-tier architecture working with a VSP) are also available; however, some feature may not work as easily or as well due to lack of cooperation from the VSP. In particular, all meta-data must be downloaded from the tier 2 server (24) which now must be accessed via the Internet. No meta-data can be inserted into the video broadcast stream, and certain additional challenges may arise, such as less consistent and reliable network access and download speeds for system content or AG content to be displayed on "virtual channels." Similarly, user usage information must be uploaded from the client device (52) to tier 2 server (24) in a more unpredictable and less reliable, and possibly lower-bandwidth environment (although in some cases the upload bandwidth can be higher via the Internet; but in that case, in the three-tier preferred embodiment the speed difference would be discovered and the Internet path (Figure 1, 36, 34) be used instead when available). Moreover, the semi-real-time meta-data required to be delivered quickly and reliably (see detailed discussion in Section (8) below) from a central real-time editing facility (Figure 1, 70) will be harder to engineer and scale without the benefit of the private VSP network and a three-tier architecture. The VSP architecture allows for the highly scalable technique of reliable multicasting from the head-end

to a large number of client devices, something not feasible over the Internet. Nevertheless, the broadband deployment of the system will have general feature parity with the three-tier preferred embodiment since even in the preferred embodiment the STB (52) is typically the location of the integration of broadcast media and EPG data with system meta-data and system-supplied video and other content.

[065] Dial-up sub-architecture: The dial-up case is quite different. Here, only small amounts of meta-data can be reliably distributed in a timely fashion. Also, possibly less than the total amount of captured usage information can be uploaded. Probably no "virtual channel" video content can be downloaded due to constrained bandwidth. Semi-real-time control data will also likely not be possible in such an environment. In this case the system would still provide a number of useful features, such as a highly personalized and customized EPG based on personal and AG data, edited playback of pre-recorded TV programs (such as re-runs), and edited playback of video-on-demand movies (which would be partially or fully cached locally on the STB/DVR/PC and played back in an edited fashion). But other features such as "virtual channels" and semi-real-time editing of live broadcasting would not be available.

4. User premises functionality: an overview

[066] The preferred embodiment of the invention provides a range and combination of new capabilities to the STB/DVR

device through the system client software and supporting infrastructure.

[067]From a content filtering and blocking perspective, the most fundamental concept is to leverage the growing power of the STB/DVR, which is now becoming a powerful computer with ample persistent storage and sufficient computational ability to modify dynamically and ephemerally all media content as desired by the user before display. This is accomplished independently from the production and distribution of the content.

[068]The customer premises equipment (CPE) approach may avoid copyright and economic control issues (such as content distribution contracts forcing the bundling of multiple unrelated media streams) that have traditionally hampered and disempowered users from getting the media content that they desire, as opposed to what the powerful producers of media content want them to consume. In the preferred embodiment of the invention content modification is accomplished after all distribution contracts have been fulfilled (the content has been delivered to the user's media device). Moreover, under copyright law the user has the right to make temporary copies for private use (for example, time-shifting), and also the right to make ephemeral modifications of playback of copyrighted content. For example, the Family Movie Act of 2005 modified the copyright laws explicitly to allow automatic, ephemeral modifications of digital content as it is played in a consumer device. Finally, the user has an obvious First

Amendment right not to watch undesired content: no one can be forced to watch portions of an entertainment program that they deem undesirable. Thus it cannot be claimed that it is illegal to use technological means to accomplish the same result that could be accomplished by more crude, manual methods (such as shutting off the TV set or turning off the sound temporarily while offensive content is being displayed).

[069]Filtering is important and uniquely accomplished in the preferred embodiment of the invention. But it only one of the important feature as far as users are concerned. While users have a strong interest in controlling, filtering, and blocking undesired content, they have an even greater, more positive interest in finding and consuming positive, desirable content. As it is today, the typical video system provides consumers with hundreds of channels and very poor tools for finding and watching or recording desirable content. The coming convergence of traditional TV distribution and Internet-based video will make the problem far worse. The user/consumer needs a trusted and knowledgeable source of information about the content that they are likely to find not only acceptable, but highly desirable and enjoyable. A happier user/consumer will be created since targeted advertising can be delivered by the system to users whose interests, preferences, habits, and even times of life (car buying time, wedding time, off to college time, and so forth) are known by the trusted content management system.

[070]An important aspect of the preferred embodiment of the system invention in a pluralistic society is the separation of technology and functional delivery of capabilities via meta-data from the non-technical opinion and judgment required to create the editorial meta-data. The user must trust the technology company building and delivering the service, but even more the user will trust the editorial judgment (both filtering and recommending) of the affinity group or groups that he/she voluntarily selects when signing up for the service, and who likely were instrumental in convincing him/her to sign up for the service in the first place.

[071]Since the technology described herein is functionally distinct from the human judgments that it is designed to deliver, and since it can deliver multiple sets of judgments embodied in different sets of relevant meta-data, the system is capable of providing what is needed in a modern democratic society.

[072]As noted above in "Background of the Invention," the system also uniquely meets the commercial needs of a highly market-driven society. A large amount of the entertainment created and consumed today—essentially all of the programming on broadcast and cable TV channels—is paid for by advertising. Yet advertisers cannot narrowly target the customers that would most likely be interested in their products but must "shotgun" their message to a broad audience hoping that enough interested people are watching to justify their investment. Moreover, some of their advertising (such as ads dealing with

male sexual dysfunction or violent movies or TV shows during family-oriented programming such as major sporting events) is actually offensive to many of their viewers, thus creating a problem far greater than disinterested viewers—hostile and unhappy viewers.

[073]The system invention provides a platform for totally custom advertising based on the system's intimate knowledge of the user's interest and activities. This information is collected directly from the user and also collected or inferred from their membership in affinity groups. That core database is supplemented by the usage of the system. As the user surfs and watches video, and surfs and utilizes the Internet, their usage of those systems is collected and stored by the system for later analysis. For example, if the user is searching for information about "Ford Explorers" on the Internet, that behavior is captured and stored. Later, this collected information can be used by the video side of the system for highly targeted advertising. The system's content editing/substitution technologies (discussed in more detail below) can also be use to dynamically substitute more desirable or more directed advertising. Finally, the level of detailed feedback available to advertisers is unprecedented: Advertisers can know whether the user watched the ad or skipped it or changed the channel, and at what point in the ad the user lost interest.

[074]Finally, from a usability perspective, while many important and real-time options are configurable using a

remote control connecting to the STB/DVR, all configuration options for the STB/DVR (and every other aspect of the system) will be available by logging into the system web site (15) from anywhere in the world using, for example, a rich web client interface. This will greatly enhance the usability of the system, which may be too complex to configure entirely from present, TV-based remote control interface.

5. User premises functionality: architecture and function

[075] The preferred embodiment of the system includes a set-top box/digital video recorder (STB/DVR) device running the system client software (STB client). Figure 3 shows additional details about the STB client. The STB/DVR (52) is, in the typical case, the junction point for unaltered television programming and advertising in the form of digital streams coming from a VSP (20, 22, 26), and a set of meta-data developed for the system by one or more meta-data creation organizations (MCOs) (70) and delivered to the system via a private or virtual private network link (79) into the system's meta-data repository (08). The editorial meta-data is then delivered via a private or virtual private network link (14) as needed to the VSP head-end/datacenter (20), and pushed out to STB clients by the tier 2 server (24) via unicast or multicast delivery (28) from the tier 2 server to the STB client.

[076] In an alternative embodiment in which the VSP is not involved with the system, all meta-data is delivered to the STB client by means of the Internet (Figure 2: 24, 18, 34, 60,

62, 52 (broadband path)), as discussed above in Sections (2) and (3).

[077] Also, in alternative embodiments the editorial meta-data is inserted into the video/audio stream up-stream of the STB client. Those options are discussed in Section (7) below.

[078] The STB client logic and architecture is as follows. First, the STB client includes a set of manager application modules that provide the framework for the real-time and semi-real-time behavior of the system. The meta-data manager (MDM) (40) receives via the wired or wireless network (86) a set of editorial meta-data that maps to electronic program guide (EPG) data and to video content. Unique identifiers (discussed below in Sections (5) and (7)) for EPG data and video program data provide a key in the meta-data database that allows for editorial meta-data to be quickly and accurately related to the standard EPG and video program data coming proximately from the VSP and ultimately from EPG producers or television channel producers such as NBC, CNN, ESPN, etc.

[079] In the case of the EPG, the editorial meta-data consists of records that reference a channel and the programs within a channel. The records contain additional ratings information about channels and programs that are much more granular and (if the MCO is doing a good job) much more accurate than the simple ratings system used by the industry (for example, TV-13 and the basis for that rating such as "sexual themes," "violence," etc.). In addition, the MCO will (optionally) provide two reviews of the program: a general one for all

users that summarizes the program and rates its quality; and a second one specifically for parents that discusses in detail any issues with the program that might make it unsuitable for children, as well as counterbalancing factors such as artistic quality and other positive elements, giving a reasoned explanation for why the MCO considers the content suitable or unsuitable for children of a certain age range, thus allowing parents to make a final decision. Finally, the meta-data coming from the MCO can also be the equivalent of completely new EPG data records which reference content and programming not found at all in the industry EPG, such as Internet-based video, or (most importantly) MCO/affinity-group created video programming. These records will have their own unique "channel" ID, standard ratings and other meta-data, and URI/URL pointing to the content source. The system will use these records to create the "virtual channels" that are an important part of the positive user experience of the system. These "virtual channels" are also important to the success of the system in that they provide one of the primary incentives for AGs to become (or contract with) MCOs to create meta-data and to add to the "network effects" that can make the system commercially viable and successful.

[080] In the case of video program data, the editorial meta-data consists of records that reference a particular television program or movie by industry-standard unique identifiers along with standard fields such as name, producer, industry-providing program rating, and so forth. The meta-data

also includes time-based information about program sub-sections ("incidents") within the video/audio streams. Each incident record includes the following data: (a) start time, end time; (b) media stream (video, audio, both, or alternate stream); (c) type of incident, rating of incident; (d) optionally, proposed incident alternative (blank screen, skip scene, obscure (portion of) video or audio, or substitute video/audio stream); and (e) pointer (URI/URL) to substitute content (along with flags to describe characteristics of alternate media such as to whether content is locally available as well as media type). As to the last item: when processing this meta-data as an incoming submission, the system meta-data server (08) will preferably use best efforts to obtain the substitute content stream and download that to the STB client along with the structured meta-data. At each step of the processing and download the characteristics flags will be updated to reflect the substitute content's local availability and state.

[081]Meanwhile, the system has already downloaded user preference information to the STB client from the user info database (10) via a private or virtual private network link (14). The MDM (40) processes and stores that user preference information in the local disk/DB (82). With both a local copy of relevant meta-data, and the requisite user information and preferences, the MDM is now ready to provide editorial decision services to the other parts of the STB client system.

[082]The first client application module requesting services from the MDM is the Home Page Manager (HPM) (44). This application can be configured to take over the screen each time the television system (or STB) is started; or it can be configured to overlay the most recently watched previous channel; or overlay a default channel; or be accessed by depressing a button, for example on a remote or the STB. The HPM provides the user a system overview, relevant news (as determined by AGs/MCOs), movie and TV recommendations (same), feedback on previous requests (such as which requested movie downloads have arrived), warnings (such as any suspicious account activity, times when the system devices have been off-line for no clear reason, or recurring attempts to access inappropriate content in the video or Internet content systems, or other system messages), and access to all account information and customization. The HPM calls the MDM to obtain the data needed to fill in these standard categories. The HPM also registers a notification mechanism with the MDM through which the HPM can receive and display real-time system alerts and events generated by the MDM (or its other system clients) at any time.

[083]The second client application module requesting services from the MDM is the customized, searchable EPG, hereafter known as the Advanced Program Guide (APG) (42). The APG utilizes standard EPG data but also merges in changes to EPG data as supplied by the MDM before displaying EPG data on-screen. The APG can be configured as a sub-screen of the HPM

or can take over the screen to maximize usability; it can easily be "zoomed" in and out of full screen mode. The APG is responsible for the customized on-screen display of industry-standard programming information, as well as user-initiated searching and filtering (such as "show all basketball games in the next week"). The APG can show channels in any order of preference as configured by the user or the user's AG(s), and/or based on past usage and predicted future usage. The views include pre-configured outline views, ratings views, recommendation-based views, filtered (search-based) views and outlines, and sorted views (such as a sort on an actor's name), all with "explorer-like" tree-structuring applied to the row and column data based on key fields and sort indices.

[084] Thus the APG provides a completely customized view into EPG data, a view determined ultimately by one or more MCOs (70) and proximately by merging EPG data with MDM-supplied meta-data. The meta-data can contain supplementary or completely new or over-riding meta-data about programs and movies, such as the detailed description of a program that is part of the standard EPG data. The APG may utilize a simple (and configurable) color-coded system to indicate which content is likely to be most desirable and which is possibly of concern to the user and his/her family (highly problematic content is not shown at all).

[085] The APG can also pull in content from local computers or can merge data from different sources (such as multiple VSPs), if the STB client is connected to the household LAN or other

input sources (such as other VSPs). The "client shim" software for the video system running on local PCs (56) can interact with the HPM and the APG such that locally-available computer-based content (movies, photos, music, and even email and instant messaging) is also available on the STB client. The HPM and APG can also manage playback of local DVD movies and other video content (whether accessible from a DVD drive built in the STB/DVR device itself, or accessible over the home LAN from a personal computer or some kind of dedicated player, networked (e.g., USB or IEEE 1394 FireWire) or otherwise connected (e.g., via HDMI cable)) using the same meta-data scheme and editorial skip/substitute model that is used for broadcast and cable content.

[086]The interaction between the HPM (44) and APG (42) on the one hand and the MDM (40) on the other involves data that is relatively static and infrequently updated. For example, other than news, most of the inputs to the HPM and APG will be based on data that is 24 or more hours old. Thus, in an alternative embodiment the merging of user information and preferences with programming information and MCO data as displayed by the HPM, and the merging of industry standard EPG data with MCO-provided meta-data as displayed by the APG, could in principle occur upstream of the STB client. For example, most HPM content could be created in the system data center (02, 04) or tier 2 server (24) and downloaded to each client on a daily basis; or the APG could be created at the first or second tiers and downloaded to each client. Or, the EPG and MCO-

provided meta-data could be processed up-stream of the EPG without considering user-specific information since there would be a relatively small number of resulting customized EPG outputs, and the user-specific customizations could then be pushed down to the STB client. However, the preferred embodiment takes advantage of the substantial local processing of the STB hardware and allows for a much more scalable system, since in effect the creation of customized and individualized content based on the merger of standardized content (both industry-standard and system/MCO-standard) happens in a massively parallel fashion when it occurs at the STB level.

[087] Turning now to the more real-time elements within the STB client system: Video content can arrive at the STB client via multiple means such as a built-in TV tuner decoding RF signals into digital data streams (84) or digital streams coming over a wired (typically Ethernet) or wireless (typically 802.11x) connection (86). Content can also be streamed from the hard disk of the device (82). In most cases hard disk content first arrived over the standard video (84) or network (86) and was looped back to the disk by via the DVR component (94). In all cases the content is normalized by the video source manager (VSM) (80) so that the remainder of the system is shielded from the details of different video formats and delivery mechanisms.

[088] The junction point between editorial meta-data for video streams and the streams themselves is the Real-Time Meta-data

Controller (RMC) (88). This module, upon receiving a command to play video from the HPM or APG, first determines what kind of stream is being requested (real-time or stored), associates the video stream ID with associated meta-data by calling the MDM interface methods, makes a fast determination about whether real-time playback is compatible with the meta-data for that program stream, and then either begins playback or else begins caching the real-time stream by calling the VSM (80) to utilize the DVR (94) to being making a stream copy.

[089]The RMC logic depends on the interaction of a number of factors. The two most important factors can be considered as two independent variables whose correlation and interaction, when combined with user preferences (including reasonable defaults) determines the output of the video display system (90, 92) when the user requests that a video program be played.

[090]The two factors are: (a) The nature of RMC's access to the video stream; specifically, whether it is available locally or "local enough" to allow skip ahead or even random access to the stream, versus a steady-state download such as the case of standard broadcast or a typical video-on-demand (VOD) service, versus a "start now" download that arrives at the STB client at an arbitrary rate. (b) The presence of various kinds of meta-data that interact with the different kinds of video streams in a complex fashion, as explained by the following definitions and table of interactions.

Video Stream Type	Explanation
Real-Time Standard Stream (RTSS)	This kind of video stream is provided by "video carousel" senders and is typical of live broadcasts and VOD systems. Datagrams containing the (for example) MPEG-2 data packets arrive in an entirely one-way fashion using their own forward error correction and are surfaced to the video playing infrastructure at a rate very similar to the playback rate. RTS streams may be of indefinite duration. Because of the nature of best-efforts digital video delivery (whether over-the-air ATSC/HD digital broadcast, or IPTV UDP unicast/ multicast), some local buffering of data must occur since data display must be coordinated by a highly accurate local clock; but buffering seldom involves more than a few seconds of data.
Start-Now Indefinite Delivery (SNID); types 1 and 2	This kind of video stream is sent on request by a "video pump" that delivers the data on a non-synchronized, best efforts basis. The send may be unreliable/datagram (typically with forward error correction), or it may be via a reliable connection-oriented protocol such as TCP/HTTP. The total size (in bytes) and playback time (in minutes/seconds) of the stream is sent as part of the initial response packet(s), but the receiver has no further control over the speed or ordering of delivered packets. Internet video delivery is typically of this sort (for example, MovieLink or YouTube downloads). A key sub-issue of SNID is the download speed, which can either be specified (as an estimate) by the sender or (more typically) measured by the receiver using a short sampling period (either an explicit sender-supported speed test, or simply by sampling the beginning of a "dumb" download). Based on the sample, the receiver can estimate the amount of time that must be buffered to ensure real-time playback. In the table below, SNID1 refers to the case where download speed is estimated to be greater than playback speed; SNID2 refers to the opposite case.
Stream Skip Forward (SSF) types 1 and 2	SSF is a variation on SNID in which the receiver can specify that certain portions of the stream be skipped. This specification can be made either at the time of stream download

	initiation or (with more sophisticated senders) at any time, to be honored on a reasonably-likely but best-efforts basis. There are no known implementations today of SSF senders, but it would be a very simple enhancement to content streaming servers currently using SNID, and it would be very helpful to the system invention, and so it is broken out as a separate case. In the table below, SSF1 refers to the case where download speed is estimated to be greater than normal (no-skipping) playback speed; SSF2 refers to the opposite case.
Stream Random Access (SRA)	SRA provides the most controlled kind of access to video streams. The receiver can "jump" at any time to any part of the video stream(s)/sub-stream(s). The canonical example is playback from local permanent storage, such as a hard disk. However, access to content over a relatively fast/reliable network (for example, a home LAN with speeds well above the streaming speed of the medium) using remote filesystem-type methods (*e.g.*, SMB/CIFS or NFS) also qualifies as SRA.

Table 1. Video Stream Types and Definitions

Meta-Data Type	Explanation
Unknown Program (UP)	The program is unknown to the MDM. By default the program will not be displayed in the APG; if displayed by user choice, the default behavior will be to reject an attempt to play. (This default, and all other defaults mentioned in this table, can be overridden in user preferences.) In the case of non-SRA delivery: by default the system will not record the program for later viewing.
Program No-Data (PND)	The program is known to the system but there is no meta-data beyond the industry-provided ratings data (if any). By default the program will not be displayed by the APG; if displayed by user choice, the default behavior will be to reject an attempt to play. However, in the case of non-SRA delivery the system will permit recording with the expectation that meta-data may be available at time of playback.
Program Rated,	The program is known to the MDM and ratings meta-data is available. No meta-data is

Insertions Unrated (PRIU)	available for "insertions" (typically advertisements or public service announcements, but the concept is generic in the system), however. By default the system will allow viewing/playback if the program rating is acceptable under current user settings (in other words, the system assumes that inserts have the same rating as programs; that assumption can be over-ridden).
Program Rated, Insertions Rated (PRIR)	The program and insertions are both known to the MDM and ratings meta-data is available for both. Both ratings will be evaluated independently and the program will be shown only if both meet or exceed current user settings.
Program Rated, Real-Time Edits, Insertions Unrated (PEIU)	The program is known to the MDM, and rated in both "raw" form and "edited" form. The meta-data includes a list of "incidents" and optionally a substitution URI/URL (which allows things like a substitute sound track without profanity) and a "short incident" hint (usage discussed below). An incident is a time-based sub-program based on begin/end data. By default unrated insertions are assumed to be acceptable if and only if the "raw" program is acceptably rated; otherwise, the entire program will not be displayed. As an alternative non-default mode, the system will attempt to detect and skip insertions (or play alternative content).
Program Rated, Real-Time Edits, Insertions Rated (PEIR)	Same as previous, with the change that insertions ratings are independently evaluated against user criteria and the program is shown or not shown accordingly. As an alternative non-default mode, the system will attempt to detect and skip insertions (or play alternative content).
Program Rated, Real-Time Edits, Insertions Real-Time Edits (PEIE)	Same as previous, with the change insertions have their own incident list and editing instructions.
Program Rated, Live Edits (PRLE)	Same as previous, with the change that in this special mode the MDM indicates to the RMC that the relevant meta-data is not locally resident, but will be arriving on a semi-real-time "as-needed" basis. This will alter the behavior of the RMC both in terms of default settings, but also in other ways discussed below.

Table 1. Meta-Data Types and Definitions

[091]The two sets of factors interact in a variety of very complicated and often subtle ways. Figures 6A&B provide a table containing an overview of those interactions, usually in terms of system defaults (most can be overridden by advanced user preference settings). The abbreviation "n/a" stands for "not applicable," which can refer either to an unreasonable usability scenario, or a technical impossibility, or both. The most common scenarios are given in bold-faced font in the "meta-data type" column, and the scenario is given a short description as well.

[092]With regard to the algorithm used by the RTC to implement the above combinations of video delivery models and meta-data scenarios: At the start of non-PRLE programs (*i.e.*, in those cases when all relevant meta-data is present at the start of program play), the RMC schedules system timer interrupts ahead of each incident event, as described by the incidents' meta-data. Upon gaining control prior to an incident start-time, the RMC queries the MDM and HPM for incident meta-data and user filtering settings and takes the appropriate filtering or blocking decision. Next, it either implements a "skip" instruction, or if given a "substitute" instruction it prepares the alternate media stream, and upon the arrival of the incident start-time, it replaces the original stream with the alternate one until the incident's end-time is reached. At this time, the original media stream is restored, and the program is completed normally. As noted, the RMC will also

need to take into consideration media streams' retrieval and buffering characteristics, which are managed by the VSM.

[093]Of course, all of this powerful and complicated technology will be rendered useless if someone in the user's household can easily by-pass the system by, for example, unplugging a co-axial cable from the back of the STB/DVR and plugging it directly into a TV set, thereby displaying all the analog channels without any controls in place. Or, a more sophisticated by-pass would be to plug the cable into a digital tuner of media center PC and view the entire set of channels available. While a simple physical by-pass cannot be prevented by the system for non-scrambled and/or non-encrypted channels, the system will certainly be able to know when its configure STB client is disconnected. The tier 2 server (24) will track its interaction with the STB (52) and will utilize (when necessary) a "heart beat" packet once every five minutes if there are no other communications packets pending to make sure that STB client is on-line and active. If the STB client cannot be contacted, this information will be noted and logged. The next time the STB client is connected, a system security warning will be placed on the system home page, indicating the loss of connectivity, and the date and duration. This information will also be surfaced to the account owner when he/she logs in via the system web server (06). These methods will ensure that if the system is physically by-passed, the account owner will be aware of the problem.

[094]Finally, it is important to note that all this powerful and complicated technology is, in the end, an *optional* way of using the video system. The system owner/user with the correct administrative password can disable any or all of its features. Moreover, the owner can configure multiple sub-accounts such that different behavior will apply depending on which family member logs in to the STB client. The owner can also use time-of-day options to customize the behavior of the system based on family viewing patterns; for example, after 10pm the system would be configured to a more "open" mode of viewing.

6. Affinity groups and Meta-data Creation Organizations

[095]The last major technical component of the preferred embodiment of the invention is discussed in the next Section (7). To understand better the technical architecture, however, is it first necessary to understand the usage model and the "human architecture" of the system, so to speak. This section will discuss those non-technical aspects of the invention.

[096]It should be clear enough what a "user" of the system is: an individual person or persons representing a family or household in their interaction with the system for video and Internet content management, aggregation, recommendation, and filtering. A large part of the power of the system, however, comes from the manner in which it is able to understand the user within a larger context or set of contexts. Those

contexts are embodied within the system as "affinity groups" (AGs).

[097] An AG is an organization known to the system representing a group of persons with whom the user self-identifies as in a member. An AG could be a church (or synagogue or mosque) group, a local community group, a parent-teacher organization, a hobby group, a sports-oriented group, or any other "civil society" group organized to help its members achieve worthwhile social goals. An "open" AG is one who allows anyone to join when becoming a member of the system; a "closed" AG is one for which the user must receive an invitation.

[098] AGs have a strong incentive to join and promote the invention system and service because it provides the leadership of the organization with a way to serve their members and to promote their own activities. For example, registered AGs are able to provide their own video content as a "virtual channel" for their members. They can also (at the user's option) receive a special place on the users' Internet web pages (since in the preferred embodiment all user Internet interaction flows through a server or other device controlled by the system).

[099] The primary value added by AGs is their provision to their members of system meta-data that either promotes desired content or blocks undesired content, or both. These two kinds of meta-data will be distinguished by referring to meta-data that promotes desired content (whether provided by the AG itself or, more commonly, by an unrelated third party) as

"associative meta-data" (AMD) and referring to meta-data that limits or blocks undesired content as "subtractive meta-data" (SMD). AMD is used to highlight and promote content; it is not used to block or alter it. Conversely, SMD is used to block or alter content, depending on user settings vis-à-vis the SMD metadata.

[0100] Organizations that directly and technically interface with the system to create and provide meta-data in standard system formats are called MCOs (Meta-data Creation Organizations). AGs need not be MCOs. However, to participate in the system they should have a relationship with one or more MCOs and delegate to those MCOs the authority to create meta-data on their behalf. MCOs need not be AGs. They can be technically-focused organizations working on behalf of one or more AGs. They can create entirely different sets of meta-data for different client AGs (although there will likely be economies of scale for an MCO servicing multiple AGs since there is likely to be significant overlap in the kinds of meta-data that need to be created). Large AGs will typically also be their own MCO, but small AGs may contract with MCOs to provide the meta-data for their particular group of users.

[0101] That said, a small AG (such as a church group) could also be an MCO or partial MCO. It could contract with a more generic MCO to provide broad and deep coverage, but also have its own team of experts or volunteers who provide editorial meta-data to supplement the base meta-data. Since all that is required to do even real-time editing is a powerful personal

computer, a fast and reliable Internet connection, and some training in the use of the editorial software, an AG could easily set up its own "cottage industry" of home-based workers doing editorial work—both recommendations and subtractive tagging—from their homes or offices. An AG could even establish a rotating team of volunteers who deal with media issues and choices every evening during "family prime time," are available to "chat" with members over the Internet, provide real-time consultation and alteration of the system behavior with respect to Internet sites (both adding and subtracting sites from the generic "black list" by changing the AG meta-data), and so forth. An AG could also ask its members aggressively to give feedback on all media consumption so that the AG-specific meta-data grows rich and powerful very rapidly based on thousands of "eyes" providing input. The only limit to the ways in which an AG can help its members to enjoy, manage, control, and screen media content is the effort the AG is willing to put into enhancing the system.

[0102] Both AGs and MCOs are well-known entities within the data model and behavior of the system. The AGs are the crucial entities with respect to user membership, behavior, and desired outcomes of total system behavior. To the user it makes no difference whether his/her AG *is* an MCO or whether the AG contracts with an MCO to provide the desired meta-data. MCOs are the crucial "behind the scenes" entities with respect to the technical functioning of the system: how meta-data is authored, replicated, distributed, and used. To the system it

makes little difference whether there is only one AG or hundreds of AGs; but the relationship with each MCO involves tight technical coupling and coordination.

[0103] The technical architecture and functioning of the meta-data system will now be described. For the reasons given, the entity discussed in the following section is primarily the MCO.

7. Editorial tools and models; meta-data creation, distribution, and usage

[0104] In its preferred embodiment, the invention includes meta-data creation organizations (MCOs) which create the meta-data that flows through the system to provide the desired outcomes to system users. Figure 4 provides and expanded view of MCOs (70, 80, 82) and their relationship to the system as a whole. An MCO is an entity that creates meta-data and feeds it into the system. In order to provide value the system must have at least one MCO. However, there is no limit to the number of MCOs that can participate in the system.

[0105] There are two basic modes of operation for an MCO: real-time and non-real-time. The non-real-time mode is typical, necessary, and relatively straightforward; the real-time mode is optional and more difficult to implement. It is possible for an MCO to provide real-time meta-data services only, but in practice such an entity would most likely need to partner with one or more non-real-time MCOs because its standalone utility to users would be limited.

[0106] The non-real-time mode operates as follows. The MCO (70) obtains video programming (either directly from networks and other publishers, or by digital recording off of broadcast or other delivery mechanisms). The video programming is stored on servers with the requisite large amount of storage (72). Human editors then review each individual program using editorial content creation workstations (ECWs) (74) and create meta-data tags associated with the program's UUID and stored in a meta-data source server (78).

[0107] For example, a trained editor (or editors; an MCO may use multiple editors and tabulate their "votes" electronically to generate a final meta-data set) reviewing a program and following the MCO's policy guidelines (as created by the MCO or a related AG) will tag a particular scene in a program with AMD (associative meta-data) tag. The tag will specify the start and end times of the scene, some general information about the scene, and provide a label from a defined list with a reason or reasons why users may want to view that particular program and scene, such as its family-friendly humor or its exhibition of patriotic values, along with a "strength" indicator to reflect the editor's judgment of the importance of the program/incident. In another part of the program the editor may tag a particular scene with an SMD (subtractive meta-data) tag. The tag will again specify start and end times, general information about the program and scene, and then a label from a defined list as to why users might find the content objectionable in some way, such as

sexual content or profane language. The severity of the incident will also be recorded according to the system's content rating scheme. This editing approach would be applied to not only program data but also "inserts" which are not permanently associated with the program (the canonical example is commercial advertising inserted throughout the program) but which may be de facto associated with it (for example, if users have recorded a broadcast program and intend to watch it played back later).

[0108] With regard to the crucial notion of "time" within the meta-data scheme: all digital video includes timing information since the delivery of packets is only approximately at the rate of playback, and so a highly accurate clock-based device (typically within the video display subsystem of the digital playback device) must use the time information transported with the digital frames to provide synchronize playback. In general, this time information is available to applications and will be used by the system and its meta-data tools and databases for universal time identifiers with respect to particular programs. The use of time codes also allows the metadata to be completely independent of the video to which it relates by sending instructions related to elapsed time, for example, of the media, instead of related to frames or sections of the video.

[0109] However, there may be cases where time information is not available in a universal or synchronized fashion. For example, a video program may be converted to analog and the

converted back to digital, with the resulting loss of the data needed for precise synchronization. Or, a video sub-program (for example, the playback on multiple nightly news programs of a single video supplied by a terrorist organization) might be best viewed as "the same" sub-program, but the timing data will be different in all manifestations of the subprogram.

[0110] In such scenarios, the preferred embodiment of the system will use best efforts to create other reliable ways of tagging video streams with both unique IDs and relative repeatable time information. One possible approach would be to transmit with the meta-data (or tag in the primary meta-data with a URI and transmit asynchronously since this part of the meta-data will be relatively large) a time sequence of image digests (mathematical "summaries" of image data). On the STB client (Figure 3, 52), the same algorithm used to generate the digests at the MCO (70, 80, 82) is calculated on the playing image stream. The resulting sequence can be used with a measurement model in a Bayesian filter to both initially estimate any timing offset between the metadata and the STB, as well as dynamically correct any temporal distortions that may arise during play. Aside from temporal distortions, the digest sequence can also be used to detect spatial distortions in playback (arising, *e.g.*, to make room for a ticker to be displayed at the bottom of the screen), which is critical to the correct rendering of spatially sensitive metadata (*e.g.*, blurring or blocking a portion of the screen). An essential feature of the digest is that it is continuous. That is, if

two images are "close" then their resulting digests are also "close." Digest functions that have this property have already been constructed and examined in the context of image watermarking. In that context, the digest is constructed from projections of the image onto a series of test images generated from a private key. In the system application of this method, there is no need for a private key, and so the system algorithm would have the freedom to choose test images to suit its purposes. As is also the case in image watermarking, the system would preprocess the images to its most essential features before calculating the digest– desaturating the image, applying the Sobel operator (to detect edges), and renormalizing, thus projecting out what might differ between images without being considered essentially different.

[0111] This digest-based approach for reconstructing "time-like" meta-data is computationally expensive, and so may not be applicable as a real-time method in the current generation of STB/DVR devices. However, as such devices gain more powerful microprocessors and programmable graphics processing units, this approach will become more and more realistic, and will be used in when either fundamental time-data is missing (rare) or (more commonly) to detect sub-programs that are "the same" even when embedded within different programs. Additionally, it may be used to verify that the program indicated is the program received (e.g., the correctly identified episode of a series), and take such actions as are

necessary, such as blocking, warning, or identifying the material.

[0112] In addition to human editors working at ECWs (74), artificial intelligence software running separately (76) can be simultaneously scanning the video programming database and creating certain kinds of meta-data, or flagging items for closer human inspection. For example, a children's TV network with generally "safe" programming that is a low priority for the MCO to review with limited human resources could be scanned by AI tools and certain usual sound or video events flagged for human review and possible tagging. Automated tools could also pre-process video streams and create tentative markers for "inserts" which will streamline later human processing. Automated tools could also create image identification markers using hashing functions to allow video scenes to be recognized by the system as somehow the "same" despite having different program ID, timecodes, etc. For example, a video tape supplied by terrorists and integrated into the nightly news by many different news organizations could be flagged with its own unique ID by automated systems using the image digest algorithms described in the previous paragraphs and elevated to a first class entity within the meta-data system. As automated tools improve, they can be promoted to creating and entering directly into the meta-data source server (78) relevant meta-data that is then distributed throughout the system.

[0113] Potentially a large number of editors working at ECWs (74) supplemented increasingly by automated creation (76) could thus create a large amount of server-based meta-data (78) associated with pre-recorded video programming. This meta-data would be used throughout the system in the following manner.

[0114] Meta-data coming from one MCO would be pushed (79) on a regular basis into the system data center (02) and the meta-data database (08) via a private or virtual private network. Other MCOs (80, 82) would do the same. The system main server (04) would then associate meta-data by program and sub-program ID with the programming available at the different VSPs (40, 42, 43) served by the system. It would also analyze which AGs are represented within a given VSP site, as well as analyze the usage patterns and content feedback coming back from system users via the system via the tier 2 server (Figure 1, 24). Based on this analysis, an appropriate subset of the meta-data would then pushed (Figure 1, 14) to tier 2 servers within the VSP private network. From the tier 2 server the meta-data would then be pushed to STB clients, where it would be utilized by the Meta-Data Manager (Figure 3, 40) to promote and/or modify content accordingly.

