8509001

THE UNITED STATES OF AMERICA

TO ALL TO WHOM THESE PRESENTS SHALL COME:

UNITED STATES DEPARTMENT OF COMMERCE

United States Patent and Trademark Office

July 22, 2024

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: ***11/623,653***
FILING DATE: ***January 16, 2007***
PATENT NUMBER: ***7716217***
ISSUE DATE: ***May 11, 2010***

Certified by

Katherine Kelly Vidal

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

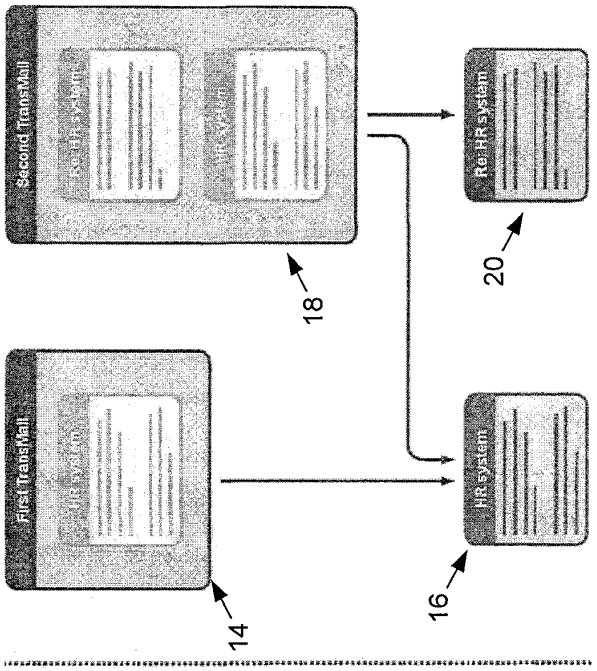


FIGURE 1B

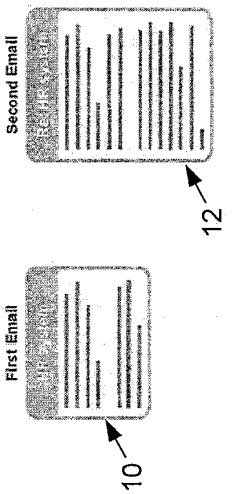


FIGURE 1A

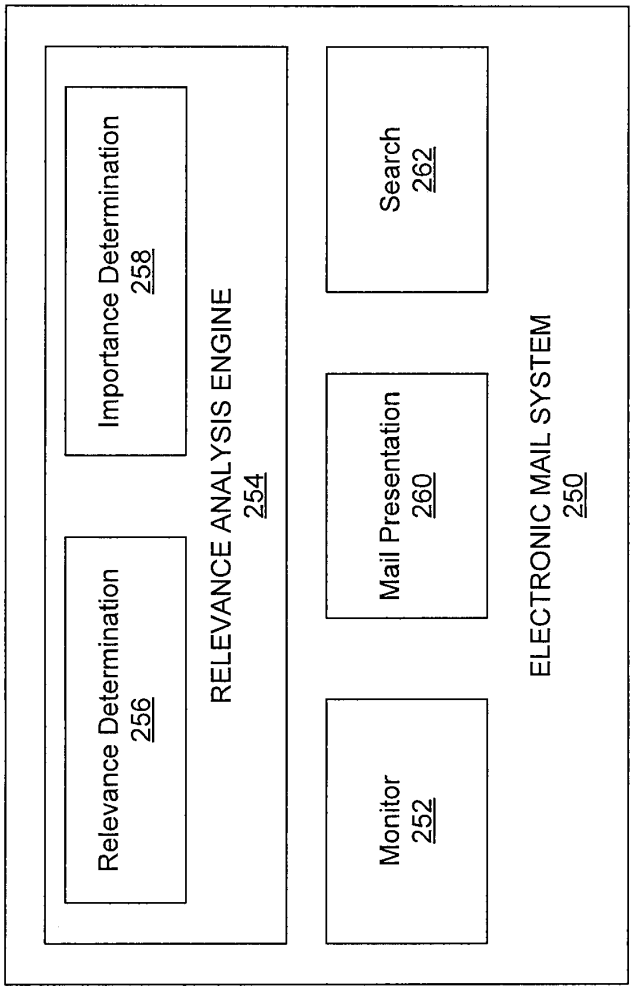

ELECTRONIC MAIL SYSTEM
250
Monitor
252
Mail Presentation
260
Search
262
RELEVANCE ANALYSIS ENGINE
254
Relevance Determination
256
Importance Determination
258

FIGURE 2

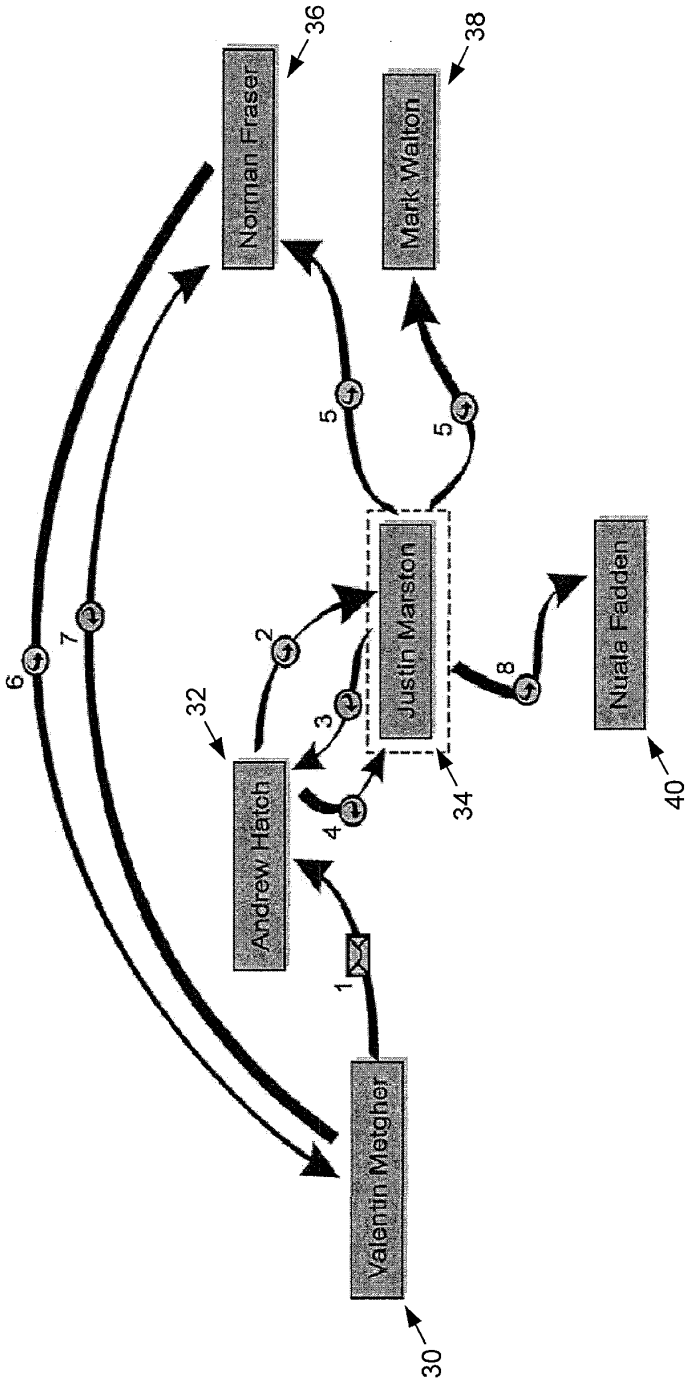

Norman Fraser
36
Mark Walton
38
Justin Marston
Andrew Hatch
32
Valentin Metgher
30
Nuala Fadden
40
34
1
2
3
4
5
5
6
7
8

FIGURE 3

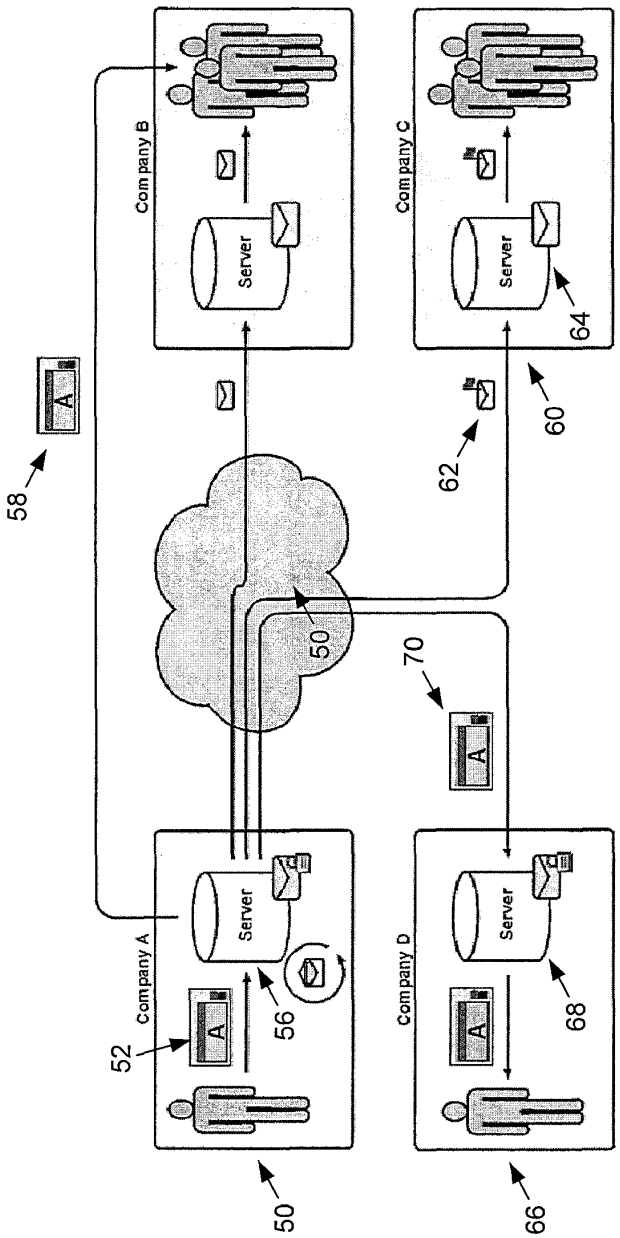

Company A
Company B
Company C
Company D
Server
A
50
52
56
58
60
62
64
66
68
70

FIGURE 4

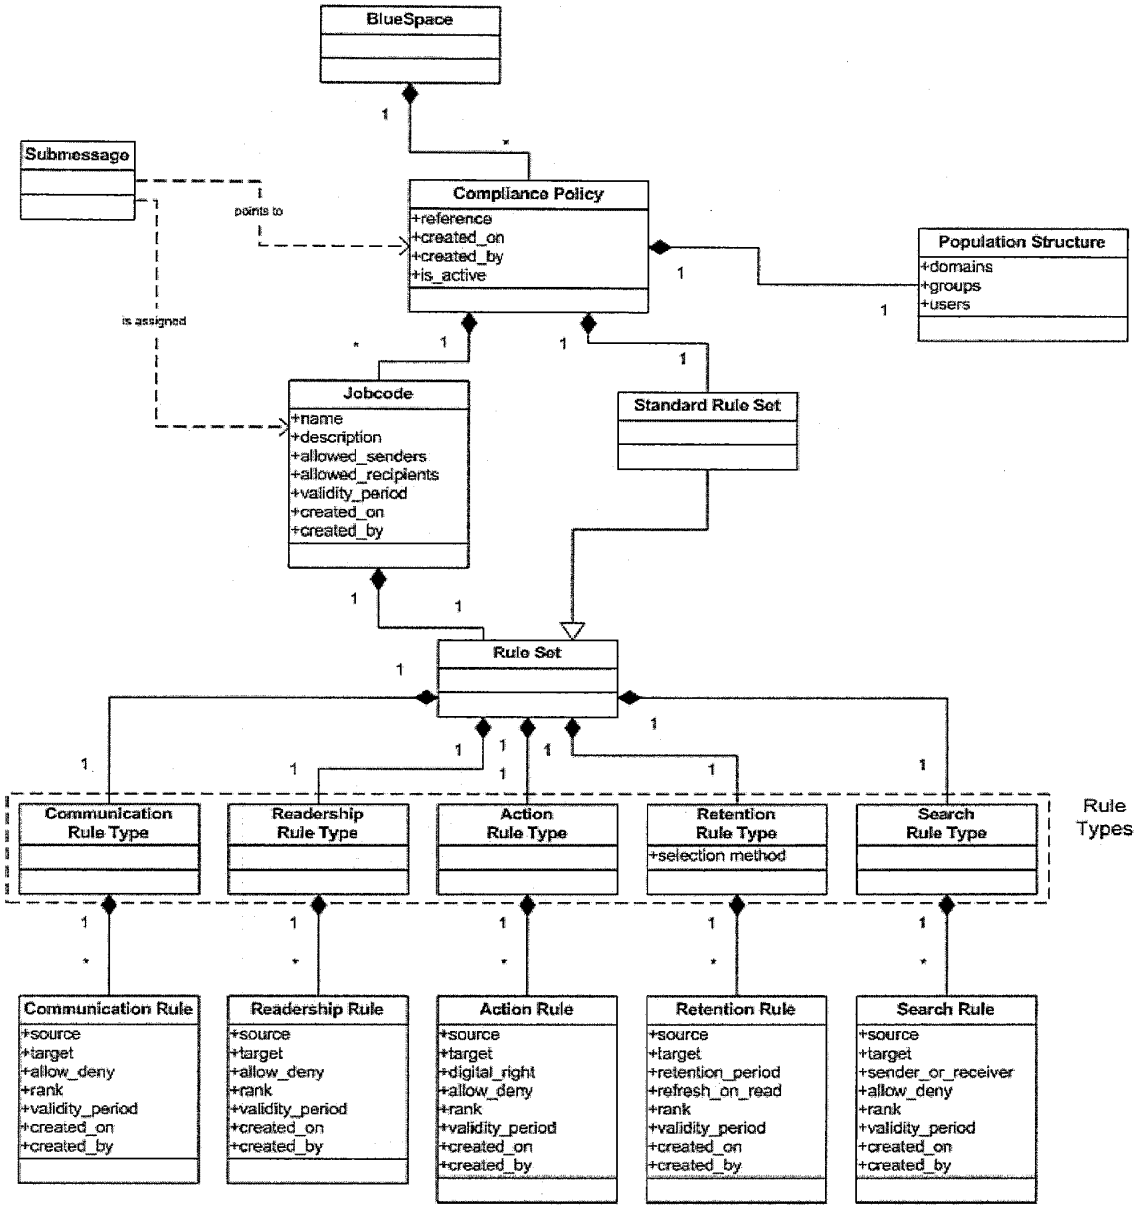

BlueSpace
Submessage
points to
is assigned
Compliance Policy
+reference
+created_on
+created_by
+is_active
Population Structure
+domains
+groups
+users
Jobcode
+name
+description
+allowed_senders
+allowed_recipients
+validity_period
+created_on
+created_by
Standard Rule Set
Rule Set
Communication Rule Type
Readership Rule Type
Action Rule Type
Retention Rule Type
+selection method
Search Rule Type
Rule Types
Communication Rule
+source
+target
+allow_deny
+rank
+validity_period
+created_on
+created_by
Readership Rule
+source
+target
+allow_deny
+rank
+validity_period
+created_on
+created_by
Action Rule
+source
+target
+digital_right
+allow_deny
+rank
+validity_period
+created_on
+created_by
Retention Rule
+source
+target
+retention_period
+refresh_on_read
+rank
+validity_period
+created_on
+created_by
Search Rule
+source
+target
+sender_or_receiver
+allow_deny
+rank
+validity_period
+created_on
+created_by

FIGURE 5

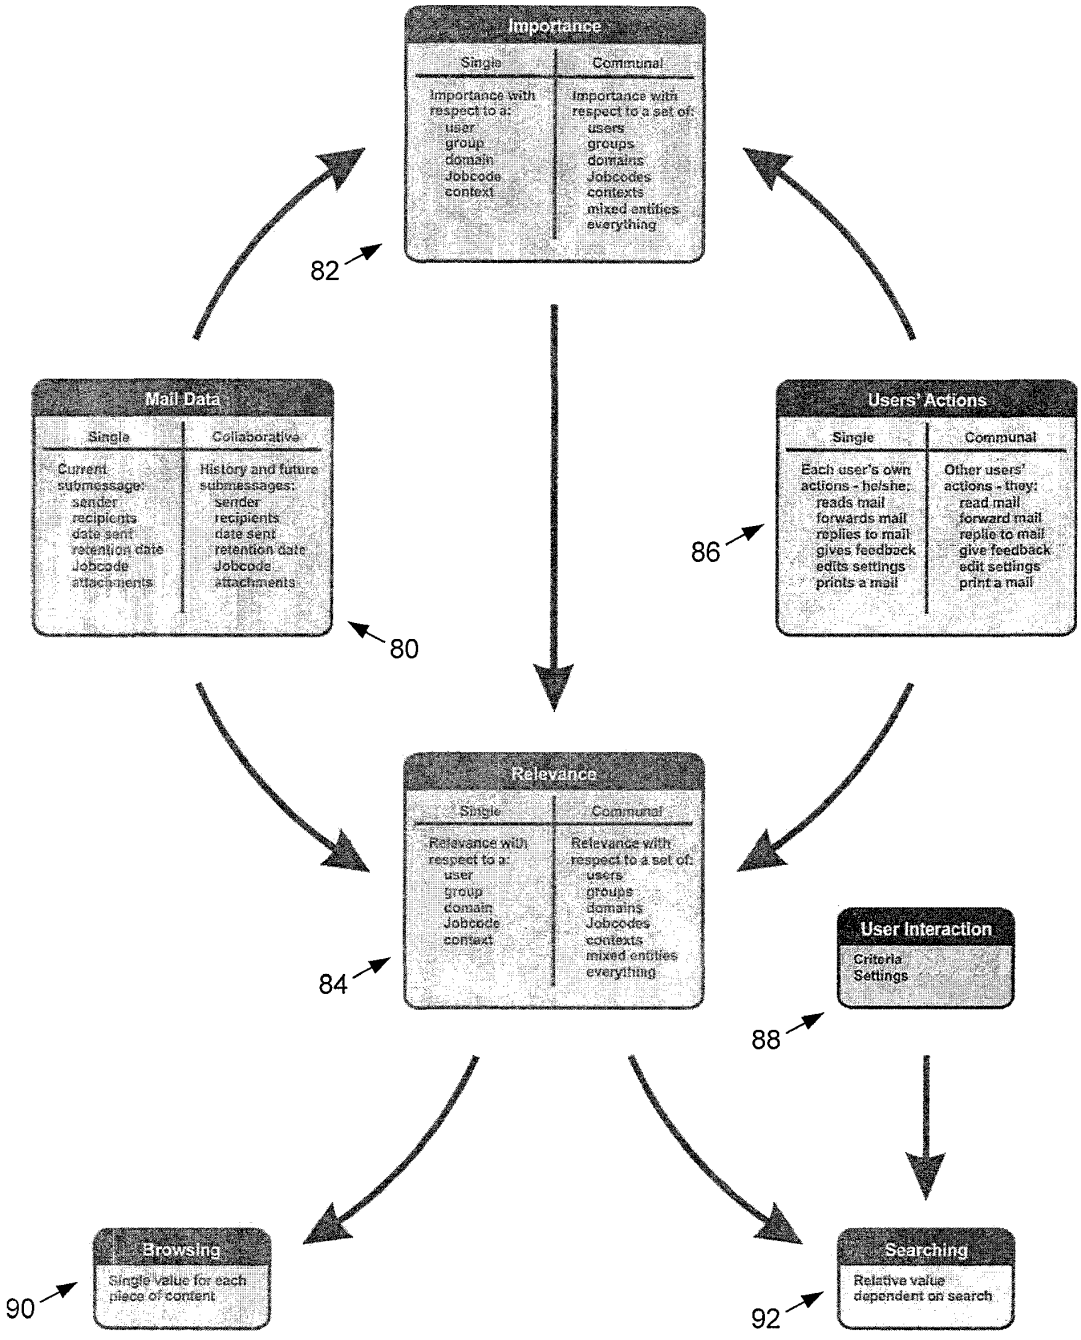

Importance
Single
Communal
Importance with respect to a: user group domain Jobcode context
Importance with respect to a set of: users groups domains Jobcodes contexts mixed entities everything
82
Mail Data
Single
Collaborative
Current submessage: sender recipients date sent retention date Jobcode attachments
History and future submessages: sender recipients date sent retention date Jobcode attachments
80
Users' Actions
Single
Communal
Each user's own actions - he/she: reads mail forwards mail replies to mail gives feedback edits settings prints a mail
Other users' actions - they: read mail forward mail replie to mail give feedback edit settings print a mail
86
Relevance
Single
Communal
Relevance with respect to a: user group domain Jobcode context
Relevance with respect to a set of: users groups domains Jobcodes contexts mixed entities everything
84
User Interaction
Criteria
Settings
88
Browsing
Single value for each piece of content
90
Searching
Relative value dependent on search
92

FIGURE 6

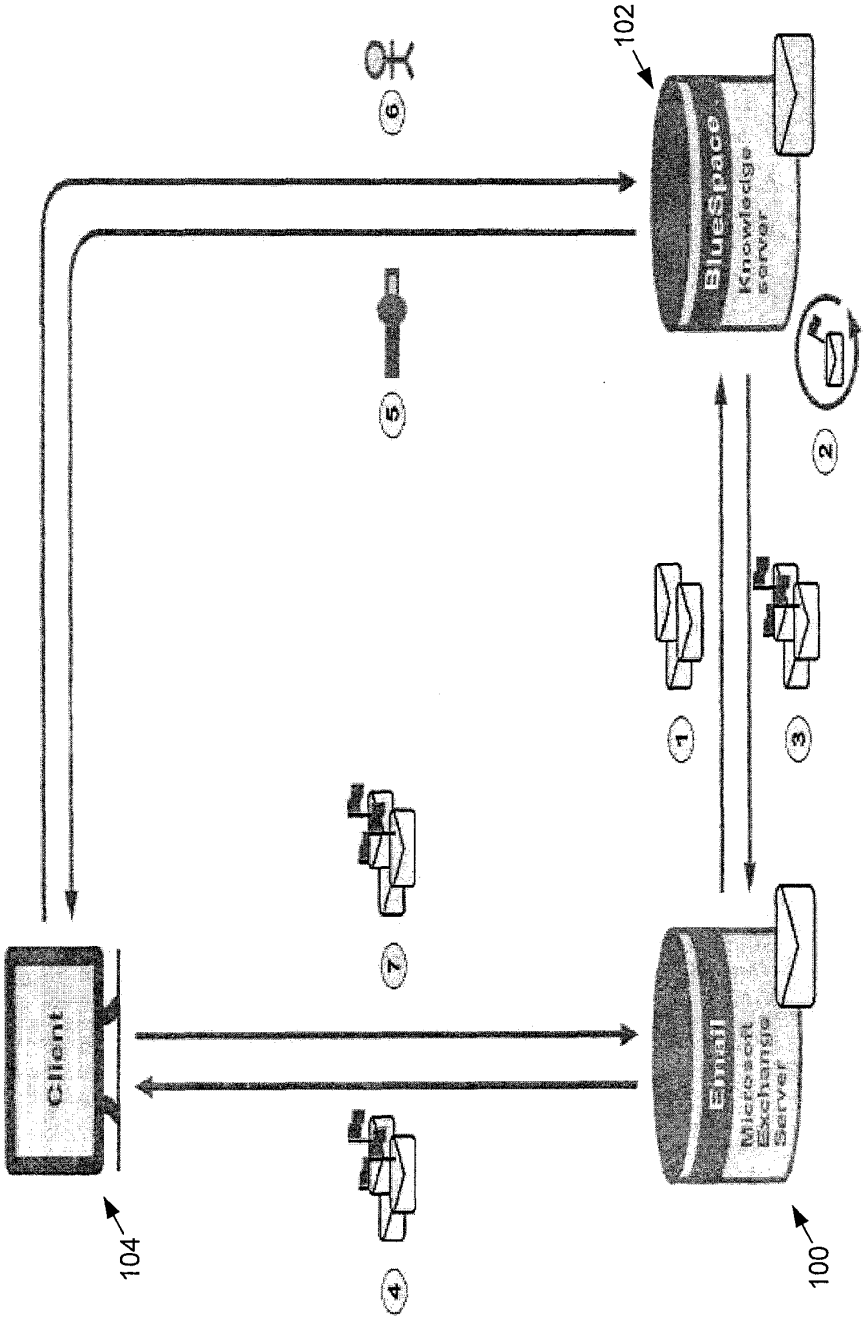

102
BlueSpace
Knowledge
server
OK
6
5
2
1
3
7
Client
Email
Microsoft
Exchange
Server
104
4
100

FIGURE 7

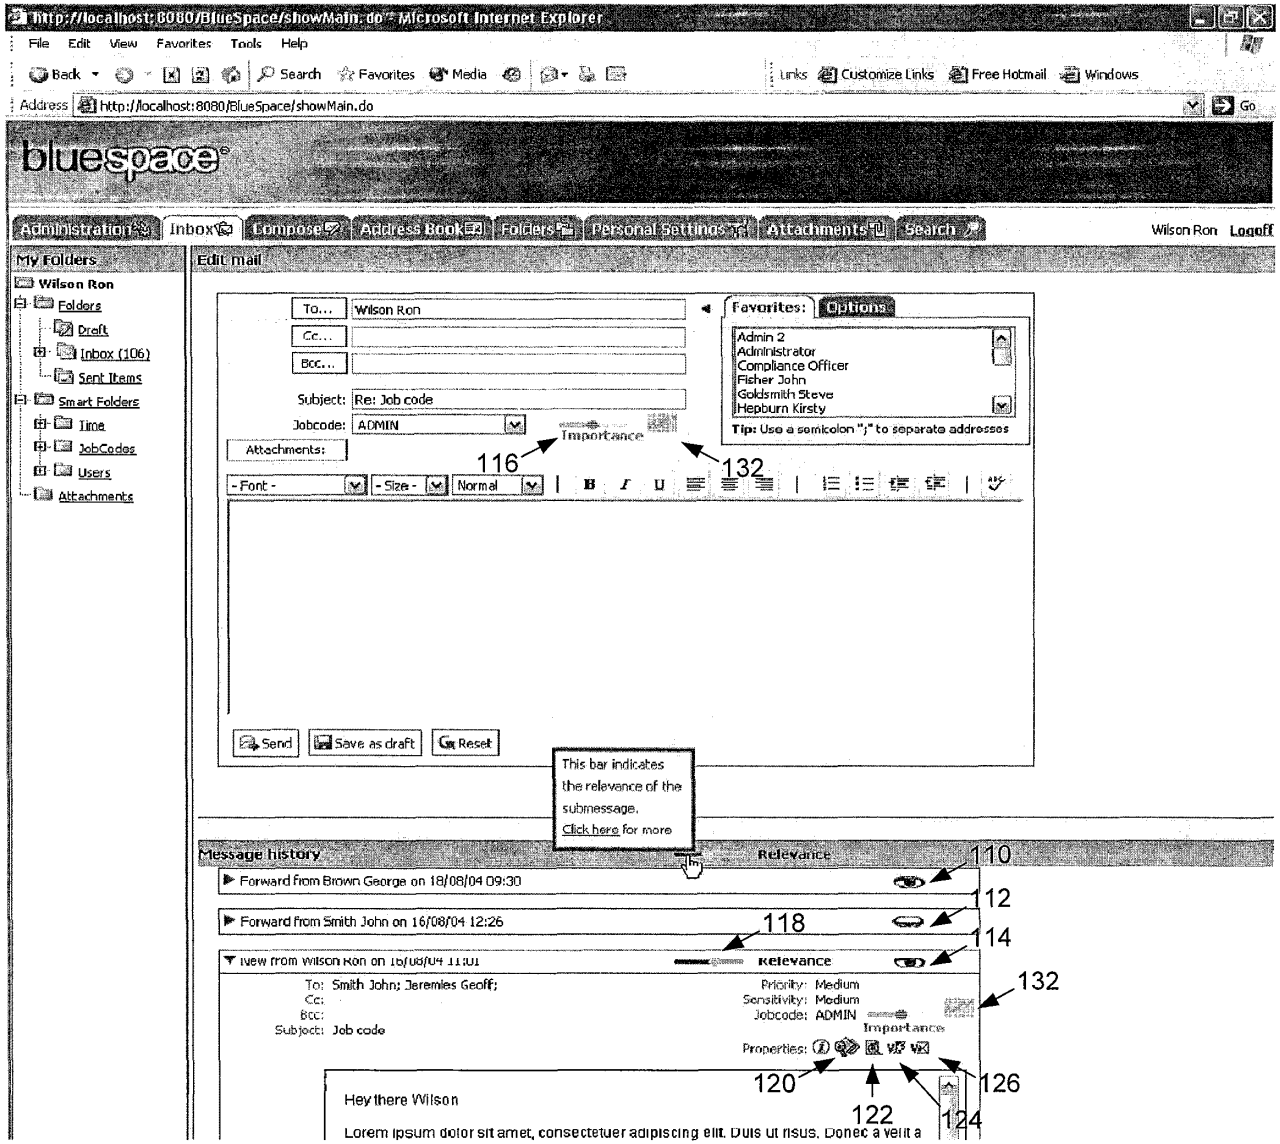


FIGURE 8

Title: DETERMINING RELEVANCE OF ELECTRONIC CONTENT
Inventor(s): Justin Marston et al.
Atty. Docket No.: 23682-12381
Sheet 8 of 22