[0115] Alternatively, due to the inherently multi-stream nature of video content (*e.g.*, MPEG-2 and MPEG-4 define a multi-logical-stream, single bit-stream transmission format), it is possible for the system to utilize the digital transmission of video as a means for distributing meta-data as

well. The meta-data can be added to the video stream itself (preferably front-loaded so that the meta-data arrives more or less completely within the first few seconds or minutes of the content streams being delivered) at some up-stream part of the video distribution system. For example, meta-data could be added to programming at the VSP head-end (Figure 1, 20, 22) by an insertion from the tier 2 server (24). One possible reason for this approach is that often there is dedicated bandwidth for video deliver that is not filled by the video content; for example, digital over-the-air broadcasting using the ATSC standards provides sufficient bandwidth for one high-definition (HD) and one standard definition (SD) encoding of all content, plus a small amount of additional bandwidth for signaling and other system uses. In the case where only HD or only SD content is being broadcast, there is ample "space" for meta-data to be added to the broadcast and rapidly delivered a potentially vast number of client systems in a reasonably reliable way (the ATSC and MPEG standards have forward-error-correction and other facilities for reliable data multi-casting in an unreliable network environment).

[0116] Turning to the meta-data usage model: It should be clear from the description thus far of the preferred embodiment of the system that a user can be associated with multiple affinity groups (AGs). In general, meta-data associated with different AGs (whether generated by one MCO or multiple MCOs) is additive: the user's meta-data is the total set of the meta-data of all AGs with which the user is

associated. The question arises: how is the different meta-data associated with different AGs applied to the user when the meta-data conflicts in some fashion? For example, one AG rates an "incident" as severity 1, another rates it as severity 2, and a third omits it altogether.

[0117] The answer has two parts. First, the user associates with AGs in prioritized list, and the priority is used to reconcile differences in SMD as well as prioritize recommendations embodied in AMD. By default, the meta-data associated with a higher priority AG takes precedence over the meta-data associated with a lower priority AG. This prioritization applies equally to the AMD/recommendation and promotion side of the system as well as the SMD/filtering and blocking side of the system. Second, with respect to SMD, the user can optionally choose either a "most restrictive" or a "least restrictive" option or a level in between (e.g., an average of the AGs). In the "most restrictive" option, the AG meta-data that identifies more incidents and/or gives incidents a higher severity is given priority over all other AG meta-data. In the "least restrictive" option the conflicting meta-data is analyzed and lower overall ratings and lower incident severity ratings are given priority. Note, however, that in the case in which one AG has a record (such as a program record or incident record) that doesn't occur at all in the meta-data associated with other AGs, that record is treated on an additive basis and thus determines the outcome, regardless of AG priority. As a result, a sub-program (for

example) that is unmarked by one AG and treated as an incident by a second AG will result in the second AG determining the outcome, since there is no meta-data conflict to resolve.

[0118] The system as so far described in this section depends on the analysis of recorded video by MCOs and its subsequent playback by the user in the presence of the relevant meta-data. This provides no benefit in the case of live TV "broadcasts" (however delivered), which is how most people still consume most video programming. In the following Section (8) the additional components and practices needed to deal with live TV within the context of the system are described. However, it is important to underscore just how valuable the system is even in the absence of any meta-data solution for live content.

[0119] (a) A unique part of the value of the system is not in the blocking and filtering features, but in the aggregation and recommendation features. All content, even if not pre-recorded and pre-screened, can potentially be part of the AMD/promotional side of the system.

[0120] (b) With regard to both AMD and SMD, the vast majority of video content is pre-recorded, and most of that pre-recorded content will be available for analysis by MCOs, either because it is a re-run of a previous live broadcast, or because it is a movie previously released in theatres or on DVD, or because through developing business relationships the participants in the system invention will be able to provide

compelling business justifications to producers to make content available for pre-screening by MCOs.

[0121] (c) Finally, with regard to both AMD and SMD, individuals and families concerned about healthy media choices are now fundamentally empowered by DVR technology to dramatically decrease their consumption of live TV and to "time-shift" most of their video viewing to later, more convenient times. When combined with the system invention and an MCO business model with sub-24-hour turnaround for the creation of meta-data for all the main broadcast channels, the result is that by simply viewing television program one or more nights later, the user of the system can greatly expand their range of possible media consumption, find better content through the recommendations of AGs as embodied in AMD, while still feeling "safe and secure" and protected from undesirable content as embodied in SMD.

[0122] In short, even with out any model or method for dealing with live TV programming, the system invention is still highly useful and valuable.

8. Live "semi-real-time" video management tools and techniques

[0123] That said, it would be more useful if, in addition to all of the capabilities of the system discussed above, there was also a mode in which live TV programming could be brought within the system umbrella. This section as illustrated by Figures 3 and 4 describes the preferred embodiment of the invention when in its live mode of operation

(the figure numbers are referenced only when relevant to the item number).

[0124] To summarize: The live video management mode takes advantage of the DVR capabilities of the client hardware and software, but instead of recording entire programs for later playback, the system uses the DVR function as a temporary cache for buffering live video and playing it back with a slight delay. During this delay period suitable meta-data will have been created and delivered to the STB client. The meta-data (received from the VSP or other sources such as Internet delivery) is then combined with the buffered live content and then played back to the screen. This mode is called the "semi-real-time mode" (SRM) mode for delivery of content to users.

[0125] SRM works as follows. Live content from a TV network is delivered to the MCO's live video servers (Figure 4, 60) while simultaneously being delivered into the user's STB client (Figure 3, 84). Because the STB client is running in "protected live" mode (as selected by the user when choosing to watch content either from the HPM (Figure 3, 44) or the APG (Figure 3, 42), the VSM (Figure 3, 80) uses the DVR component (Figure 3, 94) to make a local copy of the video stream without yet displaying it. The video is eventually displayed at a configurable interval between 60 (one minute) and 300 seconds (five minutes). For the remainder of this discussion the interval will be assumed to be 120 seconds (two minutes).

[0126] Meanwhile, one or more MCOs have implemented a parallel editorial/meta-data creation system for dealing with

live content. Human editors working at one or more live editorial workstations (LEWs) (Figure 4, 62) are watching the live broadcast. For reliability and accuracy these editors will generally work in teams of two or more, all watching the same content. (For the remainder of this description we assume that there is a team of three editors.) The LEWs have very specialized software that allow live video to be treated as though it were being played back for editing. So, for example, the editor can "pause" and "rewind" quickly and with great accuracy (to the frame level), "mark" or tag a frame as the beginning of a incident, "fast-forward" to the end of the scene/incident (if it has finished; if not to the current end of the stream) and tag again. The software has a sophisticated user interface allowing very quick application of meta-data tags and labels of the sort appropriate to live content. These tags immediately pushed into the meta-data source server (78) and treated as tentative meta-data records. The three editors are implicitly voting by all their choices and behavior. The system requires that at least two of the three editors agree on an incident before the meta-data is promoted to final state. (Automated systems (Figure 4, 64) can also be used to enhance the quality of live video management; given enough time for enhanced AI techniques and given constantly increasing computational power, automated systems could eventually become equally if not more important than human editors.)

[0127] Once the meta-data has been created (and stored for subsequent use), it is pushed to a high priority local queue on the meta-data server (78) and the pushed from the local queue to the system data center (02) and the system meta-data database server (08) via a private or virtual private network link (79). From there the data follows two paths: (a) it is stored in the long-term meta-data database like all incoming meta-data; (b) it is placed into a high priority outbound local queue and then pushed as quickly as possible via a network link (Figure 3, 14) to the tier 2 server (Figure 3, 24) and from there multicast to all STB clients where it is received into a high priority processing queue by the MDM (Figure 3, 40). The RMC (Figure 3, 88), the component in charge of coordinated, meta-data-driven playback of all video (in this case a recently-buffered video stream), then plays the delayed video stream in the normal fashion by calling the VSM (Figure 3, 80) and the MDM (Figure 3, 40) and applying the meta-data as described above against a RTSS video feed. So long as the entire process of quick editorial tagging and network distribution of meta-data takes less than 120 seconds, the result of this process is that slightly-delayed local playback on thousands or possibly millions of STB clients is essentially under the remote control of the live editorial teams at the MCOs. The length of delay will be determined by many human and technical factors, and the invention is not limited to the length of delay.

[0128] There are two important differences in SRM video that the various system components must take into account and that differentiate it from the standard non-live cases. First of all, since meta-data is arriving on a "just in time" basis, it is critical that any networking or other failures (such as an editorial team literally falling asleep) that result in the cessation of meta-data flow be evident to the STB client in a timely fashion. This requirement is met by the use of an end-to-end "heart beat" data packet based on indicators of editorial team activity that flows once every 10 seconds from the LEWs (Figure 4, 62), through the high priority queues on the meta-data source server (78) and system meta-data server (04, 08) to the tier 2 server (Figure 3, 24) and on to the MDM (Figure 3, 40) on the STB client. Failure of one packet to arrive creates an alert condition that suggests to the MDM that a system failure may have occurred. After some number of heart beat packets (in the case of a 120 delay, for example, this might be six packets or approximately 60 seconds) have failed to arrive, the STB client will "fail safe" and, depending on the current filtering level, may interrupt playback of the buffered content with an appropriate error message shown to the user (*e.g.*, "this program can no longer play in 'safe system' mode") and/or may superimpose the message in all or a portion of the screen. Of course, as an optimization when real meta-data packets are flowing they will be treated as heart beat packets as well.

[0129] Secondly, unlike the non-live case, it can be the case that an "incident start" record is created and transmitted with no corresponding "incident end" record available in time to allow the RMC (Figure 3, 88) to make a normal decision about how to merge meta-data and video source. In that case the RMC will follow a policy as configured by the system and/or AG and/or user as to what alternative content to display for a period of unknown (at start time) duration.

[0130] There are two possible models for dealing with SMD incidents that, according to user settings, result in blocking of video content. In the first and simpler model, the STB client would simply display alternative content, such as a family picture, randomized clips from home videos, a message from one of the user's AGs, and so forth. This greatly simplifies SRM playback because the live video delay (and associated details, such as local buffer size) remain constant throughout the process. The RMC would still need to be notified of the need to provide substitute content so that it could prepare the alternative media to be substituted instantaneously and seamlessly when the "start incident" timecode applies to the playback stream.

[0131] A more complicated but more seamless model works as follows. The human editors could optionally tag the "incident start" event record with a special field that implies "likely short resolution" of the incident ("short incident"); alternatively, a separate but associated "short incident" record could be pushed out through the system within 10-30

seconds of the "start incident" record. (A real-world example of when such an approach would be useful is the news item from early 2007 in which a television football producer chose to broadcast nationally an extended (five to seven second) picture of a young women in the stands wearing a t-shirt with an obscene word emblazoned across it.) The implication of these "short incident" messages to the RMC would be to plan to simply skip the entire incident and not substitute any content. But these hints are optional because alternative content would still need to be queued up by the RMC in case the "short incident" hints turn out to be wrong. What is not optional is the arrival of the "incident end" message at the STB client prior to the time that the RMC has reached the "incident start" frame from the locally-cached video playback buffer, plus some additional time such that the buffer is not completely emptied, and could respond in principle to another "start incident" request. A reasonable minimum buffer size/time delay would be 30 seconds. If the RMC receives the "incident end" message prior to 30 seconds before playing the frame pointed to by the "incident start" message, it would simply skip over the incident altogether, and continue to play video with a 30 second buffer rather than a 120 second buffer.

[0132] The skipping part is relatively easy; what remains is how to manage the small buffer size. The small buffer size means the system is much more at risk of showing undesired content (in case a "start incident" message is not generated and received within 30 seconds of the real-time start of the

incident in the broadcast TV stream). The STB client needs to rebuild the buffer to the configured size (minimum 60 seconds, default 120 seconds). There are several techniques that could be used to accomplish buffer expansion. One would be to simply slow down playback very slightly, for example, repeat a keyframe every second, thus causing a slight slowing of playback speed and gradually rebuilding the buffer. However, it would be difficult to rebuild the buffer quickly enough using slowing techniques without it being noticeable to the viewer that something was "wrong" with the video stream. Another more aggressive approach would be to insert an additional advertisement or two at the next advertising break. This would allow quick rebuilding of the desired buffer, although the opportunity for doing so may be some minutes away, and so the system will be running in a somewhat risky state in the interim. In its preferred embodiment, the system invention will use both techniques to rebuild the buffer to the desired time delay.

9. Internet content management and filtering

[0133] Internet content management and filtering is a well-known and widespread technology. In the typical case content filtering software is installed on a personal computer and from then on will block access to objectionable website and also (typically) scan other kinds of network traffic (such as instant messaging and email) looking for keywords that indicate problematic usage of the system. Web traffic is usually filtered with one of two techniques (or both). (a)

Every website requested by the user is first checked against a "black list" database of URLs compiled by the software vendor (or licensed from a third party); if the requested URL is found in the database, the software fails the request to view the site. (b) Once a site is allowed and the content downloaded from the web, the text (and sometimes graphics) of the requested page(s) is scanned dynamically for keywords and other indicators of objectionable content and the content is shown to the user only if the real-time scan indicates that the page is non-objectionable. With respect to non-web content (*e.g.*, instant messaging, email, etc.) technique (b) typically is applied to all traffic. The software will often have different log-ins available for different users so that its behavior can be tailored to (for example) an age-group-appropriate level of filtering. In almost all cases the software can be disabled or bypassed at any time by anyone with the administrative password.

[0134] There are a number of problems with typical Internet filtering software products. First of all, they must be installed and configured on each individual PC. Since parents are typically not as computer-savvy as their children, this often results in the system being compromised from the start, since the very people that the software is supposed to protect are often able to control and disable it. Second, the software is often very obtrusive and problematic in usage, since it typically alters the behavior of system rather dramatically and often conflicts with other programs running on the same

machine. Third, the software is only works on a specific set of operating systems; for example, most products in the market today work only on Microsoft Windows and provide no solution to users of Apple Macintosh or Linux computers. Fourth, and perhaps most importantly, all customization of the filtering software must occur by the user's configuration choices. The software's default behavior is "one size fits all," and the result is that it doesn't fit anyone very well. Customization by individual users to suit the needs and values of their family takes a lot of time and effort, and often users get frustrated and simply shut off the software altogether.

[0135] In the various embodiments of the system invention (Figure 5), the Internet content management and filtering technology is not based primarily on the user's computer. Instead, the system exists in the network itself, thereby greatly simplifying and enhancing the behavior of the system. Perhaps most importantly, the system's knowledge of the user through the sign-up process and chosen affinity groups (AGs) provides immediate and dynamically increasing benefit in terms of tailoring and customization of the software to meet the needs of the user and his/her family.

[0136] There are three main embodiments of the system's Internet management and filtering technology. All three modes will be utilized, and all provide similar benefits and behavior. The three embodiments/modes are:

[0137] (a) Server-based: The preferred embodiment of the system is a completely transparent system built in to the

user's Internet service provider (ISP) infrastructure. (This is the typical case for the system invention because the user typically receives both video and Internet service from the same provider, and as discussed in the preceding sections, the system is normally integrated into the VSP/ISP infrastructure.) Once the user is identified as a customer (50), all Internet traffic is routed through the system (30) and all management and filtering features are applied before Internet traffic ever reaches the home (36). An enhanced mode of operation (in which individual users and/or computers are differentiated) requires either a small additional piece of software on each computer,

[0138] (b) Gateway-based: Another, similar approach places the system intelligence in a smart gateway device ("enhanced gateway" or EG) (69). This approach is equally transparent and works even when the user's ISP has no knowledge of or relationship to the system invention and so all Internet traffic reaches the home (60) in a "raw," unfiltered fashion (61).

[0139] (c) Shim-plus-based: The final, least favored approach is based on putting a very "thin," simple piece of software known as a "shim" on the user's computer. This particular shim must have more functionality than other shims discussed in this detailed description, so it is hereafter called "shim-plus." The advantages and disadvantages of this approach are discussed in greater detail below, but the primary advantage is that the user's computer (55) is

"protected" wherever it is (50, 70) and however it connects to the Internet (36, 71). This approach thus is particularly important in the case where the user wants a portable family computer such as a laptop (55) to be "within" the system no matter where or how it connects.

[0140] In all embodiments the system's features and protections can be by-passed given the necessary security credentials. However, even the by-pass operation is enhanced in various ways, as discussed below.

[0141] The behavior of the system will first be discussed in detail in its preferred, "server-based" mode/embodiment. Next, the advantages of a group-oriented Internet content management and filtering model will be highlighted (this applies to all modes of usage). Finally, the additional two technical modes/embodiments will be discussed in terms of their differences (any features or functions not discussed remain the same).

a) Server-based mode

[0142] The system's content management and filtering capabilities begin with the use of humans and/or automated systems to scan large amounts of Internet web content, to evaluate it according to a number of industry-standard categories, and to create a URL-based meta-data database that contains this large information set. In the preferred embodiment, a system of automated "web crawlers" (84) (large and powerful computers constantly scanning the web for new content) with appropriate artificial intelligence (AI)

software scan web pages and rank the page based on a prediction of whether the page includes possibly objectionable content. This information is stored in a master URL database (82) and is periodically replicated to slave URL databases (34, 17) for usage by the system.

[0143] In the next logical step, the user's account information including family members, ages, sensitivity levels for each family member, and affinity groups (AGs) is analyzed by the Internet policy/proxy engine (32) and policy profiles are created that allow for very quick decisions about whether a given Internet usage request is permitted or denied. AGs can create their own customized "black lists" (URLs that contain questionable content) as well as "white lists" (URLs that are considered safe and desirable). The "white-listing" of a URL by an AG will override the "black-listing" of the same URL by a generic engine such as (80). To give an example of this customization: web pages based on traditional readings of the Qur'an or Bible containing criticisms of certain alternative lifestyle behavior or certain kinds of sexual behavior are often labeled as "hate speech" by generic web URL databases. An Islamic or Christian AG can either accept this labeling or may instead override the labeling with respect to certain web pages by adding those pages to their "white list." This is presumably what the user who associates himself/herself with that AG would expect and want. Thus, the system is immediately adding value to the user experience above and beyond a generic Internet filter.

[0144] In the fully protected home environment (50) all Internet traffic is routed through the Internet policy/proxy engine (IPPE) (32). Web sites on the "black list" are evaluated against user configuration and allowed or blocked. Web sites on a "dynamic sites list" (sites that have constantly changing and far-ranging kinds of information, some of which is may be quite objectionable to some users, much of which is not) are evaluated on a real-time basis. That is, each page is scanned as it is downloaded by the IPPE and passed through to the requesting client only if it meets the criteria established by the real-time scan. Beyond web usage, all other user-level protocol traffic—email, instant messaging, Internet gaming channels, etc.—is scanned in real-time.

[0145] In one common configuration, no per-user information is utilized by the IPPE; all requests coming from the household are accorded the same treatment. These household-wide policies are applied to web, instant messaging, and email traffic, and other popular user-level Internet protocols. Only when a URL is blocked is the user presented with a by-pass screen, where administrative or parental credentials can be used to by-pass the filter temporarily (be default, only for that single URL, but optionally for a short period of time).

[0146] This pure server-based approach is extremely easy to use and configure because no software need be installed on any computer in the protected home (50). It is completely automatic and safe. It can protect the user not only from

objectionable content, but also from undesirable content (such as web page pop-ups) and dangerous content (such as viruses and "phishing" schemes). (The user can even be made aware of the fact that they are using the Internet in a safe mode because the system can inject into each web page a script that displays a "safe browsing" icon on the screen of the web browser.) However, a purely server-based approach does limit some of the desirable options. For example, it would be desirable for different computers in the protected home (55, 56) to have different policies in effect depending on the user. Because in the normal case the IPPE is unable to distinguish between client computers behind a standard gateway (59), this level of customization is not possible.

[0147] The system can, however, provide per-user and per-computer levels of customization by the use of a small amount of additional software on each computer in the home. There are three possible approaches here. The simplest one from a development and deployment perspective is to utilize existing virtual private networking (VPN) technology that exists in all modern client operating systems. In particular, Windows, Macintosh, and Linux computers all include the simple VPN technology known as "point-to-point tunneling protocol" or PPTP. The identity of each computer and/or user can readily be established in a secure fashion by requiring the creation of a PPTP tunnel between the client computer (55, 56) and the IPPE (32).

[0148] There are drawbacks to the use of a technology like PPTP, however. First of all, there is some additional computational and networking overhead in the protocol, although that alone is probably not a sufficient reason not to use it. Secondly, the user model is more complicated since the user must utilized commands and utilities in the operating system that he/she is probably not familiar with. Some of that burden could be ameliorated or eliminated by the downloading through the client system web browser of a small piece of client software that automates the creation of PPTP tunnels as well as making the log-in process more like a web server log-in, as users are accustomed to doing for themselves. Third, PPTP may not be supported in the scenario where a portable computer (55) is taken outside the home and used in an arbitrary location and network (70, 78); some gateways block PPTP traffic. Finally, the PPTP solution would be very inefficient in terms of network traffic in any case where the PPTP server is not directly on the user's path to the Internet. In other words, in scenarios like protected home with an arbitrary ISP (60) or an arbitrary location/network connection (70), PPTP would route all traffic through the IPPE, which is likely not be as fast or efficient as sending only needed traffic to/through the IPPE while letting the bulk of data flow directly from the Internet (61, 71) to the home (60) or other location (70).

[0149] Thus, in its preferred embodiment the system provides one of two other possible approaches that meet the

requirements of the system without requiring the overhead and complexity of a VPN-based solution to the problem of user and system identity.

[0150] In the first preferred approach, a small layer of software called a "shim" is added to each client computer. (In function and operation, this is a different piece of software from the video-system related "client shim" layer discussed in Sections (1) and (5); and it simpler than (essentially a subset of) the "shim-plus" discussed below. But in implementation and deployment these three pieces of software would likely be combined, and almost certainly the "shim" discussed in this section would be combined with the "shim-plus" discussed below.)

[0151] In the case of the protected home (50), the only function of the shim would be to manage user and session identity. For example, when a particular computer is used for the first time after a period of inactivity, the shim would notice that no valid system session existed and would redirect the first URL request to the IPPE for a system log-in via web page supplied by the IPPE. Once the session is established (sessions are of configurable length, the default is 30 minutes), the shim would insert a session ID into each packet and the IPPE would then be able to apply the appropriate per-user policies to all Internet requests before stripping out the session ID and sending the request on to the Internet as usual. (Replies would not need to be modified or directly handled at all other than for real-time scanning (see below),

as the return path to that particular client is pre-determined by the reply IP address/port number combination managed by the gateway and the client system itself.) Sessions will be set to time-out periodically (based on a short period of inactivity (such as five minutes), and even periodically when sessions are active (such as once every 30 minutes) to protect against the case where one user starts using a computer immediately after another) so that user IDs are likely to be correct.

[0152] Another, more seamless way that the system can maintain user/session identity is by a slightly modified gateway device called a "modified gateway" (MG) (59). In this case the MG provides only one additional function (and associated setup routines), which is to provide a mapping between computer IDs on the LAN (represented by MAC address) and the use of a specific range of reply ports for Internet requests. There are approximately 63,000 dynamic port numbers available for replies to requests made according to any Internet protocol. A "network address translation" (NAT) gateway device (the type almost always used in home network environments) normally uses these port numbers to provide a private mapping between private LAN IP addresses and public IP requests. In this embodiment of the system, that mapping is utilized to provide a means by which the private ID of the computer behind the NAT can be communicated to the IPPE in an extremely efficient, seamless manner.

[0153] This part of the invention works as follows. When the MG boots and has an Internet connection, it connects over

the Internet using a secure channel to a system server in the system data center (02) at a well-known DNS name and registers itself as part of the system. If its source IP address is within a VSP/ISP network in which a system Internet Data Center (30) is present, the MG is informed of that fact and also sent the IP address of the IPPE for that network. (The IPPE is also informed asynchronously of the presence of the MG as identified by its MAC address (which functions as a UUID) as well as its current IP address.) If its source IP address is not part of a managed environment, the MG does nothing further. It makes this check each time it boots, and also once per day in case of configuration changes. (Most likely the check will succeed, as the user will not have purchased or otherwise acquired an MG unless they expect to use it within the system. However, the gateway could be taken out of the system and it will function normally.)

[0154] If the MG (59) is within an environment serviced by an IPPE (32), the two nodes work in a simple partnership that enables the IPPE easily to identify particular PCs within the private LAN (58). If the MG has previously connected to the IPPE but lost its configuration information, that information (including previously configured MAC addresses and associated reply port ranges) is downloaded at this "initial" log-in and the partnership continues as before.

[0155] The partnership between MG and IPPE is initially created as follows. The basis of the partnership is (a) the communication of internal MAC addresses to the IPPE under

special, occasional circumstances; and (b) the constant usage by the MG of a range of reply ports to uniquely identify requesting computers/devices by MAC address within the LAN.

[0156] The first part can be accomplished in a variety of ways. In the preferred embodiment, the user is requested to connect to their system account from each computer on their LAN and then to click on a special URL. The MG is always watching for a GET request to that particular URL. When it sees such a request, it incorporates the MAC address of the requesting computer (in encrypted fashion) into the request URL, and from that point forward (if not before, see next paragraph) it uses a dedicated set of reply ports for all outbound requests from that particular computer. For example, for computer X/MAC address Y, the MG will always use reply port numbers 2501-3000 (this block would have a simple identifier, such as "block 3"). (The number of ports in a block can be modified based on system experience, but it is expected that 500 ports is more than adequate for normal home computer usage, since it is extremely unlikely that any home computer will have more than 500 TCP connection or UDP sessions active at any one time.) Since there are about 63,000 port numbers available, this would allow the MG to communicate unambiguously the identity of more than 120 devices on the LAN to the IPPE simply by setting the reply port to a specified range. From that point forward, and with no further overhead or overloaded data inside packets, the IPPE can easily

determine what computer is making the request, and handle session and identity issues accordingly.

[0157] To summarize: the server-based mode of the system invention provides completely seamless, "hands-free" protection of the entire home LAN (58) without any software changes to local computers or any other technical intervention. The only thing that the user needs to do is sign up for the service via the VSP/ISP (20) or the system user interaction server (15). Once the user is known to the system, all Internet traffic from their gateway device (or, more accurately, the cable modem or DSL or fiber-to-the-home bridge) will be routed by the VSP/ISP through the system's "local" datacenter (30) and a pre-assigned IPPE (32). The system protects the home as a whole, even if someone brings in a computer that has never before had any interaction with the system, and has no software installed by or from the system. However, should the user desire further customization, such as per-user and/or per-computer settings on content management and filtering, the installation of either a client "shim" or a modified gateway (MG) (59) will allow such customizations in a simple and seamless manner.

b) Major advantages of the system over traditional Internet filters

[0158] Before discussing the two other technical system modes, and in addition to some of the technical advantages provided by the system as discussed above (such as "entire home" protection) and below (centralized management of all Internet-connected devices associated with the

user/home/account), it is important to pause to highlight the non-technical ways in which the system invention (in any mode) provides a major improvement in usability over all existing Internet content management and filtering technologies. There are two key factors that make the system much more powerful and easy to use than previous systems. One is the existence of Affinity Groups (AGs) and their relationship to the user and his/her family. The other is the connection to the user's video system.

[0159] The presence of AGs within the system allows for much more customized and desirable user experience. For example, consider a web page that is listed by the generic web filter as containing objectionable content. Sometimes web pages are listed as objectionable when they simply are not (simple false-positive). In other cases the web page would be objectionable to some members of society, but not to members of one or more AGs (context-based false positive). If a user of a standard Internet filter encounters such a page, he/she has no basis on which to know whether to unblock the page for his/her or own viewing, or his/her child's viewing. Thus, every blocked page presents a dilemma—"I think this page could be ok, but then why is it in the 'black list' database?"—without any context for resolution.

[0160] Now consider the same scenario within the system invention. First of all, as previously discussed, AGs can develop "white lists" that override the system-wide "black

list." From the start the system will be more customized to user preferences and expectations.

[0161] Secondly, the system grows more customized dynamically by simple use. The first time someone from a particular AG chooses to view a blocked page, that choice is noted in the usage database, and they are sent pop-up request asking them to rate the page and (optionally) describe why they found it useful, neutral, or harmful. Their response is logged by the system. After awhile a "collective intelligence" emerges from system usage. For example, the 100th user who encounters the blocked page will be told that "70% of members of your affinity group X clicked through this blocking page by entering the administrative password; of those, 10% were glad they did, 20% were neutral, and 70% regretted it. Click here to read their comments, sorted by most recent to oldest." In addition, aggregate behaviors will be reported to the staff of the AG and they can then take action to serve their users by either adding a URL to their "white list," or creating a special annotation for the page that explains to their users why or why not they might want to use the page or the entire site. The system will show that annotation in place of the standard block message when members of the AG encounter the questionable URL.

[0162] There is almost no limit to the kinds of user help and customization that can emerge from the system design, the aggregation of usage data, voting by users on content, and the continual feedback of usage and AG input into the system to

further refine and expand its recommendations and support for intelligent Internet and media choices by users.

[0163] The other major area of improvement brought about by the system invention comes from its ability to couple together video and Internet services into a single user experience. This begins with the fact that the user provides a set of information about himself/herself and his/her household as well as affinity group data that is shared across the two interconnected systems. The integration increases with usage, since choices made on the video system can have a positive, customizing impact on the behavior of the Internet system, and vice versa.

[0164] Integration is also key to improve reporting and feedback to the user. For example, the parents of a family can receive weekly usage reports on both video and Internet usage: how many hours the family spent watching TV, and when; the top-rated web sites and web searches, as well as the amount of time spent on instant messaging systems and web-based email, and/or the number of email messages sent through the system's integrated email system as well as arbitrary SMTP and POP/IMAP gateways. The system can also report suspicious outages: times when either the video management system and/or the Internet management system were not operational (under certain scenarios the system could be by-passed, such as a STB client connected to a traditional co-axial cable with analog channels, or the Internet system when configured to work in Enhanced Gateway or "shim-plus" mode). The system can also

signal to parents via an icon on the TV set when there are real-time problems with Internet usages, such as an unusual spike in block sites, or an unusual number of times the administrative password has been used to by-pass the filtering system.

c) Enhanced gateway and "shim-plus" modes

[0165] Returning to the technical architecture: the system has three main Internet usage modes. The first and most comprehensive mode--so-called "server-based mode"--was discussed above, along with a number of generic features and advantages of the system. Two other modes will now be discussed: enhanced gateway mode, and "shim-plus" mode. Unless noted here, all previously noted features and advantage of the system are carried over to these two modes of operation.

[0166] Server-based mode works in those circumstances in which the system is working in collaboration with the user's ISP--in other words, when all the user's Internet traffic can automatically be routed through the systems Internet Policy/Proxy Engine (IPPE). Some users of the system may not have an ISP that is working in cooperation with the system. This can happen when the user's VSP is not their ISP; or if the user is using the video system in on of its alternative modes of operation (see Sections (2) and (3) above).

[0167] In this case, the user will have the option of acquiring an enhanced gateway device, (EG) (69) which provides all the necessary functions for the system. The EG provides a number of additional functions beyond the simple computer/user

identification function provided by the modified gateway (MG) (59).

[0168] Like the MG, the EG will first contact the system data center (02) and after logging in securely, will receive the information needed to find and connect to the "closest" IPPE (32 or 16). From that point forward, the role previously played almost entirely by the IPPE will now be split between the EG and the IPPE. The EG can be thought of as some of the modules of the IPPE that have been split off and configured to run on a small computer acting as the network gateway for the home (60) that work in close partnership with the remaining modules of the IPPE on the server (32, 16).

[0169] Here are some of the functions of the IPPE now performed altogether or in part by the EG. In conjunction with the user interaction server (15) and the IPEE, the EG will manage system and user identification as well as session management. When the user interacts with the system via the web (15), the EG will transparently add information about the in-home LAN (68) as well as the connected computers indexed by MAC address. The user will be able to configure the settings for particular computer, and the EG will learn and remember the default settings for each. The EG will also provide session time-outs and request that user log in after a certain period of inactivity, or periodically even during regular activity. The EG will provide real-time scanning of email and instant messaging traffic, as well as web traffic to "dynamic" sites, and security scanning for virus, phishing schemes, and

so forth. The EG will not keep a full copy of the URL database (34), nor of all the meta-data provided by AGs. Instead, the EG will forward each URL request to the IPPE, and keep small FIFO cache of URLs and associated meta-data in-memory, with an age-out timestamp on each entry of a few hours, so that information remains fresh. Finally, the EG will upload regularly its local database of user interactions and user and system information so that the system can utilize and store this information reliably. If the EG, for example, somehow crashes or is reset to factory defaults, the system will download all of the previously "learned" configuration data to the EG when it logs in again.

[0170] The key thing to keep in mind about the EG architecture is that the EG is only making "look aside" calls to the IPPE. The bulk of Internet traffic up and down is going directly from the EG to the Internet. Video streams or other large downloads, for example, once approved by a "look aside" to the EG URL cache and, if necessary, a remote call to the IPPE, are simply streamed through the EG without any other intervention by or through the system.

[0171] The "shim-plus" mode provides a means by which a portable system such as a laptop (55, 65) can be configured from within the system to operate in a safe and integrated manner even when it is connected to the Internet from an arbitrary location (70) that is lacking both server-based and EG-based content management and filtering. Technically speaking, "shim-plus" mode is similar to EG mode, but in the

shim case the situation is simplified somewhat because the shim is tracking and managing the behavior of only a single system, not multiple systems on the LAN. But in other respects the shim acts very much like the EG, as a kind of distributed component of the IPPC.

[0172] More specifically, when installed on a PC (55, 56) the shim watches all local interactions between client applications (such as web browsers, email programs, instant-messaging programs, etc.) and the TCP/IP networking protocol stack. Some of those interactions would be intercepted temporarily and communicated to the IPPE (32), which would then return a result that blocks or permits the interaction. Others are permitted without intervention. Still others involve real-time filtering and analysis, such as the filtering of instant messages and email traffic.

[0173] In order to fit in with the remainder of the system in as seamless a mode as possible, the shim has no separate user interface. All configuration of the shim is done by configuring the machine and its users on the system web site (15), with all the configuration data stored in the user information database (10). This approach has several advantages. First of all, the account owner can track, configure, and reconfigure all computers associated with the account, regardless of their current location. For example, if a child takes a laptop to school, the parent can make configuration changes that automatically propagate to the laptop the next time it connects to the Internet. Going the

other direction, all usage information and incidents (such as blocked websites) are uploaded to the user account and integrated into a single viewing and reporting architecture for all computers associated with the account.

[0174] Finally, both the EG and "shim-plus" are designed so that if they are temporarily removed from the home or the computer, the system will recognize their absence. This "dog that didn't bark" kind of knowledge is of course ambiguous since there are many reasons why a device may have been off-line. Nevertheless, it will give to parents some assurance that the devices were in use when necessary, since a long gap in the connection times of the EG or the shim may be difficult or impossible to explain in light of other facts known to the parents.

[0175] It is to be understood that the present invention is not limited to the sole embodiment described above, but encompasses any and all embodiments within the scope of the following claims.

CLAIMS

I claim:

1. A method of customizing media content, comprising:

a video service provider broadcasting a video stream;

storing said video stream at a user premises;

at least one content editorial reviewer reviewing said video stream and indicating a portion of said video stream to be played at said user premises; and

playing back only said portions of said stored video stream at said user premises indicated by said content editorial reviewer.

2. The method of customizing media content of claim 1, wherein said content editorial review is at a location remote from said video service provider.

3. The method of customizing media content of claim 1, further comprising:

at least one affinity having a plurality of members;

each said user indicating to said video service provider an affinity group to which the user is a member, and

each said affinity group having an associated content editor indicating a portion of said video stream to be played at a premises for user belonging to said affinity group;

wherein each user premises has a playback device which receives said user indicated affinity group and plays back said portion of said video stream according to the indication of the content editor associated with the indicated affinity group.

4. The method of customizing media content of claim 1, further comprising:

receiving a program guide at the user premises indicating a first selection of video content available from the video service provider to the user premises;

displaying a portion of said program guide, wherein said displayed program guide does only displays said stored video stream at said user premises indicated by said content editorial reviewer

5. The method of customizing media content of claim 1, further comprising:

receiving and displaying a program guide in a first color at the user premises indicating a first selection of video content available from the video service provider to the user premises;

displaying a portion of said program guide indicating in a second color only said stored video stream indicated by said content editorial reviewer.

6. A method of customizing playback of video media content at a user's premises, comprising:

receiving at least one video media from a video media broadcast service and receiving a unique video media identifier for said at least one video media;

storing said at least one video media in an electronic storage device in said user premises;

receiving in said electronic storage device a video media metadata relating to a video media, said video media metadata including a unique metadata identifier identifying the unique video media identifier of the video media to which the video media metadata relates, said video media metadata further including time code instructions for displaying timed portions of the related video media on a video display or for preventing the display from displaying timed portions of the related video media on the video display;

selecting at least one video media to display on a video display at said user premises;

retrieving the unique video media identifier for said at least one video media selected for display;

retrieving a video media metadata having a metadata identifier relating to said video media identifier;

retrieving said at least one video media for display on the video display at said user premises, wherein said retrieved video media metadata causes at least a portion of said retrieved video media to not display on said video display.

7. The method of customizing playback of video media content of claim 6, wherein said video media broadcast service is selected from one of cable television and satellite television.

8. The method of customizing playback of video media content of claim 6, wherein said video media is selected from one of video on demand and MPEG.

9. The method of customizing playback of video media content of claim 6, further comprising:

providing a programming guide listing video content;

providing an indicator on said programming guide indicating whether a particular video content whether a particular video media listed in the guide contains metadata associated with said video media.

10. The method of customizing playback of video media content of claim 6, further comprising:

providing a programming guide listing video content;

providing an indicator on said programming guide indicating a particular video content includes a particular video media listed in the guide contains metadata associated with said video media, by displaying said listing in a first color;

providing an indicator on said programming guide indicating a particular video content does not include a particular video media listed in the guide contains metadata

associated with said video media, by displaying said listing in a second color different from said first color.

11. A method of targeting video content to the preferences of a viewer, comprising:

broadcasting a video segment by a first party;

reviewing the video segment by a second party and marking at least a first portion of the video segment with a first marker and marking at least a second portion of the video segment with a second marker;

receiving from a third party a first indication of a preference of the third party to view a first topic or subject matter;

receiving from the third party a second indication of a preference of the third party to not view a second topic or subject matter;

selectively playing back the video segment on a video display by a third party, where said playback skips at least the second portion of the video segment when said topic or subject matter of said second portion of the video segment matches said second indication second topic or subject matter; and

selectively indicating on the video display at least a first portion of said video display as being recommended for viewing by the third party when said topic or subject matter or said second portion of the video segment matches said first indication first topic or subject matter.

12. The method of targeting video content of claim 11, wherein said first party is a video service provider selected from the group of cable, satellite, broadcast or fiber optic television.

13. The method of targeting video content of claim 11, wherein said second party is an affinity group.

14. The method of targeting video content of claim 14, wherein said affinity group is a religious organization.

15. The method of targeting video content of claim 14, wherein said affinity group is a closed group.

16. The method of targeting video content of claim 14, wherein said video content includes at least one video segment created by said affinity group.

17. The method of targeting video content of claim 16, further comprising playing at least one affinity group created video content while said playback skips said second portion of said video segment.

18. The method of targeting video content of claim 17, wherein said played at least one affinity group created video content is the same length as said second portion of said video segment.

19. The method of targeting video content of claim 14, wherein said third party is an affinity group and said second party is an independent party contracted by said affinity group.

Fig. 1

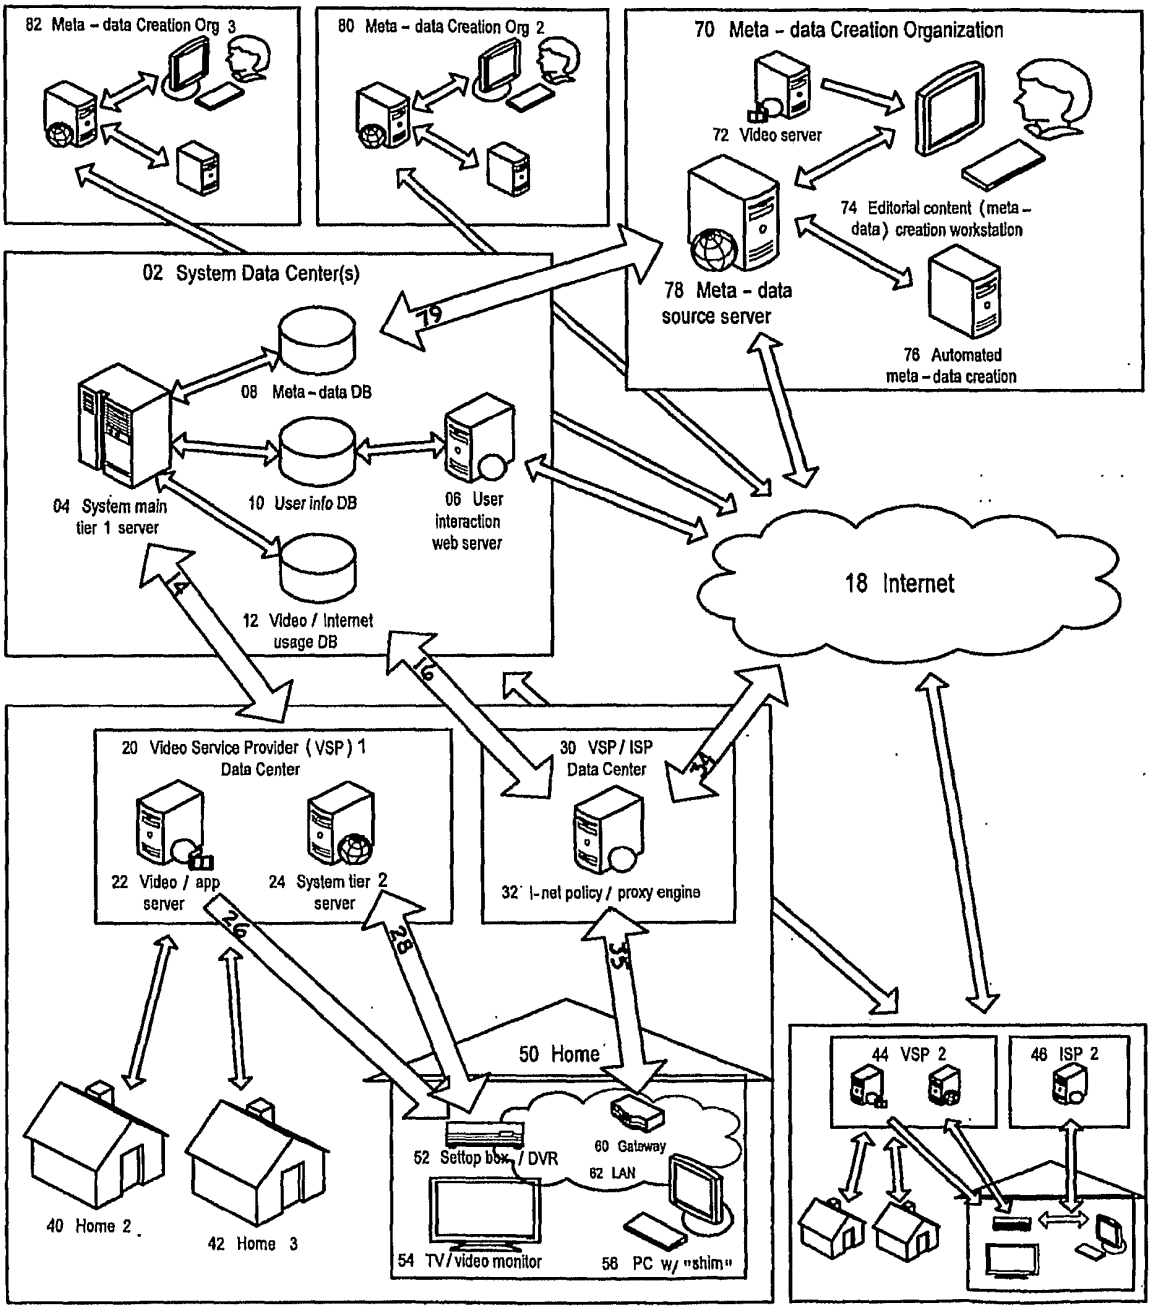

82 Meta – data Creation Org 3
80 Meta – data Creation Org 2
70 Meta – data Creation Organization
72 Video server
74 Editorial content (meta – data) creation workstation
78 Meta – data source server
76 Automated meta – data creation
02 System Data Center(s)
08 Meta – data DB
10 User info DB
04 System main tier 1 server
06 User interaction web server
12 Video / Internet usage DB
18 Internet
20 Video Service Provider (VSP) 1 Data Center
22 Video / app server
24 System tier 2 server
30 VSP / ISP Data Center
32 I-net policy / proxy engine
50 Home
52 Settop box / DVR
60 Gateway
62 LAN
54 TV / video monitor
58 PC w/ "shim"
40 Home 2
42 Home 3
44 VSP 2
46 ISP 2
79
14
16
26
28
35

Fig. 2

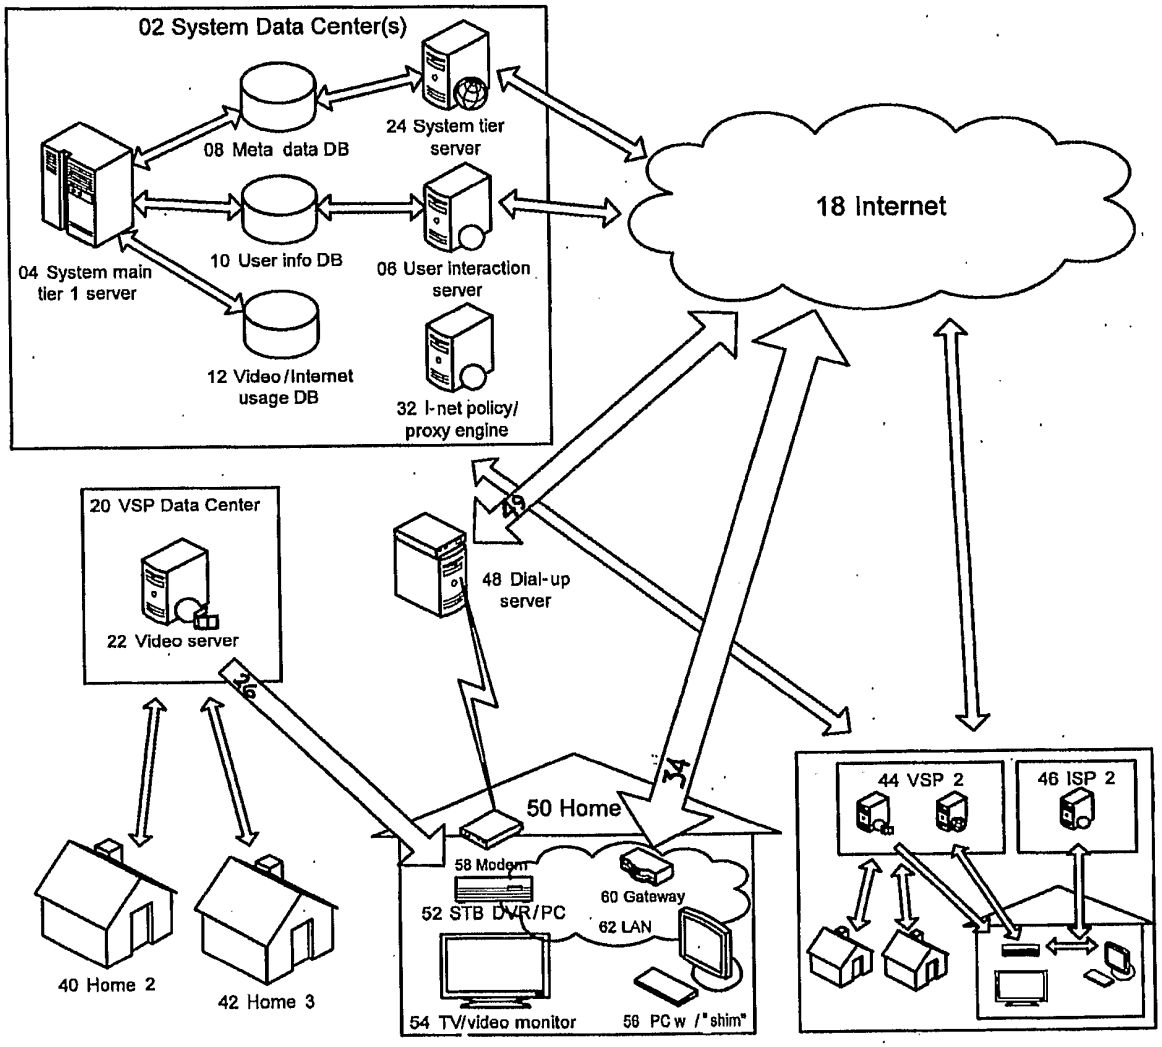

02 System Data Center(s)
08 Meta data DB
24 System tier server
10 User info DB
06 User interaction server
04 System main tier 1 server
12 Video/Internet usage DB
32 I-net policy/ proxy engine
18 Internet
20 VSP Data Center
22 Video server
48 Dial-up server
26
34
50 Home
58 Modem
52 STB DVR/PC
60 Gateway
62 LAN
54 TV/video monitor
56 PC w /"shim"
40 Home 2
42 Home 3
44 VSP 2
46 ISP 2

Fig. 3

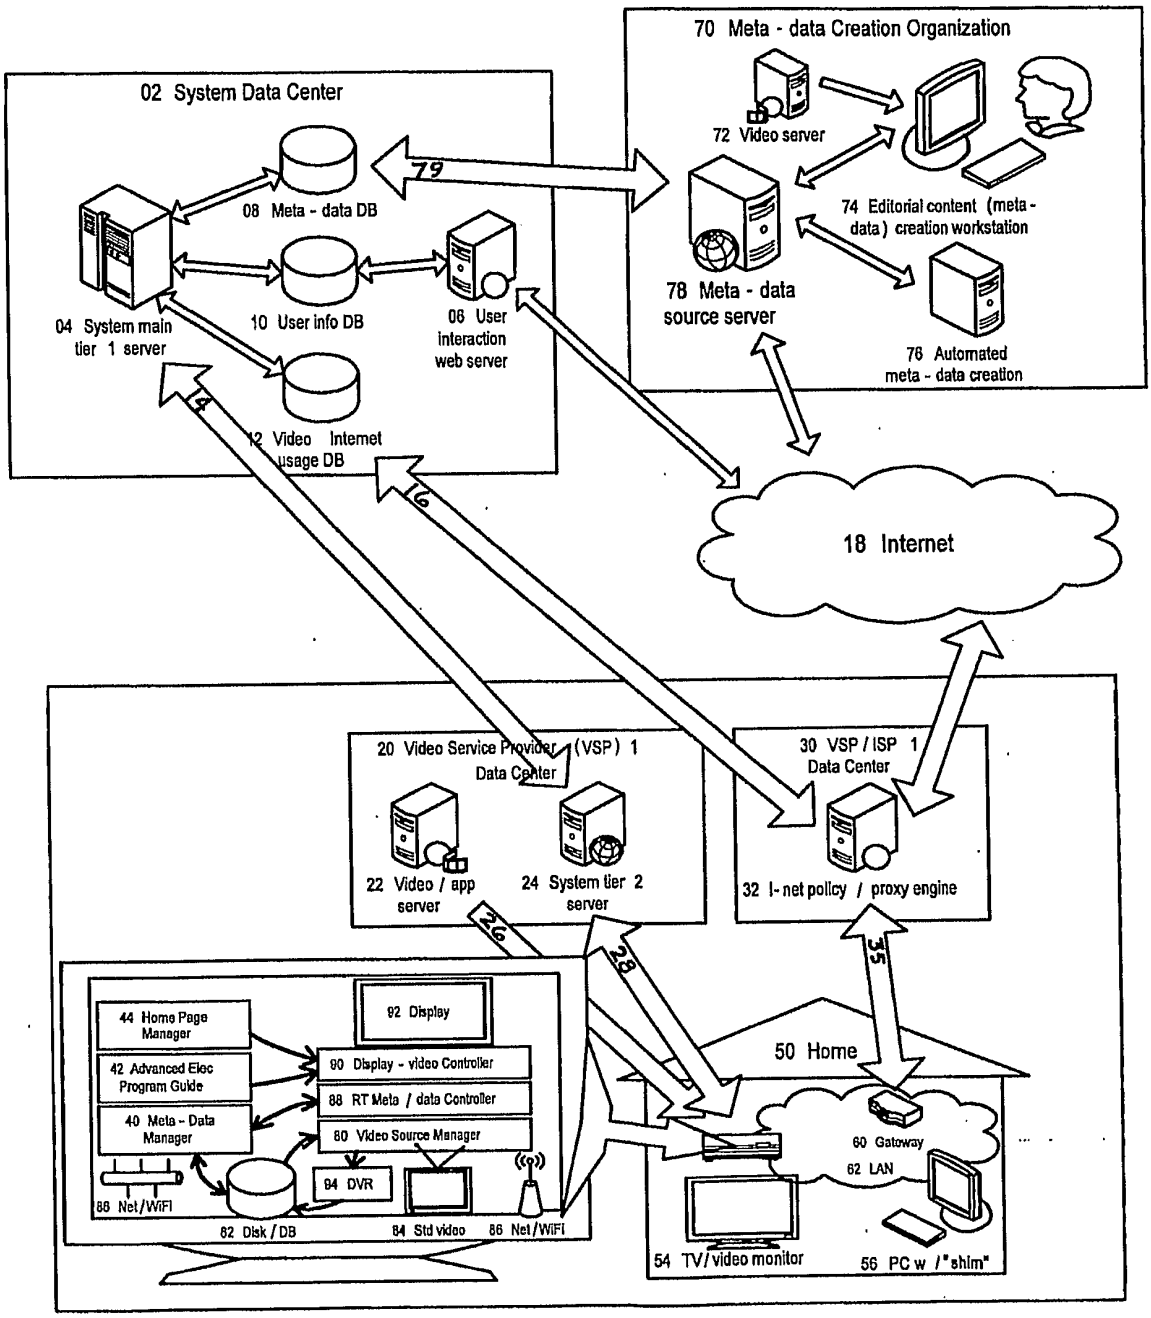

70 Meta - data Creation Organization
72 Video server
74 Editorial content (meta - data) creation workstation
78 Meta - data source server
76 Automated meta - data creation
02 System Data Center
08 Meta - data DB
79
04 System main tier 1 server
10 User info DB
06 User interaction web server
12 Video Internet usage DB
14
16
18 Internet
20 Video Service Provider (VSP) 1 Data Center
22 Video / app server
24 System tier 2 server
30 VSP / ISP 1 Data Center
32 I-net policy / proxy engine
26
28
35
44 Home Page Manager
42 Advanced Elec Program Guide
40 Meta - Data Manager
92 Display
90 Display - video Controller
88 RT Meta / data Controller
80 Video Source Manager
94 DVR
86 Net/WiFi
82 Disk / DB
84 Std video
86 Net/WiFi
50 Home
60 Gatoway
62 LAN
54 TV/video monitor
56 PC w / "shim"

Fig. 4

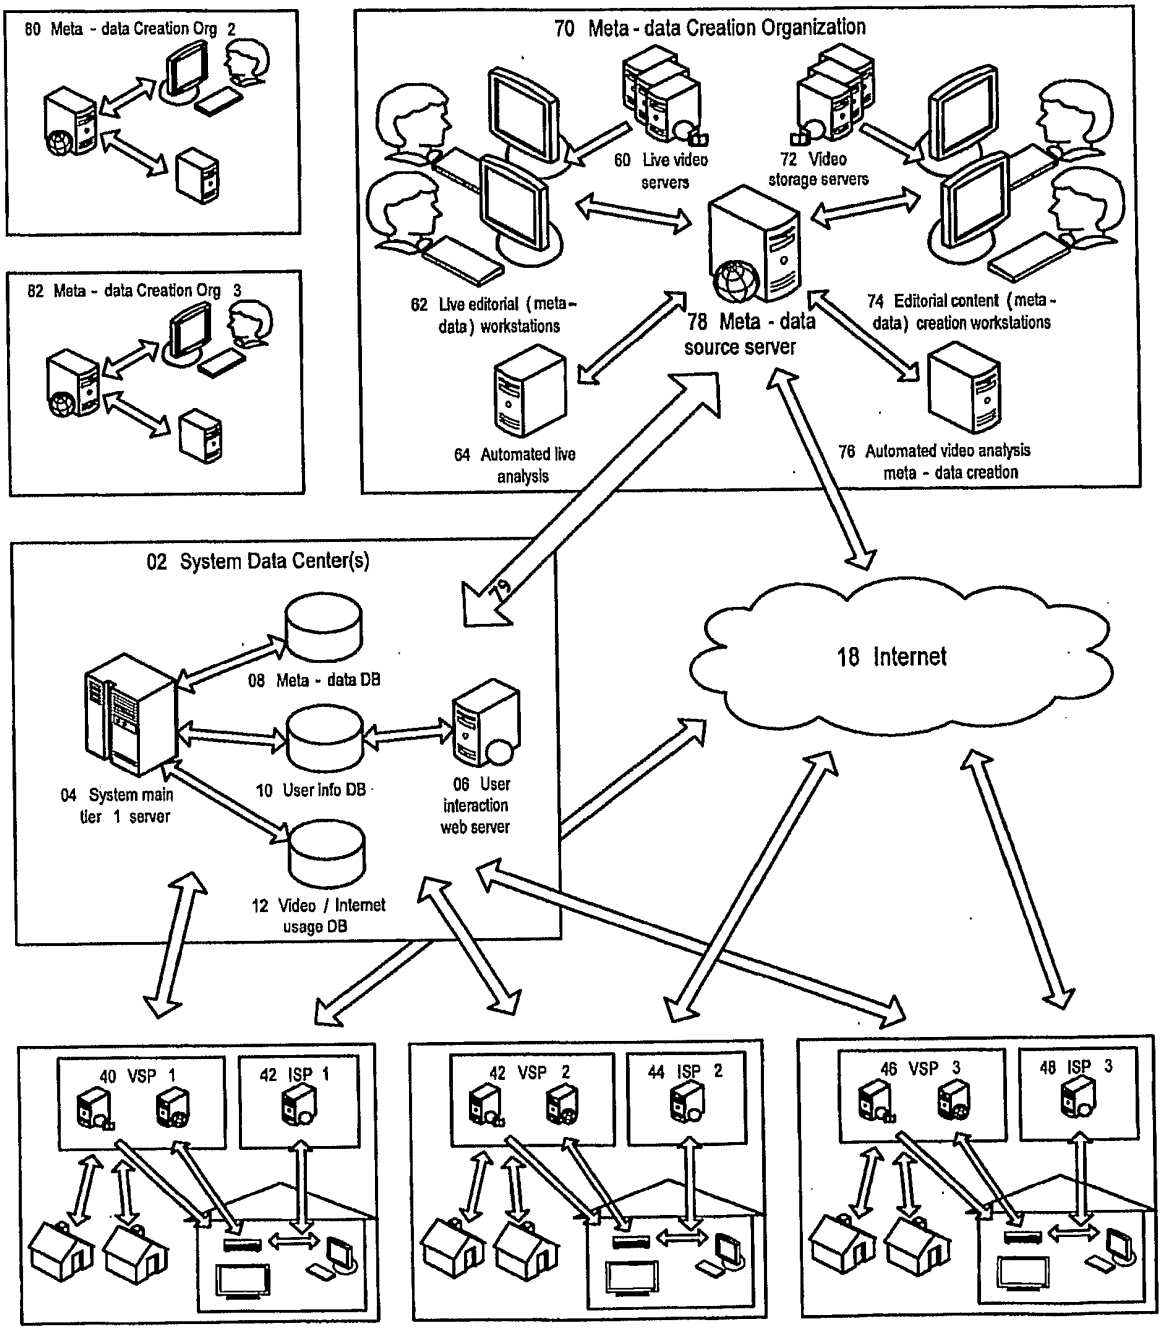
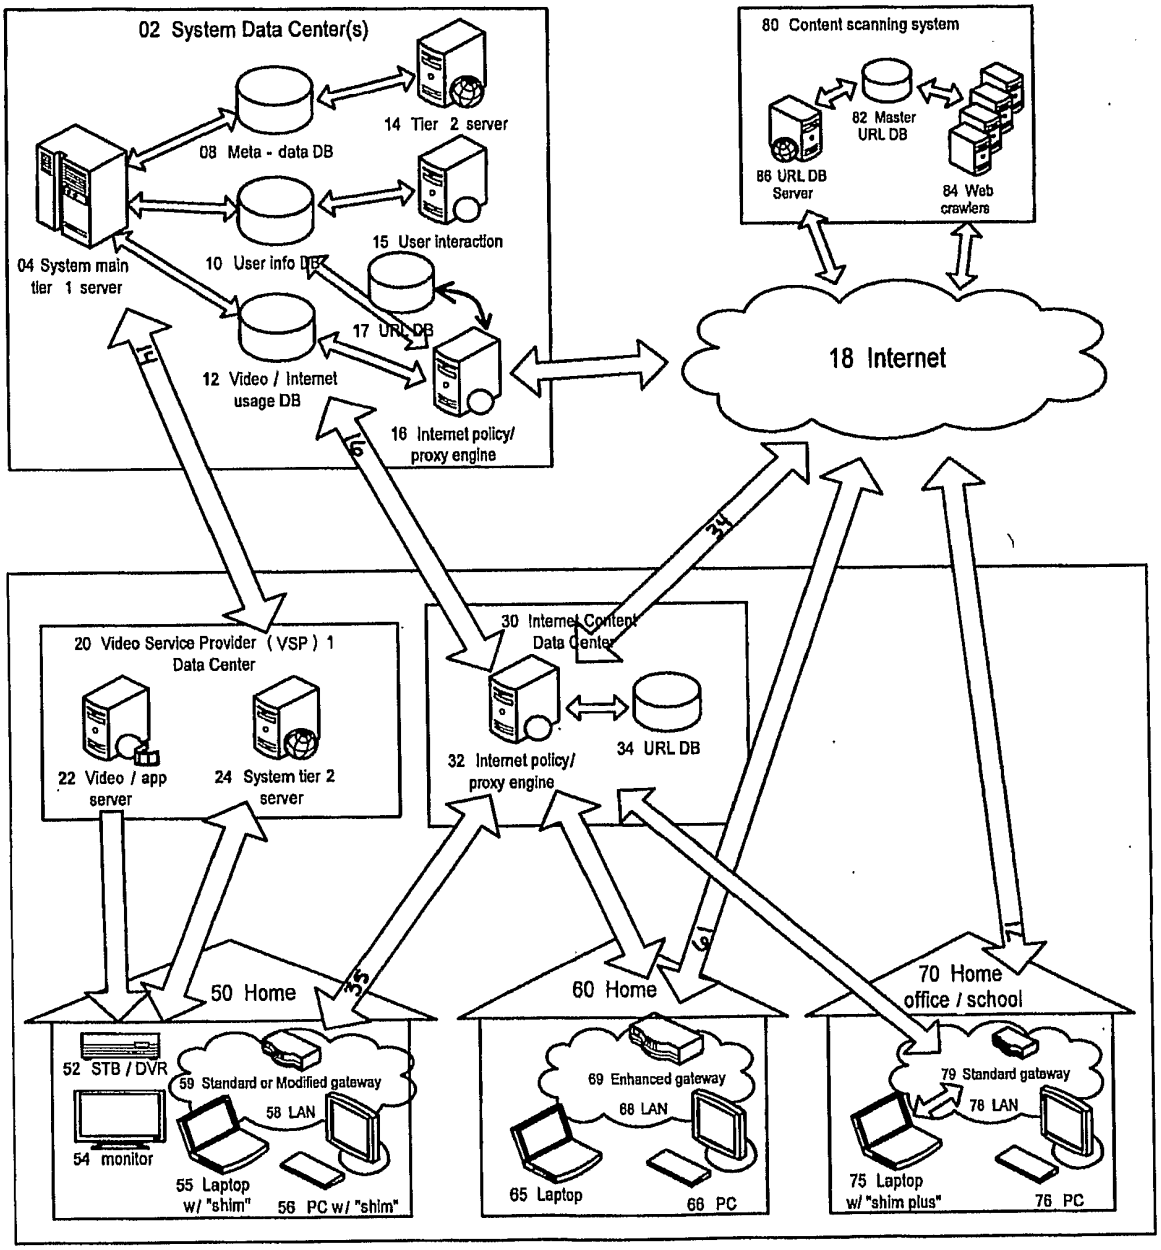

80 Meta - data Creation Org 2
82 Meta - data Creation Org 3
70 Meta - data Creation Organization
60 Live video servers
72 Video storage servers
62 Live editorial (meta-data) workstations
78 Meta - data source server
74 Editorial content (meta-data) creation workstations
64 Automated live analysis
76 Automated video analysis meta - data creation
02 System Data Center(s)
08 Meta - data DB
10 User info DB
06 User interaction web server
04 System main tier 1 server
12 Video / Internet usage DB
18 Internet
40 VSP 1
42 ISP 1
42 VSP 2
44 ISP 2
46 VSP 3
48 ISP 3

Fig. 5



SUBSTITUTE SHEET (RULE 26)

Meta-data type	RTSS	SNID1	SNID2	SSF1 & SSF2	SRA
UP	In the canonical case, an unknown broadcast or VOD; default is not to display; or, if displayed, not record	Unknown non-RTSS video stream; same	Same	Same	Same, except recording is n/a; and playback not allowed so long as no meta-data available
PND Recording unknown content for later playback	Known broadcast or VOD; default is not to display; if displayed, recording allowed	Known non-RTSS video stream; same	Same	Same	Same, except recording is n/a; and playback not allowed so long as no meta-data available
PRIU Content with overall rating but no detailed meta-data; no insertion meta-data	Rated broadcast; displayed and played based on user preferences (by default insertions deemed compatible); VOD typically does not have inserts, but if so treated the same as broadcast	Known non-RTSS video stream; playback (if allowed) begins soon after user selection	Same; playback (if allowed) begins after a calculated delay	Same, depending on download speed	Same as RTSS case
PRIR	Rated broadcast and rated insertions, both must be compatible with user preferences	Rated non-RTSS and rated insertions; playback (if allowed) begins soon after user selection	Same; playback (if allowed) begins after a calculated delay	Same, depending on download speed	Same as RTSS case
PEIU Re-runs of previously broadcast TV programs; also, all VOD without inserts	By default play not allowed, since program requires edits and insertions are unknown (can be overridden by user preferences, the result is then like the PEIR case)	Same; playback (if allowed) begins soon after user selection	Same; playback (if allowed) begins after a calculated delay	Same, depending on download speed (but see next row)	Same as RTSS case, except that all editing behavior readily implemented in real-time playback

FIGURE 6A

Meta-data type	RTSS	SNID1	SNID2	SSF1 & SSF2	SRA
PEIR	Edited broadcast and rated insertions (or in the case of typical VOD, edited stream with no insertions) must meet user settings; if there is a "skip-required" field, the start of the program will be delayed; if "skip optional," the system will delay the start of viewing to allow for smoother edited playback (else customized or standard substitution will be used); system may attempt to rebuild buffer after skip by slight slowing of playback	Edited non-RTSS and rated insertions must meet user settings; if there are "skip" fields, those will be used as part of start-time calculation	Same (longer calculated delay)	Same, except that "skip" fields can be implemented on the stream itself (at some level of granularity) by instructions to SSF server, allowing quicker start-up in both cases	Same as RTSS case, except that all editing behavior readily implemented in real-time playback
PEIE Recorded programs and inserts played back after meta-data creation	Same as above, but now insertions as well as program data enter in to the various calculations; this is a the canonical scenario for "off the air" program recording that are played back later after meta-data has been created	Same as above and left	Same (longer calculated delay)	Same as above	Same as RTSS case, except that all editing behavior readily implemented in real-time playback; the canonical example of this mode is a broadcast TV show with commercials previously recorded to disk, reviewed soon after by editors to create the PEIE meta-data, and then played back (for example) 24 hours later
PRLE Live programming	Similar to PEIE case, but special considerations apply: meta-data is being created on-the-fly by editorial team (see Section (8) below), so delay calculations must be adjusted accordingly; meta-data is not all present at beginning, so will utilize fail-safe mode (if no meta-data received in a timely fashion, stop playback).	N/a	Likely n/a, but theoretically possible; longer start delay for possible incidents and "skips" must be estimated	SSF1 n/a; SSF2 unlikely, but theoretically possible; longer start delay must be estimated (meta-data arriving in semi-real-time could be used for skip-ahead)	N/a (becomes PEIE case on playback, assuming complete set of meta-data has been captured and distributed throughout the system)

FIGURE 6B

Electronic Acknowledgement Receipt

EFS ID:	47121030
Application Number:	17714965
International Application Number:	
Confirmation Number:	7004
Title of Invention:	Targeted television advertisements based on online behavior
First Named Inventor/Applicant Name:	Roy Shkedi
Customer Number:	26362
Filer:	Louis J. Hoffman/Michelle Lisa
Filer Authorized By:	Louis J. Hoffman
Attorney Docket Number:	14-TV1-CON7
Receipt Date:	30-NOV-2022
Filing Date:	06-APR-2022
Time Stamp:	20:59:34
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	22-11-30-IIQ-TV1-CON7-Supplemental-IDS.pdf	446882 433d003c38ffb78a3478f008cb5fefd779c7dc48	no	4

Warnings:

Information:					
This is not an USPTO supplied IDS fillable form					
2	Foreign Reference	WO2007117613A2-1.pdf	3911740 5794f609f272c60cbc0cf09d7fc6bfd04834d54b	no	105
Warnings:					
Information:					
			Total Files Size (in bytes):	4358622	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

Doc Code: DIST.E.FILE **Document Description: Electronic Terminal Disclaimer - Filed**	PTO/SB/26 U.S. Patent and Trademark Office Department of Commerce

Electronic Petition Request	**TERMINAL DISCLAIMER TO OBVIATE A DOUBLE PATENTING REJECTION OVER A "PRIOR" PATENT**
Application Number	17714965
Filing Date	06-Apr-2022
First Named Inventor	Roy Shkedi
Attorney Docket Number	14-TV1-CON7
Title of Invention	Targeted television advertisements based on online behavior

☒ Filing of terminal disclaimer does not obviate requirement for response under 37 CFR 1.111 to outstanding Office Action

☒ This electronic Terminal Disclaimer is not being used for a Joint Research Agreement.

Owner	Percent Interest
Intent IQ, LLC	100%

The owner(s) with percent interest listed above in the instant application hereby disclaims, except as provided below, the terminal part of the statutory term of any patent granted on the instant application which would extend beyond the expiration date of the full statutory term of prior patent number(s)

10715878

as the term of said prior patent is presently shortened by any terminal disclaimer. The owner hereby agrees that any patent so granted on the instant application shall be enforceable only for and during such period that it and the prior patent are commonly owned. This agreement runs with any patent granted on the instant application and is binding upon the grantee, its successors or assigns.