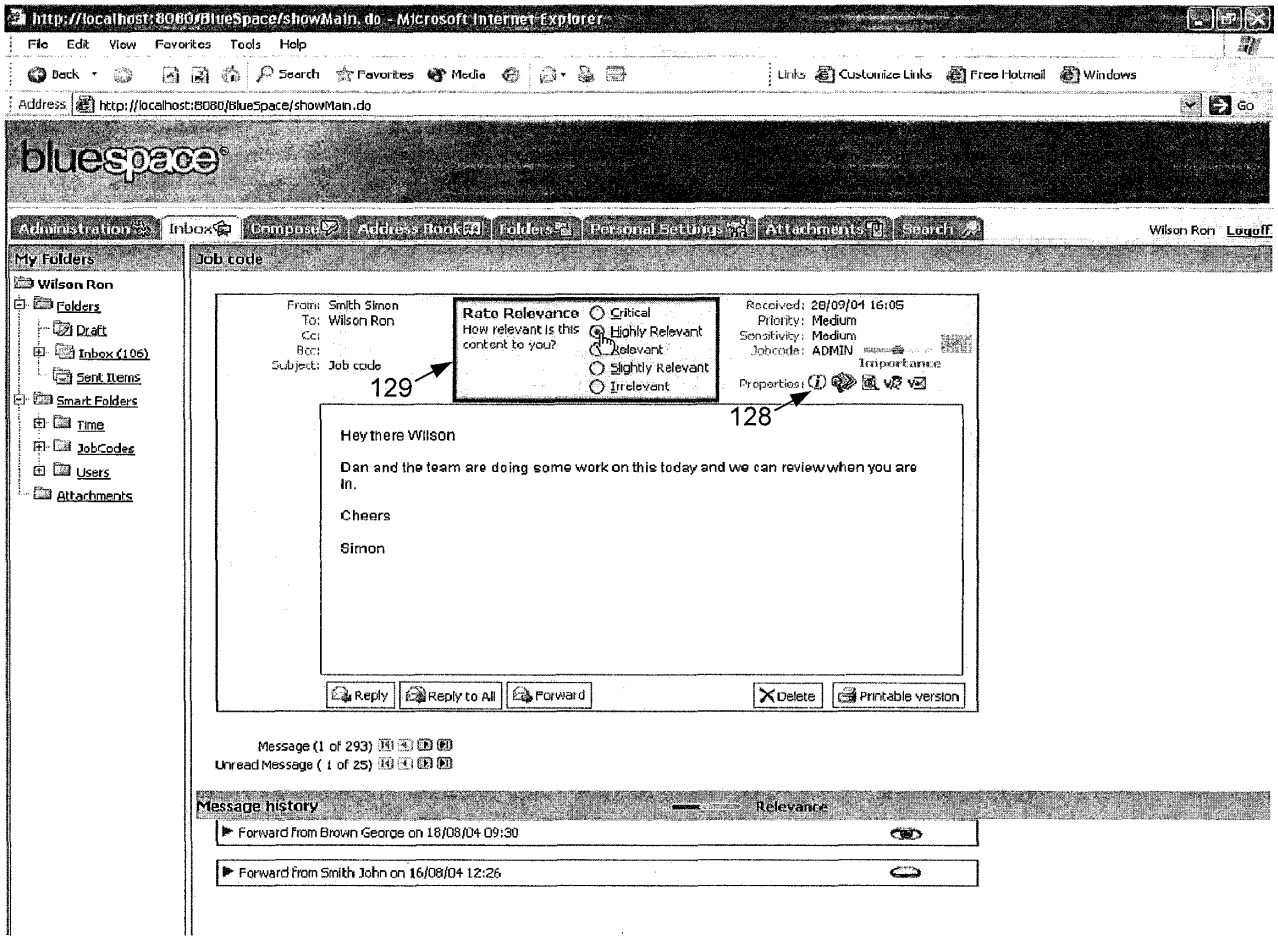

http://localhost:8080/BlueSpace/showMain.do - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Address http://localhost:8080/BlueSpace/showMain.do
bluespace
Administration Inbox Compose Address Book Folders Personal Settings Attachments Search
Wilson Ron Logoff
My Folders
Wilson Ron
Folders
Draft
Inbox (106)
Sent Items
Smart Folders
Time
JobCodes
Users
Attachments
Job code
From: Smith Simon
To: Wilson Ron
Cc:
Bcc:
Subject: Job code
Rate Relevance
How relevant is this content to you?
Critical
Highly Relevant
Relevant
Slightly Relevant
Irrelevant
129
Received: 28/09/04 16:05
Priority: Medium
Sensitivity: Medium
Jobcode: ADMIN
Importance
Properties:
128
Hey there Wilson
Dan and the team are doing some work on this today and we can review when you are in.
Cheers
Simon
Reply
Reply to All
Forward
Delete
Printable version
Message (1 of 293)
Unread Message (1 of 25)
Message history
Relevance
Forward from Brown George on 18/08/04 09:30
Forward from Smith John on 16/08/04 12:26

FIGURE 9

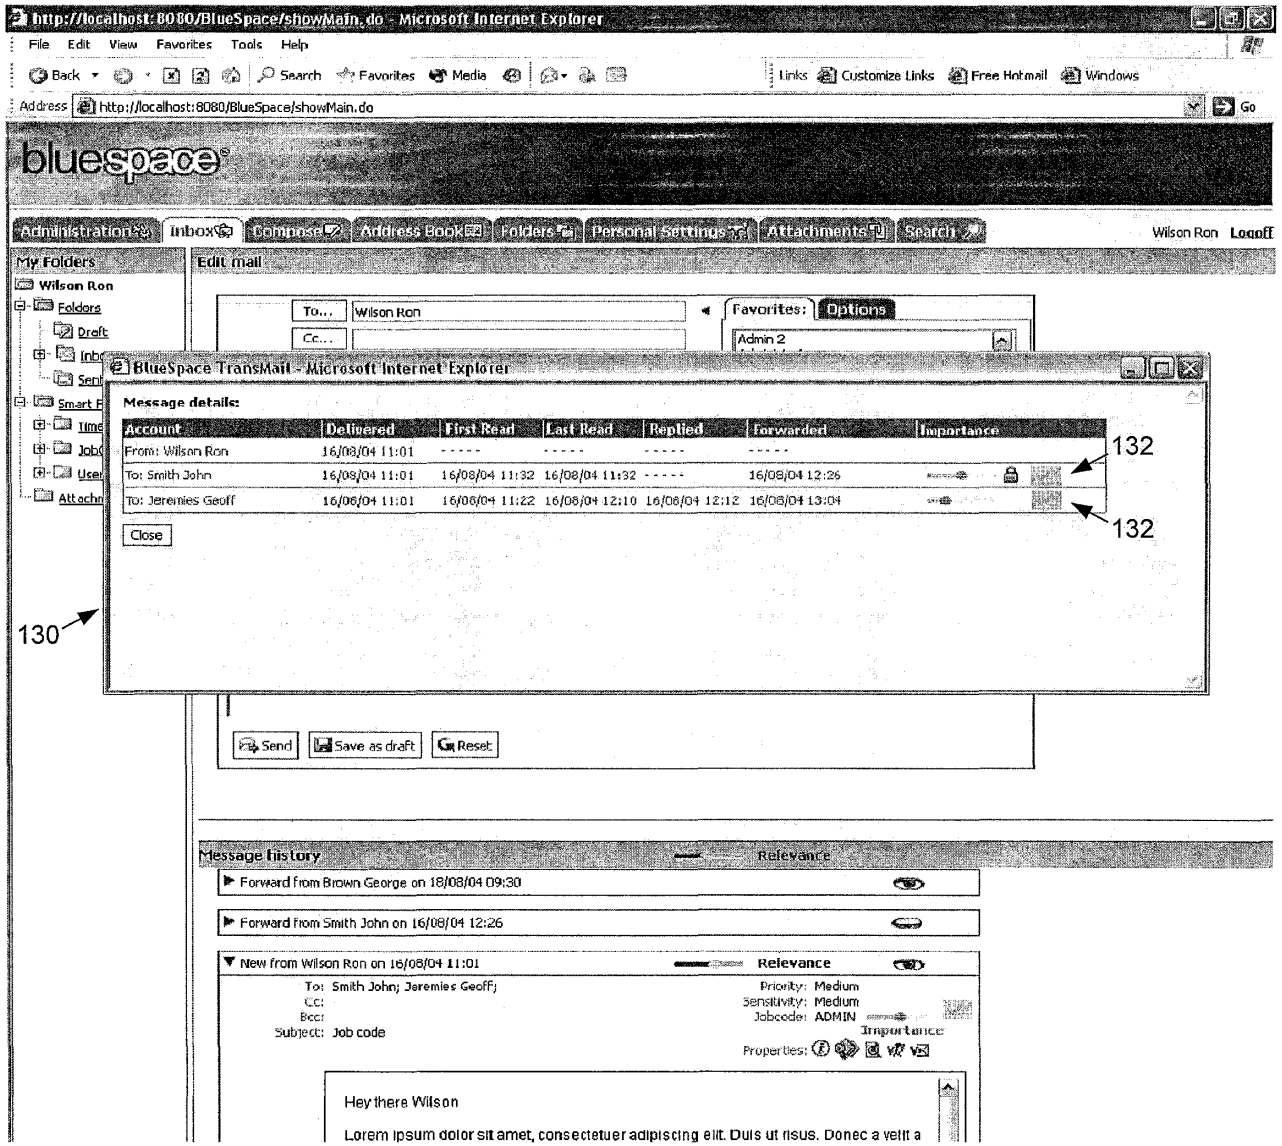

http://localhost:8080/BlueSpace/showMain.do - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Links Customize Links Free Hotmail Windows
Address http://localhost:8080/BlueSpace/showMain.do
Go
bluespace
Administration
Inbox
Compose
Address Book
Folders
Personal Settings
Attachments
Search
Wilson Ron Logoff
My Folders
Wilson Ron
Folders
Draft
Edit mail
To...
Wilson Ron
Cc...
Favorites:
Options
Admin 2
BlueSpace TransMail - Microsoft Internet Explorer
Message details:
Account
Delivered
First Read
Last Read
Replied
Forwarded
Importance
From: Wilson Ron 16/08/04 11:01 ----- ----- ----- -----
To: Smith John 16/08/04 11:01 16/08/04 11:32 16/08/04 11:32 ----- 16/08/04 12:26
To: Jeremies Geoff 16/08/04 11:01 16/08/04 11:22 16/08/04 12:10 16/08/04 12:12 16/08/04 13:04
Close
132
132
130
Send
Save as draft
Reset
Message history
Relevance
Forward from Brown George on 18/08/04 09:30
Forward from Smith John on 16/08/04 12:26
New from Wilson Ron on 16/08/04 11:01
Relevance
To: Smith John; Jeremies Geoff;
Cc:
Bcc:
Subject: Job code
Priority: Medium
Sensitivity: Medium
Jobcode: ADMIN
Importance:
Properties:
Hey there Wilson
Lorem ipsum dolor sit amet, consectetuer adipiscing elit. Duis ut risus. Donec a velit a