In making the above disclaimer, the owner does not disclaim the terminal part of the term of any patent granted on the instant application that would extend to the expiration date of the full statutory term of the prior patent, "as the term of said prior patent is presently shortened by any terminal disclaimer," in the event that said prior patent later:
- expires for failure to pay a maintenance fee;
- is held unenforceable;
- is found invalid by a court of competent jurisdiction;
- is statutorily disclaimed in whole or terminally disclaimed under 37 CFR 1.321;
- has all claims canceled by a reexamination certificate;
- is reissued; or
- is in any manner terminated prior to the expiration of its full statutory term as presently shortened by any terminal disclaimer.

◉ Terminal disclaimer fee under 37 CFR 1.20(d) is included with Electronic Terminal Disclaimer request.

○ I certify, in accordance with 37 CFR 1.4(d)(4), that the terminal disclaimer fee under 37 CFR 1.20(d) required for this terminal disclaimer has already been paid in the above-identified application.	
Applicant claims the following fee status: ○ Small Entity ○ Micro Entity ◉ Regular Undiscounted I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.	
THIS PORTION MUST BE COMPLETED BY THE SIGNATORY OR SIGNATORIES I certify, in accordance with 37 CFR 1.4(d)(4) that I am: ◉ An attorney or agent registered to practice before the Patent and Trademark Office who is of record in this application Registration Number 38918 ○ A sole inventor ○ A joint inventor; I certify that I am authorized to sign this submission on behalf of all of the inventors as evidenced by the power of attorney in the application ○ A joint inventor; all of whom are signing this request	
Signature	/Louis J. Hoffman/
Name	Louis Hoffman

*Statement under 37 CFR 3.73(b) is required if terminal disclaimer is signed by the assignee (owner).
Form PTO/SB/96 may be used for making this certification. See MPEP § 324.

Electronic Patent Application Fee Transmittal

Application Number:	17714965
Filing Date:	06-Apr-2022
Title of Invention:	Targeted television advertisements based on online behavior
First Named Inventor/Applicant Name:	Roy Shkedi
Filer:	Louis J. Hoffman
Attorney Docket Number:	14-TV1-CON7

Filed as Large Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
STATUTORY OR TERMINAL DISCLAIMER	1814	1	170	170
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Extension-of-Time:				
Miscellaneous:				
			Total in USD ($)	**170**

Doc Code: DISQ.E.FILE
Document Description: Electronic Terminal Disclaimer – Approved

Application No.: 17714965

Filing Date: 06-Apr-2022

Applicant/Patent under Reexamination: Shkedi

Electronic Terminal Disclaimer filed on December 1, 2022

☒ APPROVED

This patent is subject to a terminal disclaimer

☐ DISAPPROVED

Approved/Disapproved by: Electronic Terminal Disclaimer automatically approved by EFS-Web

U.S. Patent and Trademark Office

Electronic Acknowledgement Receipt	
EFS ID:	47124976
Application Number:	17714965
International Application Number:	
Confirmation Number:	7004
Title of Invention:	Targeted television advertisements based on online behavior
First Named Inventor/Applicant Name:	Roy Shkedi
Customer Number:	26362
Filer:	Louis J. Hoffman
Filer Authorized By:	
Attorney Docket Number:	14-TV1-CON7
Receipt Date:	01-DEC-2022
Filing Date:	06-APR-2022
Time Stamp:	13:27:06
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	CARD
Payment was successfully received in RAM	$170
RAM confirmation Number	E2022B1D27022591
Deposit Account	
Authorized User	
The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:	

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Terminal Disclaimer-Filed (Electronic)	eTerminal-Disclaimer.pdf	41278 1d5db08c9054c0f0beebb4d6fb81024856bc496a	no	2
Warnings:					
Information:					
2	Fee Worksheet (SB06)	fee-info.pdf	37974 68519cb47eca32c975447b6f090e35239b465790	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	79252		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.
National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.
New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor	:	Shkedi, Roy	Art Unit	:	2425
Serial No.	:	17/714,965	Examiner	:	Ekpo, Nnenna
Filing Date	:	04/06/2022	Conf. No.	:	7004
Title	:	Targeted television advertisements based on online behavior			

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Filed via EFS-Web on December 1, 2022

INTERVIEW REPORT

Dear Sir:

1. The examiner telephoned the undersigned this morning to report that the case was in condition for allowance provided that applicant filed a terminal disclaimer over co-owned parent patent USP 10,715,878. No rejection has been made for obviousness-type double-patenting, and the examiner did not provide any explanation of why the claims were thought to be obvious over the prior patent. Nevertheless, in the interview, applicant's attorney indicated that the examiner's request should not be an issue.

Parent U.S. Patent 10,715,878 is commonly owned, and a patent issuing from the instant application is expected to expire on the same day as that parent patent. In view of those facts, it does not seem necessary to consider the relationship between the claim sets. Assignee is filing concurrently an eTerminal Disclaimer, for automatic processing and immediate approval through EFS-Web. By filing a terminal disclaimer to simplify and expedite allowance, assignee does not concede that the rejected claims are patentably indistinct from the claims of the parent patent or that they contain parallel, inherent, or related limitations. *See* MPEP 804.02(II) ("The filing of a terminal disclaimer to obviate a rejection based on nonstatutory double patenting is not an admission of the propriety of the rejection.") (citing *Quad Envtl. Techs. Corp. v. Union Sanitary Dist.*, 946

F.2d 870, 874 (Fed. Cir. 1991) (the "filing of a terminal disclaimer simply serves the statutory function of removing the rejection of double patenting, and raises neither a presumption nor estoppel on the merits of the rejection.")).

2. In the same call, the undersigned also brought to the examiner's attention that applicant had filed yesterday two Information Disclosure Statements. One was the replacement form PTO-1449 discussed previously (*see* interview report filed yesterday), which contains references that the examiner had reviewed in parent applications. The other, the undersigned told the examiner, was a Supplemental Information Disclosure Statement (*see* PTO form PTO/SB/08a filed yesterday), which contains three new references. The undersigned indicated to the examiner that the three new references should be considered not only in connection with this application but also in connection with copending application Serial No. 17/877,870, because the '870 application contains counterpart method claims and because this application is a parent of the '870 application. *See* MPEP 609.02 ("The examiner of the continuing application will consider information which has been considered by the Office in the parent application."); MPEP 707 ("In all continuation … applications, the parent applications should be reviewed for pertinent prior art"); MPEP 904 ("In all continuing applications, the parent applications must be reviewed by the examiner for pertinent prior art. The examiner must consider prior art which was cited and considered in the parent application."); MPEP 2001.06(b) ("If the application under examination is identified as a continuation … of an earlier application, the examiner will consider the prior art properly cited in the earlier application.… Accordingly, no separate citation of the same prior art need be made in the later application, unless applicant wants a listing of the prior art printed on the face of the patent."). Given that the '870 application has been allowed, applicant does not wish to delay issuance just to have these three references printed on the face of the patent that will issue from the '870 application.

The examiner indicated that she would study the two IDS filings and the references cited therein.

The undersigned and applicant thank the examiner for the careful examination and consideration of this application, as well as for the fast action. Please feel free to telephone the undersigned if it would in any way advance completion of this application.

Respectfully submitted,

INTENT IQ, LLC

by its attorney

Dated: December 1, 2022

/Louis J. Hoffman/
Louis J. Hoffman
Reg. No. 38,918

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, Arizona 85260
(480) 948-3295

Electronic Acknowledgement Receipt	
EFS ID:	47125064
Application Number:	17714965
International Application Number:	
Confirmation Number:	7004
Title of Invention:	Targeted television advertisements based on online behavior
First Named Inventor/Applicant Name:	Roy Shkedi
Customer Number:	26362
Filer:	Louis J. Hoffman
Filer Authorized By:	
Attorney Docket Number:	14-TV1-CON7
Receipt Date:	01-DEC-2022
Filing Date:	06-APR-2022
Time Stamp:	13:29:20
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Applicant summary of interview with examiner	22-12-01-IIQ-TV1_CON7-Interview_Report2.pdf	96093 9b2fdf5e56089612f6c547f7fd6e84f0f4834c0505	no	3

Warnings:

Information:	
Total Files Size (in bytes):	96093

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

NOTICE OF ALLOWANCE AND FEE(S) DUE

26362 7590 12/15/2022

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, AZ 85260

EXAMINER
EKPO, NNENNA NGOZI

ART UNIT	PAPER NUMBER
2425	

DATE MAILED: 12/15/2022

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
17/714,965	04/06/2022	Roy Shkedi	14-TV1-CON7	7004

TITLE OF INVENTION: Targeted television advertisements based on online behavior

APPLN. TYPE	ENTITY STATUS	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	UNDISCOUNTED	$1200	$0.00	$0.00	$1200	03/15/2023

THE APPLICATION IDENTIFIED ABOVE HAS BEEN EXAMINED AND IS ALLOWED FOR ISSUANCE AS A PATENT. PROSECUTION ON THE MERITS IS CLOSED. THIS NOTICE OF ALLOWANCE IS NOT A GRANT OF PATENT RIGHTS. THIS APPLICATION IS SUBJECT TO WITHDRAWAL FROM ISSUE AT THE INITIATIVE OF THE OFFICE OR UPON PETITION BY THE APPLICANT. SEE 37 CFR 1.313 AND MPEP 1308.

THE ISSUE FEE AND PUBLICATION FEE (IF REQUIRED) MUST BE PAID WITHIN THREE MONTHS FROM THE MAILING DATE OF THIS NOTICE OR THIS APPLICATION SHALL BE REGARDED AS ABANDONED. THIS STATUTORY PERIOD CANNOT BE EXTENDED. SEE 35 U.S.C. 151. THE ISSUE FEE DUE INDICATED ABOVE DOES NOT REFLECT A CREDIT FOR ANY PREVIOUSLY PAID ISSUE FEE IN THIS APPLICATION. IF AN ISSUE FEE HAS PREVIOUSLY BEEN PAID IN THIS APPLICATION (AS SHOWN ABOVE), THE RETURN OF PART B OF THIS FORM WILL BE CONSIDERED A REQUEST TO REAPPLY THE PREVIOUSLY PAID ISSUE FEE TOWARD THE ISSUE FEE NOW DUE.

HOW TO REPLY TO THIS NOTICE:

I. Review the ENTITY STATUS shown above. If the ENTITY STATUS is shown as SMALL or MICRO, verify whether entitlement to that entity status still applies.

If the ENTITY STATUS is the same as shown above, pay the TOTAL FEE(S) DUE shown above.

If the ENTITY STATUS is changed from that shown above, on PART B - FEE(S) TRANSMITTAL, complete section number 5 titled "Change in Entity Status (from status indicated above)".

For purposes of this notice, small entity fees are 1/2 the amount of undiscounted fees, and micro entity fees are 1/2 the amount of small entity fees.

II. PART B - FEE(S) TRANSMITTAL, or its equivalent, must be completed and returned to the United States Patent and Trademark Office (USPTO) with your ISSUE FEE and PUBLICATION FEE (if required). If you are charging the fee(s) to your deposit account, section "4b" of Part B - Fee(s) Transmittal should be completed. If an equivalent of Part B is filed, a request to reapply a previously paid issue fee must be clearly made, and delays in processing may occur due to the difficulty in recognizing the paper as an equivalent of Part B.

III. All communications regarding this application must give the application number. Please direct all communications prior to issuance to Mail Stop ISSUE FEE unless advised to the contrary.

IMPORTANT REMINDER: Maintenance fees are due in utility patents issuing on applications filed on or after Dec. 12, 1980. It is patentee's responsibility to ensure timely payment of maintenance fees when due. More information is available at www.uspto.gov/PatentMaintenanceFees.

PART B - FEE(S) TRANSMITTAL

Complete and send this form, together with applicable fee(s), by mail or fax, or via EFS-Web.

By mail, send to: Mail Stop ISSUE FEE
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

By fax, send to: (571)-273-2885

INSTRUCTIONS: This form should be used for transmitting the ISSUE FEE and PUBLICATION FEE (if required). Blocks 1 through 5 should be completed where appropriate. All further correspondence including the Patent, advance orders and notification of maintenance fees will be mailed to the current correspondence address as indicated unless corrected below or directed otherwise in Block 1, by (a) specifying a new correspondence address; and/or (b) indicating a separate "FEE ADDRESS" for maintenance fee notifications.

CURRENT CORRESPONDENCE ADDRESS (Note: Use Block 1 for any change of address)

26362 7590 12/15/2022

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, AZ 85260

Note: A certificate of mailing can only be used for domestic mailings of the Fee(s) Transmittal. This certificate cannot be used for any other accompanying papers. Each additional paper, such as an assignment or formal drawing, must have its own certificate of mailing or transmission.

Certificate of Mailing or Transmission

I hereby certify that this Fee(s) Transmittal is being deposited with the United States Postal Service with sufficient postage for first class mail in an envelope addressed to the Mail Stop ISSUE FEE address above, or being transmitted to the USPTO via EFS-Web or by facsimile to (571) 273-2885, on the date below.

(Typed or printed name)
(Signature)
(Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
17/714,965	04/06/2022	Roy Shkedi	14-TV1-CON7	7004

TITLE OF INVENTION: Targeted television advertisements based on online behavior

APPLN. TYPE	ENTITY STATUS	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	UNDISCOUNTED	$1200	$0.00	$0.00	$1200	03/15/2023

EXAMINER	ART UNIT	CLASS-SUBCLASS
EKPO, NNENNA NGOZI	2425	725-034000

1. Change of correspondence address or indication of "Fee Address" (37 CFR 1.363).

❑ Change of correspondence address (or Change of Correspondence Address form PTO/AIA/122 or PTO/SB/122) attached.

❑ "Fee Address" indication (or "Fee Address" Indication form PTO/AIA/47 or PTO/SB/47; Rev 03-02 or more recent) attached. **Use of a Customer Number is required.**

2. For printing on the patent front page, list
(1) The names of up to 3 registered patent attorneys or agents OR, alternatively,
(2) The name of a single firm (having as a member a registered attorney or agent) and the names of up to 2 registered patent attorneys or agents. If no name is listed, no name will be printed.

1 _____
2 _____
3 _____

3. ASSIGNEE NAME AND RESIDENCE DATA TO BE PRINTED ON THE PATENT (print or type)

PLEASE NOTE: Unless an assignee is identified below, no assignee data will appear on the patent. If an assignee is identified below, the document must have been previously recorded, or filed for recordation, as set forth in 37 CFR 3.11 and 37 CFR 3.81(a). Completion of this form is NOT a substitute for filing an assignment.

(A) NAME OF ASSIGNEE (B) RESIDENCE: (CITY and STATE OR COUNTRY)

Please check the appropriate assignee category or categories (will not be printed on the patent) : ❑ Individual ❑ Corporation or other private group entity ❑ Government

4a. Fees submitted: ❑ Issue Fee ❑ Publication Fee (if required) ❑ Advance Order - # of Copies _____

4b. Method of Payment: *(Please first reapply any previously paid fee shown above)*

❑ Electronic Payment via EFS-Web ❑ Enclosed check ❑ Non-electronic payment by credit card (Attach form PTO-2038)

❑ The Director is hereby authorized to charge the required fee(s), any deficiency, or credit any overpayment to Deposit Account No. _____

5. **Change in Entity Status** (from status indicated above)

❑ Applicant certifying micro entity status. See 37 CFR 1.29

❑ Applicant asserting small entity status. See 37 CFR 1.27

❑ Applicant changing to regular undiscounted fee status.

NOTE: Absent a valid certification of Micro Entity Status (see forms PTO/SB/15A and 15B), issue fee payment in the micro entity amount will not be accepted at the risk of application abandonment.
NOTE: If the application was previously under micro entity status, checking this box will be taken to be a notification of loss of entitlement to micro entity status.
NOTE: Checking this box will be taken to be a notification of loss of entitlement to small or micro entity status, as applicable.

NOTE: This form must be signed in accordance with 37 CFR 1.31 and 1.33. See 37 CFR 1.4 for signature requirements and certifications.

Authorized Signature _____ Date _____

Typed or printed name _____ Registration No. _____

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
17/714,965	04/06/2022	Roy Shkedi	14-TV1-CON7	7004

26362 7590 12/15/2022

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, AZ 85260

EXAMINER	
EKPO, NNENNA NGOZI	
ART UNIT	PAPER NUMBER
2425	

DATE MAILED: 12/15/2022

Determination of Patent Term Adjustment under 35 U.S.C. 154 (b)
(Applications filed on or after May 29, 2000)

The Office has discontinued providing a Patent Term Adjustment (PTA) calculation with the Notice of Allowance.

Section 1(h)(2) of the AIA Technical Corrections Act amended 35 U.S.C. 154(b)(3)(B)(i) to eliminate the requirement that the Office provide a patent term adjustment determination with the notice of allowance. See Revisions to Patent Term Adjustment, 78 Fed. Reg. 19416, 19417 (Apr. 1, 2013). Therefore, the Office is no longer providing an initial patent term adjustment determination with the notice of allowance. The Office will continue to provide a patent term adjustment determination with the Issue Notification Letter that is mailed to applicant approximately three weeks prior to the issue date of the patent, and will include the patent term adjustment on the patent. Any request for reconsideration of the patent term adjustment determination (or reinstatement of patent term adjustment) should follow the process outlined in 37 CFR 1.705.

Any questions regarding the Patent Term Extension or Adjustment determination should be directed to the Office of Patent Legal Administration at (571)-272-7702. Questions relating to issue and publication fee payments should be directed to the Customer Service Center of the Office of Patent Publication at 1-(888)-786-0101 or (571)-272-4200.

OMB Clearance and PRA Burden Statement for PTOL-85 Part B

The Paperwork Reduction Act (PRA) of 1995 requires Federal agencies to obtain Office of Management and Budget approval before requesting most types of information from the public. When OMB approves an agency request to collect information from the public, OMB (i) provides a valid OMB Control Number and expiration date for the agency to display on the instrument that will be used to collect the information and (ii) requires the agency to inform the public about the OMB Control Number's legal significance in accordance with 5 CFR 1320.5(b).

The information collected by PTOL-85 Part B is required by 37 CFR 1.311. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 30 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, Virginia 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, Virginia 22313-1450. Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Notice of Allowability	Application No. 17/714,965	Applicant(s) Shkedi, Roy	
	Examiner NNENNA N EKPO	Art Unit 2425	AIA (FITF) Status No

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address--

All claims being allowable, PROSECUTION ON THE MERITS IS (OR REMAINS) CLOSED in this application. If not included herewith (or previously mailed), a Notice of Allowance (PTOL-85) or other appropriate communication will be mailed in due course. **THIS NOTICE OF ALLOWABILITY IS NOT A GRANT OF PATENT RIGHTS.** This application is subject to withdrawal from issue at the initiative of the Office or upon petition by the applicant. See 37 CFR 1.313 and MPEP 1308.

1. ☐ This communication is responsive to _____.
   ☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____.

2. ☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on _____; the restriction requirement and election have been incorporated into this action.

3. ☑ The allowed claim(s) is/are 52-77 . As a result of the allowed claim(s), you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information, please see **http://www.uspto.gov/patents/init_events/pph/index.jsp** or send an inquiry to **PPHfeedback@uspto.gov.**

4. ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

   Certified copies:

   a) ☐ All b) ☐ Some* c) ☐ None of the:

   1. ☐ Certified copies of the priority documents have been received.
   2. ☐ Certified copies of the priority documents have been received in Application No. _____ .
   3. ☐ Copies of the certified copies of the priority documents have been received in this national stage application from the International Bureau (PCT Rule 17.2(a)).

   * Certified copies not received: _____ .

Applicant has THREE MONTHS FROM THE "MAILING DATE" of this communication to file a reply complying with the requirements noted below. Failure to timely comply will result in ABANDONMENT of this application.
THIS THREE-MONTH PERIOD IS NOT EXTENDABLE.

5. ☐ CORRECTED DRAWINGS (as "replacement sheets") must be submitted.
   ☐ including changes required by the attached Examiner's Amendment / Comment or in the Office action of Paper No./Mail Date _____ .

   Identifying indicia such as the application number (see 37 CFR 1.84(c)) should be written on the drawings in the front (not the back) of each sheet. Replacement sheet(s) should be labeled as such in the header according to 37 CFR 1.121(d).

6. ☐ DEPOSIT OF and/or INFORMATION about the deposit of BIOLOGICAL MATERIAL must be submitted. Note the attached Examiner's comment regarding REQUIREMENT FOR THE DEPOSIT OF BIOLOGICAL MATERIAL.

Attachment(s)

1. ☑ Notice of References Cited (PTO-892)
2. ☑ Information Disclosure Statements (PTO/SB/08), Paper No./Mail Date See Continuation Sheet.
3. ☐ Examiner's Comment Regarding Requirement for Deposit of Biological Material _____.
4. ☐ Interview Summary (PTO-413), Paper No./Mail Date. _____.
5. ☐ Examiner's Amendment/Comment
6. ☑ Examiner's Statement of Reasons for Allowance
7. ☐ Other _____.

/NNENNA N EKPO/
Primary Examiner, Art Unit 2425

Continuation of Attachment(s) 2. Information Disclosure Statements (PTO/SB/08), Paper No./Mail Date: 04/08/22 (2 pages), 08/02/2022, 09/16/2022, 11/30/2022

DETAILED ACTION

Notice of Pre-AIA or AIA Status

1. The present application is being examined under the pre-AIA first to invent provisions.

Information Disclosure Statement

2. The references listed in the Information Disclosure Statement filed on April 08, 2022 (2 pages), 08/02/2022, September 16, 2022 and November 30, 2022 has been considered by the examiner (see attached PTO-1449 form).

Terminal Disclaimer

3. The terminal disclaimer filed on 12/01/2022 disclaiming the terminal portion of any patent granted on this application which would extend beyond the expiration date of **U.S. Patent No. 10,715,878** has been reviewed and is accepted. The terminal disclaimer has been recorded.

Allowable Subject Matter

4. **Claims 52-77** are allowed.

5. The following is an examiner's statement of reasons for allowance: The prior art fails to teach or fairly disclose "A computer system comprising at least one hardware server computer, wherein the computer system is programmed and connected to

perform a method of using common IP addresses to identify associations among multiple electronic devices and using an identified association to cause cross-device action, the method comprising:... (B) an IP address through which the one of the devices accessed the Internet, and (C) a timestamp identifying when the one of the devices accessed the Internet; (ii) at least some of the associations indicate association between at least two devices of the two or more Internet-accessing electronic devices, which association is determined by using the stored data to determine that both of the at least two devices have accessed the Internet via a common IP address at respective times satisfying a predetermined time-relation criterion; (iii) at least some of the associations indicate association between at least two devices of the two or more Internet-accessing electronic devices, which association is determined without use of personally identifiable information; and (iv) the computer system determining the associations accesses the Internet without using any of the common IP addresses of the associations identified in part (ii); and (b) when the computer system and a first device identified in a first association are in communication with one another, with the computer system using a device identifier of the first device and user profile information associated with a device identifier of a second device identified in the first association to cause an automatic action to be taken with respect to the first device, wherein: (i) the first association is one of the associations identified in part (a)(ii); and (ii) the first association was stored in the electronic database described in part (a) prior to the computer system and the first device being in communication with one another as described in part (b)".

Citation of Pertinent Prior Art

6. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

Estugrul et al.	U.S. Pub. No. 2008/0222283
Blasko	U. S. Pub. No. 2001/0049620
Howcroft et al.	U.S. Pub. No. 2008/0040742
Haeuser et al.	U.S. Pub. No. 2007/0283384
Lee	U.S. Pub. No. 2002/0010653
Conkwright et al.	U.S. Pub. No. 2002/0133490

7. Any comments considered necessary by applicant must be submitted no later than the payment of the issue fee and, to avoid processing delays, should preferably accompany the issue fee. Such submissions should be clearly labeled "Comments on Statement of Reasons for Allowance."

Conclusion

8. Any inquiry concerning this communication or earlier communications from the examiner should be directed to NNENNA NGOZI EKPO whose telephone number is (571)270-1663. The examiner can normally be reached M-W 10:00am - 6:30pm, TH-F 8:00am - 4:30pm.

Examiner interviews are available via telephone, in-person, and video conferencing using a USPTO supplied web-based collaboration tool. To schedule an

interview, applicant is encouraged to use the USPTO Automated Interview Request (AIR) at http://www.uspto.gov/interviewpractice.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Brian Pendleton can be reached on 571-272-7527. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of published or unpublished applications may be obtained from Patent Center. Unpublished application information in Patent Center is available to registered users. To file and manage patent submissions in Patent Center, visit: https://patentcenter.uspto.gov. Visit https://www.uspto.gov/patents/apply/patent-center for more information about Patent Center and https://www.uspto.gov/patents/docx for information about filing in DOCX format. For additional questions, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

NNENNA EKPO
Primary Examiner
Art Unit 2425

/NNENNA N EKPO/
Primary Examiner, Art Unit 2425
November 28, 2022.

Notice of References Cited	Application/Control No. 17/714,965	Applicant(s)/Patent Under Reexamination Shkedi, Roy	
	Examiner NNENNA N EKPO	Art Unit 2425	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	CPC Classification	US Classification
*	A	US-20080222283-A1	09-2008	Ertugrul; Kent Thomas	H04L67/535	709/224
*	B	US-20010049620-A1	12-2001	Blasko, John P.	G06Q30/0207	705/14.1
*	C	US-20080040742-A1	02-2008	Howcroft; Jerald Robert	H04N21/6582	725/35
*	D	US-20070283384-A1	12-2007	Haeuser; Anthony A.	H04N7/165	725/35
*	E	US-20020010653-A1	01-2002	Lee, Ki Kuk	G06Q30/06	709/206
*	F	US-20020133490-A1	09-2002	Conkwright, G. Colby	G06Q30/0277	707/999.009
	G					
	H					
	I					
	J					
	K					
	L					
	M					

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	CPC Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	
	V	
	W	
	X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	17/714,965	Shkedi, Roy
	Examiner	**Art Unit**
	NNENNA N EKPO	2425

CPC - Searched*		
Symbol	**Date**	**Examiner**
H04N21/812, G06Q30/0256, H04L61/5007, G06Q30/0251	11/28/2022	Nnenna Ekpo
G06Q30/0241, H04N21/6125, H04N21/64322	11/28/2022	Nnenna Ekpo
H04N21/25891, H04N21/44222, H04N21/4532	11/28/2022	Nnenna Ekpo
H04N21/44224, G06Q30/0255, G06F16/951, G06Q30/0269	11/28/2022	Nnenna Ekpo
H04N21/42684, G06F16/248, H04L67/306, H04L12/66	11/28/2022	Nnenna Ekpo

CPC Combination Sets - Searched*		
Symbol	**Date**	**Examiner**

US Classification - Searched*			
Class	**Subclass**	**Date**	**Examiner**

* See search history printout included with this form or the SEARCH NOTES box below to determine the scope of the search.

Search Notes		
Search Notes	**Date**	**Examiner**
Inventor Name Search	11/28/2022	Nnenna Ekpo
NPL Search: Google	11/28/2022	Nnenna Ekpo
PE2E Search	11/28/2022	Nnenna Ekpo
CPC Search further limited with text string search	11/28/2022	Nnenna Ekpo
IPR2022-01315, IPR2022-01505 - Reviewed Petition for Inter Partes Review of U.S. Patent No. 10,715,878 on 11/28/2022	11/28/2022	Nnenna Ekpo

/NNENNA N EKPO/ Primary Examiner, Art Unit 2425	

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	17/714,965	Shkedi, Roy
	Examiner	Art Unit
	NNENNA N EKPO	2425

Interference Search			
US Class/CPC Symbol	**US Subclass/CPC Group**	**Date**	**Examiner**
H04N21	812, 6125, 64322, 25891, 44222, 4532, 44224, 42684	11/28/2022	Nnenna Ekpo
G06Q30	0256, 0251, 0241, 0255, 0269	11/28/2022	Nnenna Ekpo
H04L61	5007	11/28/2022	Nnenna Ekpo
H04L67	306	11/28/2022	Nnenna Ekpo
H04L12	66	11/28/2022	Nnenna Ekpo
G06F16	951, 248	11/28/2022	Nnenna Ekpo

/NNENNA N EKPO/ Primary Examiner, Art Unit 2425	

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	17/714,965	Shkedi, Roy
	Examiner	**Art Unit**
	NNENNA N EKPO	2425

CPC						
Symbol					**Type**	**Version**
H04N	/	21	/	812	F	2013-01-01
G06F	/	16	/	951	I	2019-01-01
G06Q	/	30	/	0241	I	2013-01-01
G06Q	/	30	/	0251	I	2013-01-01
G06Q	/	30	/	0255	I	2013-01-01
G06Q	/	30	/	0256	I	2013-01-01
G06Q	/	30	/	0269	I	2013-01-01
H04L	/	12	/	66	I	2013-01-01
H04L	/	67	/	306	I	2013-01-01
H04N	/	21	/	44224	I	2020-08-01
H04N	/	21	/	44222	I	2013-01-01
H04L	/	61	/	5007	I	2022-05-01
H04N	/	21	/	25891	I	2013-01-01
H04N	/	21	/	42684	I	2013-01-01
H04N	/	21	/	4532	I	2013-01-01
H04N	/	21	/	6125	I	2013-01-01
H04N	/	21	/	64322	I	2013-01-01
G06F	/	16	/	248	I	2019-01-01

CPC Combination Sets							
Symbol					**Type**	**Set**	**Ranking**	**Version**
	/		/					

NONE		**Total Claims Allowed:**	
(Assistant Examiner)	(Date)	26	
/NNENNA N EKPO/ Primary Examiner, Art Unit 2425	28 November 2022	O.G. Print Claim(s)	O.G. Print Figure
(Primary Examiner)	(Date)	1	6A

U.S. Patent and Trademark Office — Part of Paper No.: 20221128

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	17/714,965	Shkedi, Roy
	Examiner	**Art Unit**
	NNENNA N EKPO	2425

INTERNATIONAL CLASSIFICATION				
CLAIMED				
H04N	/	21	/	81
G06Q	/	30	/	02
H04L	/	61	/	5007
G06Q	/	30	/	02
H04N	/	21	/	61
H04N	/	21	/	643
H04N	/	21	/	258
H04N	/	21	/	442
H04N	/	21	/	45
G06F	/	16	/	951
H04N	/	21	/	426
G06F	/	16	/	248
H04L	/	67	/	306
H04L	/	12	/	66

NON-CLAIMED				
	/		/	

US ORIGINAL CLASSIFICATION	
CLASS	**SUBCLASS**

CROSS REFERENCES(S)						
CLASS	**SUBCLASS (ONE SUBCLASS PER BLOCK)**					

NONE		Total Claims Allowed:	
(Assistant Examiner)	(Date)	26	
/NNENNA N EKPO/ Primary Examiner, Art Unit 2425	28 November 2022	O.G. Print Claim(s)	O.G. Print Figure
(Primary Examiner)	(Date)	1	6A

U.S. Patent and Trademark Office — Part of Paper No.: 20221128

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	17/714,965	Shkedi, Roy
	Examiner	**Art Unit**
	NNENNA N EKPO	2425

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIMS

Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original
-	1	-	11	-	21	-	31	-	41	-	51	10	61	20	71
-	2	-	12	-	22	-	32	-	42	1	52	11	62	21	72
-	3	-	13	-	23	-	33	-	43	2	53	12	63	22	73
-	4	-	14	-	24	-	34	-	44	3	54	13	64	23	74
-	5	-	15	-	25	-	35	-	45	4	55	14	65	24	75
-	6	-	16	-	26	-	36	-	46	5	56	15	66	25	76
-	7	-	17	-	27	-	37	-	47	6	57	16	67	26	77
-	8	-	18	-	28	-	38	-	48	7	58	17	68		
-	9	-	19	-	29	-	39	-	49	8	59	18	69		
-	10	-	20	-	30	-	40	-	50	9	60	19	70		

NONE		**Total Claims Allowed:**	
(Assistant Examiner)	(Date)	26	
/NNENNA N EKPO/ Primary Examiner, Art Unit 2425	28 November 2022	O.G. Print Claim(s)	O.G. Print Figure
(Primary Examiner)	(Date)	1	6A

U.S. Patent and Trademark Office — Part of Paper No.: 20221128

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	17/714,965	Shkedi, Roy
	Examiner	Art Unit
	NNENNA N EKPO	2425

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

CLAIMS

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	09/23/2022								
-	1	-								
-	2	-								
-	3	-								
-	4	-								
-	5	-								
-	6	-								
-	7	-								
-	8	-								
-	9	-								
-	10	-								
-	11	-								
-	12	-								
-	13	-								
-	14	-								
-	15	-								
-	16	-								
-	17	-								
-	18	-								
-	19	-								
-	20	-								
-	21	-								
-	22	-								
-	23	-								
-	24	-								
-	25	-								
-	26	-								
-	27	-								
-	28	-								
-	29	-								
-	30	-								
-	31	-								
-	32	-								
-	33	-								
-	34	-								
-	35	-								
-	36	-								
-	37	-								
-	38	-								
-	39	-								
-	40	-								
-	41	-								
-	42	-								

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	17/714,965	Shkedi, Roy
	Examiner	**Art Unit**
	NNENNA N EKPO	2425

CLAIM		DATE								
Final	Original	09/23/2022								
-	43	-								
-	44	-								
-	45	-								
-	46	-								
-	47	-								
-	48	-								
-	49	-								
-	50	-								
-	51	-								
1	52	=								
2	53	=								
3	54	=								
4	55	=								
5	56	=								
6	57	=								
7	58	=								
8	59	=								
9	60	=								
10	61	=								
11	62	=								
12	63	=								
13	64	=								
14	65	=								
15	66	=								
16	67	=								
17	68	=								
18	69	=								
19	70	=								
20	71	=								
21	72	=								
22	73	=								
23	74	=								
24	75	=								
25	76	=								
26	77	=								

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor	:	Shkedi, Roy	Art Unit	:	3693 (in parent)
Serial No.	:	17/714,965	Examiner	:	Ekpo, Nnenna (in parent)
Filing Date	:	04/06/2022	Conf. No.	:	7004
Title	:	Targeted television advertisements based on online behavior			

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Filed via EFS-Web on August 2, 2022

INFORMATION DISCLOSURE STATEMENT

Dear Sir:

Assignee discloses the references listed on attached form PTO/SB/08. Assignee submits herewith copies of all NPL references that were not previously cited in this application.

Also, assignee wishes to direct the examiner to the co-pending applications and proceedings in the table below, as well as the references cited in those proceedings. All listed applications and the listed patent share the same first inventor:

Serial No.	Filing Date	Earliest Parent Filing Date
17/714,936	04/06/2022	06/14/2007
17/877,870	07/29/2022	04/17/2007
17/843,311	06/17/2022	08/03/2011 (prov.)
IPR2022-00953	05/24/2022	N/A - re: USP 8,566,164 (same family as SN 17/401,078, previously cited)
IPR2022-00959	05/13/2022	N/A - re: USP 8,595,069 (same family as SN 17/401,078, previously cited)
IPR2022-00960	05/13/2022	N/A - re: USP 10,321,198 (same family as SN 17/401,078, previously cited)
IPR2022-00773	03/28/2022	Re: parent USP 8,677,398

CONSIDERED: /N.N.E/

IPR2022-01236	07/05/2022	Re: parent USP 8,677,398
IPR2022-01315	07/27/2022	Re: parent USP 10,715,878

If the Office has any questions, please feel free to contact assignee's undersigned attorney of record.

Respectfully submitted,

INTENT IQ, LLC.

by its attorney

Dated: August 2, 2022

/Louis J. Hoffman/
Louis J. Hoffman
Reg. No. 38,918

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, Arizona 85260
(480) 948-3295

/NNENNA N EKPO/ 11/28/2022

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For
Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

U.S. PATENT DOCUMENTS

Examiner Initials	Cite No.	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	1	5105184	04/14/1992	Pirani et al.
	2	5220655	06/15/1993	Tsutsui
	3	5359508	10/25/1994	Rossides
	4	5442633	08/15/1995	Perkins et al.
	5	5572643	11/05/1996	Judson
	6	5710884	01/20/1998	Dedrick
	7	5724424	03/03/1998	Gifford
	8	5724521	03/03/1998	Dedrick
	9	5751956	05/12/1998	Kirsch
	10	5774670	06/30/1998	Montulli
	11	5794210	08/11/1998	Goldhaber et al.
	12	5809242	09/15/1998	Shaw et al.
	13	5819092	10/06/1998	Ferguson et al.
	14	5819285	10/06/1998	Damico et al.
	15	5835718	11/10/1998	Blewett
	16	5838314	11/17/1998	Neel et al.
	17	5848396	12/08/1998	Gerace
	18	5855008	12/29/1998	Goldhaber et al.
	19	5870546	02/09/1999	Kirsch
	20	5870550	02/09/1999	Wesinger Jr. et al.
	21	5901287	05/04/1999	Bull et al.
	22	5918014	06/29/1999	Robinson
	23	5933811	08/03/1999	Angles et al.
	24	5937390	08/10/1999	Hyodo
	25	5937392	08/10/1999	Alberts
	26	5948061	09/07/1999	Merriman et al.
	27	5959623	09/28/1999	van Hoff et al.
	28	5963915	10/05/1999	Kirsch
	29	5978013	11/02/1999	Jones et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

_____	30	5995943	11/30/1999	Bull et al.
_____	31	5999912	12/07/1999	Wodarz et al.
_____	32	6002394	12/14/1999	Schein et al.
_____	33	6014638	01/11/2000	Burge et al.
_____	34	6016504	01/18/2000	Arnold et al.
_____	35	6018768	01/25/2000	Ullman et al.
_____	36	6026368	02/15/2000	Brown et al.
_____	37	6026369	02/15/2000	Capek
_____	38	6026374	02/15/2000	Chess
_____	39	6047310	04/04/2000	Kamakura et al.
_____	40	6049835	04/11/2000	Gagnon
_____	41	6055573	04/25/2000	Gardenswartz et al.
_____	42	6073241	06/06/2000	Rosenberg et al.
_____	43	6128663	10/03/2000	Thomas
_____	44	6133912	10/17/2000	Montero
_____	45	6134532	10/17/2000	Lazarus et al.
_____	46	6157946	12/05/2000	Itakura et al.
_____	47	6160989	12/12/2000	Hendricks et al.
_____	48	6177931	01/23/2001	Alexander et al.
_____	49	6188398	02/13/2001	Collins-Rector et al.
_____	50	6189008	02/13/2001	Easty et al.
_____	51	6205432	03/20/2001	Gabbard et al.
_____	52	6216112	04/10/2001	Fuller et al.
_____	53	6216129	04/10/2001	Eldering
_____	54	6285987	09/04/2001	Roth et al.
_____	55	6289318	09/11/2001	Barber
_____	56	6308202	10/23/2001	Cohn et al.
_____	57	6327574	12/04/2001	Kramer et al.
_____	58	6366298	04/02/2002	Haitsuka et al.
_____	59	6377986	04/23/2002	Philyaw
_____	60	6385592	05/07/2002	Angles et al.
_____	61	6401075	06/04/2002	Mason et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

_____	62	6408278	06/18/2002	Carney et al.
_____	63	6415270	07/02/2002	Rackson et al.
_____	64	6446261	09/03/2002	Rosser
_____	65	6449657	09/10/2002	Stanbach Jr. et al.
_____	66	6463533	10/08/2002	Calamera et al.
_____	67	6463585	10/08/2002	Hendricks et al.
_____	68	6490587	12/03/2002	Easty et al.
_____	69	6536041	03/18/2003	Knudson et al.
_____	70	6633850	10/14/2003	Gabbard et al.
_____	71	6654725	11/25/2003	Langheinrich et al.
_____	72	6684194	01/27/2004	Eldering et al.
_____	73	6718551	04/06/2004	Swix et al.
_____	74	6771290	08/03/2004	Hoyle
_____	75	6810527	10/26/2004	Conrad et al.
_____	76	6837436	01/04/2005	Swartz et al.
_____	77	6845396	01/18/2005	Kanojia et al.
_____	78	6847992	01/25/2005	Haitsuka et al.
_____	79	6868392	03/15/2005	Ogasawara
_____	80	6907566	06/14/2005	McElfresh et al.
_____	81	6985882	01/10/2006	Del Sesto
_____	82	7013290	03/14/2006	Ananian
_____	83	7038637	05/02/2006	Eller et al.
_____	84	7051351	05/23/2006	Goldman et al.
_____	85	7092926	08/15/2006	Cerrato
_____	86	7100183	08/29/2006	Kunkel et al.
_____	87	7146329	12/05/2006	Conkwright et al.
_____	88	7152237	12/19/2006	Flickinger et al.
_____	89	7158943	01/02/2007	Van der Riet
_____	90	7188085	03/06/2007	Pelletier
_____	91	7260823	08/21/2007	Schlack et al.
_____	92	7272853	09/18/2007	Goodman et al.
_____	93	7299195	11/20/2007	Tawakol et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

_____	94	7302419	11/27/2007	Conkwright et al.
_____	95	7139723	01/13/2000	Conkwright et al.
_____	96	7305459	12/04/2007	Klemba et al.
_____	97	7580855	08/25/2009	Fernandes
_____	98	7640214	12/29/2009	Black et al.
_____	99	7756744	07/13/2010	Schiff et al.
_____	100	7805332	09/28/2010	Wilson
_____	101	7806329	10/05/2010	Dmitriev et al.
_____	102	7856372	12/21/2010	Ullah
_____	103	7856373	12/21/2010	Ullah
_____	104	7913275	03/22/2011	Hendricks et al.
_____	105	7930252	04/19/2011	Bender et al.
_____	106	7937383	05/03/2011	Hintze et al.
_____	107	7966647	06/21/2011	Igoe et al.
_____	108	7971260	06/28/2011	Spalink et al.
_____	109	8015184	09/06/2011	Zito et al.
_____	110	8015286	09/06/2011	Jenkins
_____	111	8019881	09/13/2011	Sandhu et al.
_____	112	8024765	09/20/2011	Ramanathan et al.
_____	113	8079048	12/13/2011	Howcroft
_____	114	8086491	12/27/2011	Matz et al.
_____	115	8180674	05/15/2012	Madhavan
_____	116	8700464	04/15/2014	McAllister et al.
_____	117	8725571	05/13/2014	McAllister et al.
_____	118	8762201	06/24/2014	Noonan
_____	119	8880677	11/04/2014	Julia et al.
_____	120	20010034646	10/25/2001	Hoyt et al.
_____	121	20010042002	11/15/2001	Koopersmitch
_____	122	20010049620	12/06/2001	Blasko
_____	123	20020010653	01/24/2002	Lee
_____	124	20020010928	01/24/2002	Sahota
_____	125	20020013943	01/31/2002	Haberman et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For
Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

_____	126	20020019769	02/14/2002	Barrits et al.
_____	127	20020032603	03/14/2002	Yeiser
_____	128	20020046105	04/18/2002	Gardenswartz et al.
_____	129	20020049968	04/25/2002	Wilson et al.
_____	130	20020059094	05/16/2002	Hosea et al.
_____	131	20020059602	05/16/2002	Macrae et al.
_____	132	20020078444	06/20/2002	Krewin et al.
_____	133	20020082910	06/27/2002	Kontogouris
_____	134	20020082923	06/27/2002	Merriman et al.
_____	135	20020083043	06/27/2002	Hoshi et al.
_____	136	20020103850	08/01/2002	Moyer et al.
_____	137	20020104086	08/01/2002	Tomsen et al.
_____	138	20020112013	08/15/2002	Walsh
_____	139	20020120933	08/29/2002	Knudson et al.
_____	140	20020120935	08/29/2002	Huber et al.
_____	141	20020123928	09/05/2002	Eldering et al.
_____	142	20020124249	09/05/2002	Shintani
_____	143	20020124253	09/05/2002	Eyer et al.
_____	144	20020129362	09/12/2002	Chang et al.
_____	145	20020129368	09/12/2002	Schlack et al.
_____	146	20020147645	10/02/2002	Alao et al.
_____	147	20020194058	12/19/2002	Eldering
_____	148	20020194604	12/19/2002	Sanchez et al.
_____	149	20030023481	01/30/2003	Calvert et al.
_____	150	20030041156	02/27/2003	Pickover et al.
_____	151	20030050863	03/13/2003	Radwin
_____	152	20030051242	03/13/2003	Donnelly
_____	153	20030066078	04/03/2003	Bjorgan et al.
_____	154	20030093311	05/15/2003	Knowlson
_____	155	20030126597	07/03/2003	Darby et al.
_____	156	20030131355	07/10/2003	Berenson et al.
_____	157	20030145323	07/31/2003	Hendricks et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

_____	158	20030149975	08/07/2003	Eldering et al.
_____	159	20030187726	10/02/2003	Bull et al.
_____	160	20030187949	10/02/2003	Bhatt et al.
_____	161	20030208756	11/06/2003	Macrae et al.
_____	162	20030208758	11/06/2003	Schein et al.
_____	163	20030226141	12/04/2003	Krasnow et al.
_____	164	20030229900	12/11/2003	Reisman
_____	165	20030231207	12/18/2003	Huang
_____	166	20040002938	01/01/2004	Deguchi
_____	167	20040003406	01/01/2004	Billmaier
_____	168	20040010546	01/15/2004	Klug et al.
_____	169	20040031058	02/12/2004	Reisman
_____	170	20040044566	03/04/2004	Bostelmann et al.
_____	171	20040073915	04/15/2004	Dureau
_____	172	20040078809	04/22/2004	Drazin
_____	173	20040107437	06/03/2004	Reichardt et al.
_____	174	20040117827	06/17/2004	Karaoguz et al.
_____	175	20040123314	06/24/2004	Bova
_____	176	20040163101	08/19/2004	Swix et al.
_____	177	20040172650	09/02/2004	Hawkins et al.
_____	178	20050021387	01/27/2005	Gottfurcht
_____	179	20050049957	03/03/2005	Vakili et al.
_____	180	20050076393	04/07/2005	Sussman
_____	181	20050125289	06/06/2005	Beyda et al.
_____	182	20050125290	06/09/2005	Beyda et al.
_____	183	20050132405	06/16/2005	AbiEzzi et al.
_____	184	20050165638	07/28/2005	Piller
_____	185	20050165643	07/28/2005	Wilson et al.
_____	186	20050165644	07/28/2005	Beyda et al.
_____	187	20050177420	08/11/2005	Tanahashi
_____	188	20050187823	08/25/2005	Howes
_____	189	20060018466	01/26/2006	Adelstein et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For
Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

_____	190	20060116924	06/01/2006	Angles et al.
_____	191	20060137015	06/22/2006	Fahrny et al.
_____	192	20060184989	08/17/2006	Slothouber
_____	193	20060229930	10/12/2006	Gottfurcht
_____	194	20060253323	11/09/2006	Phan et al.
_____	195	20060259357	11/16/2006	Chiu
_____	196	20060271690	11/30/2006	Banga et al.
_____	197	20060271953	11/30/2006	Jacoby et al.
_____	198	20060271959	11/30/2006	Jacoby et al.
_____	199	20060287916	12/21/2006	Starr et al.
_____	200	20070027901	02/01/2007	Chan et al.
_____	201	20070038516	02/15/2007	Apple et al.
_____	202	20070039020	02/15/2007	Cansler et al.
_____	203	20070073585	03/29/2007	Apple et al.
_____	204	20070078774	04/05/2007	Brown
_____	205	20070078849	04/05/2007	Slothouber
_____	206	20070088603	04/19/2007	Jouppi et al.
_____	207	20070089127	04/19/2007	Flickinger et al.
_____	208	20070100690	05/03/2007	Hopkins
_____	209	20070130015	06/07/2007	Starr et al.
_____	210	20070157228	07/05/2007	Bayer et al.
_____	211	20070198327	08/23/2007	Yazdani et al.
_____	212	20070208619	09/06/2007	Branam et al.
_____	213	20070208789	09/06/2007	Reichman
_____	214	20070214048	09/13/2007	Chan et al.
_____	215	20070220553	09/20/2007	Branam
_____	216	20070233571	10/04/2007	Eldering et al.
_____	217	20070244750	10/18/2007	Grannan et al.
_____	218	20070266403	11/15/2007	Ou et al.
_____	219	20070277220	11/29/2007	Shikuma et al.
_____	220	20070283384	12/06/2007	Haeuser et al.
_____	221	20070294721	12/20/2007	Haeuser et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For
Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

_____	222	20080016540	01/17/2008	Savoor et al.
_____	223	20080040742	02/14/2008	Howcroft et al.
_____	224	20080046942	02/21/2008	Merlin
_____	225	20080052392	02/28/2008	Webster et al.
_____	226	20080082413	04/03/2008	Madhavan
_____	227	20080109306	05/08/2008	Maigret et al.
_____	228	20080109376	05/08/2008	Walsh et al.
_____	229	20080109843	05/08/2008	Ullah
_____	230	20080155613	06/26/2008	Benya et al
_____	231	20080181225	07/31/2008	Zampiello
_____	232	20080183573	07/31/2008	Muschetto
_____	233	20080189360	08/07/2008	Kiley et al.
_____	234	20080195664	08/14/2008	Maharajh et al.
_____	235	20080201772	08/21/2008	Mondaeev et al.
_____	236	20080222283	09/11/2008	Ertugrul et al.
_____	237	20080243822	10/02/2008	Campbell et al.
_____	238	20080244076	10/02/2008	Shah et al.
_____	239	20080255944	10/16/2008	Shah et al.
_____	240	20080262901	10/23/2008	Banga et al.
_____	241	20080263633	10/23/2008	Banga et al.
_____	242	20080275785	11/06/2008	Altberg et al.
_____	243	20080307112	12/11/2008	Andreasson et al.
_____	244	20090037949	02/05/2009	Birch
_____	245	20090049469	02/19/2009	Small et al.
_____	246	20090106100	04/23/2009	Mashinsky
_____	247	20090171780	07/02/2009	Aldrey et al.
_____	248	20090204706	08/13/2009	Ertugrul et al.
_____	249	20100076848	03/25/2010	Stefanik et al.
_____	250	20110125592	05/26/2011	Mesaros
_____	251	20110167043	07/07/2011	Hintze et al.
_____	252	20110258049	10/20/2011	Ramer et al.
_____	253	20110307548	12/15/2011	Fisk et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

_____	254	20130254035	09/26/2013	Ramer et al.
_____	255	6832207	12/14/2004	Shkedi
_____	256	6925440	08/02/2005	Shkedi
_____	257	6973436	12/06/2005	Shkedi
_____	258	7072853	07/04/2006	Shkedi
_____	259	7428493	09/23/2008	Shkedi
_____	260	7454364	11/18/2008	Shkedi
_____	261	7747745	06/29/2010	Shkedi
_____	262	7822637	10/26/2010	Shkedi
_____	263	7822639	10/26/2010	Shkedi
_____	264	7861260	12/28/2010	Shkedi
_____	265	7890609	02/15/2011	Shkedi
_____	266	7979307	07/12/2011	Shkedi
_____	267	8051444	11/01/2011	Shkedi
_____	268	8200822	06/12/2012	Shkedi
_____	269	8204783	06/19/2012	Shkedi
_____	270	8204965	06/19/2012	Shkedi
_____	271	8239264	08/07/2012	Shkedi
_____	272	8244574	08/14/2012	Shkedi
_____	273	8244582	08/14/2012	Shkedi
_____	274	8244583	08/14/2012	Shkedi
_____	275	8244586	08/14/2012	Shkedi
_____	276	8280758	10/02/2012	Shkedi
_____	277	8494904	07/23/2013	Shkedi
_____	278	8566164	10/22/2013	Shkedi et al.
_____	279	8589210	11/19/2013	Shkedi
_____	280	8595069	11/26/2013	Shkedi et al.
_____	281	8600815	12/03/2013	Shkedi
_____	282	8607267	12/10/2013	Shkedi
_____	283	8683502	03/25/2014	Shkedi et al.
_____	284	8713600	04/29/2014	Shkedi
_____	285	8775249	07/08/2014	Shkedi

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For
Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

_____	286	8959146	02/17/2015	Shkedi
_____	287	8997138	03/31/2015	Shkedi
_____	288	9071886	06/30/2015	Shkedi
_____	289	9078035	07/07/2015	Shkedi et al.
_____	290	9083853	07/14/2015	Shkedi
_____	291	9131282	09/08/2015	Shkedi
_____	292	9208514	12/08/2015	Shkedi
_____	293	9226019	12/29/2015	Shkedi
_____	294	9271024	02/23/2016	Shkedi et al.
_____	295	9351053	05/24/2016	Shkedi
_____	296	9369779	06/14/2016	Shkedi
_____	297	9508089	11/29/2016	Shkedi
_____	298	9571865	02/14/2017	Shkedi
_____	299	9591380	03/07/2017	Shkedi et al.
_____	300	9756372	09/05/2017	Shkedi
_____	301	9800917	10/24/2017	Shkedi
_____	302	9813778	11/07/2017	Shkedi
_____	303	9830615	11/28/2017	Shkedi
_____	304	10009636	06/26/2018	Shkedi
_____	305	10026100	07/17/2018	Shkedi
_____	306	10134054	11/20/2018	Shkedi
_____	307	10178442	01/08/2019	Shkedi et al.
_____	308	10306282	05/28/2019	Shkedi
_____	309	10321198	6/11/2019	Shkedi et al.
_____	310	10405058	09/03/2019	Shkedi
_____	311	10475073	11/12/2019	Shkedi
_____	312	10628857	04/21/2020	Shkedi
_____	313	10645438	05/05/2020	Shkedi
_____	314	10715878	07/14/2020	Shkedi
_____	315	10771860	09/08/2020	Shkedi et al.
_____	316	10839423	11/17/2020	Shkedi
_____	317	20050086112	04/21/2005	Shkedi

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

_____	318	20110035256	02/10/2011	Shkedi et al.
_____	319	20110040607	02/17/2011	Shkedi
_____	320	20090172723	07/02/2009	Shkedi et al.

FOREIGN PATENT DOCUMENTS

Examiner Initials		Document Number	Publication Date MM/DD/YYYY	Name of Patentee	Translation Yes / No
_____	321	WO 2005/046237	05/19/2005	Gutta et al.	No
_____	322	KR-10-0793513	01/03/2008	ICU Research & Industrial Cooperation Group (ETRI)	Yes
_____	323	KR 10-2009-0113456	08/14/2012	SK Planet Co., Ltd.	Yes
_____	324	JP2007-274246	10/18/2007	Casio Computer Co. Ltd.	Yes
_____	325	JP 2006-004441	01/05/2006	Yokogawa Electric Corp.	Yes
_____	326	WO0184440	11/08/2001	PSM Global, Inc.	No
_____	327	WO2004086680	10/07/2004	Thomson Licensing S.A.	Yes
_____	328	JP2003-331183	11/21/2003	NTT Communications KK	Yes
_____	329	WO1997021183A1	06/12/1997	Bell Communications Research Inc.	No
_____	330	WO1999057660A1	11/11/1999	IWEB LTD.	No
_____	331	WO00/004709	01/27/2000	Ellis	No
_____	332	WO 01/054034 A1	07/26/2001	Phillips et al.	No

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	333	Engage Technologies, Inc.; Form S-1A; filed with the SEC; pp. 36-50; July 19, 1999.
_____	334	Chen, Andy; "Behavioral Matchmaking, Part 3: GPS and Behavioral Targeting"; from website http://www.clickz.com; November 24, 2004.
_____	335	Battelle, John; "The Search - How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture"; pp. 167-171; 2005.
_____	336	"BSkyB in Google Link-up"; Financial Times (FT.com); December 6, 2006.
_____	337	"Google moves into TV ads with BSkyB Deal"; Guardian Unlimited; December 7, 2006.
_____	338	Chen, Andy; "The New Behavioral Wunder"; from webpage http://www. clickz.com; April 11, 2007.
_____	339	"NetZero Launches Ad-Backed Access Service"; from website www.clickz. com; October 20, 1998.
_____	340	"Nissan, BellSouth, Music Blvd. Among Latest Advertisers to Sign Up With NetZero"; press release from www.irconnect.com; November 16, 1998.
_____	341	"LookSmart and NetZero Extend and Enhance Winning Partnership"; press release found on www.shareholder.com; February 9, 2000.
_____	342	"NetZero Guarantees Increase in Web Site Traffic for New Adverstisers"; press release from www.irconnect.com; circa 2000.
_____	343	Saunders, Christopher; "NetZero to Launch Mobile Ad Platform"; from website www.clickz.com; March 29, 2001.
_____	344	"180Solutions Launches All-In-One Online Ad Solution - Metricsdirect Outperforms Conventional Cost-Per-Click Models"; from www.zango.com; March 30, 2004.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For
Information Disclosure Statement

Page 13 of 33

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	345	Keegan, Paul; "The Man Who Can Save Advertising"; from http://money.cnn.com; November 1, 2004.
_____	346	"Visible World Targets Advertising Industry"; from http://informitv.com; February 20, 2005.
_____	347	"Venture Market Summary"; VentureWire Alert; March 29, 2005.
_____	348	Olsen, S.; "TiVo is in talks with Internet search giants Google and Yahoo over possible deal ... television and the Web..."; www.news.com; April 18, 2005.
_____	349	"This Ad's for You -- Just You"; BusinessWeek online; June 28, 2005.
_____	350	"Cable's Big Bet On Hyper-Targeting. Time Warner will test new software that sends different ads to different viewers."; BusinessWire online; July 4,2005.
_____	351	"Comcast, Cox Buy Interactive TV Company"; from Yahoo.com Financial News; July 13, 2005.
_____	352	"TiVo Launches New Interactive Advertising Technology"; press release from www.tivo.com; July 18, 2005.
_____	353	Grant, Peter; "To Ward Off New Competitors, Comcast Builds a Mini Internet"; Wall Street Journal online; October 13, 2005.
_____	354	"TiVo Announces First Advertising Search Product for Television"; press release from www.tivo.com; November 28, 2005.
_____	355	"TiVo Begins Rollout of Online Services Now Accessible Directly on the TV"; press release from www.tivo.com; December 1, 2005.
_____	356	"Comcast, Time Warner Cable hopeful on network DVRs"; from www.reuters.com; March 30, 2006.
_____	357	"TiVo Launches Television's New Advertising Search Product"; press release from www.tivo.com; May 8, 2006.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Page 14 of 33

Serial No.:	17/714,965
Applicant:	Shkedi, Roy
Filed:	04/06/2022
Art Unit:	2425
Examiner:	Ekpo, Nnenna
Conf. No.:	7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	358	"TiVo and Brightcove Partner to Deliver Broadband Content to TiVo Subscribers; press release from www.tivo.com; May 10, 2006.
_____	359	"TiVo Debuts New Lexus Car Configuration Application"; press release from www.tivo.com; May 11, 2006.
_____	360	Newcomb, Kevin; "BMW Uses TiVo Technology to Improve Product Placement"; from www.clickz.com; June 21, 2006.
_____	361	Wingfield, Nick; "New TiVo Feature Lets Families Share Videos On TV Via the Web"; The Wall Street Journal Online; November 14, 2006.
_____	362	Kaye, Kate; "TiVo Broadens CBS Offerings, Adds Web-to-TV Services"; from www.clickz.com; November 15, 2006.
_____	363	Grant, Peter and Marr, Merissa; "Comcast Nears Pact With Disney To Offer ABC Shows On Demand"; The Wall Street Journal Online; November 21, 2006.
_____	364	"TiVo to Insert Ads At End of Probrams"; The Wall Street Journal Online; November 28, 2006.
_____	365	FAQ from Zango Ad Services; http://adservices.zango.com; 2007.
_____	366	"NetZero and Juno's sophisticated targeting capabilities enable advertisers to hit an exact audience of qualified consumers"; from www.untd.com; 2007.
_____	367	Shukla, Anuradha; "Visible World Offers 'Any Screen' Solution for Customized Video Advertising"; from www.tmcnet.com; February 6, 2007.
_____	368	Tam, Pui-Wing; "TiVo, Amazon to Deliver Web Video to TV Sets"; The Wall Stree Journal Online; February 7, 2007.
_____	369	Story, Louise; "The Web drives an advertising boom without the need for agencies"; International Herald Tribune (iht.com); February 8, 2007.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

CONSIDERED: /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For
Information Disclosure Statement

Serial No.:	17/714,965
Applicant:	Shkedi, Roy
Filed:	04/06/2022
Art Unit:	2425
Examiner:	Ekpo, Nnenna
Conf. No.:	7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	370	Todd, Daniel; "The New World of Desktop Advertising"; DMNews.com; March 28, 2007.
_____	371	"Medio Systems to Launch Mobile Performance Ad Network - Search and recommendation technologies propel advancements in mobile advertising", article from Mobilestreams website, March 19, 2007.
_____	372	HUNT, TIMOTHY J., "Moving Target", Financial Post, October 1, 2000.
_____	373	Luening E.; "Free ISP NetZero beefs up ad services"; web page at http://www.news.com/Free-ISP-NetZero-beefs-up-ad-services/2100-1023_3-241067.html; January 2, 2002.
_____	374	BERGSTEIN B.; "Ad-targeted System Monitors Your Interests with ISP's Help"; from website www.siliconvalley.com; December 10, 2007.
_____	375	WHITE B.; "Watching What You See on the Web"; Wall Street Journal Online; December 6, 2007.
_____	376	"What's New: The Latest On Technology Deals"; from Dow Jones VentureWire; as published on Wall Street Journal Online; October 22, 2007.
_____	377	"DoubleClick Launches New Marketing Solutions"; from website www.clickz.com; October 6, 1998.
_____	378	Poikselktä, Miikka, et al.; "IMS Concepts"; The IMS: IP 1-15 Multimedia Concepts and Services, Second Edition; John Wiley & Sons, Ltd.; 2006; pp. 43-116.
_____	379	Baseview.com website (ClassManagerPro) from www.archive.org, archived January 16, 1997
_____	380	S-1/A SEC Filing, filed by FlyCast Communications Corp. on April 30, 1999
_____	381	24/7 RealMedia, The Science of Digital Marketing, Company Overview (2007)
_____	382	aaddzz.com webpages from www.archive.org, archived January 30, 1998

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Page 16 of 33

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	383	AdForce, Inc., User Guide, A Complete Guide to AdForce, Version 2.6 (1998)
_____	384	Fawcett, Adrienne W., Publishers thankful but nervous about boom, Magazinedotcom, Hot spending boosts …, adage.com, November 1, 1999
_____	385	Ad-Star.com website archive from www.archive.org (archived on April 12, 1997 and February 1, 1997)
_____	386	Chang et al., "Goodies" in exchange for Consumer Information on the Internet, The Economics and Issues, Proceedings of the Thirty-First Annual Hawaii International Conference on System Sciences, 1998
_____	387	avenuea.com webpage, FAQ, www.avenuea.com/corpinfo/faq.html, archived October 2, 1999
_____	388	Baek et al., Implementation of Smart Home Control using LabVIEW and PDA," IEEE (2004)
_____	389	Bardram et al., Experiences from Real-World Deployment of Context-Aware Technologies in a Hospital Environment, Center for Pervasive Healthcare (2006)
_____	390	Advertising-A Unified Approach to Adverstising Management, Baseview Products, Inc. brochure, October 14, 1999
_____	391	Ribeiro-Neto et al., Extracting Semi-Structured Data Through Examples, Proceedings of the Eighth international Conference on Information and Knowledge Management (1999)
_____	392	Raskutti and Beitz, Acquiring Customer Preferences for Information Filtering, a Heuristic-Statistical Approach, Fifth Annual Conference on User Modeling (1996)
_____	393	AdKnowledge, Campaign Manager-Quick Reference Card (1996)
_____	394	Emert, 2 Web Ad Firms Decide to Marry, Palo Alto's Focalink, ClickOver Merging, San Francisco Chronicle, November 17, 1997

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	395	Castro et al., A Probabilistic Room Location Service for Wireless Networked Environments, Ubicomp 2001, LNCS 2201, pp. 18-34 (2001)
_____	396	Chiu et al., LiteMinutes, An Internet-Based System for Multimedia Meeting Minutes," WWW10, May 2-5, 2001
_____	397	Cox et al., IntelliBadge, Towards Providing Location-Aware Value-Added Services at Academic Conferences," UbiComp 2003, LNCS 2864, pp. 264-280 (2003)
_____	398	Cuzzocrea et al., Enriching Conceptual Modeling of XML-Aware Adaptive Web Systems with Object-Oriented Constructs and UML Diagrams," EC-WEB 2004, LNCS 3182, pp. 248-258 (2004)
_____	399	Elderbrock & Borwankar, Building Successful Internet Businesses, IDG Worldwide, 1996
_____	400	David W. Schumann & Ester Thorson, Advertising and the World Wide Web, Lawrence Erlbaum Associates, 1999
_____	401	Williamson, Debra Aho, Web ads mark 2nd birthday with decisive issues ahead, adage.com, October 21, 1996
_____	402	Williamson, Debra Aho, Targeting distinguishes AdForce from the pack, adage.com, March 1, 1999
_____	403	Doubleclick.com webpages from www.archived.org, archived February 5, 1998
_____	404	Tittel et al., Foundations of World Wide Web Programming with HTML & CGI, IDG Books Worldwide (1995)
_____	405	Genco et al., An Augmented Campus Design for Context-aware Service Provision, SIGUCCS'05 (November 6-9, 2005
_____	406	Genco, et al. An Agent-Based Service Network for Personal Mobile Devices, IEEE (2006)

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Page 18 of 33

Serial No.:	17/714,965
Applicant:	Shkedi, Roy
Filed:	04/06/2022
Art Unit:	2425
Examiner:	Ekpo, Nnenna
Conf. No.:	7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	407	Hess et al., A Context-Aware Data Management System for Ubiquitous Computing Applications, Proceedings of the 23rd Intl. Conf. on Distributed Computing Systems (2003)
_____	408	Graham, Ian S., HTML Sourcebook, A Complete Guide to HTML, John Wiley & Sons (1995)
_____	409	Aaddzz Brochure, Information Access Technologies, Inc. (1997)
_____	410	Iwasaki et al., Azim, Direction Based Service using Azimuth Based Position Estimation, ICDCS (2004)
_____	411	Gilbert, SPECIAL REPORT, Agencies centralize Web ad serving, adage.com, Mar 1, 1999
_____	412	Palmer & Eriksen, Digital Newspapers Explore Marketing on the Internet, Communications of the ACM, Vol. 42, No. 9, September 1999
_____	413	Angwin, Julia, Focalink Will Monitor World of Online Ads, San Francisco Chronicle, February 7, 1996
_____	414	Gallagher & Parsons, A Framework for Targeting Banner Advertising on the Internet, Proceedings of the Thirtieth Hawaii International Conference on System Sciences, January 1997
_____	415	Cleland, Kim, Accipiter latest entrant in race to manage ads; company sets sights on helping content providers target audience, Adage.com, September 9, 1996
_____	416	Cleland, Kim, New Player joins fray of Web ad targeters, adage.com, February 24, 1997
_____	417	Bayne, Kim M., AdKnowledge rolls out Web ad evaluation tool, adage.com, July 8, 1999
_____	418	Langheinrich et al., Unintrusive Customization Techniques for Web Advertising, COMPUTER NETWORKS 31, 1259–1272, May 1999

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.:	17/714,965
Applicant:	Shkedi, Roy
Filed:	04/06/2022
Art Unit:	2425
Examiner:	Ekpo, Nnenna
Conf. No.:	7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	419	Maekawa et al., A Java-Based Information Browsing System in a Remote Display Environment, IEEE (2004)
_____	420	Meeker, Mary, The Internet Advertising Report, Harper Business, 1997
_____	421	McCarthy et al., UniCast, OutCast & GroupCast, Three Steps Toward Ubiquitous, Peripheral Displays, Ubicomp 2001, LNCS 2201, pp. 332-345 (2001)
_____	422	Nakajima et al., Universal Interaction with Networked Home Appliances, ICDCS'02, IEEE (2002)
_____	423	Nakazawa et al., A Bridging Framework for Universal Interoperability in Pervasive Systems, ICDCS'06 IEEE (2006)
_____	424	Yeager & McGrath, Web Server Technology, The Advance Guide for World Wide Web Information Providers, Morgan Kaufmann Publishers (1996)
_____	425	Nawyn et al., Embedding Behavior Modification Strategies into a Consumer Electronic Device, A Case Study, Ubicomp 2006, LNCS 4206, pp. 297-314 (2006)
_____	426	net.Genesis & Devra Hall, Build a Website, The Programmer's Guide to Creating, Building, and Maintaining a Web Presence, Prima Publishing (1995)
_____	427	NetGravity Launches AdServer 3.5 with Beefed Up Targeting, clickz.com, September 10, 1998
_____	428	Novak & Hoffman, New Metrics for New Media, Toward the Development of Web Measurement Standards, Project 2000, elab.vanderbilt.edu, September 26, 1996
_____	429	Ponnekanti et al., ICrafter, A Service Framework for Ubiquitous Computing Environments," Ubicomp 2001, LNCS 2201, pp. 56-75 (2001)
_____	430	RealMedia.com webpages, www.archive.org, archived on May 15, 19977 and June 7, 1997

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	431	Bruner, Rick E., Interactive, NetGravity to launch new ad service bureau, adage.com, December 1, 1997
_____	432	Bruner, Rick E., Small networks chase per-click ad business, Aaddzz, ValueClick compete, AdvertisingAge, September 8, 1997
_____	433	Zeff & Aronson, Advertising on the Internet, 1st Edition, John Wiley & Sons (1997)
_____	434	Zeff & Aronson, Advertising on the Internet, 2nd Edition, John Wiley & Sons (1999)
_____	435	Ruta et al., Ontology driven resource discovery in a Bluetooth based m-marketplace, CEC/EEE'06, IEEE (2006)
_____	436	Smetters et al., Instant Matchmaking, Simple and Secure Integrated Ubiquitous Computing Environments, Ubicomp 2006, LNCS 4206, pp. 477-794 (2006)
_____	437	Sogo et al., LightHolder, A Platform for Real-World Mobile Commerce, CEC/EEE'06, IEEE (2006)
_____	438	Sumi et al., Digital Assistant for Supporting Conference Participants, An Attempt to Combine Mobile, Ubiquitous and Web Computing, Ubicomp 2001, LNCS 2201, pp. 156-175 (2001)
_____	439	Singh & Dalal, Web Home Pages as Advertisements, Communications of the ACM, August 1999
_____	440	The ClickZ Network, Netscape Selects Real Media's Open AdStream, January 13, 1998
_____	441	The ClickZ Network, Real Media Signs More Media Companies, December 30, 1998
_____	442	Trevor et al., Issues in Personalizing Shared Ubiquitous Devices, UbiComp 2002, LNCS 2498, pp. 56-72 (2002)