FIGURE 10

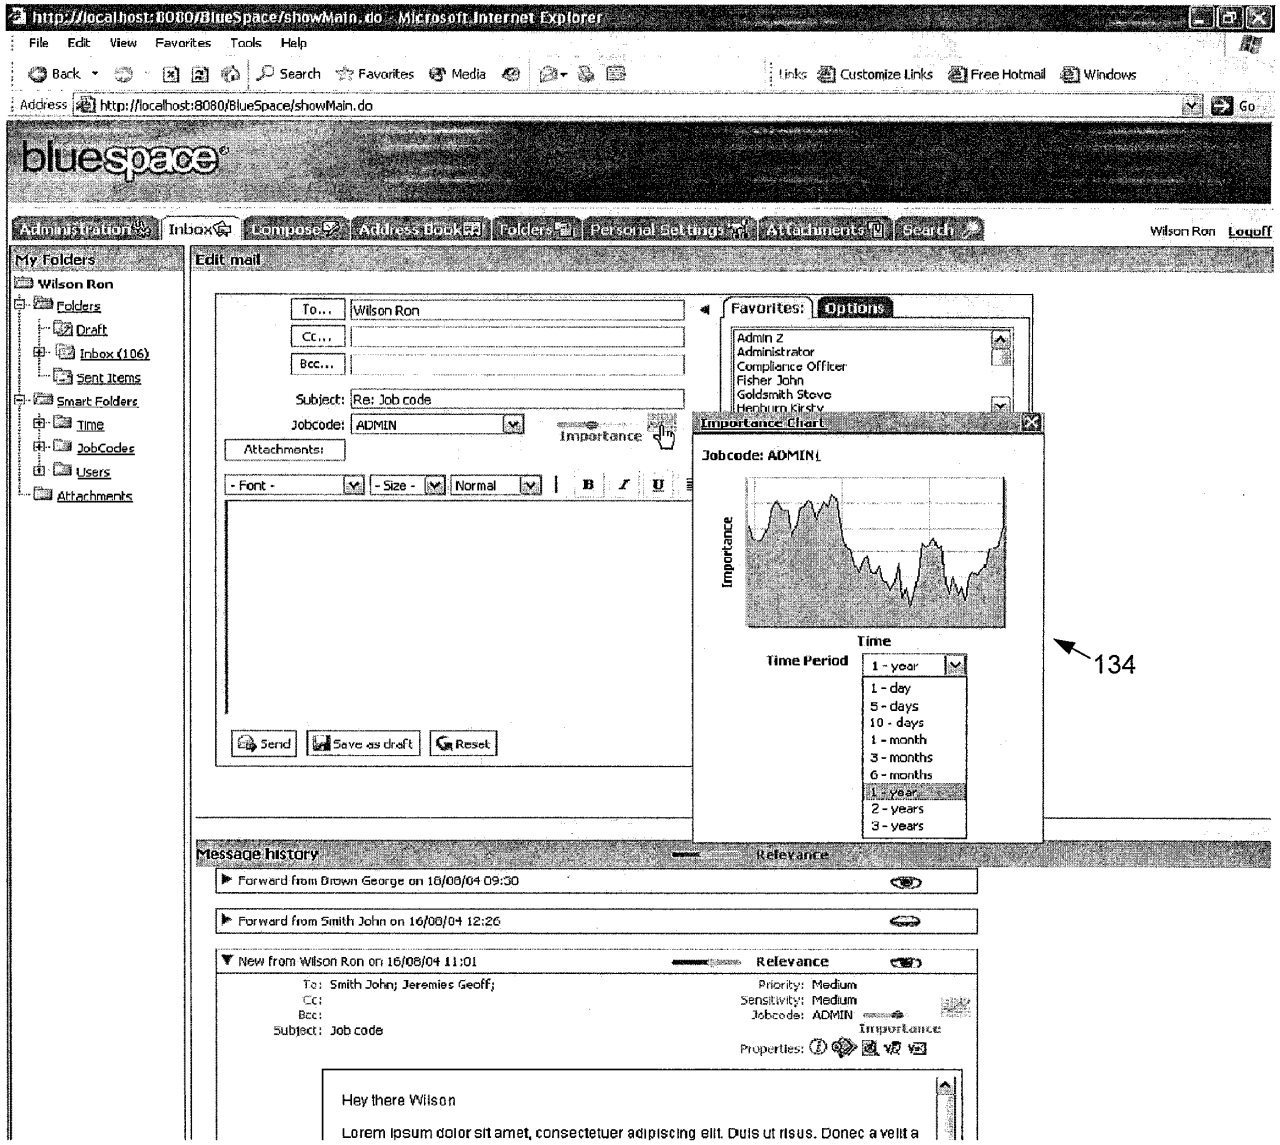

http://localhost:8080/BlueSpace/showMain.do - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Links Customize Links Free Hotmail Windows
Address http://localhost:8080/BlueSpace/showMain.do
Go
bluespace
Administration Inbox Compose Address Book Folder Personal Settings Attachments Search
Wilson Ron Logoff
My Folders
Wilson Ron
Folders
Draft
Inbox (106)
Sent Items
Smart Folders
Time
JobCodes
Users
Attachments
Edit mail
To... Wilson Ron
Cc...
Bcc...
Subject: Re: Job code
Jobcode: ADMIN
Importance
Attachments:
- Font - - Size - Normal
Send Save as draft Reset
Favorites: Options
Admin Z
Administrator
Compliance Officer
Fisher John
Goldsmith Steve
Henburn Kirsty
Importance Chart
Jobcode: ADMIN
Importance
Time
Time Period 1 - year
1 - day
5 - days
10 - days
1 - month
3 - months
6 - months
1 - year
2 - years
3 - years
134
Message history
Relevance
Forward from Brown George on 16/08/04 09:30
Forward from Smith John on 16/08/04 12:26
New from Wilson Ron on 16/08/04 11:01
Relevance
To: Smith John; Jeremies Geoff;
Cc:
Bcc:
Subject: Job code
Priority: Medium
Sensitivity: Medium
Jobcode: ADMIN
Importance
Properties:
Hey there Wilson
Lorem ipsum dolor sit amet, consectetuer adipiscing elit. Duis ut risus. Donec a velit a

FIGURE 11


http://localhost:8000/BlueSpace/showMain.do - Microsoft Internet Explorer
bluespace
Inbox
137
136
From
Subject
Relevance
Jobcode
Received
First Read
Last Read
Sensitivity
Priority
brandonm
Re: Job Codes
Sara Ball
Test
governance patent
Re: more (ignore test)...
search patent
WTEVA
Medium

FIGURE 12

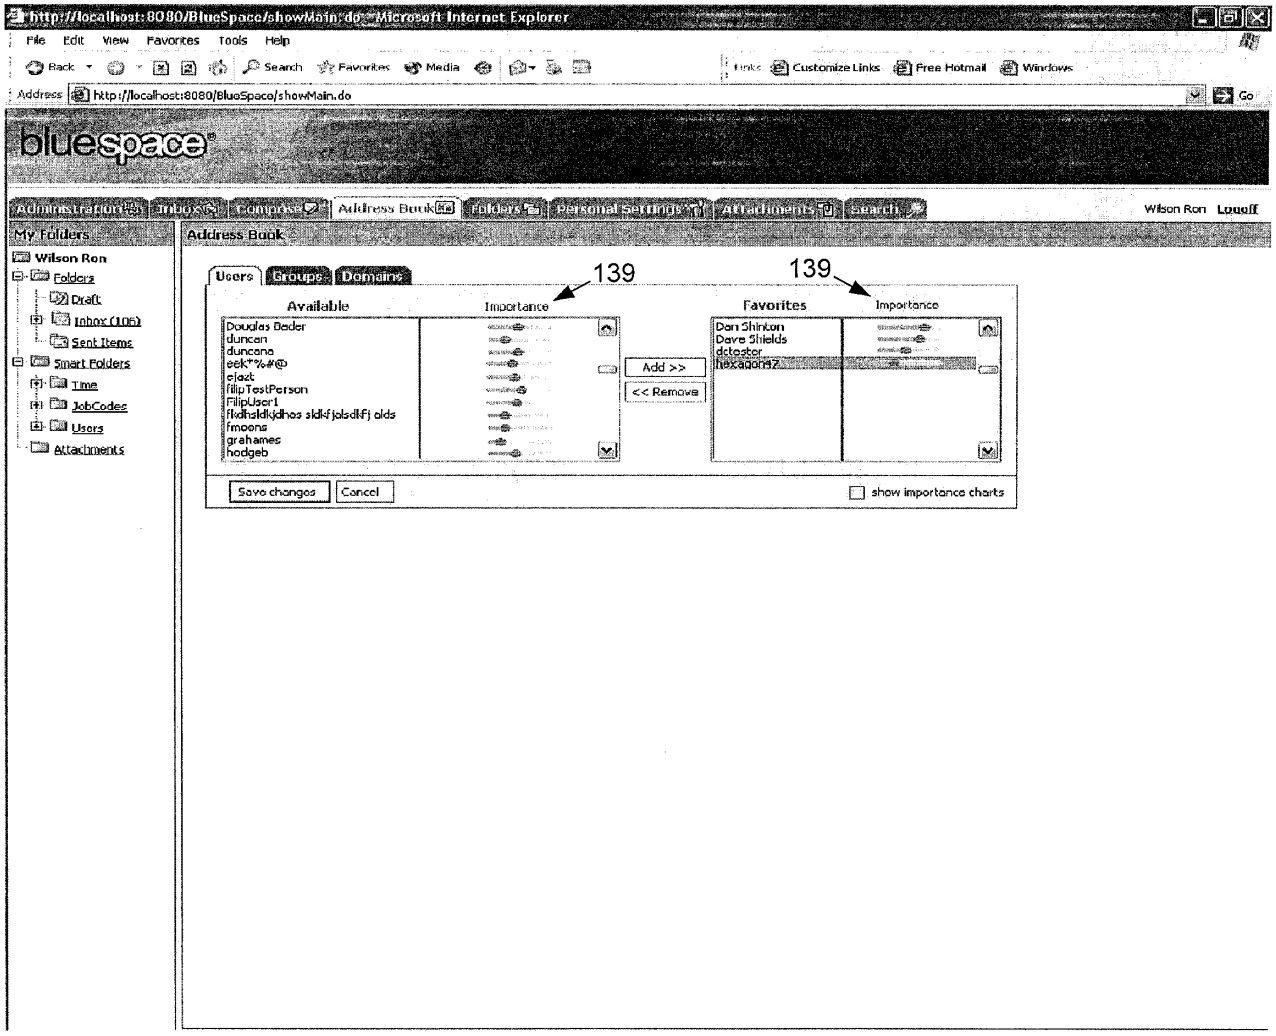

http://localhost:8080/BlueSpace/showMain.do - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://localhost:8080/BlueSpace/showMain.do
Go
bluespace
Administration
Inbox
Compose
Address Book
Folders
Personal Settings
Attachments
Search
Wilson Ron Logoff
My Folders
Wilson Ron
Folders
Draft
Inbox (105)
Sent Items
Smart Folders
Time
JobCodes
Users
Attachments
Address Book
Users
Groups
Domains
139
139
Available
Importance
Favorites
Importance
Douglas Beder
duncan
duncana
eek*%#@
elodt
FilipTestPerson
FilipUser1
fkdhsldkjdhas sldkfjalsdlkfj alds
fmoons
grahames
hodgeb
Add >>
<< Remove
Dan Shinton
Dave Shields
dctaster
hexagon7
Save changes
Cancel
show importance charts

FIGURE 13

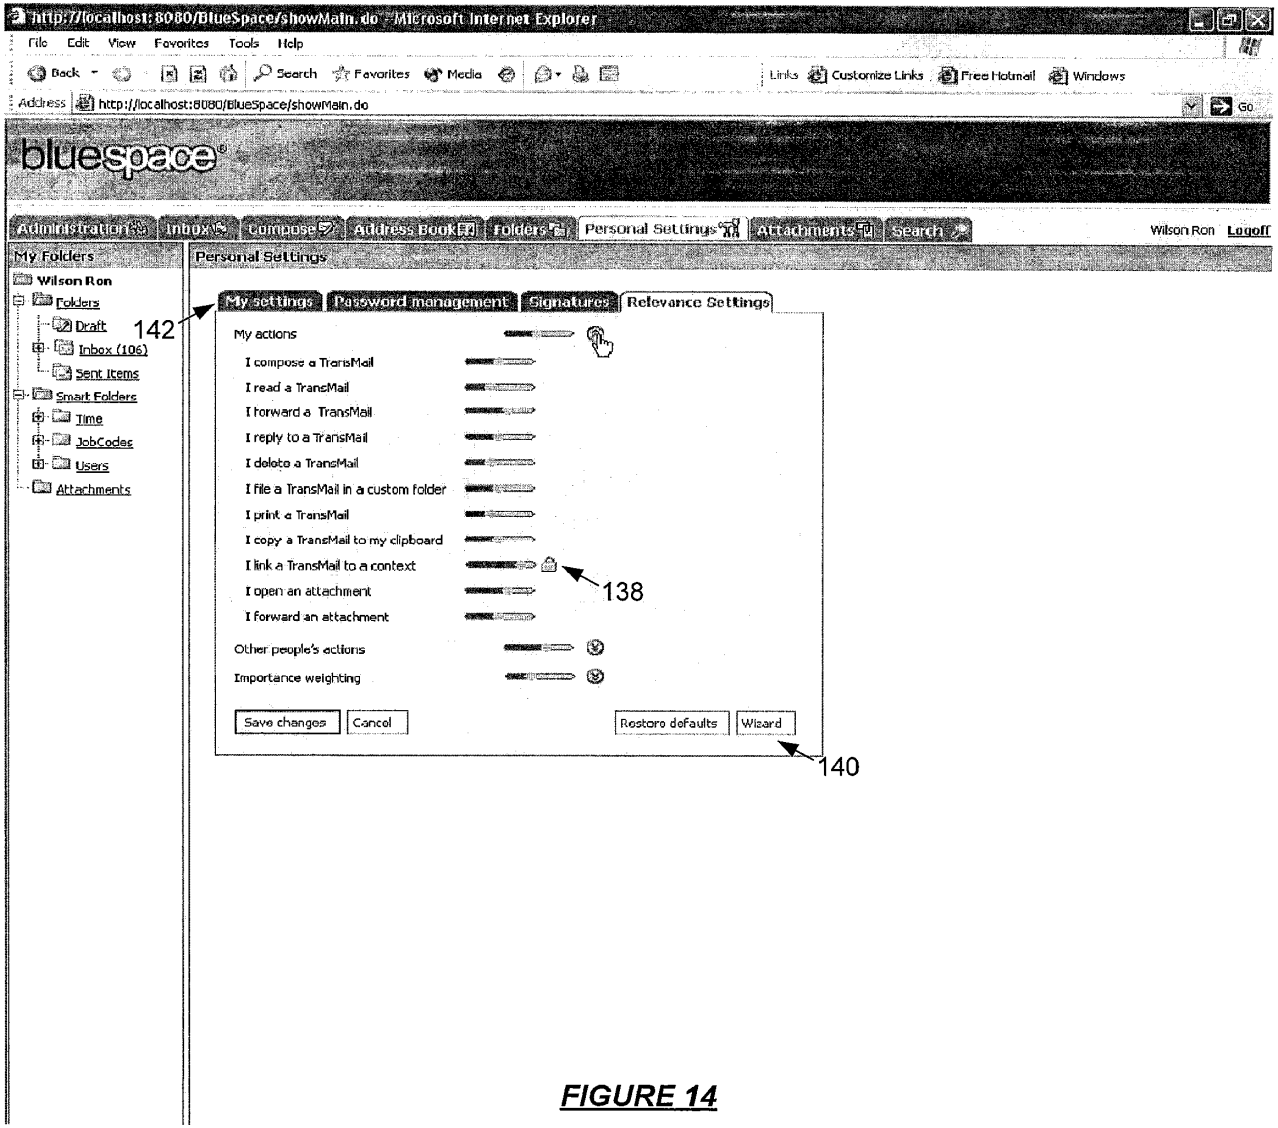

http://localhost:8080/BlueSpace/showMain.do - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Address http://localhost:8080/BlueSpace/showMain.do
bluespace
Administration Inbox Compose Address Book Folders Personal Settings Attachments Search
Wilson Ron Logoff
My Folders
Wilson Ron
Folders
Draft
Inbox (106)
Sent Items
Smart Folders
Time
JobCodes
Users
Attachments
142
Personal Settings
My settings Password management Signatures Relevance Settings
My actions
I compose a TransMail
I read a TransMail
I forward a TransMail
I reply to a TransMail
I delete a TransMail
I file a TransMail in a custom folder
I print a TransMail
I copy a TransMail to my clipboard
I link a TransMail to a context
I open an attachment
I forward an attachment
138
Other people's actions
Importance weighting
Save changes
Cancel
Restore defaults
Wizard
140

FIGURE 14

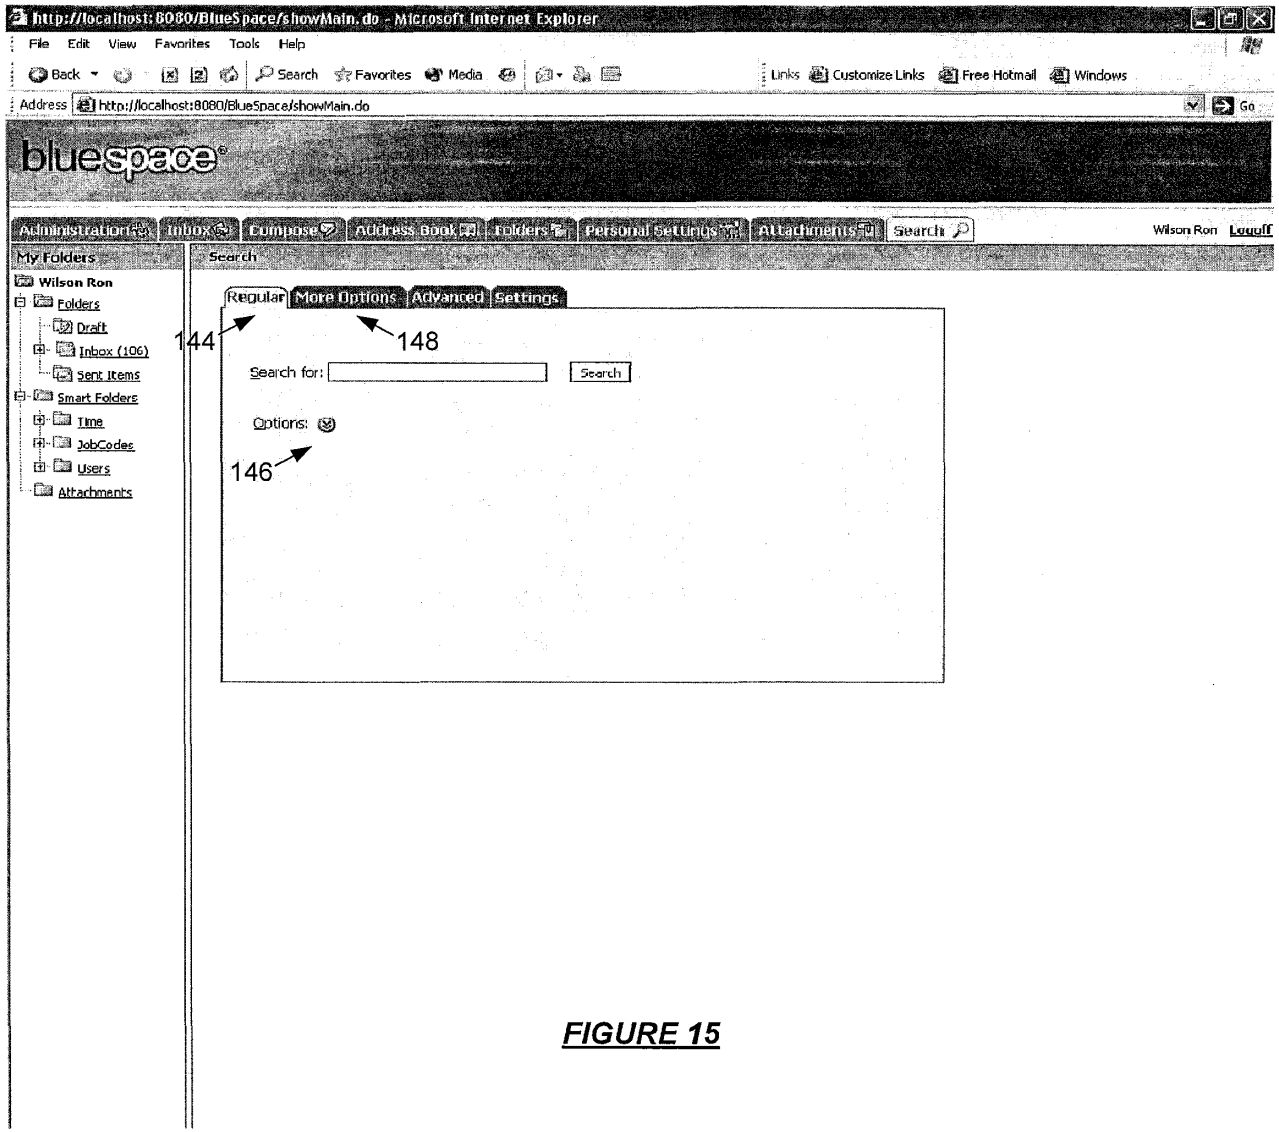

http://localhost:8080/BlueSpace/showMain.do - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Links Customize Links Free Hotmail Windows
Address http://localhost:8080/BlueSpace/showMain.do
Go
bluespace
Administration Inbox Compose Address Book Folders Personal Settings Attachments Search
Wilson Ron Logoff
My Folders
Search
Wilson Ron
Folders
Draft
Inbox (106)
Sent Items
Smart Folders
Time
JobCodes
Users
Attachments
Regular More Options Advanced Settings
144
148
Search for:
Search
Options:
146

FIGURE 15

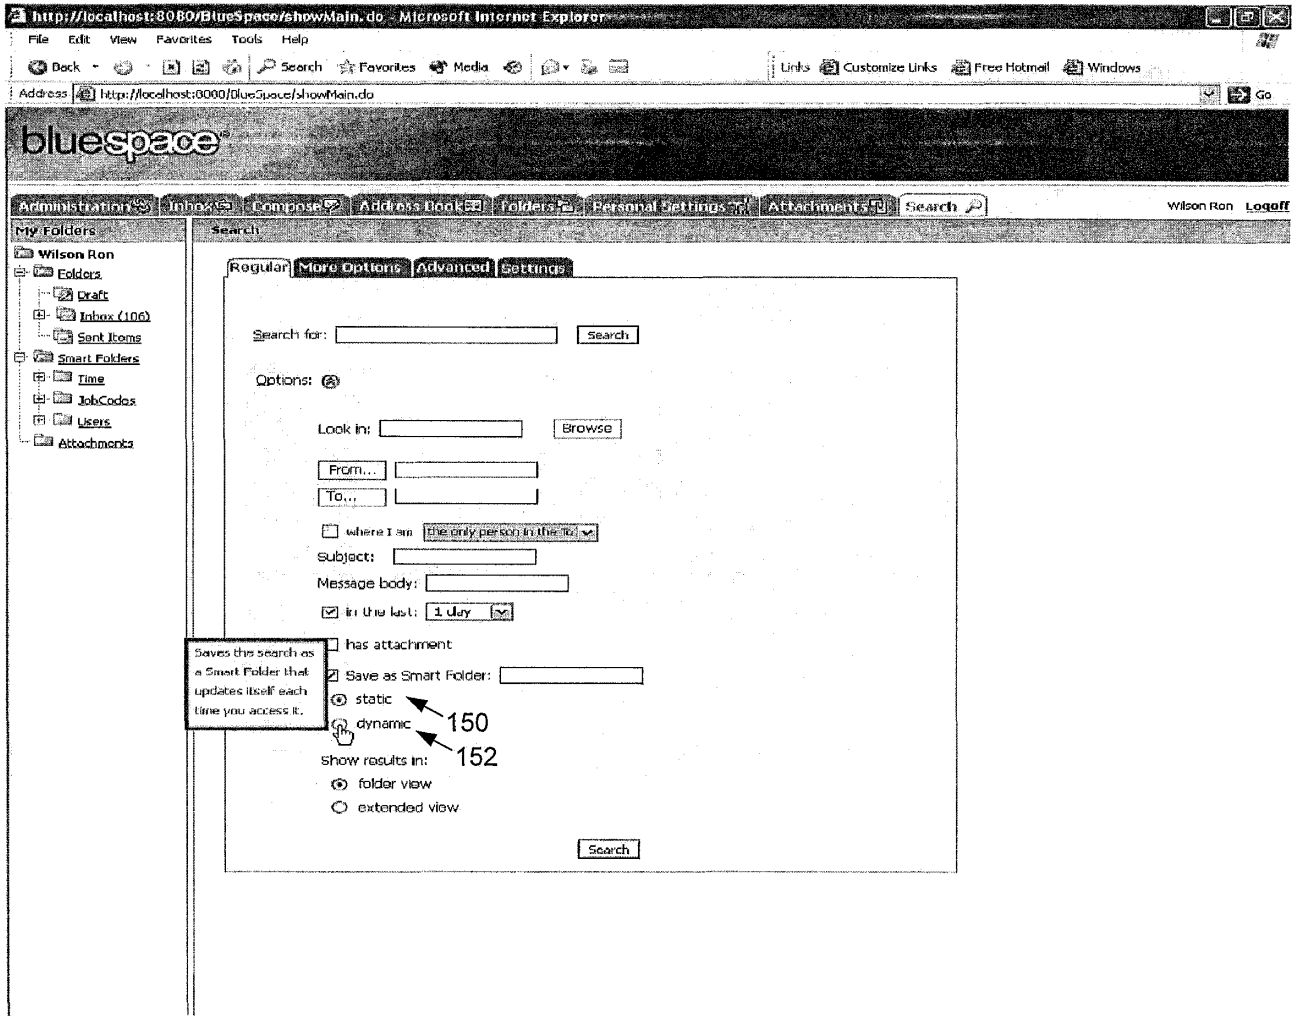


FIGURE 16

Title: DETERMINING RELEVANCE OF ELECTRONIC CONTENT
Inventor(s): Justin Marston et al.
Atty. Docket No.: 23682-12381
Sheet 16 of 22

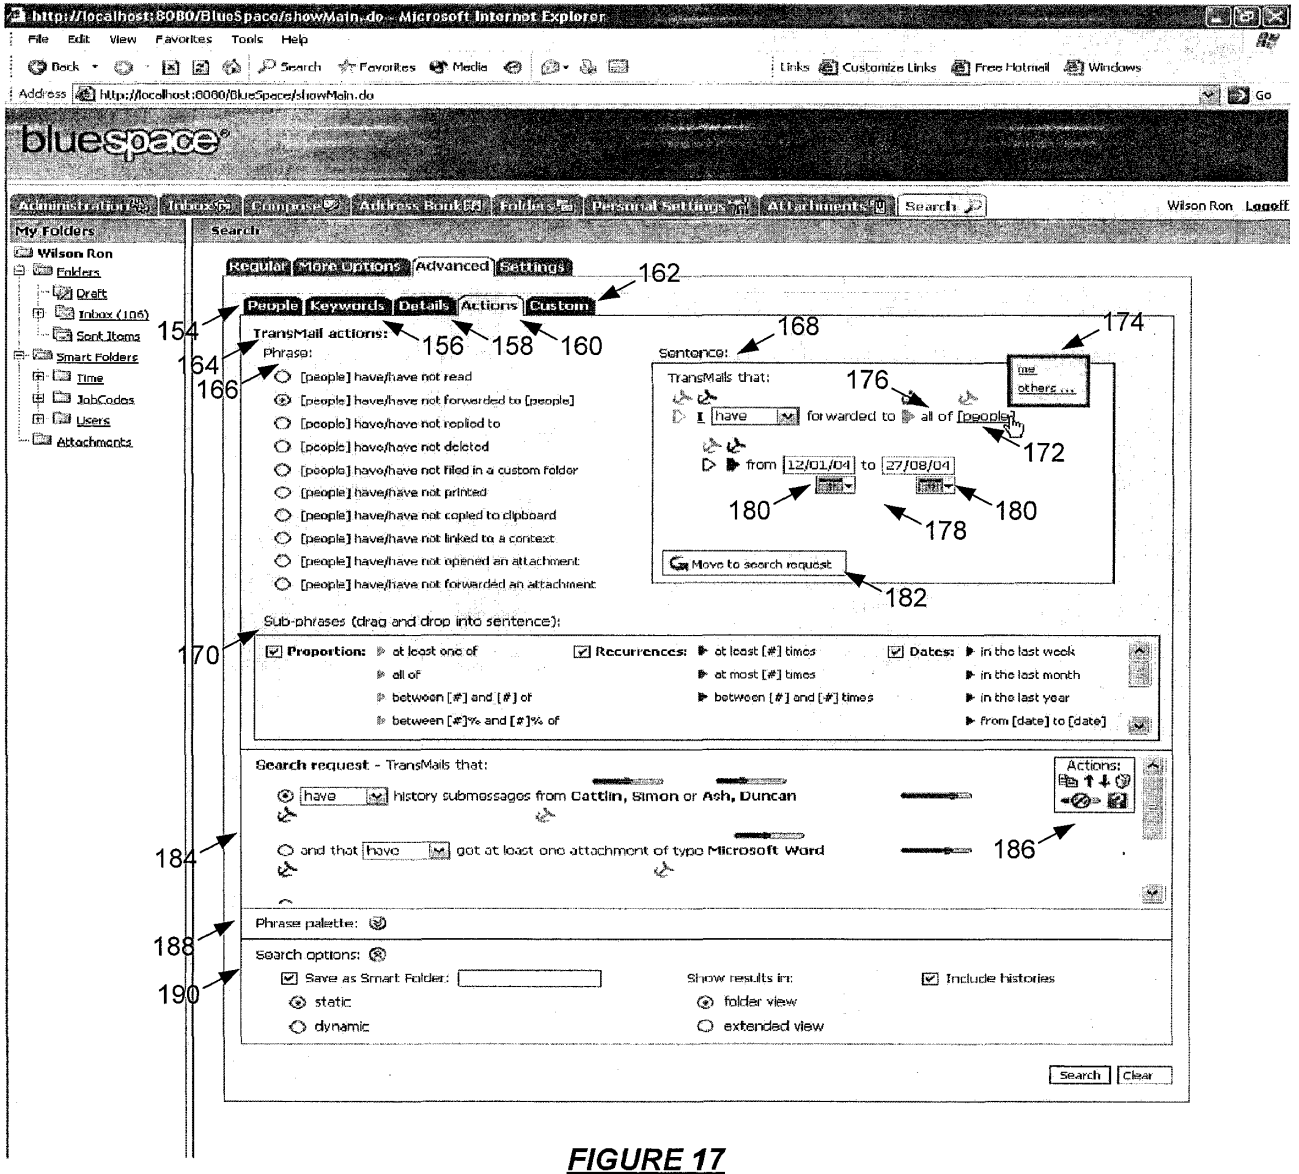


FIGURE 17

Title: DETERMINING RELEVANCE OF ELECTRONIC CONTENT
Inventor(s): Justin Marston et al.
Atty. Docket No.: 23682-12381
Sheet 17 of 22