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For
Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	443	www.Spinbox.com webpages, www.archive.org, archived beginning Nov. 11, 1998
_____	444	Bakos, The Emerging Role of Electronic Marketplaces on the Internet, Communications of the ACM, Vol. 41, No. 8, August 1998
_____	445	Youiji Kohda & Susumu Endo, Ubiquitous Advertising on the WWW, Merging Advertisement on the Browser, Fifth International World Wide Web Conference (May 1996)
_____	446	Zhang et al., Smart Alarm Clock – A Networked Home Appliance with Bluetooth Connection, IEEE (2004)
_____	447	Baseview.com webpages from www.archive.org, archived January 16, 1997
_____	448	AdForce, Inc., Quick Reference Guide, Version 3.0 (1999)
_____	449	AdManagerPro; Ad Entry Manual; User's Manual V.1.05; 1997
_____	450	AdManagerPro; Ad Entry Manual; User's Manual V.2.0; 1997
_____	451	AdManagerPro; Administration and Receivables Manual; Administrator's Manual Version 1.0.5; November 1997
_____	452	AdManagerPro; Administration and Receivables Manual; Administrator's Manual Version 2.0; June 1998
_____	453	Commercialization of the World Wide Web, The Role of Cookies, vanderbilt.edu web pages, archive.org, February 1, 1997
_____	454	Crabtree et al., Adaptive Personal Agents, Personal Technologies (1998)
_____	455	Dalton, Gregory, OPS, Answer to Cookies?, InformationWeek, October 13, 1997
_____	456	Dunn et al.; Privacy and Profiling on the Web; Submitted to W3C on June 2, 1997

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	457	firefly.com web pages, archive.org, archived 1998-2000
_____	458	Gilbert, Jennifer, CMGI's Web dealmaking hits big-time, Advertising Age, October 4, 1999
_____	459	Hersley et al., Implementation of OPS Over HTTP, Submitted to W3C on June 2, 1997
_____	460	Hersley et al., Proposal for an Open Profiling Standard, Submitted to W3C on June 2, 1997
_____	461	Hersley et al., Standard Practices of OPS Systems, Submitted to W3C on June 2, 1997
_____	462	junkbusters.com web pages; archive.org; archived April 1997 through November 2000
_____	463	Karpinski, Richard; ICE application to revolutionize e-commerce; Business Marketing, March 1998
_____	464	matchlogic.com web pages, archive.org, April 13, 2000
_____	465	Roscheisen, R. Martin; A Network-Centric Design for Relationship-Based Rights Management (Dissertation); December 1997
_____	466	BARRY M. LEINER ET AL.; Brief History of the Internet; Internet Society; October 15, 2012
_____	467	KRISTOL, D. & MONTULLI, L.; HTTP State Management Mechanism; Internet Engineering Task Force; available at https://tools.ietf.org/html/rfc2109; February 1997
_____	468	TUNGATE, MARK; Adland: A Global History of Advertising, 2d Ed. (2013)

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

CONSIDERED: /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.:	17/714,965
Applicant:	Shkedi, Roy
Filed:	04/06/2022
Art Unit:	2425
Examiner:	Ekpo, Nnenna
Conf. No.:	7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	469	JOHNSON, BRADLEY; Interactive: Planet Oasis Gives Web Sites Promotion Clout; Advertising Age; available at http://adage.com/article/news/interactive-planet-oasis-web-sites-promotion-clout/78467; July 8, 1996
_____	470	JONES, DUNCAN; The History of Pay Per Click (PPC) Marketing; available at https://www.duncanjonesnz.com/history-of-ppc-marketing; September 23, 2014
_____	471	D.A.R.P.A., Internet Protocol; DARPA Internet Program Protocol Specification; Internet Engineering Task Force; available at https://tools.ietf.org/html/rfc791; September 1981
_____	472	IEEE Standard 802.3; IEEE Standard for Information technology-Telecommunications and information exchange between systems-Local and metropolitan area networks-Specific requirements; Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications; 2002 Edition
_____	473	IEEE Standard 802.11; IEEE Standard for Information technology-Telecommunications and information exchange between systems-Local and metropolitan area networks-Specific requirements; Part 11: Wireless Medium Access Control (MAC) and Physical Layer Specifications; 1999 Edition
_____	474	LO, Shih-Hsiang et al.; Controlling Digital TV Set-Top Box with Mobile Devices via an IP Network; National Taiwan University; pp. 935-942; May 20, 2006
_____	475	FORD, Bryan et al.; Persistent Personal Names for Globally Connected Mobile Devices; OSDI '06 Paper; pp. 233-248; Nov. 2006
_____	476	FALLAHKHAIR, S. et al.; Development of a cross-platform ubiquitous language learning service via mobile phone and interactive television; Journal of Computer Assisted Learning, 23, pp. 312-325; 2007
_____	477	LYU, Jihye et al.; Design of Open APIs for Personalized IPTV Service; ICACT 2007; February 12-14, 2007

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	478	GEIER, Eric; The Business Side of Hosting a Hotspot; http://www.ciscopress.com/articles/printerfriendly/680821; January 12, 2007
_____	479	MOON, Aekyung et al.; Context-Aware Active Services in Ubiquitous Computing Environments; ETRI Journal, Volume 29, Number 2; April 2007
_____	480	The Gresham Hotel, http://ambassadordublin.com/the-gresham-hotel; February 24, 2007
_____	481	Time Warner Cable, Using BIAP Systems' Interactive Technology, Expands Availability of 'Fantasy Football Tracker'; http://www.timewarner.com/newsroom/press-releases/2006/09/19/time-warner-cable-using-biap-systems-interactive-technology; September 19, 2006
_____	482	YANG, Herong; 'SendMail' Mail Transfer Agent Developed by Eric Allman, http://www.herongyang.com/Computer-History/Sendmail-Mail-Transfer-Agent-by-Eric-Allman-1983.html; 1983
_____	483	DEDRICK, Rick; Interactive Electronic Advertising; IEEE; 1994
_____	484	DEDRICK, Rick; A Consumption Model for Targeted Electronic Advertising; IEEE Multimedia; 1995
_____	485	ROBERTSON, Scott et al.; Dual Device User Interface Design: PDAs and Interactive Television; Applied Research, US WEST Advanced Technologies; 1996
_____	486	GORDON, Mary Ellen et al.; Consumer attitudes towards Internet; International Marketing Review, Vol. 14, Issue 5; pp. 362-375; 1997
_____	487	BEIGL, Michael et al.; The UbicompBrowser; Telecooperation Office, University of Karlsruhe; 1998
_____	488	RISDEN, Kirsten et al; Interactive Advertising: Patterns of Use and Effectiveness; CHI '98, 18-23; pp. 219-224; April 1998

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For
Information Disclosure Statement

Page 25 of 33

Serial No.:	17/714,965
Applicant:	Shkedi, Roy
Filed:	04/06/2022
Art Unit:	2425
Examiner:	Ekpo, Nnenna
Conf. No.:	7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	489	DREZE, Xavier and ZUFRYDEN, Fred; Is Internet Advertising ready For Prime Time?; April 5, 1998
_____	490	McCANDLESS, Michael; Web Advertising; IEEE Intelligent Systems; May/June 1998
_____	491	AGGARWAL, Charu C. et al.; A Framework for the Optimizing of WWW Advertising; June 3, 1998
_____	492	CHAN, M.C. and WOO, T.Y.; Next-generation wireless data services: Architecture and experience; Personal Commuincations, IEEE, 6(1); 1999
_____	493	SAMAR, Vipin; Single Sign-on Using Cookies for Web Applications; IEEE 8th International Workshops on Enabling Technologies: Infrastucture for Collaborative Enterprises (WET ICE '99); June 16-18, 1999
_____	494	SHEN, Andrew; Online Profiling Project – Comment, P994809; November 8, 1999
_____	495	SIEW, David and YI, Xun; Agent-mediated Internet Advertising; IEEE; 2000
_____	496	PHAM, Thai-Lai et al; A Situated Computing Framework for Mobile and Ubiquitous Multimedia Access using Small Screen and Composite Devices; January 2000
_____	497	TOMLIN, John A.; An Entropy Approach To Unintrusive Targeted Advertising On The Web; IBM Research Divison; June 2000
_____	498	SMYTH, Barry and COTTER, Paul; A Personalized Television Listings Service; Communications of the ACM; Vol. 43, No. 8, pp. 107-111; August 2000
_____	499	GOLDEN, G. Richard III; Service Advertisement and Discovery: Enabling Universal Device Cooperation; IEEE Internet Computing; September-October 2000
_____	500	BOLCER, Gregory Alan et al.; Peer-to-Peer Architectures and the Magi™ Open-Source Infrastructure; Endeavors Technology, Inc.; December 6, 2000

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

CONSIDERED: /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	501	LEKAKOS, George et al.; Information Systems In The Living Room: A Case Study Of Personalized Interactive TV Design; Global Co-Operation in the New Millenium; The 9th European Conference on Information Systems; Bled, Slovenia, June 27-29, 2001
_____	502	JUELS, Ari; Targeted Advertising… And Privacy Too; RSA Laboratories; 2001
_____	503	LEKAKOS, Georgios et al.; An Integrated Approach to Interactive and Personalized TV Advertising; eL TRUN- The eBusiness Center, Department of Management Science & Technology, Athens University of Economics and Business; 2001
_____	504	PRAMATARIS, K. et al.; Personalized Interactive TV Advertising: The iMEDIA Business Model; Electronic Markets; Vol. 11 (1): 1-9; 2001
_____	505	BOZIOS, Theodoros et al.; Advanced Techniques for Personalized Advertising in a Digital TV Environment: The iMEDIA System; INTRACOM S.A., R&D Division, Development Programmes Department; 2001
_____	506	KOUTSORODI, Artemis et al.; A Management Architecture For Enabling Wireless System Operation In The B3g Context And Exploitation Of The Abc Concept; National Technical University of Athens, School of Electrical and Computer Engineering; 2001
_____	507	NGAI, E.W.T.; Selection of web sites for online advertising using the AHP; Information & Management 40; pp. 233-242; 2001
_____	508	BERNSTEIN, D.J.; E-mail re: Internet host SMTP server survey; http://cr.yp.to/surveys/smtpsoftware.txt; November 29, 1996
_____	509	GERSHMAN, Anatole; Ubiquitous Commerce – Always On, Always Aware, Always Pro-active; The 2002 International Symposium on Applications and the Internet (SAINT 2002); January 30, 2002

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	510	WANG, Helen J. and KATZ, Randy H.; Mobility Support in Unified Communication Networks; Computer Science Department, University of California, Berkeley; January 31, 2001
_____	511	NG, Vincent and HO, Mok Kwan; An Intelligent Agent For Web Advertisements; International Journal of Foundations of Computer Science; Vol. 13, No. 4; pp. 531-554; 2002
_____	512	CHORIANOPOULOS, Kostas et al.; Metrics For Advertisement Effectiveness Measurement In The Interactive TV Environment: The iMEDIA Case; 14th Bled Electronic Commerce Conference; June 25-26, 2001
_____	513	MISHRA, Pinki; Coming to a TV Set Near You; Catalog Age, 18, 8; ProQuest Central, pg. 61; July 2001
_____	514	BONCELLA, Robert; Internet Privacy – At Home and At Work; Communications of the Association for Information Systems; Volume 7, Article 14; September 2001
_____	515	KRAVETS; Robin et al.; A Cooperative Approach to User Mobility; Department of Computer Science, University of Illinois; October 2001
_____	516	ADYA, Atul et al.; Analyzing the Browse Patterns of Mobile Clients; Microsoft Research; November 1, 2001
_____	517	HASHMAN, Steve and KNUDSEN, Steven; The Application of Jini™ Technology to Enhance the Delivery of Mobile Services; PsiNaptic; December 2001
_____	518	Equal Access Networks for real broadband access; Cisco Systems white paper; 2002
_____	519	HOFER, Thomas et al.; Context-Awareness on Mobile Devices - the Hydrogen Approach; Proceedings of the 36th Hawaii International Conference on System Sciences; 2002

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	520	CARTER, Casey and KRAVETS, Robin; User Devices Cooperating to Support Resource Aggregation; Department of Computer Science, University of Illinois; 2002
_____	521	LEKAKOS, Georgios and GIAGLIS, George M.; Delivering Personalized Advertisements in Digital Television: A Methodology and Empirical Evaluation; January 2002
_____	522	PRASAD, Ashutosh et al.; Advertising versus pay-per-view in electronic media; International Journal of Research in Marketing 20, 13-30; 2003
_____	523	PODNAR, Ivana et al.; Mobile Push: Delivering Content to Mobile Users; July 2, 2002
_____	524	Coactive TV: CoTVTM Enables Powerful Interactive TV Now; press release at http://www.teleshuttle.com/cotv/CoTVPR902.htm; September 19, 2002
_____	525	Coactive TV, A new way to enable and profit from Interactive TV now; archived at https://web.archive.org/web/20020920101402/http://www.teleshuttle.com/cotv/default.htm; September 20, 2002
_____	526	REISMAN, Richard R.; Coactive Television Interactive TV, the Web, and 'Co-vergence': A User and Content-Centric Re-visioning; white paper at: http://www.teleshuttle.com/cotv/CoTVWP2.htm; September 2002
_____	527	REKIMOTO; Jun et al.; Proximal Interactions: A Direct Manipulation Technique for Wireless Networking; Interaction Laboratory, Sony Computer Sciences Laboratories, Inc.; 2003
_____	528	BAE, Sung Min et al.; Fuzzy Web Ad Selector Based on Web Usage Mining; IEEE Intelligent Systems; 2003
_____	529	RATSIMOR, Olga et al.; eNcentive: A Framework for Intelligent Marketing in Mobile Peer-To-Peer Environments; Department of Computer Science and Electrical Engineering, University of Maryland; 2003

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	530	LEE, Jeongkeun et al.; Hybrid gateway advertisement scheme for connecting mobile ad hoc networks to the Internet; 2003
_____	531	YUAN, Soe-Tsyr and TSAO, Y.W.; A recommendation mechanism for contextualized mobile advertising; Expert Systems with Applications 24, pp. 399-414; 2003
_____	532	BILCHEV, G. and MARSTON, D.; Personalised Advertising – exploiting the distributed user profile; BT Technology Journal, Vol. 21, No. 1; January 2003
_____	533	HANSEN, Lars M.; Blocking spam with Sendmail; http://www.hansenonline.net/networking/sendmail-spam.html; February 2003
_____	534	VANEM, Erik et al.; Managing Heterogeneous Services And Devices With The Device Unifying Service; Telenor R&D; March 24, 2003
_____	535	ITO, Masaki et al.; Smart Furniture: Improvising Ubiquitous Hot- spot Environment; Graduate School of Media and Governance, Keio University; May 19, 2003
_____	536	MAHLER, Alwin and STEINFIELD, Charles; The Evolving Hot Spot Market for Broadband Access; October 2003
_____	537	SCHULZRINNE, Henning et al.; Ubiquitous Computing in Home Networks; IEEE Communications Magazine; November 2003
_____	538	LEE, Choonhwa et al.; Konark: A system and protocols for device independent, peer-to-peer discovery and delivery of mobile services; IEEE Xplore Document; November 2003
_____	539	SPANGLER, William E. et al.; Using Data Mining To Profile Tv Viewers; Communications of the ACH; Vol. 46, No. 12; pp. 67-72; December 2003
_____	540	PRATA, Alcina et al.; iTV Enhanced System for Generating Multi- Device Personalized Online Learning Environments; 2004

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.:	17/714,965
Applicant:	Shkedi, Roy
Filed:	04/06/2022
Art Unit:	2425
Examiner:	Ekpo, Nnenna
Conf. No.:	7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	541	THAWANI, Amit et al.; Viewing Characteristics based Personalized Ad Streaming in an Interactive TV Environment; Applied Research Group, Satyam Computer Services, Ltd.; 2004
_____	542	HRISTOVA, Nataliya and O'HARE, G.M.P.; Ad-me: Wireless Advertising Adapted to the User Location, Device and Emotions; Proceedings of the 37th Hawaii International Conference on System Sciences; 2004
_____	543	BRAUN, Elmar et al.; Accessing Web Applications With Multiple Context-Aware Devices; Department of Computer Science, Darmstadt University of Technology; 2004
_____	544	THAWANI, Amit et al.; Context Aware Personalized Ad Insertion in an Interactive TV Environment; Applied Research Group, Satyam Computer Services Ltd.; January 2004
_____	545	BAKER, Loren; Behavioral Targeting and Contextual Advertizing; Search Engine Journal; September 1, 2004
_____	546	PHELPS, Joseph E. et al.; Viral Marketing or Electronic Word-of-Mouth Advertising: Examining Consumer Responses and Motivations to Pass Along Email; Journal of Advertising Research; Vol. 44, No. 4; December 2004
_____	547	FALLAHKHAIR, Sanaz; Dual Device User Interface Design for Ubiquitous Language Learning: Mobile Phone and Interactive Television (iTV); Proceedings of the 2005 IEEE International Workshop on Wireless and Mobile Technologies in Education; 2005
_____	548	BELLMAN, Steven et al.; Duane; Interactive Television Advertising: A Research Agenda; ANZMAC 2005 Conference: Advertising/Marketing Communications Issues; 2005
_____	549	YU, Zhiwen et al.; User Preference Learning for Multimedia Personalization in Pervasive Computing Environment; KES 2005, LNAI 3682; pp. 236-242; 2005

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

CONSIDERED: /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	550	KIM, Munjo et al.; Target Advertisement Service Using TV Viewers' Profile Inference; ; PCM 2005, Part T, LNCS 3767, pp. 202-211; 2005
_____	551	HARRIS, Amy and IRELAND, Greg; Enabling IPTV: What Carriers Need to Know to Succeed; white paper from IDC Analyze the Future; May 2005
_____	552	ATIS IPTV Exploratory Group Report and Recommendation to the TOPS Council; Alliance for Telecommunications Industry Solutions; July 2005
_____	553	HARTE, Lawrence; 10 Things You Need to Know About IP Television; iptvmagazine.com; July 2005
_____	554	RIBEIRO-NETO, Berthier et al.; Impedance Coupling in Content-targeted Advertising; SIGIR '05; August 15-19, 2005
_____	555	OpenTV® IPTV Solutions; white paper; OpenTV, Inc.; September 2005
_____	556	WAGNER, Ben; Driving IPTV Growth: The Challenges and Perspectives; Texas Instruments; September 2005
_____	557	JENSEN, Jens F.; Interactive Television: New Genres, New Format, New Content; Proceedings of the Second Australasian Conference on Interactive Entertainment; pp. 89-96; November 23, 2005
_____	558	EPIC – Privacy and Human Rights Report 2006; http://www.worldlii.org/int/journals/EPICPrivHR/2006/PHR2006-Behavior.html; 2006
_____	559	DRAKE, Eddie C.; Next-Generation Architecture for IPTV Middleware; Myrio Corporation; 2006
_____	560	MATE; Sujeet et al.; Movable-Multimedia: Session Mobility in Ubiquitous Computing Ecosystem; MUM '06; December 4-6, 2006
_____	561	GAL-OR, Esther et al.; Targeted Advertising Strategies on Television; Management Science 52(5); pp. 713-725; 2006

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Page 32 of 33

Serial No.:	17/714,965
Applicant:	Shkedi, Roy
Filed:	04/06/2022
Art Unit:	2425
Examiner:	Ekpo, Nnenna
Conf. No.:	7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	562	YOON, Jeong Hyun et al.; TV-Anytime based personalized bi-directional metadata service system; 2006
_____	563	KIM, Jong Woo and DU, Stephen; Design for an Interactive Television Advertising System; Proceedings of the 39th Hawaii International Conference on System Sciences; 2006
_____	564	DE CASTRO, Josephine E. and SHIMAKAWA, Hiromitsu; Mobile Advertisement System Utilizing User's Contextual Information; Proceedings of the 7th International Conference on Mobile Data Management; 2006
_____	565	CAGENIUS, Torbjörn et al.; Evolving the TV experience: Anytime, anywhere, any device; Ericsson Review No. 3; pp. 107-111; 2006
_____	566	TASKIRAN, Cuneyt et al.; Seamless Content Consumption Across Media Delivery Channels and Devices Using Media Bundles; SCTE Conference on Emerging Technologies; January 23-25, 2007
_____	567	MERISAVO, Marko and VESANEN, Jari; The effectiveness of targeted mobile advertising in selling mobile services: an empirical study; Int. J. Mobile Communications, Vol. 4, No. 2, pp. 119-127; 2006
_____	568	FERNÁNDEZ, Yolanda Blanco et al.; AVATAR: An Improved Solution for Personalized TV based on Semantic Inference; Department of Telematics Engineering, University of Vigo; January 11, 2006
_____	569	Front Porch WiFi Messaging and Advertising Solution; Front Porch; February 24, 2006
_____	570	YANG, Wan-Shiou et al.; Mining Social Networks for Targeted Advertising; Proceedings of the 39th Hawaii International Conference on System Sciences; February 24, 2006
_____	571	HARWOOD, Susie; The Revolution Masterclass on cross-media planning; Revolution; London; pp. 62-64; May 2006

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Page 33 of 33

Serial No.:	17/714,965
Applicant:	Shkedi, Roy
Filed:	04/06/2022
Art Unit:	2425
Examiner:	Ekpo, Nnenna
Conf. No.:	7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	572	SPANGLER, William E. et al.; Exploring The Privacy Implications Of Addressable Advertising And Viewer Profiling; Communications of the ACM, Vol. 49, No. 5; May 2006
_____	573	Yahoo's Invalidity Contentions (pages related to parent patents in this family); AlmondNet, Inc. et al. v. Yahoo! Inc., Civil Action No. 16-cv-1557-ILG-SMG, U.S. District Court for the Eastern District of New York; June 30, 2017
_____	574	Exhibits E-1 through E-14 (as revised) of Yahoo's Invalidity Contentions cited above; June 30, 2017 and August 22, 2017

/NNENNA N EKPO/

Examiner: Date Considered: 12/02/2022

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Form PTO-1449 (Modified)
List of Patents and Publications For Information Disclosure Statement

Page 1 of 21

Serial No.:	17/714,965
Applicant:	Shkedi, Roy
Filed:	04/06/2022
Art Unit:	2425 (in parent)
Examiner:	Ekpo, Nnenna (in parent)
Conf. No.:	7004

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	1	5220655	06/15/1993	Tsutsui
	2	5359508	10/25/1994	Rossides
	3	5442633	08/15/1995	Perkins et al.
	4	5572643	11/05/1996	Judson
	5	5636346	06/03/1997	Saxe
	6	5710884	01/20/1998	Dedrick
	7	5717923	02/10/1998	Dedrick
	8	5724424	03/03/1998	Gifford
	9	5724521	03/03/1998	Dedrick
	10	5751956	05/12/1998	Kirsch
	11	5752238	05/12/1998	Dedrick
	12	5774670	06/30/1998	Montulli
	13	5794210	08/11/1998	Goldhaber, et al.
	14	5819092	10/06/1998	Ferguson et al.
	15	5819285	10/06/1998	Damico et al.
	16	5835718	11/10/1998	Blewett
	17	5848396	12/08/1998	Gerace
	18	5855008	12/29/1998	Goldhaber, et al.
	19	5870546	02/09/1999	Kirsch
	20	5870550	02/09/1999	Wesinger Jr. et al.
	21	5901287	05/04/1999	Bull et al.
	22	5918014	06/29/1999	Robinson
	23	5933811	08/03/1999	Angles et al.
	24	5937390	08/10/1999	Hyodo
	25	5937392	08/10/1999	Alberts
	26	5948061	09/07/1999	Merriman, et al.
	27	5959623	09/28/1999	van Hoff et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	28	5963915	10/05/1999	Kirsch
	29	5999912	12/07/1999	Wodarz et al.
	30	6014638	01/11/2000	Burge et al.
	31	6016504	01/18/2000	Arnold et al.
	32	6026368	02/15/2000	Brown et al.
	33	6026369	02/15/2000	Capek
	34	6026374	02/15/2000	Chess
	35	6044376	03/28/2000	Kurtzman, II
	36	6047310	04/04/2000	Kamakura et al.
	37	6049835	04/11/2000	Gagnon
	38	6055513	04/25/2000	Katz, et al.
	39	6073241	06/06/2000	Rosenberg et al.
	40	6128663	10/03/2000	Thomas
	41	6133912	10/17/2000	Montero
	42	6134532	10/17/2000	Lazarus et al.
	43	6141010	10/31/2000	Hoyle
	44	6177931	01/23/2001	Alexander et al.
	45	6188398	02/13/2001	Collins-Rector et al.
	46	6205432	03/20/2001	Gabbard et al.
	47	6216112	04/10/2001	Fuller et al.
	48	6216129	04/10/2001	Eldering
	49	6269361	07/31/2001	Davis et al.
	50	6285987	09/04/2001	Roth et al.
	51	6285991	09/04/2001	Powar
	52	6289318	09/11/2001	Barber
	53	6308202	10/23/2001	Cohn et al.
	54	6324519	11/27/2001	Eldering

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	55	6366298	04/02/2002	Haitsuka et al.
	56	6377986	04/23/2002	Philyaw et al.
	57	6385592	05/07/2002	Angles, et al.
	58	6401075	06/04/2002	Mason et al.
	59	6408278	06/18/2002	Carney et al.
	60	6415270	07/02/2002	Rackson et al.
	61	6415322	07/02/2002	Jaye
	62	6449657	09/10/2002	Stanbach Jr. et al.
	63	6484492	11/26/2002	Meholic, et al.
	64	6487538	11/26/2002	Gupta, et al.
	65	6505169	01/07/2003	Bhagavath et al.
	66	6615247	09/02/2003	Murphy
	67	6633850	10/14/2003	Gabbard et al.
	68	6523041	02/18/2003	Morgan et al.
	69	6654725	11/25/2003	Langheinrich et al.
	70	6766327	07/20/2004	Morgan et al.
	71	6810527	10/26/2004	Conrad et al.
	72	6850988	02/01/2005	Reed
	73	6868392	03/15/2005	Ogasawara
	74	6868525	03/15/2005	Szabo
	75	6907566	06/14/2005	McElfresh et al.
	76	6983379	01/03/2006	Spalink et al.
	77	6985882	01/10/2006	Del Sesto
	78	7007074	02/28/2006	Radwin
	79	7038637	05/02/2006	Eller et al.
	80	7039594	05/02/2006	Gersting
	81	7039599	05/02/2006	Merriman et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	82	7043531	05/09/2006	Seibel et al.
	83	7058593	06/06/2006	Merritt
	84	7065550	06/20/2006	Raghunandan
	85	7085682	08/01/2006	Heller et al.
	86	7155508	12/26/2006	Sankuratripati, et al.
	87	7165105	01/16/2007	Reiner et al.
	88	7181438	02/20/2007	Szabo
	89	7188076	03/06/2007	Bensemana
	90	7194424	03/20/2007	Greer et al.
	91	7246150	07/17/2007	Donoho et al.
	92	7299195	11/20/2007	Tawakol et al.
	93	7337133	02/26/2008	Bezos et al.
	94	7360251	04/15/2008	Spalink et al.
	95	7386599	06/10/2008	Piersol et al.
	96	7530020	05/05/2009	Szabo
	97	7552069	06/23/2009	Kepecs
	98	7558795	07/07/2009	Malik et al.
	99	7580855	08/25/2009	Fernandes
	100	7606810	10/20/2009	Jeavons
	101	7698422	04/13/2010	Vanderhook et al.
	102	7739594	06/15/2010	Vasilik
	103	7756744	07/13/2010	Schiff et al.
	104	7797274	09/14/2010	Strathearn et al.
	105	7822636	10/26/2010	Ferber et al.
	106	7899801	03/01/2011	Bronson et al.
	107	7930252	04/19/2011	Bender et al.
	108	7949553	05/24/2011	Wisniewski et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	109	8015184	09/06/2011	Zito et al.
	110	8019881	09/13/2011	Sandhu et al.
	111	8036929	10/11/2011	Reisman
	112	8050970	11/01/2011	Barry et al.
	113	8082332	12/20/2011	Gilbert
	114	8095450	01/10/2012	Walker et al.
	115	8117068	02/14/2012	Shuster et al.
	116	8135799	03/13/2012	Mekikian
	117	8175908	05/08/2012	Anderson
	118	8364540	01/29/2013	Soroca et al.
	119	8510174	08/13/2013	Walker et al.
	120	8527352	09/03/2013	Chatwin et al.
	121	8554940	10/08/2013	Jennings et al.
	122	8751310	06/10/2014	van Datta et al.
	123	8768766	07/01/2014	Ellis et al.
	124	8880677	11/04/2014	Julia et al.
	125	8930204	01/06/2015	Igoe et al.
	126	8996406	03/31/2015	Jain
	127	9117217	08/25/2015	Wilson et al.
	128	20010014868	08/16/2001	Herz et al.
	129	20010020242	09/06/2001	Gupta, et al.
	130	20010034637	10/25/2001	Lin et al.
	131	20010034646	10/25/2001	Hoyt et al.
	132	20010039515	11/08/2001	Mayadas
	133	20010042002	11/15/2001	Koopersmitch
	134	20010044743	11/22/2001	McKinley et al.
	135	20020002541	01/03/2002	Williams

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	136	20020026360	02/28/2002	McGregor et al.
	137	20020038350	03/28/2002	Lambert et al.
	138	20020046286	04/18/2002	Caldwell et al.
	139	20020072965	06/13/2002	Merriman et al.
	140	20020072971	06/13/2002	DeBusk, et al.
	141	20020077891	06/20/2002	Castle et al.
	142	20020082910	06/27/2002	Kontogouris
	143	20020082923	06/27/2002	Merriman, et al.
	144	20020091568	07/11/2002	Kraft et al.
	145	20020099603	07/25/2002	Bandura
	146	20020112013	08/15/2002	Walsh
	147	20020112171	08/15/2002	Ginter et al.
	148	20020123928	09/05/2002	Eldering et al.
	149	20020133399	09/19/2002	Main
	150	20020143632	10/03/2002	Walter et al.
	151	20020161633	10/31/2002	Jacob et al.
	152	20020161636	10/31/2002	Takahashi et al.
	153	20020184088	12/05/2002	Rosenberg
	154	20020194058	12/19/2002	Eldering
	155	20030023481	01/30/2003	Calvert et al.
	156	20030028451	02/06/2003	Ananian
	157	20030036944	02/20/2003	Lesandrini et al.
	158	20030046156	03/06/2003	Cromer, et al.
	159	20030046689	03/06/2003	Gaos
	160	20030050827	03/13/2003	Hennessey
	161	20030050863	03/13/2003	Radwin
	162	20030065558	04/03/2003	Shaw et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	163	20030065805	04/03/2003	Barnes Jr.
	164	20030083931	05/01/2003	Lang
	165	20030083938	05/01/2003	Smith et al.
	166	20030101126	05/29/2003	Cheung et al.
	167	20030110497	06/12/2003	Yassin, et al.
	168	20030135413	07/17/2003	Nishi, et al.
	169	20030135853	07/17/2003	Goldman, et al.
	170	20030149938	08/07/2003	McElfresh, et al.
	171	20030154126	08/14/2003	Gehlot et al.
	172	20030171990	09/11/2003	Rao, et al.
	173	20030208404	11/06/2003	Michie
	174	20030220837	11/27/2003	Asayama
	175	20030229542	12/11/2003	Morrisroe
	176	20040024632	02/05/2004	Perry
	177	20040044566	03/04/2004	Bostelmann et al.
	178	20040059632	03/25/2004	Kang et al.
	179	20040068436	04/08/2004	Boubek et al.
	180	20040117322	06/17/2004	Bjorksten et al.
	181	20040176995	09/09/2004	Fusz
	182	20040186776	09/23/2004	Llach
	183	20040193488	09/30/2004	Khoo et al.
	184	20040260605	12/23/2004	McIntyre et al.
	185	20050010472	01/13/2005	Quatse et al.
	186	20050021397	01/27/2005	Cui, et al.
	187	20050028188	02/03/2005	Latona, et al.
	188	20050038698	02/17/2005	Lukose et al.
	189	20050038702	02/17/2005	Merriman et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	190	20050080772	04/14/2005	Bem
	191	20050120045	06/02/2005	Klawon
	192	20050120091	06/02/2005	Casais
	193	20050131762	06/16/2005	Bharat et al.
	194	20050131762	06/16/2005	Bharat et al.
	195	20050165644	07/28/2005	Beyda, et al.
	196	20050177613	08/11/2005	Dresden
	197	20050187818	08/25/2005	Zito et al.
	198	20050192863	09/01/2005	Mohan
	199	20050216336	09/29/2005	Roberts et al.
	200	20050216823	09/29/2005	Petersen, et al.
	201	20050222905	10/06/2005	Wills
	202	20050278443	12/15/2005	Winner, et al.
	203	20050289590	12/29/2005	Cheok et al.
	204	20060020508	01/26/2006	Gorti et al.
	205	20060020510	01/26/2006	Vest
	206	20060026061	02/02/2006	Collins
	207	20060026061	02/02/2006	Collins
	208	20060026064	02/02/2006	Collins
	209	20060029296	02/09/2006	King et al.
	210	20060053047	03/09/2006	Garcia et al.
	211	20060064351	03/23/2006	Hamor, et al.
	212	20060069614	03/30/2006	Agarwal, et al.
	213	20060095322	05/04/2006	Dierks
	214	20060098899	05/11/2006	King et al.
	215	20060106710	05/18/2006	Meek, et al.
	216	20060116924	06/01/2006	Angles, et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	217	20060122879	06/08/2006	O'Kelley
	218	20060129455	06/15/2006	Shah
	219	20060161553	07/20/2006	Woo
	220	20060183467	08/17/2006	Stewart
	221	20060200377	09/07/2006	Wolfe et al.
	222	20060212349	09/21/2006	Brady
	223	20060212350	09/21/2006	Ellis et al.
	224	20060224496	10/05/2006	Sandholm et al.
	225	20060253432	11/09/2006	Eagle et al.
	226	20060271690	11/30/2006	Banga et al.
	227	20060282327	12/14/2006	Neal et al.
	228	20060287916	12/21/2006	Starr et al.
	229	20070027850	02/01/2007	Chan et al.
	230	20070027901	02/01/2007	Chan et al.
	231	20070038516	02/15/2007	Apple et al.
	232	20070067297	03/22/2007	Kublickis
	233	20070079331	04/05/2007	van Datta et al.
	234	20070088606	04/19/2007	Jouppi et al.
	235	20070121843	05/31/2007	Atazky et al.
	236	20070130015	06/07/2007	Starr et al.
	237	20070130070	06/07/2007	Williams
	238	20070150234	06/28/2007	Wicks
	239	20070150457	06/28/2007	Goldman et al.
	240	20070150468	06/28/2007	Goldman et al.
	241	20070150483	06/28/2007	Goldman et al.
	242	20070150721	06/28/2007	Goldman et al.
	243	20070157228	07/05/2007	Bayer et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	244	20070157245	07/05/2007	Collins
	245	20070162396	07/12/2007	Goldman et al.
	246	20070179845	08/02/2007	Jain
	247	20070179846	08/02/2007	Jain et al.
	248	20070179848	08/02/2007	Jain
	249	20070179849	08/02/2007	Jain
	250	20070179853	08/02/2007	Feige et al.
	251	20070192166	08/16/2007	Van Luchene
	252	20070192168	08/16/2007	Van Luchene
	253	20070198385	08/23/2007	McGill et al.
	254	20070203794	08/30/2007	Hajdukiewicz et al.
	255	20070214037	09/13/2007	Shubert et al.
	256	20070214048	09/13/2007	Chan et al.
	257	20070239517	10/11/2007	Chung et al.
	258	20070239518	10/11/2007	Chung et al.
	259	20070260624	11/08/2007	Chung et al.
	260	20080052392	02/28/2008	Webster et al.
	261	20080109306	05/08/2008	Maigret et al.
	262	20080262901	10/23/2008	Banga et al.
	263	20080263633	10/23/2008	Banga et al.
	264	20080319850	12/25/2008	Shaul et al.
	265	20090106100	04/23/2009	Mashinsky
	266	20090222329	09/03/2009	Ramer et al.
	267	20090234711	09/17/2009	Ramer et al.
	268	20090234861	09/17/2009	Ramer et al.
	269	20090254971	10/08/2009	Herz et al.
	270	20090281954	11/12/2009	Waelbroeck et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	271	20090299857	12/03/2009	Brubaker
	272	20100257053	10/07/2010	Ferber et al.
	273	20100268605	10/21/2010	Waelbroeck et al.
	274	20100293057	11/18/2010	Haveliwala et al.
	275	20100293057	11/18/2010	Haveliwala et al.
	276	20100332583	12/30/2010	Szabo
	277	20110112903	05/12/2011	Goldman et al.
	278	20110191140	08/04/2011	Newman et al.
	279	20110258049	10/20/2011	Ramer et al.
	280	20120209705	08/16/2012	Ramer et al.
	281	20120209706	08/16/2012	Ramer et al.
	282	20120209707	08/16/2012	Ramer et al.
	283	20120209708	08/16/2012	Ramer et al.
	284	20120209709	08/16/2012	Ramer et al.
	285	20120209710	08/16/2012	Ramer et al.
	286	20120215635	08/23/2012	Ramer et al.
	287	20120215639	08/23/2012	Ramer et al.
	288	20130097019	04/18/2013	Ferber et al.
	289	20130097026	04/18/2013	Ferber et al.
	290	20130232001	09/05/2013	van Datta et al.
	291	20130254035	09/26/2013	Ramer et al.
	292	20140324576	10/30/2014	van Datta et al.
	293	6832207	12/14/2004	Shkedi
	294	6925440	08/02/2005	Shkedi
	295	6973436	12/06/2005	Shkedi
	296	7072853	07/04/2006	Shkedi
	297	7428493	09/23/2008	Shkedi

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	298	7454364	11/18/2008	Shkedi
	299	7822637	10/26/2010	Shkedi
	300	7822639	10/26/2010	Shkedi
	301	7861260	12/28/2010	Shkedi
	302	7890609	02/15/2011	Shkedi
	303	7979307	07/12/2011	Shkedi
	304	8051444	11/01/2011	Shkedi
	305	8239264	08/07/2012	Shkedi
	306	8244574	08/14/2012	Shkedi
	307	8244582	08/14/2012	Shkedi
	308	8244583	08/14/2012	Shkedi
	309	8244586	08/14/2012	Shkedi
	310	8280758	10/02/2012	Shkedi
	311	8281336	10/02/2012	Shkedi
	312	8494904	07/23/2013	Shkedi
	313	8566164	10/22/2013	Shkedi et al.
	314	8589210	11/19/2013	Shkedi
	315	8595069	11/26/2013	Shkedi et al.
	316	8600815	12/03/2013	Shkedi
	317	8607267	12/10/2013	Shkedi
	318	8677398	03/18/2014	Shkedi
	319	8683502	03/25/2014	Shkedi et al.
	320	8695032	04/08/2014	Shkedi
	321	8713600	04/29/2014	Shkedi
	322	8775249	07/08/2014	Shkedi
	323	8997138	03/31/2015	Shkedi
	324	9071886	06/30/2015	Shkedi

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	325	9078035	07/07/2015	Shkedi et al.
	326	9083853	07/14/2015	Shkedi
	327	9131282	09/08/2015	Shkedi
	328	9226019	12/29/2015	Shkedi
	329	9271024	02/23/2016	Shkedi et al.
	330	9351053	05/24/2016	Shkedi
	331	9369779	06/14/2016	Shkedi
	332	9571865	02/14/2017	Shkedi
	333	9591380	03/07/2017	Shkedi et al.
	334	9756372	09/05/2017	Shkedi
	335	9800917	10/24/2017	Shkedi
	336	9813778	11/07/2017	Shkedi
	337	10009636	06/26/2018	Shkedi
	338	10026100	07/17/2018	Shkedi
	339	10070200	09/04/2018	Shkedi
	340	10178442	01/08/2019	Shkedi
	341	10306282	05/28/2019	Shkedi
	342	10321198	06/11/2019	Shkedi et al.
	343	10405058	09/03/2019	Shkedi et al.
	344	10628857	04/21/2020	Shkedi
	345	10645438	05/05/2020	Shkedi
	346	10715878	07/14/2020	Shkedi
	347	10771860	09/08/2020	Shkedi et al.
	348	10839423	11/17/2020	Shkedi
	349	10984445	4/20/2021	Shkedi
	350	11082753	4/8/2021	Shkedi et al.
	351	11093970	9/12/2019	Shkedi

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

U.S. PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee or Applicant of Cited Document
	352	11095952	8/17/2021	Shkedi et al.
	353	11282105	03/03/2022	Shkedi et al.
	354	20050086112	04/21/2005	Shkedi
	355	20090049468	02/19/2009	Shkedi
	356	201100035256	2/10/2011	Shkedi et al.
	357	20110040607	2/17/2011	Shkedi
	358	20090172723	07/02/2009	Shkedi et al.