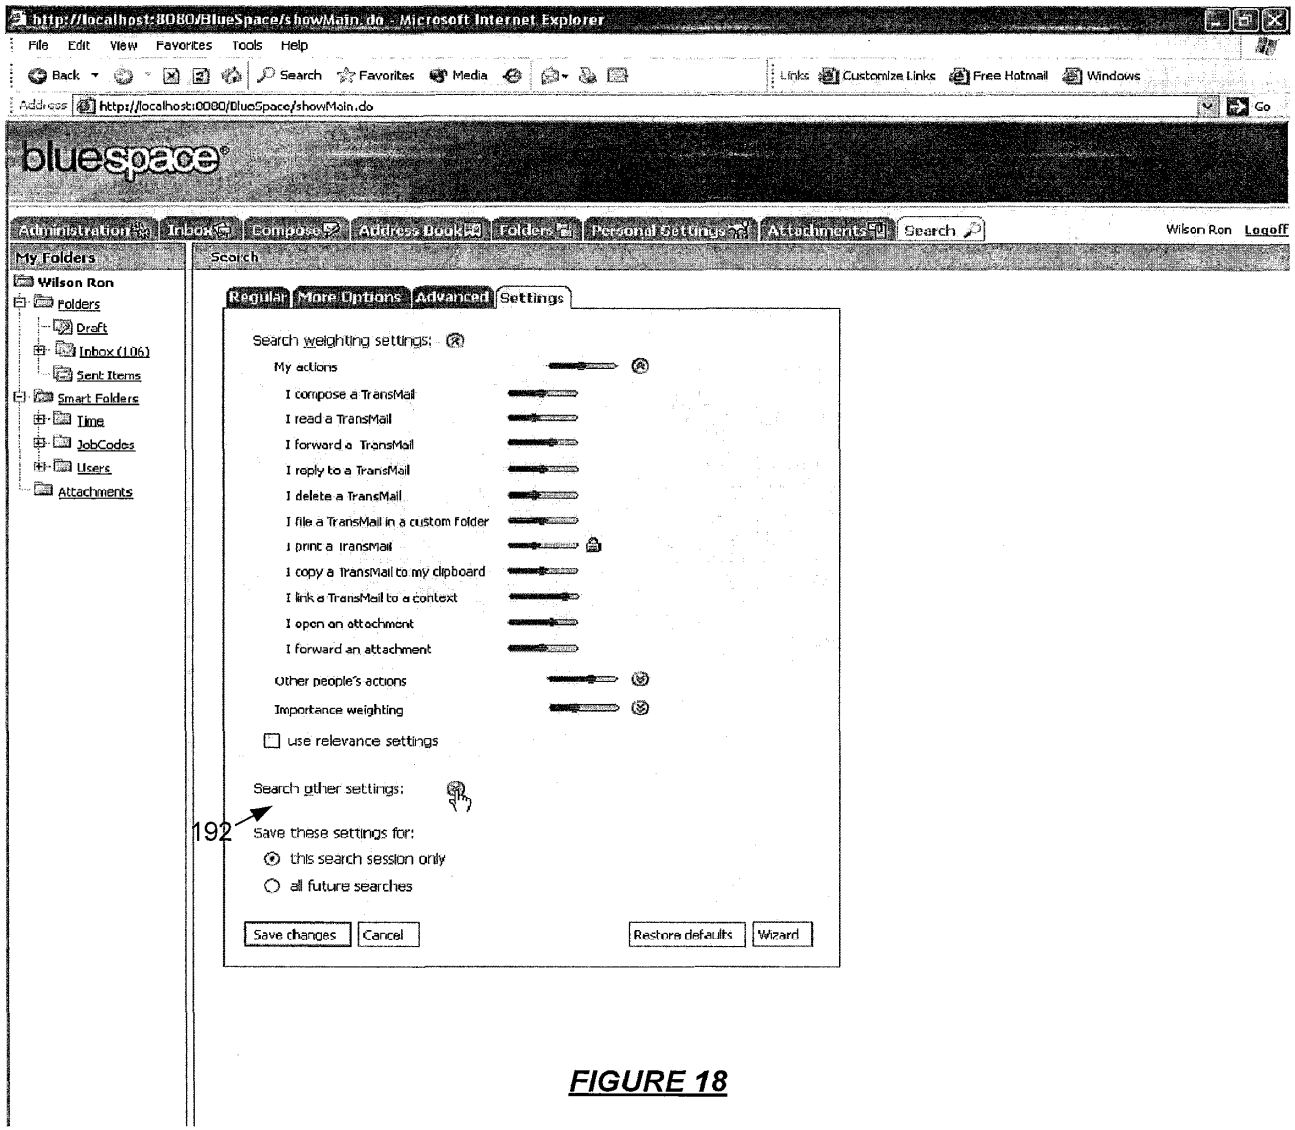


FIGURE 18

Title: DETERMINING RELEVANCE OF ELECTRONIC CONTENT
Inventor(s): Justin Marston et al.
Atty. Docket No.: 23682-12381
Sheet 18 of 22

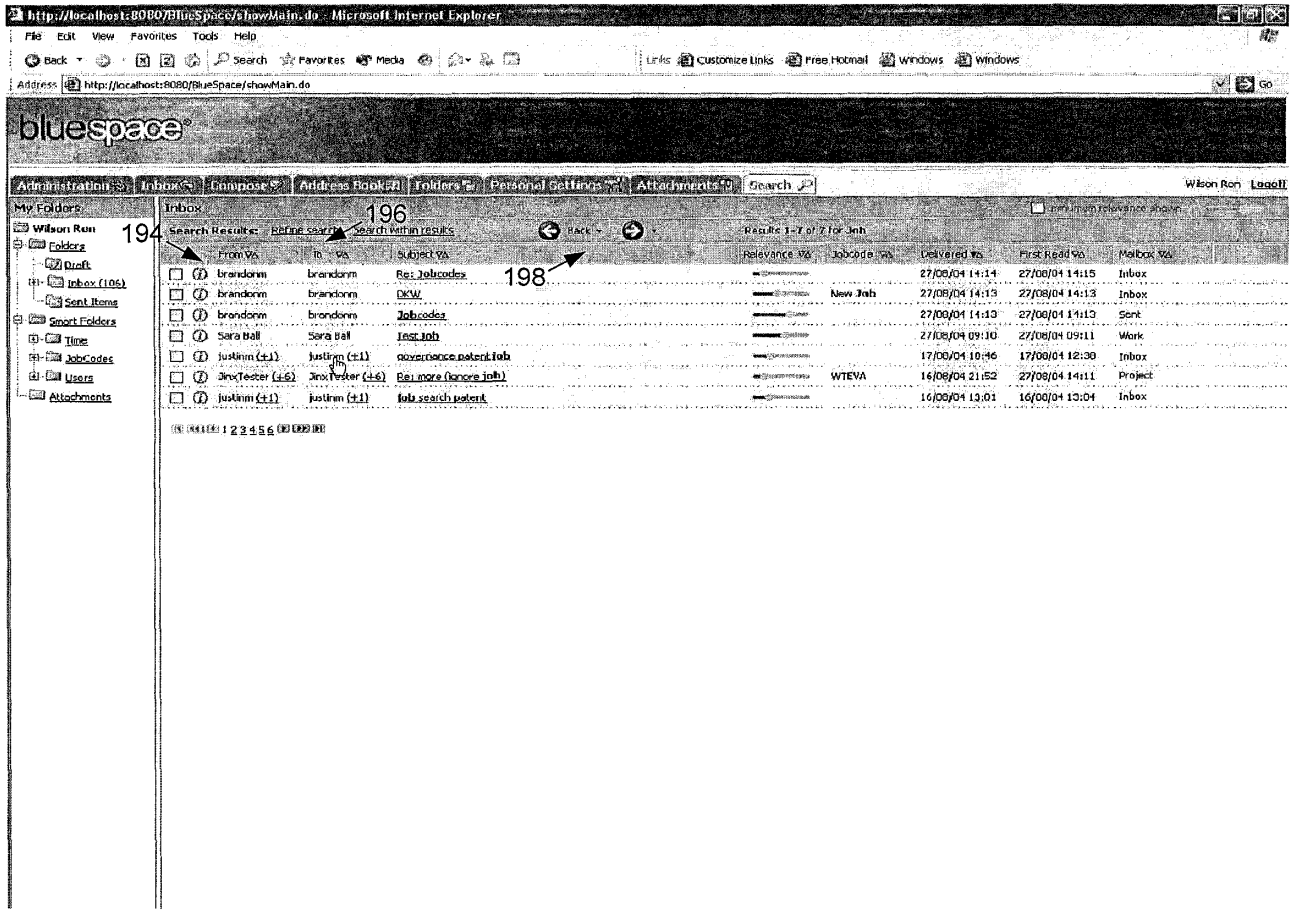

http://localhost:8080/BlueSpace/showMain.do - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Address http://localhost:8080/BlueSpace/showMain.do
bluespace
Administration Inbox Compose Address Book Folders Personal Settings Attachments Search
Wilson Ron Logoff
My Folders
Wilson Ron
Folders
Draft
Inbox (106)
Sent Items
Smart Folders
Time
JobCodes
Users
Attachments
Inbox
Search Results: Refine search Search within results
194
196
198
From
To
Subject
Relevance
Jobcode
Delivered
First Read
Mailbox
brandonm brandonm Re: Jobcodes 27/08/04 14:14 27/08/04 14:15 Inbox
brandonm brandonm DKW New Job 27/08/04 14:13 27/08/04 14:13 Inbox
brandonm brandonm Jobcodes 27/08/04 14:13 27/08/04 14:13 Sent
Sara Ball Sara Ball test job 27/08/04 09:10 27/08/04 09:11 Work
justinm (+1) justinm (+1) governance patent job 17/08/04 10:46 17/08/04 12:30 Inbox
JinxTester (+6) JinxTester (+6) Re: more (ignore job) WTEVA 16/08/04 21:52 27/08/04 14:11 Project
justinm (+1) justinm (+1) job search patent 16/08/04 13:01 16/08/04 13:04 Inbox
1 2 3 4 5 6

FIGURE 19

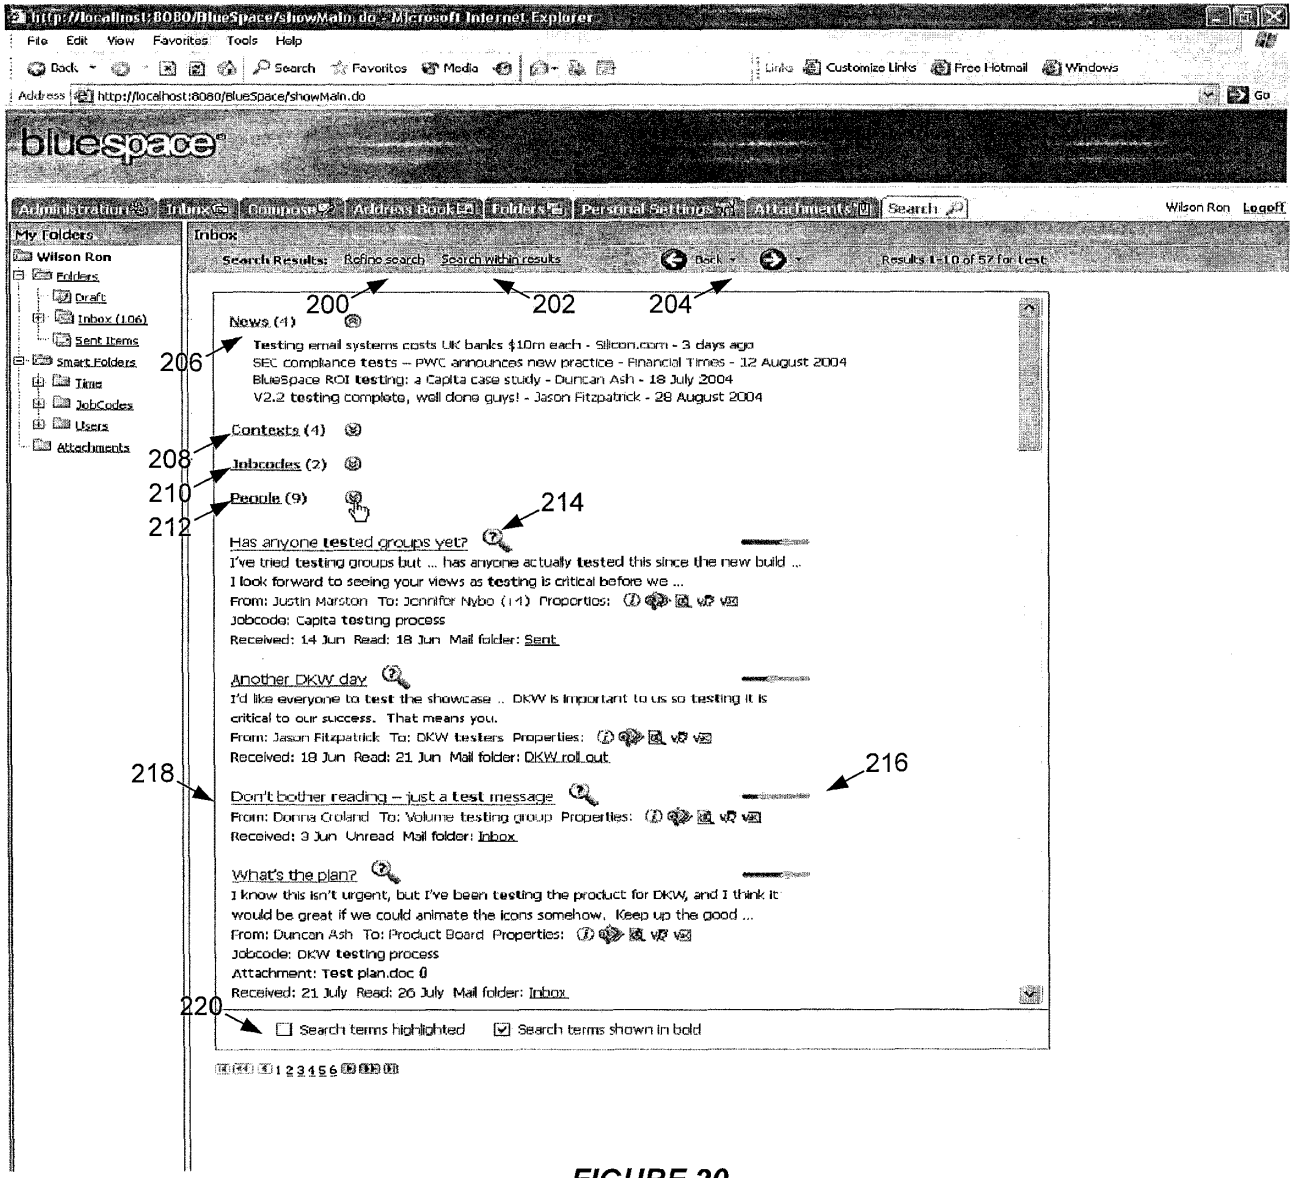


FIGURE 20

Title: DETERMINING RELEVANCE OF ELECTRONIC CONTENT
Inventor(s): Justin Marston et al.
Atty. Docket No.: 23682-12381
Sheet 20 of 22