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

FOREIGN PATENT DOCUMENTS

Examiner Initials	Item	Document Number	Publication Date MM/DD/YYYY	Name of Patentee	Translation Yes / No
_____	359	WO 2005/046237	05/19/2005	Gutta et al.	No
_____	360	WO 2004/084109	09/30/2004	NHN Corp.	No
_____	361	KR 10-2004-0099222	11/26/2004	Echo Marketing Co., Ltd.	Yes
_____	362	KR 10-2001-0096343	11/07/2001	MyOne Card Corp.	Yes
_____	363	JP 2003-331183	11/21/2003	Masuda et al	Yes
_____	364	KR 20020037411	05/21/2002	Sin	Yes
_____	365	JP 2012-073990	04/12/2012	Takada	Yes
_____	366	WO1997021183A1	06/12/1997	Bell Communications Research, Inc.	No
_____	367	WO1999057660A1	11/11/1999	IWEB LTD.	No
_____	368	CA 2246746	03/12/1998	Blumenau	No
_____	369	WO 99/44159	09/02/1999	Jaye	No
_____	370	WO 01/54034	07/26/2001	Phillips et al.	No
	371	WO2005022887A2	03/10/2005	Gottfurcht	No

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	372	Engage Technologies, Inc., Form S-1A, filed with the SEC, pp. 36-50, July 19, 1999
_____	373	The e-Marketing Report, Morgan Stanley Dean Whitter, May, 2000, portions
_____	374	ClickOver, Inc. introduces ClickWise, the first intelligent, high-performance ad management system for electronic advertising, Business Wire, March 24, 1997
_____	375	Advertiser Solutions, https://web.archive.org/web/20050325042542/http://www.premiumnetwork.com/advertisers.shtml, archived March 25, 2005
_____	376	baseview.com website (ClassManagerPro) from www.archive.org, archived January 16, 1997
_____	377	Media Kit – Marketers, RevenueScience, https://web.archive.org/web/20051124105455/http://www.revenuescience.com/datafile/Mediakit_marketers_02.pdf, archived November 24, 2005

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
		Publisher Solutions, https://web.archive.org/web/20050325061154/http://www.premiumnetwork.com/publishers.shtml, archived March 25, 2005
	378	
	379	S-1/A SEC Filing, filed by FlyCast Communications Corporation on April 30, 1999
	380	RevenueScience (for Publishers), Behavioral Targeting, Delivering on the Promise of One-to-One Marketing (2005)
	381	24/7 RealMedia, The Science of Digital Marketing, Company Overview (2007)
	382	aaddzz.com webpages from www.archive.org, archived January 30, 1998
	383	AdForce, Inc., User Guide, A Complete Guide to AdForce, Version 2.6 (1998)
	384	FAWCETT, Adrienne W., Publishers thankful but nervous about boom, Magazinedotcom, Hot spending boosts ..., adage.com, November 1, 1999
	385	Ad-Star.com website archive from www.archive.org (archived on April 12, 1997 and February 1, 1997)
	386	CHANG et al., "Goodies" in exchange for Consumer Information on the Internet, The Economics and Issues, Proceedings of the Thirty-First Annual Hawaii International Conference on System Sciences, 1998
	387	avenuea.com webpage, FAQ, www.avenuea.com/corpinfo/faq.html, archived October 2, 1999
	388	Advertising-A Unified Approach to Adverstising Management, Baseview Products, Inc. brochure, October 14, 1999
	389	RIBEIRO-NETO et al., Extracting Semi-Structured Data Through Examples, Proceedings of the Eighth international Conference on Information and Knowledge Management (1999)
	390	RASKUTTI and BEITZ, Acquiring Customer Preferences for Information Filtering, a Heuristic-Statistical Approach, Fifth Annual Conference on User Modeling (1996)
	391	AdKnowledge, Campaign Manager-Quick Reference Card (1996)
	392	EMERT, 2 Web Ad Firms Decide to Marry, Palo Alto's Focalink, ClickOver Merging, San Francisco Chronicle, November 17, 1997
	393	ELDERBROCK & BORWANKAR, Building Successful Internet Businesses, IDG Worldwide (1996)
	394	SCHUMANN, David W. & THORSON, Ester, Advertising and the World Wide Web, Lawrence Erlbaum Associates (1999)
	395	WILLIAMSON, Debra Aho, Web ads mark 2nd birthday with decisive issues ahead, adage.com, October 21, 1996

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
	396	WILLIAMSON, Debra Aho, Targeting distinguishes AdForce from the pack, adage.com, March 1, 1999
	397	Doubleclick.com webpages from www.archived.org, archived February 5, 1998
	398	TITTEL et al., Foundations of World Wide Web Programming with HTML & CGI, IDG Books Worldwide (1995)
	399	HARRISON et al., An Emerging Marketplace for Digital Advertising Based on Amalgamated Digital Signage Networks, IEEE International Conference on E-Commerce, June 24-27, 2003
	400	GRAHAM, Ian S., HTML Sourcebook, A Complete Guide to HTML, John Wiley & Sons (1995)
	401	Aaddzz Brochure, Information Access Technologies, Inc. (1997)
	402	GILBERT, SPECIAL REPORT, Agencies centralize Web ad serving, adage.com, Mar 1, 1999
	403	PALMER & ERIKSEN, Digital Newspapers Explore Marketing on the Internet, Communications of the ACM, Vol. 42, No. 9, September 1999
	404	ANGWIN, Julia, Focalink Will Monitor World of Online Ads, San Francisco Chronicle, February 7, 1996
	405	GALLAGHER & PARSONS, A Framework for Targeting Banner Advertising on the Internet, Proceedings of the Thirtieth Hawaii International Conference on System Sciences, January 1997
	406	KAZIENKO & ADAMSKI, Personalized Web Advertising Methods, Third International Conference, AH 2004, August 23-26, 2004
	407	CLELAND, Kim, Accipiter latest entrant in race to manage ads, company sets sights on helping content providers target audience, Adage.com, September 9, 1996
	408	CLELAND, Kim, New Player joins fray of Web ad targeters, adage.com, February 24, 1997
	409	BAYNE, Kim M., AdKnowledge rolls out Web ad evaluation tool, adage.com, July 8, 1999
	410	LANGHEINRICH et al., Unintrusive Customization Techniques for Web Advertising, COMPUTER NETWORKS 31, 1259–1272, May 1999
	411	MAEKAWA et al., A Java-Based Information Browsing System in a Remote Display Environment, IEEE (2004)
	412	MEEKER, Mary, The Internet Advertising Report, Harper Business (1997)
	413	YEAGER & MCGRATH, Web Server Technology, The Advance Guide for World Wide Web Information Providers, Morgan Kaufmann Publishers (1996)

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
	414	net.Genesis & Devra Hall, Build a Website, The Programmer's Guide to Creating, Building, and Maintaining a Web Presence, Prima Publishing (1995)
	415	NetGravity Launches AdServer 3.5 with Beefed Up Targeting, clickz.com, September 10, 1998
	416	NOVAK & HOFFMAN, New Metrics for New Media, Toward the Development of Web Measurement Standards, Project 2000, elab.vanderbilt.edu, September 26, 1996
	417	RealMedia.com webpages, www.archive.org, archived on May 15, 19977 and June 7, 1997
	418	BRUNER, Rick E., Interactive, NetGravity to launch new ad service bureau, adage.com, December 1, 1997
	419	BRUNER, Rick E., Small networks chase per-click ad business, Aaddzz, ValueClick compete, AdvertisingAge, September 8, 1997
	420	ZEFF & ARONSON, Advertising on the Internet, 1st Edition, John Wiley & Sons (1997)
	421	ZEFF & ARONSON, Advertising on the Internet, 2nd Edition, John Wiley & Sons (1999)
	422	SINGH & DALAL, Web Home Pages as Advertisements, Communications of the ACM, August 1999
	423	The ClickZ Network, Netscape Selects Real Media's Open AdStream, January 13, 1998
	424	The ClickZ Network, Real Media Signs More Media Companies, December 30, 1998
	425	www.Spinbox.com webpages, www.archive.org, archived beginning Nov. 11, 1998
	426	BAKOS, The Emerging Role of Electronic Marketplaces on the Internet, Communications of the ACM, Vol. 41, No. 8, August 1998
	427	KOHDA, Youiji & ENDO, Susumu, Ubiquitous Advertising on the WWW, Merging Advertisement on the Browser, Fifth International World Wide Web Conference, May 1996
	428	Baseview.com webpages from www.archive.org, archived January 16, 1997
	429	AdForce, Inc., Quick Reference Guide, Version 3.0 (1999)
	430	AdManagerPro, Ad Entry Manual, User's Manual V.1.05, 1997
	431	AdManagerPro, Ad Entry Manual, User's Manual V.2.0, 1997
	432	AdManagerPro, Administration and Receivables Manual, Administrator's Manual Version 1.0.5, November 1997
	433	AdManagerPro, Administration and Receivables Manual, Administrator's Manual Version 2.0, June 1998

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
		Commercialization of the World Wide Web, The Role of Cookies, vanderbilt.edu
	434	web pages, archive.org, February 1, 1997
	435	CRABTREE et al., Adaptive Personal Agents, Personal Technologies (1998)
	436	DALTON, Gregory, OPS, Answer to Cookies?, InformationWeek, October 13, 1997
	437	DUNN et al., Privacy and Profiling on the Web, Submitted to W3C on June 2, 1997
	438	firefly.com web pages, archive.org, archived 1998-2000
		GILBERT, Jennifer, CMGI's Web dealmaking hits big-time, Advertising Age, October
	439	4, 1999
		HERSLEY et al., Implementation of OPS Over HTTP, Submitted to W3C on June 2,
	440	1997
		HERSLEY et al., Proposal for an Open Profiling Standard, Submitted to W3C on June
	441	2, 1997
		HERSLEY et al., Standard Practices of OPS Systems, Submitted to W3C on June 2,
	442	1997
		junkbusters.com web pages, archive.org, archived April 1997 through November
	443	2000
		KARPINSKI, Richard, ICE application to revolutionize e-commerce, Business
	444	Marketing, March 1998
	445	matchlogic.com web pages, archive.org, April 13, 2000
		ROSCHEISEN, R. Martin, A Network-Centric Design for Relationship-Based Rights
	446	Management (Dissertation), December 1997
		KRISTOL, D. & MONTULLI, L., HTTP State Management Mechanism, Internet
	447	Engineering Task Force, available at https://tools.ietf.org/html/rfc2109, February 1997
		BERNERS-LEE, TIM, Information Management: A Proposal, CERN, available at
	448	http://www.w3.org/History/1989/proposal.html, May 1990
		BERNERS-LEE, TIM, Uniform Resource Locators (URL), A Syntax for the Expression
		of Access Information of Objects on the Network, Internet Engineering Task Force,
	449	available at http://www.w3.org/Addressing/URL/url-spec.txt, March 21, 1994
		BERNERS-LEE, TIM, The Original HTTP as defined in 1991, W3C, available at
	450	http://www.w3.org/Protocols/HTTP/AsImplemented.html (1991)
		BERNERS-LEE, TIM, et al., Hypertext Transfer Protocol – HTTP/1.0, Network
		Working Group Request for Comments: 1945, Internet Engineering Task Force,
	451	available at https://tools.ietf.org/html/rfc1945, May 1996
		LEINER, BARRY M. ET AL., Brief History of the Internet, Internet Society, October
	452	15, 2012

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
	453	BERNERS-LEE, T. & CONNOLLY, D., Hypertext Markup Language – 2.0, Network Working Group Request for Comments: 1866, Internet Engineering Task Force, available at https://tools.ietf.org/html/rfc1866, November 1995
	454	TUNGATE, MARK, Adland-A Global History of Advertising, 2d Ed. (2013)
	455	JOHNSON, BRADLEY, Interactive, Planet Oasis Gives Web Sites Promotion Clout, Advertising Age, available at http://adage.com/article/news/interactive-planet-oasis-web-sites-promotion-clout/78467, July 8, 1996
	456	JONES, DUNCAN, The History of Pay Per Click (PPC) Marketing, Online Marketing Bog (www.duncanjonesnz.com) available at https://www.duncanjonesnz.com/history-of-ppc-marketing, September 23, 2014
	457	MOUKHEIBER, Z., DoubleClick is watching you, Forbes Magazine, November 4, 1996
	458	FIELDING, R. et al., Hypertext Transfer Protocol – HTTP/1.1, Network Working Group, RFC 2068, January 1997
	459	GLENN, A. and MILLMAN, D., Access Management of Web-based Services, Ariel Glenn, D-Lib Magazine, September 1998
	460	SAMAR, V., Single Sign-on Using Cookies for Web Applications, IEEE 8th Intl. Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises, June 16-18, 1999
	461	KIYOMITSU, H., An E-Commerce Technology: Web Personalization on the Patron Web, Kobe University Economic Review 46 (2000)
	462	WARKENTIN, M. and BAPNA, R., Intelligent Agent-Based Data Mining in Electronic Markets, AMCIS 2000 Proceedings (2000)
	463	BONCELLA, R., Internet Privacy – At Home and At Work, Communications of AIS, Vol. 7, Article 14, September 2001
	464	New Engage ECHO Remarketing Solution Aimed at Decreasing Customer Acquisition Costs, Improving Customer Retention, BusinessWire, September 25, 2000
	465	BILCHEV, G. and MARSTON, D., Personalized advertising – exploiting the distributed user profile, BT Technology Journal, Vol. 21, No. 1, January 2003
	466	PAN, Y. and ZINKHAN, G.M., Netvertising characteristics, challenges and opportunities: a research agenda, Int. J. Internet Marketing and Advertising, Vol. 1, No. 3, 2004

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425 (in parent)
Examiner: Ekpo, Nnenna (in parent)
Conf. No.: 7004

OTHER ART

Examiner Initials	Item	AUTHOR, Title, Date, Pertinent Pages, Etc.
_____	467	Google Launches Self-Service Advertising Program, http://googlepress.blogspot.com/2000/10/google-launches-self-service.html, October 23, 2000
_____	468	Google Introduces New Pricing For Popular Self-Service Online Advertising Program, http://googlepress.blogspot.com/2002/02/google-introduces-new-pricing-for.html, February 20, 2002
_____	469	AdWords Separate Content Bidding, http://googlepress.blogspot.com/2005/11/adwords-separate-content-bidding_22.html, November 22, 2005
_____	470	WashingtonPost.Newsweek Interactive Selects Revenue Science Audience SearchTM for Behavioral Targeting Services, https://www.audiencescience.com/washingtonpost-newsweek-interactive-selects-revenue-science-audience-search-for-behavioral-targeting-services/, September 20, 2004
_____	471	Revenue Science Hosts Behavioral Targeting Workshop at Interactive Advertising World, https://www.audiencescience.com/revenue-science-hosts-behavioral-targeting-workshop-at-interactive-advertising-world/, September 17, 2004
_____	472	Online Advertising and Privacy Survey Shows Consumers Hold Strong Preference for Targeted Advertising, https://www.audiencescience.com/online-advertising-and-privacy-survey-shows-consumers-hold-strong-preference-for-targeted-advertising/, September 9, 2004
_____	473	Yahoo's Invalidity Contentions (pages related to parent patents in this family); AlmondNet, Inc. et al. v. Yahoo! Inc., Civil Action No. 16-cv-1557-ILG-SMG, U.S. District Court for the Eastern District of New York: June 30, 2017
_____	474	Exhibits C-1 through C-12 (as revised) of Yahoo's Invalidity Contentions cited above; June 30, 2017 and August 22, 2017

Examiner: _____ Date Considered: _____

EXAMINER: Initial if reference considered whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

Doc code: IDS
Doc description: Information Disclosure Statement (IDS) Filed

PTO/SB/08a (01-22)
Approved for use through 05/31/2024. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (**Not for submission under 37 CFR 1.99)**		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Kekpo, Nnenna
	Attorney Docket Number	TV1-CON7

U.S.PATENTS						
Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	7843923		2010-11-30	Baum	
	2	7899862		2011-03-11	Appelman et al.	
	3	6034689		2000-03-07	White et al.	
	4	7958224		2011-06-07	Amjadi	
	5	8131271		2012-03-06	Ramer et al.	
	6	6915271		07/05/2005	Meyer et al.	
	7	7251628		07/31/2007	Edlund et al.	
	8	5835896		11/10/1998	Fisher et al.	

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Kekpo, Nnenna
	Attorney Docket Number	TV1-CON7

	9	7356045		2008-04-08	Satapati et al.	
	10	8180674		2012-05-15	Madhavan	
	11	9495682		2016-11-15	Vallier	
	12	7908645		2011-03-15	Varghese	
	13	7299195		2007-11-20	Tawakol et al.	
	14	7805332		2010-09-28	Wilson	
	15	7894448		2011-02-22	Lillibridge et al.	
	16	8296437		2012-10-23	Pankratov	
	17	7853255		2010-12-14	Karaoguz et al.	
	18	8040819		2011-10-18	Harrington et al.	
	19	7039699		2006-05-02	Narin et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Kekpo, Nnenna
	Attorney Docket Number	TV1-CON7

	20	11301898		2022-04-12	Shkedi	
	21	11303973		2022-04-12	Shkedi et al.	
	22	11368763		2022-06-21	Shkedi et al.	

If you wish to add additional U.S. Patent citation information please click the Add button.

U.S.PATENT APPLICATION PUBLICATIONS

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20070157247		07/05/2005	Cordray et al.	
	2	20080113674		2008-05-15	Baig	
	3	20070067459		2007-03-22	Baal-Haness	
	4	20030229893		2003-12-11	Sgaraglino	
	5	20080281697		2008-11-13	Whitehead	
	6	20060069616		2006-03-30	Bau	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Kekpo, Nnenna
	Attorney Docket Number	TV1-CON7

	7	20040215515		2004-10-28	Perry	
	8	20050086109		2005-04-21	McFadden et al.	
	9	20070061331		2007-03-15	Ramer et al.	
	10	20040147265		2004-07-29	Kelley et al.	
	11	20010020242		2001-09-06	Gupta et al.	
	12	20070260514		2007-11-08	Burdick et al.	
	13	20060242294		2006-10-26	Damick et al.	
	14	20040107287		2004-06-03	Ananda et al	
	15	20030207685		2003-11-06	Rankin	
	16	20040133689		2004-07-08	Vasisht	
	17	20020116313		2002-08-22	Detering	

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Kekpo, Nnenna
	Attorney Docket Number	TV1-CON7

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] ¡	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button

NON-PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	Fletcher et al., Practical Web Traffic Analysis: Standard, Privacy, Techniques, Results (2002); Chapters 1 and 3	☐
	2	Microsoft Computer Dictionary (5th ed. 2002); pp. vii-viii, 444	☐
	3	ZWICKY, Building Internet Firewalls (2d ed. 2000): table of contents, ch. 5	☐
	4	Hahn, The Internet Complete Reference (2d ed. 1996); chapters 10, 25, 26	☐
	5	P. Srisuresh et al., Request for Comments (RFC) 2663, IP Network Translator (NAT) Terminology and Considerations, August 1999 (“RFC 2663”)	☐
	6	Request for Comments (RFC) 793, Transmission Control Protocol, DARPA Internet Program Protocol Specification, September 1981 (“RFC 793”)	☐

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Kekpo, Nnenna
	Attorney Docket Number	TV1-CON7

	7	Robbin Zeff & Brad Aronson, Advertising on the Internet (2d ed. 1999) ("Zeff"); chapters 1, 2, 4, 6, 7, 9, and 10	☐
	8	RFC 950, "Internet Standard Subnetting Procedure," August 1985; pp. 1-18	☐
	9	RFC 1531, "Dynamic Host Configuration Protocol," October 1993; pp. 1-39	☐
	10	Definition of "modem," H. Newton, Newton's Telecom Dictionary, CMP Books, 2000	☐

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature	/NNENNA N EKPO/	Date Considered	11/28/2022

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Kekpo, Nnenna
	Attorney Docket Number	TV1-CON7

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Louis J. Hoffman/	Date (YYYY-MM-DD)	2022-08-02
Name/Print	Louis J. Hoffman	Registration Number	38918

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

PE2E SEARCH - Search History (Prior Art)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	British Equivalents	Time Stamp
L1	152	Shkedi-Roy$.in.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/08 09:30 AM
L2	2	"20190124412"	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/08 09:34 AM
L3	21416	(H04N21/812, /64322, /42684, /4532, /6125, /25891, /44222$.CPC.)	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/08 11:07 AM
L4	12905	(G06F16/248$.CPC.)	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/08 11:08 AM
L5	13540	(G06Q30/0241, /0269, /0256, /0255, /0251$.CPC.)	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/08 11:10 AM
L6	8242	(H04L61/2007$.CPC.)	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/08 11:11 AM
L7	34062	(H04L67/306$.CPC.)	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/08 11:11 AM
L8	15397	(H04L12/66$.CPC.)	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/08 11:11 AM
L9	31332	(G06F16/951$.CPC.)	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/08 11:12 AM
L10	130454	L3 OR L4 OR L5 OR L6 OR L7 OR L8 OR L9	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/08 11:13 AM
L11	6979	L10 and ((mobile cellular cell portable handheld) near2 (device)) same ("STB" settopbox$2 "STT" set$top$terminal or receiver)	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/08 11:16 AM
L12	6901	L11 and ("LAN" local area network)	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/08 11:17 AM
L13	4404	L11 and ("LAN" local$area$network)	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/08 11:18 AM
L14	366	L10 and ((mobile cellular cell portable handheld) near2 (device)) same ("STB" settopbox$2 "STT" set$top$terminal or	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/08 11:18 AM

		receiver) same ("LAN" local$area$network)					
L15	2	"20090217319"	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/08 11:50 AM
L16	20	("2009/0217319").URPN.	(USPAT)	OR	OFF	OFF	2020/01/08 11:51 AM
L17	35	("2008/0259906").URPN.	(USPAT)	OR	OFF	OFF	2020/01/08 01:07 PM
L18	140	("20010003184" | "20020013943" | "20020032603" | "20020078444" | "20020082910" | "20020120933" | "20020123928" | "20020124249" | "20020124253" | "20020129368" | "20020194058" | "20020194604" | "20030051242" | "20030066078" | "20030135853" | "20030145323" | "20030149975" | "20030208756" | "20030208758" | "20030221191" | "20040003406" | "20040010546" | "20040073915" | "20040078809" | "20040107437" | "20040117827" | "20040163101" | "20040172650" | "20050165638" | "20050165644" | "20050246736" | "20050283796" | "20060116924" | "20060212900" | "20060248570" | "20060253323" | "20060271953" | "20060277569" | "20070022032" | "20070089127" | "20070100690" | "20070180469" | "20070204297" | "20070233571" | "20070244750" | "20070266400" | "20070266403" | "20070277220" |	(US-PGPUB; USPAT; USOCR)	OR	OFF	OFF	2020/01/08 01:21 PM

		"20070283384" | "20070294721" | "20080016540" | "20080040742" | "20080243822" | "20080259906" | "20080313194" | "20090049468" | "20090094119" | "20090172723" | "20100192172" | "5848396" | "5855008" | "5933811" | "5948061" | "6160989" | "6216129" | "6366298" | "6385592" | "6446261" | "6536041" | "6684194" | "6718551" | "6771290" | "6845396" | "6925440" | "6973436" | "7072853" | "7092926" | "7100183" | "7152237" | "7260823" | "7328448" | "7428493" | "7454364").PN. OR ("8051444").URPN.					
L19	1	"10178442".pn.	(US-PGPUB; USPAT; USOCR)	OR	OFF	OFF	2020/01/08 05:54 PM
L20	4	"10178442".pn.	(US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/01/08 05:55 PM
L21	4039	725/34.ccls.	(US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/01/08 07:09 PM
L22	57	(US-5105184-$ | US-5220655-$ | US-5359508-$ | US-5442633-$ | US-5572643-$ | US-5710884-$ | US-5724424-$ | US-5724521-$ | US-5751956-$ | US-5774670-$ | US-5794210-$ | US-5809242-$ | US-5819092-$ | US-5819285-$ | US-5835718-$ | US-5838314-$ | US-5848396-$ | US-5855008-$ | US-5870546-$ | US-5870550-$ | US-5901287-$ | US-5918014-$ | US-5933811-$ | US-5937390-$ | US-	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/09 09:32 AM

		5937392-$ | US-5948061-$ | US-5959623-$ | US-5963915-$ | US-5978013-$).DID.					
L23	58	(US-5995943-$ | US-5999912-$ | US-6002394-$ | US-6014638-$ | US-6016504-$ | US-6018768-$ | US-6026368-$ | US-6026369-$ | US-6026374-$ | US-6047310-$ | US-6049835-$ | US-6055573-$ | US-6073241-$ | US-6128663-$ | US-6133912-$ | US-6134532-$ | US-6157946-$ | US-6160989-$ | US-6177931-$ | US-6188398-$ | US-6189008-$ | US-6205432-$ | US-6216112-$ | US-6216129-$ | US-6285987-$ | US-6289318-$ | US-6308202-$ | US-6327574-$ | US-6366298-$).DID.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/09 09:33 AM
L24	59	(US-6377986-$ | US-6385592-$ | US-6401075-$ | US-6408278-$ | US-6415270-$ | US-6446261-$ | US-6449657-$ | US-6463533-$ | US-6463585-$ | US-6490587-$ | US-6536041-$ | US-6633850-$ | US-6654725-$ | US-6684194-$ | US-6718551-$ | US-6771290-$ | US-6810527-$ | US-6837436-$ | US-6845396-$ | US-6847992-$ | US-6868392-$ | US-6907566-$ | US-6985882-$ | US-	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/09 09:35 AM

		7013290-$ | US-7038637-$ | US-7051351-$ | US-7092926-$ | US-7100183-$ | US-7146329-$).DID.					
L25	57	(US-7152237-$ | US-7158943-$ | US-7188085-$ | US-7260823-$ | US-7272853-$ | US-7299195-$ | US-7302419-$ | US-7305459-$ | US-7580855-$ | US-7640214-$ | US-7756744-$ | US-7805332-$ | US-7806329-$ | US-7856372-$ | US-7856373-$ | US-7913275-$ | US-7930252-$ | US-7937383-$ | US-7966647-$ | US-7971260-$ | US-8015184-$ | US-8015286-$ | US-8024765-$ | US-8079048-$ | US-8086491-$ | US-8180674-$ | US-8700464-$ | US-8725571-$ | US-8762201-$).DID.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/09 09:36 AM
L26	57	(US-7152237-$ | US-7158943-$ | US-7188085-$ | US-7260823-$ | US-7272853-$ | US-7299195-$ | US-7302419-$ | US-7305459-$ | US-7580855-$ | US-7640214-$ | US-7756744-$ | US-7805332-$ | US-7806329-$ | US-7856372-$ | US-7856373-$ | US-7913275-$ | US-7930252-$ | US-7937383-$ | US-7966647-$ | US-7971260-$ | US-8015184-$ | US-8015286-$ | US-	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/09 09:37 AM

		8024765-$ | US-8079048-$ | US-8086491-$ | US-8180674-$ | US-8700464-$ | US-8725571-$ | US-8762201-$).DID.					
L27	58	(US-20010034646-$ | US-20010042002-$ | US-20010049620-$ | US-20020010653-$ | US-20020010928-$ | US-20020013943-$ | US-20020019769-$ | US-20020032603-$ | US-20020046105-$ | US-20020059094-$ | US-20020059602-$ | US-20020078444-$ | US-20020082910-$ | US-20020082923-$ | US-20020083043-$ | US-20020103850-$ | US-20020104086-$ | US-20020112013-$ | US-20020120933-$ | US-20020120935-$ | US-20020123928-$ | US-20020124249-$ | US-20020124253-$ | US-20020129362-$ | US-20020129368-$ | US-20020194058-$ | US-20020194604-$ | US-20030023481-$ | US-8880677-$).DID.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/09 09:38 AM
L28	57	(US-20030041156-$ | US-20030051242-$ | US-20030066078-$ | US-20030093311-$ | US-20030126597-$ | US-20030131355-$ | US-20030145323-$ | US-20030149975-$ | US-20030187726-$ | US-20030187949-$ | US-20030208756-$ | US-20030208758-$ | US-20030226141-$ | US-20030229900-$ | US-20030231207-$ | US-20040002938-$ | US-20040003406-$ | US-20040010546-$ | US-20040031058-$ | US-20040073915-$ | US-20040078809-$ |	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/09 09:39 AM

		US-20040107437-$ | US-20040117827-$ | US-20040123314-$ | US-20040163101-$ | US-20040172650-$ | US-20050049957-$ | US-20050076393-$ | US-20050125289-$).DID.					
L29	57	(US-20050125290-$ | US-20050132405-$ | US-20050165638-$ | US-20050165643-$ | US-20050165644-$ | US-20050177420-$ | US-20050187823-$ | US-20060018466-$ | US-20060116924-$ | US-20060137015-$ | US-20060184989-$ | US-20060253323-$ | US-20060259357-$ | US-20060271690-$ | US-20060271953-$ | US-20060271959-$ | US-20060287916-$ | US-20070027901-$ | US-20070038516-$ | US-20070073585-$ | US-20070078774-$ | US-20070088603-$ | US-20070089127-$ | US-20070100690-$ | US-20070130015-$ | US-20070157228-$ | US-20070198327-$ | US-20070208619-$ | US-20070214048-$).DID.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/09 09:41 AM
L30	58	(US-20070220553-$ | US-20070233571-$ | US-20070244750-$ | US-20070266403-$ | US-20070277220-$ | US-20070283384-$ | US-20070294721-$ | US-20080016540-$ | US-20080040742-$ | US-20080052392-$ | US-20080082413-$ | US-20080109306-$ | US-20080109376-$ | US-20080109843-$ | US-20080155613-$ | US-20080181225-$ | US-20080183573-$ | US-20080189360-$ |	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/09 09:42 AM

		US-20080201772-$ | US-20080222283-$ | US-20080243822-$ | US-20080244076-$ | US-20080255944-$ | US-20080262901-$ | US-20080275785-$ | US-20090037949-$ | US-20090049469-$ | US-20090106100-$ | US-20090171780-$).DID.					
L31	47	(US-20090204706-$ | US-20100076848-$ | US-20110125592-$ | US-20110167043-$ | US-20110258049-$ | US-20110307548-$ | US-20130254035-$ | US-6832207-$ | US-6925440-$ | US-6973436-$ | US-7072853-$ | US-7428493-$ | US-7454364-$ | US-7747745-$ | US-7822637-$ | US-7822639-$ | US-7890609-$ | US-7979307-$ | US-8051444-$ | US-8200822-$ | US-8204783-$ | US-8204965-$ | US-8239264-$ | US-8244574-$ | US-8244582-$ | US-8244583-$ | US-8244586-$ | US-8280758-$ | US-8494904-$).DID.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/09 09:42 AM
L32	51	(US-10009636-$ | US-10026100-$ | US-10134054-$ | US-20050086112-$ | US-20110035256-$ | US-8566164-$ | US-8589210-$ | US-8595069-$ | US-8600815-$ | US-8607267-$ | US-8683502-$ | US-8713600-$ | US-8775249-$ | US-8959146-$ | US-8997138-$ | US-9071886-$ | US-	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/09 09:44 AM

		9078035-$ | US-9083853-$ | US-9131282-$ | US-9208514-$ | US-9226019-$ | US-9271024-$ | US-9351053-$ | US-9508089-$ | US-9571865-$ | US-9591380-$ | US-9756372-$ | US-9800917-$ | US-9830615-$).DID.					
L33	4	(US-20090113456-$ | US-0113456-$ | US-0793513-$).DID.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/09 09:46 AM
L34	2	"20190124412"	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/01/09 11:25 AM
L35	1	"16228479"	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/04/26 03:11 PM
L36	2	"7861260".pn.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/04/26 03:14 PM
L39	5	"9510025".pn.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/04/27 07:33 AM
L40	4	"10116971".pn.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2020/04/27 07:33 AM
L41	3	"5105184".pn.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/06/08 01:17 PM
L42	3	"5220655".pn.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/06/08 01:17 PM
L43	3	"5359508".pn.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/06/08 01:17 PM
L44	3	"5442633".pn.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/06/08 01:17 PM
L45	3	"5572643".pn.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/06/08 01:17 PM
L46	3	"5710884".pn.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/06/08 01:17 PM
L47	57	(US-5105184-$ | US-5220655-$ | US-5359508-$ | US-5442633-$ | US-5572643-$ | US-	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/06/08 01:29 PM

		5710884-$ | US- 5724424-$ | US- 5724521-$ | US- 5751956-$ | US- 5774670-$ | US- 5794210-$ | US- 5809242-$ | US- 5819092-$ | US- 5819285-$ | US- 5835718-$ | US- 5838314-$ | US- 5848396-$ | US- 5855008-$ | US- 5870546-$ | US- 5870550-$ | US- 5901287-$ | US- 5918014-$ | US- 5933811-$ | US- 5937390-$ | US- 5937392-$ | US- 5948061-$ | US- 5959623-$ | US- 5963915-$ | US- 5978013-$).DID.					
L48	58	(US-5995943-$ | US- 5999912-$ | US- 6002394-$ | US- 6014638-$ | US- 6016504-$ | US- 6018768-$ | US- 6026368-$ | US- 6026369-$ | US- 6026374-$ | US- 6047310-$ | US- 6049835-$ | US- 6055573-$ | US- 6073241-$ | US- 6128663-$ | US- 6133912-$ | US- 6134532-$ | US- 6157946-$ | US- 6160989-$ | US- 6177931-$ | US- 6188398-$ | US- 6189008-$ | US- 6205432-$ | US- 6216112-$ | US- 6216129-$ | US- 6285987-$ | US- 6289318-$ | US- 6308202-$ | US- 6327574-$ | US- 6366298-$).DID.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/06/08 01:31 PM
L49	59	(US-6377986-$ | US- 6385592-$ | US- 6401075-$ | US- 6408278-$ | US-	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/06/08 01:32 PM

		6415270-$ | US-6446261-$ | US-6449657-$ | US-6463533-$ | US-6463585-$ | US-6490587-$ | US-6536041-$ | US-6633850-$ | US-6654725-$ | US-6684194-$ | US-6718551-$ | US-6771290-$ | US-6810527-$ | US-6837436-$ | US-6845396-$ | US-6847992-$ | US-6868392-$ | US-6907566-$ | US-6985882-$ | US-7013290-$ | US-7038637-$ | US-7051351-$ | US-7092926-$ | US-7100183-$ | US-7146329-$).DID.					
L50	57	(US-7152237-$ | US-7158943-$ | US-7188085-$ | US-7260823-$ | US-7272853-$ | US-7299195-$ | US-7302419-$ | US-7305459-$ | US-7580855-$ | US-7640214-$ | US-7756744-$ | US-7805332-$ | US-7806329-$ | US-7856372-$ | US-7856373-$ | US-7913275-$ | US-7930252-$ | US-7937383-$ | US-7966647-$ | US-7971260-$ | US-8015184-$ | US-8015286-$ | US-8024765-$ | US-8079048-$ | US-8086491-$ | US-8180674-$ | US-8700464-$ | US-8725571-$ | US-8762201-$).DID.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/06/08 01:33 PM
L51	58	(US-20010034646-$ | US-20010042002-$ | US-20010049620-$ |	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/06/08 01:35 PM

		US-20020010653-$ | US-20020010928-$ | US-20020013943-$ | US-20020019769-$ | US-20020032603-$ | US-20020046105-$ | US-20020059094-$ | US-20020059602-$ | US-20020078444-$ | US-20020082910-$ | US-20020082923-$ | US-20020083043-$ | US-20020103850-$ | US-20020104086-$ | US-20020112013-$ | US-20020120933-$ | US-20020120935-$ | US-20020123928-$ | US-20020124249-$ | US-20020124253-$ | US-20020129362-$ | US-20020129368-$ | US-20020194058-$ | US-20020194604-$ | US-20030023481-$ | US-8880677-$).DID.					
L52	57	(US-20030041156-$ | US-20030051242-$ | US-20030066078-$ | US-20030093311-$ | US-20030126597-$ | US-20030131355-$ | US-20030145323-$ | US-20030149975-$ | US-20030187726-$ | US-20030187949-$ | US-20030208756-$ | US-20030208758-$ | US-20030226141-$ | US-20030229900-$ | US-20030231207-$ | US-20040002938-$ | US-20040003406-$ | US-20040010546-$ | US-20040031058-$ | US-20040073915-$ | US-20040078809-$ | US-20040107437-$ | US-20040117827-$ | US-20040123314-$ | US-20040163101-$ | US-20040172650-$ | US-20050049957-$ | US-20050076393-$ | US-20050125289-$).DID.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/06/08 01:36 PM
L53	57	(US-20050125290-$ |	(US-PGPUB; USPAT;	OR	OFF	OFF	2020/06/08

		US-20050132405-$ | US-20050165638-$ | US-20050165643-$ | US-20050165644-$ | US-20050177420-$ | US-20050187823-$ | US-20060018466-$ | US-20060116924-$ | US-20060137015-$ | US-20060184989-$ | US-20060253323-$ | US-20060259357-$ | US-20060271690-$ | US-20060271953-$ | US-20060271959-$ | US-20060287916-$ | US-20070027901-$ | US-20070038516-$ | US-20070073585-$ | US-20070078774-$ | US-20070088603-$ | US-20070089127-$ | US-20070100690-$ | US-20070130015-$ | US-20070157228-$ | US-20070198327-$ | US-20070208619-$ | US-20070214048-$).DID.	EPO; JPO; DERWENT)				01:37 PM
L54	58	(US-20070220553-$ | US-20070233571-$ | US-20070244750-$ | US-20070266403-$ | US-20070277220-$ | US-20070283384-$ | US-20070294721-$ | US-20080016540-$ | US-20080040742-$ | US-20080052392-$ | US-20080082413-$ | US-20080109306-$ | US-20080109376-$ | US-20080109843-$ | US-20080155613-$ | US-20080181225-$ | US-20080183573-$ | US-20080189360-$ | US-20080201772-$ | US-20080222283-$ | US-20080243822-$ | US-20080244076-$ | US-20080255944-$ | US-20080262901-$ | US-20080275785-$ | US-20090037949-$ | US-20090049469-$ | US-20090106100-$ | US-20090171780-	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/06/08 01:37 PM

		$).DID.					
L55	47	(US-20090204706-$ | US-20100076848-$ | US-20110125592-$ | US-20110167043-$ | US-20110258049-$ | US-20110307548-$ | US-20130254035-$ | US-6832207-$ | US-6925440-$ | US-6973436-$ | US-7072853-$ | US-7428493-$ | US-7454364-$ | US-7747745-$ | US-7822637-$ | US-7822639-$ | US-7890609-$ | US-7979307-$ | US-8051444-$ | US-8200822-$ | US-8204783-$ | US-8204965-$ | US-8239264-$ | US-8244574-$ | US-8244582-$ | US-8244583-$ | US-8244586-$ | US-8280758-$ | US-8494904-$).DID.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/06/08 01:38 PM
L56	51	(US-10009636-$ | US-10026100-$ | US-10134054-$ | US-20050086112-$ | US-20110035256-$ | US-8566164-$ | US-8589210-$ | US-8595069-$ | US-8600815-$ | US-8607267-$ | US-8683502-$ | US-8713600-$ | US-8775249-$ | US-8959146-$ | US-8997138-$ | US-9071886-$ | US-9078035-$ | US-9083853-$ | US-9131282-$ | US-9208514-$ | US-9226019-$ | US-9271024-$ | US-9351053-$ | US-9508089-$ | US-9571865-$ | US-9591380-$ | US-9756372-$ | US-	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/06/08 01:39 PM

		9800917-$ | US-9830615-$).DID.					
L57	8	(US-20090113456-$ | US-0113456-$ | US-0793513-$ | US-20090172723-$ | US-20110040607-$).DID.	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	OFF	OFF	2020/06/08 01:40 PM
L58	1	"10715878".pn.	(US-PGPUB; USPAT)	OR	ON	ON	2022/09/23 05:04 PM
L59	16248	(common "same" identical) WITH ("IP" internet ADJ protocol) near3 (address$3) WITH (device$1)	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2022/09/26 09:58 AM
L60	1257	(common "same" identical) WITH ("IP" internet ADJ protocol) near3 (address$3) WITH (multiple plural several many) near3(device$1)	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2022/09/26 10:09 AM
L61	1257	(common "same" identical) WITH ("IP" internet ADJ protocol) near3 (address$3) WITH (multiple plural several many multitude) near3(device$1)	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2022/09/26 10:12 AM
L62	3033	(time) ADJ (stamp period duration) WITH (access$3 use using) WITH (internet)	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2022/09/26 10:16 AM
L63	6	L61 AND (time) ADJ (stamp period duration) WITH (access$3 use using) WITH (internet)	(US-PGPUB; USPAT; EPO; JPO; DERWENT)	OR	ON	ON	2022/09/26 10:16 AM
L64	3219	(device terminal receiver "STB" settopbox "STT" settopterminal) WITH (associat$3 relat$3 link$3) SAME ("PII" personally ADJ identifiable ADJ information)	(US-PGPUB; USPAT)	OR	ON	ON	2022/09/26 10:28 AM
L65	2	L62 AND L64	(US-PGPUB; USPAT)	OR	ON	ON	2022/09/26 10:29 AM
L66	150	(device terminal receiver "STB" settopbox "STT" settopterminal) WITH (associat$3 relat$3 link$3) SAME (non) near2("PII" personally ADJ identifiable ADJ	(US-PGPUB; USPAT)	OR	ON	ON	2022/09/26 10:29 AM

		information)					
L68	144836	(H04N21/812 OR G06F16/951 OR G06F16/248 OR H04N21/44224 OR H04L61/5007 OR G06Q/0241 OR G06Q/0251 OR G06Q30/0255 OR G06Q30/0255 OR G06Q30/0256 OR G06Q30/0269 OR H04L12/66 OR H04L67/306 OR H04N21/25891 OR H04N21/42684 OR H04N21/44222, H04N21/4532 OR H04N21/6125 OR H04N21/64322).CPC.	(US-PGPUB; USPAT)	OR	ON	ON	2022/09/26 10:40 AM
L69	78	(H04N21/812 OR G06F16/951 OR G06F16/248 OR H04N21/44224 OR H04L61/5007 OR G06Q/0241 OR G06Q/0251 OR G06Q30/0255 OR G06Q30/0255 OR G06Q30/0256 OR G06Q30/0269 OR H04L12/66 OR H04L67/306 OR H04N21/25891 OR H04N21/42684 OR H04N21/44222, H04N21/4532 OR H04N21/6125 OR H04N21/64322).CPC. AND L66	(US-PGPUB; USPAT)	OR	ON	ON	2022/09/26 10:40 AM

PE2E SEARCH - Search History (Interference)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	British Equivalents	Time Stamp
N37	7616	(H04N21/812, /64322, /42684, /4532, /6125, /25891, /44222$.CPC.)	(USPAT)	OR	ON	ON	2020/04/26 03:27 PM
N38	2854	(G06Q30/0241, /0269, /0256, /0255, /0251$.CPC.)	(USPAT)	OR	ON	ON	2020/04/26 03:31 PM
N39	6	L61 AND (time) ADJ (stamp period duration) WITH (access$3 use using) WITH (internet)	(US-PGPUB; USPAT)	OR	ON	ON	2022/09/26 10:17 AM
N40	150	(device terminal	(US-PGPUB; USPAT)	OR	ON	ON	2022/09/26

		receiver "STB" settopbox "STT" settopterminal) WITH (associat$3 relat$3 link$3) SAME (non) near2("PII" personally ADJ identifiable ADJ information)					10:30 AM
N41	144836	(H04N21/812 OR G06F16/951 OR G06F16/248 OR H04N21/44224 OR H04L61/5007 OR G06Q/0241 OR G06Q/0251 OR G06Q30/0255 OR G06Q30/0255 OR G06Q30/0256 OR G06Q30/0269 OR H04L12/66 OR H04L67/306 OR H04N21/25891 OR H04N21/42684 OR H04N21/44222, H04N21/4532 OR H04N21/6125 OR H04N21/64322).CPC.	(US-PGPUB; USPAT)	OR	ON	ON	2022/09/26 10:39 AM
N42	78	(H04N21/812 OR G06F16/951 OR G06F16/248 OR H04N21/44224 OR H04L61/5007 OR G06Q/0241 OR G06Q/0251 OR G06Q30/0255 OR G06Q30/0255 OR G06Q30/0256 OR G06Q30/0269 OR H04L12/66 OR H04L67/306 OR H04N21/25891 OR H04N21/42684 OR H04N21/44222, H04N21/4532 OR H04N21/6125 OR H04N21/64322).CPC. AND L66	(US-PGPUB; USPAT)	OR	ON	ON	2022/09/26 10:40 AM
N43	510	L64 AND L68	(US-PGPUB; USPAT)	OR	ON	ON	2022/09/26 10:41 AM

Doc code: IDS
Doc description: Information Disclosure Statement (IDS) Filed

PTO/SB/08a (01-22)
Approved for use through 05/31/2024. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (**Not for submission under 37 CFR 1.99)**		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Ekpo, Nnenna
	Attorney Docket Number	TV1-CON7

U.S.PATENTS

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	9110903		08/18/2005	Martinez et al.	

If you wish to add additional U.S. Patent citation information please click the Add button.

U.S.PATENT APPLICATION PUBLICATIONS

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20050278731		12/15/2005	Cameron et al.	
	2	20080120178		05/22/2008	Martinez et al.	
	3	20050160002		07/21/2005	Roetter et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Ekpo, Nnenna
	Attorney Docket Number	TV1-CON7

If you wish to add additional Foreign Patent Document citation information please click the Add button

NON-PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature	/NNENNA N EKPO/	Date Considered	11/28/2022

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Ekpo, Nnenna
	Attorney Docket Number	TV1-CON7

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Louis J. Hoffman/	Date (YYYY-MM-DD)	2022-09-16
Name/Print	Louis J. Hoffman	Registration Number	38918

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

BIB DATA SHEET

CONFIRMATION NO. 7004

SERIAL NUMBER	FILING or 371(c) DATE RULE	CLASS	GROUP ART UNIT	ATTORNEY DOCKET NO.
17/714,965	04/06/2022	725	2425	14-TV1-CON7

APPLICANTS
Intent IQ, LLC, Long Island City, NY;

INVENTORS
Roy Shkedi, Forest Hills, NY;

**** CONTINUING DATA ***************************
This application is a CON of 16/925,655 07/10/2020 PAT 11303973
which is a CON of 16/228,479 12/20/2018 PAT 10715878
which is a CON of 15/803,631 11/03/2017 PAT 10178442
which is a CON of 15/179,356 06/10/2016 PAT 9813778
which is a CON of 14/246,525 04/07/2014 PAT 9369779
which is a DIV of 13/097,885 04/29/2011 PAT 8695032
which is a DIV of 12/860,666 08/20/2010 PAT 8281336
which is a CON of 11/736,544 04/17/2007 PAT 7861260

**** FOREIGN APPLICATIONS ***************************

**** IF REQUIRED, FOREIGN FILING LICENSE GRANTED ****
04/11/2022

Foreign Priority claimed ❑ Yes ☑ No 35 USC 119(a-d) conditions met ❑ Yes ☑ No Verified and Acknowledged /NNENNA NGOZI EKPO/ Examiner's Signature	❑ Met after Allowance Initials	STATE OR COUNTRY	SHEETS DRAWINGS	TOTAL CLAIMS	INDEPENDENT CLAIMS
		NY	18	26	2

ADDRESS

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, AZ 85260
UNITED STATES

TITLE

Targeted television advertisements based on online behavior

FILING FEE RECEIVED		
2580	FEES: Authority has been given in Paper No._____ to charge/credit DEPOSIT ACCOUNT No._____ for following:	❑ All Fees ❑ 1.16 Fees (Filing) ❑ 1.17 Fees (Processing Ext. of time) ❑ 1.18 Fees (Issue) ❑ Other _____ ❑ Credit

BIB (Rev. 05/07).

OK TO ENTER: /N.N.E/

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor	:	Shkedi, Roy	Art Unit	:	2425 (in parent)
Serial No.	:	17/714,965	Examiner	:	Ekpo, Nnenna (in parent)
Filing Date	:	04/06/2022	Conf. No.	:	7004
Title	:	Targeted television advertisements based on online behavior			

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Filed via EFS-Web on April 8, 2022

INFORMATION DISCLOSURE STATEMENT

Dear Sir:

Assignee discloses the references listed on form PTO-1449 attached to this document. All foreign and non-patent literature references listed were cited in parent application Serial No. 16/925,655, therefore, pursuant to Office rules, copies of those references are not included here.

Assignee also wishes to direct Examiner to the applications listed in the table below (which share the same first inventor) as well as the references cited therein:

Serial No.	Filing Date	Earliest Parent Filing Date
17/401,078	8/12/2021	12/31/2007
17/097,867	11/13/2020	06/16/2006
17/092,651	11/09/2020	06/13/2013
17/390,211	7/30/2021	08/03/2011 (prov.)
17/401,052	08/12/2021	06/19/2006 (prov.)
17/699,297	3/1/2022	08/28/2020 (prov.)

If the Office has any questions, please feel free to contact assignee's undersigned attorney of record.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

OK TO ENTER: /N.N.E/

Respectfully submitted,

INTENT IQ, LLC.

by its attorney

Dated: April 8, 2022

/Louis J. Hoffman/
Louis J. Hoffman
Reg. No. 38,918

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, Arizona 85260
(480) 948-3295

/NNENNA N EKPO/ 11/28/2022

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Doc code: IDS
Doc description: Information Disclosure Statement (IDS) Filed

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (**Not for submission under 37 CFR 1.99**)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Ekpo, Nnenna
	Attorney Docket Number	TV1-CON7

U.S.PATENTS

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button.

U.S.PATENT APPLICATION PUBLICATIONS

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20070061838		2007-03-15	Grubbs et al	
	2	20050108091		2005-05-19	Sotak et al	

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1	WO 2007/117613			2007-10-18	Ferguson		☐

If you wish to add additional Foreign Patent Document citation information please click the Add button

NON-PATENT LITERATURE DOCUMENTS

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Ekpo, Nnenna
	Attorney Docket Number	TV1-CON7

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature	/NNENNA N EKPO/	Date Considered	12/02/2022

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Ekpo, Nnenna
	Attorney Docket Number	TV1-CON7

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Louis J. Hoffman/	Date (YYYY-MM-DD)	2022-11-30
Name/Print	Louis J. Hoffman	Registration Number	38918

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor	:	Shkedi, Roy	Art Unit	:	3693 (in parent)
Serial No.	:	17/714,965	Examiner	:	Ekpo, Nnenna (in parent)
Filing Date	:	04/06/2022	Conf. No.	:	7004
Title	:	Targeted television advertisements based on online behavior			

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Filed via EFS-Web on September 16, 2022

INFORMATION DISCLOSURE STATEMENT

Dear Sir:

Assignee discloses the references listed on attached form PTO/SB/08.

Also, assignee wishes to direct the examiner to the co-pending proceedings in the table below, as well as the references cited in those proceedings. All listed applications and the listed patent share the same first inventor:

Serial No.	Filing Date	Patent Challenged
IPR2022-01507	9/7/2022	Re: USP 10,321,198 (same family as SN 17/401,078, previously cited)
IPR2022-01506	9/7/2022	Re: USP 8,566,164 (same family as SN 17/401,078, previously cited)
IPR2022-01420	8/12/2022	Re: parent USP 8,677,398
IPR2022-01505	9/7/2022	Re: parent USP 10,715,878

If the Office has any questions, please feel free to contact assignee's undersigned attorney of record.

Respectfully submitted,

INTENT IQ, LLC.

by its attorney

Dated: September 16, 2022

/Louis J. Hoffman/
Louis J. Hoffman
Reg. No. 38,918

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, Arizona 85260
(480) 948-3295

/NNENNA N EKPO/

11/28/2022

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

To: Louis@valuablepatents.com,donald@valuablepatents.com,dsalavi@northwestpatent.com
From: PAIR_eOfficeAction@uspto.gov
Cc: PAIR_eOfficeAction@uspto.gov
Subject: Private PAIR Correspondence Notification for Customer Number 26362

Dec 15, 2022 04:37:42 AM

Dear PAIR Customer:

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, AZ 85260
UNITED STATES

The following USPTO patent application(s) associated with your Customer Number, 26362 , have new outgoing correspondence. This correspondence is now available for viewing in Private PAIR.

The official date of notification of the outgoing correspondence will be indicated on the form PTOL-90 accompanying the correspondence.

Disclaimer:
The list of documents shown below is provided as a courtesy and is not part of the official file wrapper. The content of the images shown in PAIR is the official record.

Application	Document	Mailroom Date	Attorney Docket No.
17714965	NOA	12/15/2022	14-TV1-CON7
	892	12/15/2022	14-TV1-CON7
	1449	12/15/2022	14-TV1-CON7
	1449	12/15/2022	14-TV1-CON7
	1449	12/15/2022	14-TV1-CON7
	1449	12/15/2022	14-TV1-CON7
	1449	12/15/2022	14-TV1-CON7
	1449	12/15/2022	14-TV1-CON7
	1449	12/15/2022	14-TV1-CON7
	1449	12/15/2022	14-TV1-CON7

To view your correspondence online or update your email addresses, please visit us anytime at https://sportal.uspto.gov/secure/myportal/privatepair.

If you have any questions, please email the Electronic Business Center (EBC) at EBC@uspto.gov with 'e-Office Action' on the subject line or call 1-866-217-9197 during the following hours:

Monday - Friday 6:00 a.m. to 12:00 a.m.

Thank you for prompt attention to this notice,

UNITED STATES PATENT AND TRADEMARK OFFICE

PATENT APPLICATION INFORMATION RETRIEVAL SYSTEM

Issue Fee Transmittal Form

Application Number	Filing Date	First Named Inventor	Atty. Docket No.	Confirmation No.
17714965	06-Apr-2022	Roy Shkedi	14-TV1-CON7	7004

TITLE OF INVENTION :

Targeted television advertisements based on online behavior

Entity Status	Application Type	Art Unit	Class - Subclass	EXAMINER
Regular Undiscounted	Utility under 35 USC 111(a)	2425	034000	NNENNA EKPO

Issue Fee Due	Publication Due	Total Fee(s) Due	Date Due	Prev. Paid Fee
$1200	$0	$1200	15-Mar-2023	$0

1.Change of Correspondence Address and/or Indication Of Fee Address (37 CFR 1.33 & 1.363)

Current Correspondence Address:	Current Indicated Fee Address :
26362 LOUIS J. HOFFMAN, P.C. 7689 East Paradise Lane Suite 2 Scottsdale AZ 85260 UNITED STATES 480-948-3295 Louis@valuablepatents.com	
☐ Change of correspondence address requested, system generated AIA/122-EFS form attached	☐ Fee Address indication requested, system generated SB/47-EFS form attached

2.Entity Status

Change in Entity Status

○ Applicant certifying micro entity status; system generated Micro Entity certification form attached. See 37 CFR 1.29.
Note: Absent a valid certification of micro entity status, issue fee payment in the micro entity amount will not be accepted at the risk of application abandonment. If this box is checked, you will be prompted to choose a micro entity status on the gross income basis (37 CFR 1.29(a)) or the institution of higher education basis (37 CFR 1.29(d)), and make the applicable certification online.

○ Applicant asserting small entity status. See 37 CFR 1.27.
Note: If the application was previously under micro entity status, checking this box will be taken to be a notification of loss of entitlement to micro entity status.

◉ Applicant changing to regular undiscounted fee status.
Note: Checking this box will be taken to be a notification of loss of entitlement to small or micro entity status, as applicable.

3.The Following Fee(s) Are Submitted:

☒ Issue Fee	☐ I authorize USPTO to apply my previously paid issue fee to the current fees due
☐ Publication Fee	☐ The Director is hereby authorized to apply my previously paid issue fee to the current fee due and to charge deficient fees to Deposit Account Number _____
☐ Advance Order - # of copies _____	☐ If **in addition to** the payment of the issue fee amount submitted with this form, there are any discrepancies in any amount(s) due, the Director is authorized to charge any deficiency, or credit any overpayment, to Deposit Account Number _____ . The **issue fee must be submitted** with this form. **If payment of the issue fee does not accompany this form, checking this box and providing a deposit account number will NOT be effective to satisfy full payment of the fee(s) due.**

4.Firm and/or Attorney Names To Be Printed

NOTE: If no name is listed, no name will be printed
For printing on the patent front page, list to be displayed as entered

1. Louis J. Hoffman
2. David S. Alavi
3.

5.Assignee Name(s) and Residence Data To Be Printed

PLEASE NOTE: Unless an assignee is identified below, no assignee data will appear on the patent. If an assignee is identified below, the document has been filed for recordation as set forth in 37 CFR 3.11. Completion of this form is NOT a substitute for filing an assignment.

Name	City	State	Country	Category
Intent IQ, LLC	Long Island City	NEW YORK	UNITED STATES	other

6.Signature

I certify, in accordance with 37 CFR 1.4(d)(4) that I am an attorney or agent registered to practice before the Patent and Trademark Office who has filed and has been granted power of attorney in this application. I also certify that this Fee(s) Transmittal form is being transmitted to the USPTO via EFS-WEB on the date indicated below.

Signature	/Louis J. Hoffman/	**Date**	12-16-2022
Name	Louis J. Hoffman	**Registration Number**	38918

Electronic Patent Application Fee Transmittal

Application Number:	17714965
Filing Date:	06-Apr-2022
Title of Invention:	Targeted television advertisements based on online behavior
First Named Inventor/Applicant Name:	Roy Shkedi
Filer:	Louis J. Hoffman
Attorney Docket Number:	14-TV1-CON7

Filed as Large Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
UTILITY APPL ISSUE FEE	1501	1	1200	1200
PUBL. FEE- EARLY, VOLUNTARY, OR NORMAL	1504	1	0	0
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				
Miscellaneous:				
			Total in USD ($)	**1200**

Electronic Acknowledgement Receipt	
EFS ID:	47227868
Application Number:	17714965
International Application Number:	
Confirmation Number:	7004
Title of Invention:	Targeted television advertisements based on online behavior
First Named Inventor/Applicant Name:	Roy Shkedi
Customer Number:	26362
Filer:	Louis J. Hoffman
Filer Authorized By:	
Attorney Docket Number:	14-TV1-CON7
Receipt Date:	16-DEC-2022
Filing Date:	06-APR-2022
Time Stamp:	19:37:01
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	CARD
Payment was successfully received in RAM	$1200
RAM confirmation Number	E2022BFJ36582838
Deposit Account	
Authorized User	
The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:	

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Issue Fee Payment (PTO-85B)	Web85b.pdf	62488 5540aaf1501c71a920783f1b4d001f543429 55a3	no	2
Warnings:					
Information:					
2	Fee Worksheet (SB06)	fee-info.pdf	40163 84e3bfc11571f1fb04fc9ee94bcd37ff7877a 381	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	102651		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.
National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.
New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	17714965
	Filing Date	2022-04-06
	First Named Inventor	Shkedi, Roy
	Art Unit	2425
	Examiner Name	Kekpo, Nnenna
	Attorney Docket Number	TV1-CON7

	20	11301898		2022-04-12	Shkedi	
	21	11303973		2022-04-12	Shkedi et al.	
	22	11368763		2022-06-21	Shkedi et al.	

If you wish to add additional U.S. Patent citation information please click the Add button.

U.S.PATENT APPLICATION PUBLICATIONS

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
Change(s) applied to document. /P.P./ 12/19/2022	1	20070157247		~~07/05/2005~~ 07/2007	Cordray et al.	
	2	20080113674		2008-05-15	Baig	
	3	20070067459		2007-03-22	Baal-Haness	
	4	20030229893		2003-12-11	Sgaraglino	
	5	20080281697		2008-11-13	Whitehead	
	6	20060069616		2006-03-30	Bau	

Form PTO-1449 (Modified)
List of Patents and Publications For
Information Disclosure Statement

Serial No.: 17/714,965
Applicant: Shkedi, Roy
Filed: 04/06/2022
Art Unit: 2425
Examiner: Ekpo, Nnenna
Conf. No.: 7004

_____	94	7302419	11/27/2007	Conkwright et al.
Change(s) applied _____	95	7139723	~~01/13/2000~~ 11/2006	Conkwright et al.
to document, _____	96	7305459	12/04/2007	Klemba et al.
/P.P./ _____	97	7580855	08/25/2009	Fernandes
12/19/2022 _____	98	7640214	12/29/2009	Black et al.
_____	99	7756744	07/13/2010	Schiff et al.
_____	100	7805332	09/28/2010	Wilson
_____	101	7806329	10/05/2010	Dmitriev et al.
_____	102	7856372	12/21/2010	Ullah
_____	103	7856373	12/21/2010	Ullah
_____	104	7913275	03/22/2011	Hendricks et al.
_____	105	7930252	04/19/2011	Bender et al.
_____	106	7937383	05/03/2011	Hintze et al.
_____	107	7966647	06/21/2011	Igoe et al.
_____	108	7971260	06/28/2011	Spalink et al.
_____	109	8015184	09/06/2011	Zito et al.
_____	110	8015286	09/06/2011	Jenkins
_____	111	8019881	09/13/2011	Sandhu et al.
_____	112	8024765	09/20/2011	Ramanathan et al.
_____	113	8079048	12/13/2011	Howcroft
_____	114	8086491	12/27/2011	Matz et al.
_____	115	8180674	05/15/2012	Madhavan
_____	116	8700464	04/15/2014	McAllister et al.
_____	117	8725571	05/13/2014	McAllister et al.
_____	118	8762201	06/24/2014	Noonan
_____	119	8880677	11/04/2014	Julia et al.
_____	120	20010034646	10/25/2001	Hoyt et al.
_____	121	20010042002	11/15/2001	Koopersmitch
_____	122	20010049620	12/06/2001	Blasko
_____	123	20020010653	01/24/2002	Lee
_____	124	20020010928	01/24/2002	Sahota
_____	125	20020013943	01/31/2002	Haberman et al.

Examiner: Date Considered: _____

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance *and* not considered. Include copy of this form with next communication to applicant.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /N.N.E/

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	ISSUE DATE	PATENT NO.	ATTORNEY DOCKET NO.	CONFIRMATION NO.
17/714,965	01/24/2023	11564015	14-TV1-CON7	7004

26362 7590 01/04/2023

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, AZ 85260

ISSUE NOTIFICATION

The projected patent number and issue date are specified above.

Determination of Patent Term Adjustment under 35 U.S.C. 154 (b)
(application filed on or after May 29, 2000)

The Patent Term Adjustment is 0 day(s). Any patent to issue from the above-identified application will include an indication of the adjustment on the front page.

If a Continued Prosecution Application (CPA) was filed in the above-identified application, the filing date that determines Patent Term Adjustment is the filing date of the most recent CPA.

Applicant will be able to obtain more detailed information by accessing the Patent Application Information Retrieval (PAIR) WEB site (http://pair.uspto.gov).

Any questions regarding the Patent Term Extension or Adjustment determination should be directed to the Office of Patent Legal Administration at (571)-272-7702. Questions relating to issue and publication fee payments should be directed to the Application Assistance Unit (AAU) of the Office of Patents Stakeholder Experience (OPSE), Stakeholder Support Division (SSD) at (571)-272-4200.

INVENTOR(s) (Please see PAIR WEB site http://pair.uspto.gov for additional inventors):

Roy Shkedi, Forest Hills, NY;

APPLICANT(s) (Please see PAIR WEB site http://pair.uspto.gov for additional applicants):

Intent IQ, LLC, Long Island City, NY;

The United States represents the largest, most dynamic marketplace in the world and is an unparalleled location for business investment, innovation, and commercialization of new technologies. The USA offers tremendous resources and advantages for those who invest and manufacture goods here. Through SelectUSA, our nation works to encourage and facilitate business investment. To learn more about why the USA is the best country in the world to develop technology, manufacture products, and grow your business, visit SelectUSA.gov.

To: Louis@valuablepatents.com,donald@valuablepatents.com,dsalavi@northwestpatent.com
From: PAIR_eOfficeAction@uspto.gov
Cc: PAIR_eOfficeAction@uspto.gov
Subject: Private PAIR Correspondence Notification for Customer Number 26362

Jan 05, 2023 08:11:37 AM

Dear PAIR Customer:

LOUIS J. HOFFMAN, P.C.
7689 East Paradise Lane
Suite 2
Scottsdale, AZ 85260
UNITED STATES

The following USPTO patent application(s) associated with your Customer Number, 26362 , have new outgoing correspondence. This correspondence is now available for viewing in Private PAIR.

The official date of notification of the outgoing correspondence will be indicated on the form PTOL-90 accompanying the correspondence.

Disclaimer:
The list of documents shown below is provided as a courtesy and is not part of the official file wrapper. The content of the images shown in PAIR is the official record.

Application	Document	Mailroom Date	Attorney Docket No.
17714965	ISSUE.NTF	01/04/2023	14-TV1-CON7

To view your correspondence online or update your email addresses, please visit us anytime at https://sportal.uspto.gov/secure/myportal/privatepair.

If you have any questions, please email the Electronic Business Center (EBC) at EBC@uspto.gov with 'e-Office Action' on the subject line or call 1-866-217-9197 during the following hours:

Monday - Friday 6:00 a.m. to 12:00 a.m.

Thank you for prompt attention to this notice,

UNITED STATES PATENT AND TRADEMARK OFFICE

PATENT APPLICATION INFORMATION RETRIEVAL SYSTEM

Disclaimer

7,979,307 B2 - Roy Shkedi, New York, NY (US). METHOD AND STORED PROGRAM FOR ACCUMULATING DESCRIPTIVE PROFILE DATA ALONG WITH SOURCE INFORMATION FOR USE IN TARGETING THIRD-PARTY ADVERTISEMENTS. Patent date July 12, 2011. Disclaimer filed September 6, 2017, by the inventor.

I hereby disclaim the following complete claims 8 and 10 of said patent.

(Official Gazette February 21, 2023)