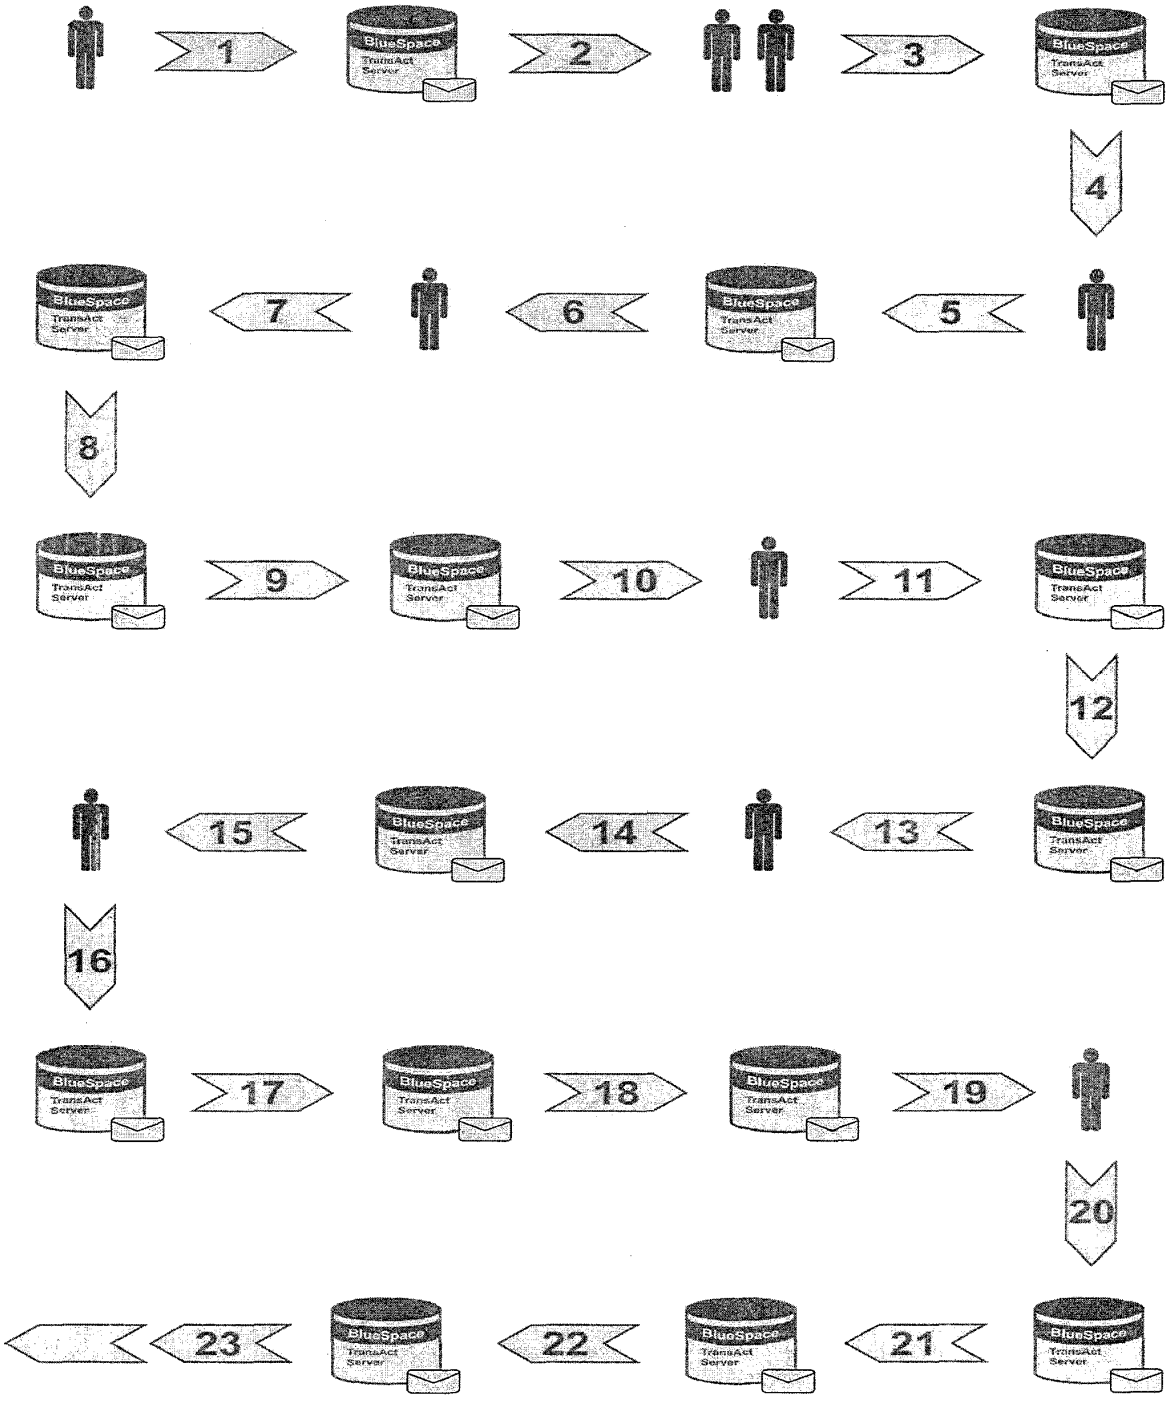
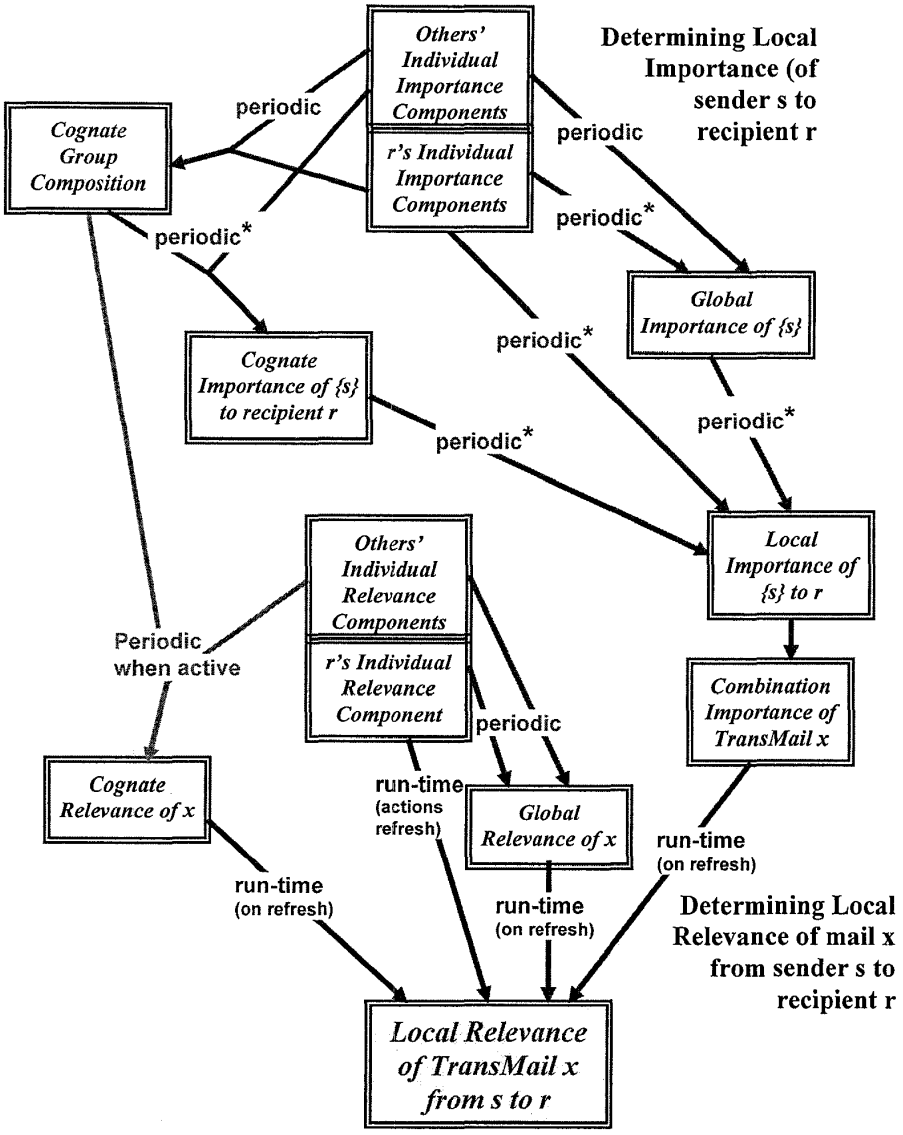

1
BlueSpace TransAct Server
2
3
BlueSpace TransAct Server
4
BlueSpace TransAct Server
7
6
BlueSpace TransAct Server
5
8
BlueSpace TransAct Server
9
BlueSpace TransAct Server
10
11
BlueSpace TransAct Server
12
15
BlueSpace TransAct Server
14
13
BlueSpace TransAct Server
16
BlueSpace TransAct Server
17
BlueSpace TransAct Server
18
BlueSpace TransAct Server
19
20
23
BlueSpace TransAct Server
22
BlueSpace TransAct Server
21
BlueSpace TransAct Server

FIGURE 21



Note: If this is the **first** contact between *s* and *r* then the arrows marked * will be absent and the arrival of mail will trigger the calculations of arrows marked * (which are usually done only on periodic review) because no local importance figure will previously exist.

FIGURE 22

PTO/SB/14 (08-05)
Approved for use through 07/31/2006. 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	23682-12381
	Application Number	
Title of Invention	DETERMINING RELEVANCE OF ELECTRONIC CONTENT	

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.)

Applicant Information:

Applicant 1

Applicant Authority ◉ Inventor ○ Legal Representative under 35 U.S.C. 117 ○ Party of Interest under 35 U.S.C. 118

Prefix	Given Name	Middle Name	Family Name	Suffix
Mr.	Justin		Marston	

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

City	Austin	State/Province	TX	Country of Residence [i]	US

Citizenship under 37 CFR 1.41(b) [i] UK

Mailing Address of Applicant:

Address 1	P.O. Box 30570		
Address 2			
City	Austin	State/Province	TX
Postal Code	78755	Country [i]	US

Applicant 2

Applicant Authority ◉ Inventor ○ Legal Representative under 35 U.S.C. 117 ○ Party of Interest under 35 U.S.C. 118

Prefix	Given Name	Middle Name	Family Name	Suffix
Dr.	Paul		Marston	

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

City	Lancashire	Country Of Residence [i]	UK

Citizenship under 37 CFR 1.41(b) [i] UK

Mailing Address of Applicant:

Address 1	143 Longmeanygate, Midge Hall		
Address 2	Preston		
City	Lancashire	State/Province	
Postal Code	PR26 6TD	Country [i]	UK

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).

PTO/SB/14 (08-05)
Approved for use through 07/31/2006. 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	23682-12381
	Application Number	
Title of Invention	DETERMINING RELEVANCE OF ELECTRONIC CONTENT	

☐ **An Address is being provided for the correspondence Information of this application.**			
Customer Number	00758		
Email Address	bhoffman@fenwick.com	Add Email	Remove Email

Application Information:

Title of the Invention	DETERMINING RELEVANCE OF ELECTRONIC CONTENT		
Attorney Docket Number	23682-12381	**Small Entity Status Claimed**	☑
Application Type	Nonprovisional		
Subject Matter	Utility		
Suggested Class (if any)		**Sub Class (if any)**	
Suggested Technology Center (if any)			
Total Number of Drawing Sheets (if any)	22	**Suggested Figure for Publication (if any)**	
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 been and will not be the subject of an application filed in another country, or under a multilateral 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).
Enter either 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	○ US Representative (37 CFR 11.9)
Customer Number	00758		

Domestic Priority Information:

This section allows for the applicant to claim benefit under 35 U.S.C. 119(e), 120, 121, or 365(c). Providing this 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(a)(2) or CFR 1.78(a)(4), and need not otherwise be made part of the specification.

Prior Application Status	Pending		Remove
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)
	non provisional of	60758828	2006-01-13

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

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	23682-12381
	Application Number	
Title of Invention	DETERMINING RELEVANCE OF ELECTRONIC CONTENT	

Foreign Priority Information:

This section allows for the applicant to claim benefit of foreign priority and to identify any prior foreign application for which priority is not claimed. 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(a).

Remove

Application Number	Country [i]	Parent Filing Date (YYYY-MM-DD)	Priority Claimed
			◉ Yes ○ No

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

Assignee Information:

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

Assignee 1

If the Assignee is an Organization check here. ☑

Organization Name	BlueSpace Software Corporation		
Mailing Address Information:			
Address 1	P.O. Box 30570		
Address 2			
City	Austin	**State/Province**	TX
Country [i]	US	Postal Code	78755
Phone Number		Fax Number	
Email Address			

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

Signature:

A signature of the applicant or representative is required in accordance with 37 CFR 1.33 and 10.18. Please see 37 CFR 1.4(d) for the form of the signature.

Signature	/Brian Hoffman/			Date (YYYY-MM-DD)	2007-01-16
First Name	Brian	Last Name	Hoffman	Registration Number	39713

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 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.

DETERMINING RELEVANCE OF ELECTRONIC CONTENT

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Application No. 60/758,828, filed January 13, 2006, the entirety of which is hereby incorporated by reference.

BACKGROUND

[0002] Electronic mail is perhaps the most common form of modern communication among individuals and entities. Virtually anyone with access to a computer network (e.g., the Internet) can send and receive email. To a large degree, the use of email has replaced reliance on more traditional forms of communication, such as telephone and postal mail. In fact, for many corporations and organizations, email is now a mission-critical tool, the failure of which can greatly hinder, or altogether stop, the operations of the corporation or organization. Further, although email is already a primary form of communication as described above, the volume of email continues to considerably grow at a fast rate and merits the close attention of email administrators.

[0003] Given the heavy use of electronic mail, many email users receive hundreds of different emails a day. Oftentimes, much of the email received by an individual is irrelevant (or more colloquially, is "junk") and serves as a nuisance that the individual must navigate through. In other situations, while a certain set of received emails may all be at least somewhat relevant to the recipient, the recipient must still take the time to go

through those emails to determine which ones are more important than others.

[0004] As email is received and sent over time, the recipient's "inbox" and "sent folder" becomes a vast store of information. For many corporations and organizations, email might be the most comprehensive and up-to-date repository of knowledge possessed by the corporation or organization. However, despite the potential richness of information in stored email, typical email applications are resigned to providing unintelligent mechanisms (e.g., keyword, sender, recipient, and/or date range searches) for mining through email.

SUMMARY

[0005] According to at least one aspect of one or more embodiments of the present invention, a system and computer program product include: a monitoring module configured to track actions by a plurality of users associated with an application for managing electronic content; a relevance analysis module coupled to the monitoring module and configured to determine a relevance to a user of a piece of electronic content sent from a sender to the user, where the relevance is determined dependent on a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content, and a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other recipient of the piece of electronic content.

[0006] According to another aspect of one or more embodiments of the present invention, a method comprises tracking actions by a plurality of users associated with an

application for managing electronic content, and determining a relevance to a user of a piece of electronic content sent from a sender to the user. The relevance is determined dependent on: a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content, and a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content.

[0007] The features and advantages described herein are not all inclusive, and, in particular, many additional features and advantages will be apparent to those skilled in the art in view of the following description. Moreover, it should be noted that the language used herein has been principally selected for readability and instructional purposes and may not have been selected to circumscribe the present invention.

BRIEF DESCRIPTION OF DRAWINGS

[0008] Figure 1A shows a traditional prior art email system.

[0009] Figure 1B shows a relational electronic mail system in accordance with one or more embodiments of the present invention.

[0010] Figure 2 shows an electronic mail system in accordance with an embodiment of the present invention.

[0011] Figure 3 shows an example of a series of electronic mail transactions in accordance with one or more embodiments of the present invention.

[0012] Figure 4 shows different techniques as to how an electronic mail object may

be sent to external recipients in accordance with one or more embodiments of the present invention.

[0013] Figure 5 shows a compliance model in accordance with one or more embodiments of the present invention.

[0014] Figure 6 shows an overview of information flow in an electronic mail system in accordance with one or more embodiments of the present invention.

[0015] Figure 7 shows an electronic mail system integration in accordance with one or more embodiments of the present invention.

[0016] Figure 8 shows a screenshot of an example of a mail view web page in accordance with one or more embodiments of the present invention.

[0017] Figure 9 shows a screenshot of an example of a mail view web page in accordance with one or more embodiments of the present invention.

[0018] Figure 10 shows a screenshot of an example of a mail view web page in accordance with one or more embodiments of the present invention.

[0019] Figure 11 shows a screenshot of an example of a web page with an importance chart in accordance with one or more embodiments of the present invention.

[0020] Figure 12 shows a screenshot of an example of a folder view web page in accordance with one or more embodiments of the present invention.

[0021] Figure 13 shows a screenshot of an example of an address book web page in accordance with one or more embodiments of the present invention.

[0022] Figure 14 shows a screenshot of an example of a preference setting web page in accordance with one or more embodiments of the present invention.

[0023] Figure 15 shows a screenshot of an example of a search interface web page in accordance with one or more embodiments of the present invention.

[0024] Figure 16 shows a screenshot of an example of a search interface web page in accordance with one or more embodiments of the present invention.

[0025] Figure 17 shows a screenshot of an example of an advanced search interface web page in accordance with one or more embodiments of the present invention.

[0026] Figure 18 shows a screenshot of an example of a search settings web page in accordance with one or more embodiments of the present invention.

[0027] Figure 19 shows a screenshot of an example of a folder view search results web page in accordance with one or more embodiments of the present invention.

[0028] Figure 20 shows a screenshot of an example of an extended view search results web page in accordance with one or more embodiments of the present invention.

[0029] Figure 21 shows an example flow process of a relevance mechanism flow in accordance with one or more embodiments of the present invention.

[0030] Figure 22 shows a relevance determination mapping and flow in accordance with one or more embodiments of the present invention.

[0031] Each of the figures referenced above depict an embodiment of the present invention for purposes of illustration only. Those skilled in the art will readily recognize

from the following description that one or more other embodiments of the structures, methods, and systems illustrated herein may be used without departing from the principles of the present invention.

DETAILED DESCRIPTION

[0032] In the following description of embodiments of the present invention, numerous specific details are set forth in order to provide a more thorough understanding of the present invention. However, it will be apparent to one skilled in the art that the present invention may be practiced without one or more of these specific details. In other instances, well-known features have not been described in detail to avoid unnecessarily complicating the description.

[0033] In general, embodiments of the present invention relate to a technique for determining the relative relevances of electronic mail messages. Particularly, one or more embodiments, in addition to possibly having traditional email search capabilities, may determine the relevance and/or importance of an instance of electronic mail based on, for example, user actions (e.g., what did the user do with that piece of electronic mail), other user actions (e.g., what did other users do with that piece of electronic mail), and/or administrative settings.

[0034] It is noted that although one or more embodiments are described herein with respect to traditional electronic mail, one or more embodiments may be applicable to other forms of electronic messages and content. For example, one or more embodiments may be used for determining relevance in document management applications. In

another example, one or more embodiments may be used in conjunction with instant message (IM) communications and/or short message service (SMS) messages.

[0035] Some forms of electronic mail, such as "TransMail" developed by BLUESPACE SOFTWARE CORP., use a relational architecture. In such a relational mail system, a piece of mail written by a user is managed as a single instance object during its life in the mail system. To illustrate the use of a relational mail system, Figures 1A and 1B show a comparison between traditional e-mail and email and electronic mail in a relational architecture in accordance with one or more embodiments. With traditional email, a first email **10** is replied to with a second email **12**, in which the content of the first email **10** is duplicated – there is no link between the first email **10** and the second email **12**. In a relational mail system according to one or more embodiments, a first electronic mail object **14** points to a first message **16**, and a second electronic mail object **18**, representing the reply to the first electronic mail **14**, points to the first message **16** and a second message **20**, where the second message **20** represents the reply content. Thus, in a relational mail system, electronic mail is transacted through the use of links, or "pointers," among electronic mail objects as they are created over time.

[0036] To further illustrate the use of links in a relational mail system, Figure 3 shows an example of a series of electronic mail transactions (transactions numbered in Figure 3) in accordance with one or more embodiments. In step 1, user **30** sends an electronic mail object to user **32** containing submessage A. In step 2, user **32** forwards the electronic mail object to user **34**, thereby creating a new electronic mail object that contains B|A (a current submessage B, with history portion submessage A). Then, in step

3, user **34** replies to user **32**, creating an electronic mail object that contains C|BA (a current submessage C, with history submessages B and A). In step 4, user **32** replies to user **34**, creating an electronic mail object that contains D|CBA (a current submessage D, with history submessages C, B, and A). Then, in step 5, user **34** forwards the electronic mail object containing D|CBA to users **36**, **38**, thereby creating new electronic mail objects containing E|DCBA (a current submessage E, with history submessages D, C, B, and A). In step 6, user **36**, seeking clarification from user **30**, sends an electronic mail object to user **30**, thereby creating an electronic mail object containing F|EDCBA (a current submessage F, with history submessages E, D, C, B, and A). User **30**, in step 7, replies to user **36**, thereby creating an electronic mail object that contains G|FEDCBA (a current submessage G, with history submessages F, E, D, C, B, and A). At some later time, in step 8, user **34** forwards the electronic mail object received from user **32** to user **40**, thereby creating an electronic mail object containing H|DCBA (a current submessage H, with history submessages D, C, B, and A).

[0037] As apparent from the above description with reference to Figure 3, a "current" submessage is a submessage at the top of the chain for a given electronic mail object. The current submessage will have been written by the user who created the electronic mail object. For example, in Figure 1B, submessage **16** is the current submessage of electronic mail object **14** and submessage **20** is the current submessage of electronic mail object **18**. In step 5 of Figure 3, for example, submessage E is the current submessage.

[0038] The previous submessage in the chain of a given electronic mail object are

referred to herein as "history" submessages. History submessages are content created before the current submessage and its respective electronic mail object. For example, in Figure 1B, submessage **16** is a history submessage of electronic mail object **18**. In step 5 of Figure 3, for example, submessages D, C, B, and A are the history submessages. Moreover, it is noted that in one or more embodiments, a submessage may only have one history path.

[0039] Further, for a given submessage, submessages that appear in later electronic mail objects that reference the original submessage are referred to herein as "future" submessages. For example, in Figure 1B, submessage **20** is a future submessage with respect to submessage **16**. In step 5 of Figure 3, for example, submessages F and G are future submessages, but submessage H is not because the electronic mail object created in step 8 does not include submessage E. Rather, submessage H is a future submessage in step 4 of Figure 3. Moreover, it is noted that in one or more embodiments, a submessage may have a plurality of future paths, as it may be included in several branching forward and reply structures.

[0040] An electronic mail object in accordance with one or more embodiments may be sent to a destination external to a relational mail system in which the electronic mail object is created. For example, now with reference to Figure 4, it shows different techniques as to how an electronic mail object may be sent to external recipients in accordance with one or more embodiments. In a first technique, a user in company **50** sends an electronic mail object **52** in accordance with one or more embodiments to recipients in company **54**. This occurs by a relational mail server **56** in company **50**

sending an electronic mail notification (with, for example, a uniform resource locator (URL) to the electronic mail object **52**) **58** to the recipients in company **54**. The recipients in company **54** may then select the URL to read the electronic mail object **52** using a web browser.

[0041] Still referring to Figure 4, in another technique, a user in company **50** sends an electronic mail object **52** in accordance with one or more embodiments to recipients in company **60**. In this case, the relational mail server **56** wraps the electronic mail object **52** in a simple mail transfer protocol (SMTP) envelope **62** and sends it like a regular email to an email server **64** in company **60**. The recipients in company **60** may then view the email as they normally would, without having to select a web link.

[0042] In another technique, a user in company **50** sends an electronic mail object **52** in accordance with one or more embodiments to recipients in company **66**. In this case, the relational mail server **56** may determine that a mail server **68** in company **66** is a relational-type mail server, and thus, may send an electronic mail object **70** using a protocol based on web services.

[0043] In one or more embodiments, electronic mail content may be managed according to rules specified in a compliance model. Such a model may be used to determine, for example, how long content is kept for, who can access content, where content should reside, and other allowable actions to be carried out by users. Figure 5 shows an example, in unified modeling language (UML) representation, of a compliance model in accordance with one or more embodiments. Using such a compliance model, search rules may be layered over a content system, determining, for example, who has the

right to search different pieces of content in the system. For example, in the case of an investment bank, an analyst might have the right to search all content related to analyst processes written by other analysts; however, the same analyst may not have the ability to search content that does not pertain to his/her role in the bank.

[0044] Turning now to a more detailed description of one or more embodiments, an electronic mail relevance application helps users make sense of vast quantities of electronic mail by inferring a user's interest in content and other users based on monitoring that user's actions with in the application. Such relevance may be thought of as having at least two components: relevance, which applies to content (e.g., submessages); and importance, which applies to entities (e.g., users and job codes).

[0045] A "relevance" component may be used to rate electronic mail content in an order reflecting a user's likely interest in the content. For example, a high relevance value may indicate that the content is likely to be of significant interest to the user, while a low relevance value may indicate that the content is likely to be of little interest to the user.

[0046] An "importance" component may used to rate an entity in order to reflect a user's likely interest in content involving that entity. For example, a high importance value may indicates that a given job code is likely to be of significant interest to the user in question, while a low importance value may indicates that a job code is likely to be of little interest to the user. As used herein, a "job code" may be thought of as metadata used to track the process to which a particular electronic mail object relates.

[0047] In one or more embodiments, relevance and importance values may be determined based on monitoring users' actions. Such monitoring may occur as part of an automatic, background process. As will be apparent from the description herein, users' actions may indicate valuable information about both content and entities in the system. For example, if a user receives an electronic mail object, saves the electronic mail in a custom folder and subsequently forwards it to three other users, then the electronic mail is likely to be much more relevant to those users than if the sending user had just deleted the electronic mail as soon as he or she had finished reading it. In another example of associating relevance and importance in an electronic mail system, it is likely that a user's direct manager will be important to the user because the user will typically read messages from their manager in a timely manner, will often reply to those messages, and will send their manager a significant number of electronic mail objects. As will be described in detail further below, by monitoring a collection of user actions and then performing an analysis of those actions, valuable information may be learned for subsequent use in an electronic mail application in accordance with one or more embodiments.

[0048] As described above, a relational electronic mail system in accordance with one or more embodiments separates electronic mail objects into its constituent message objects, likely written by different authors (see, for example, Figure 1B). The use of such constituent parts promotes a level of granularity that may used to support a relevance analysis engine that can be applied to the electronic mail system. In such a manner, behavioral information may be efficiently gleaned from the mail environment and used to

help improve content relevance as presented to one or more users.

[0049] Now referring to Figure 2, it shows an example representation of an electronic mail system **250** in accordance with one or more embodiments. As shown in Figure 2, the electronic mail system **250** has a plurality of modules, where a "module" may be any program, circuitry, logic, and/or functionality implemented in either software and/or hardware.

[0050] The electronic mail system **250** has a monitoring module **252**, which generally monitors and registers user actions and those of users associated or corresponding with a particular user. Such behavioral information may then serve as inputs to a relevance analysis engine **254**, which is formed of a relevance determination module **256** and an importance determination module **258**. The relevance determination module **256** assesses user action data provided by the monitoring module **252** and calculates relevance scores for objects accessible to the user. The importance determination module **258** assesses user action data provided by the monitoring module **252** and calculates importance scores for objects accessible to the user. It is noted that operations of the relevance analysis engine **254** may not be limited to applying to any one user, but may instead, for example, be used to determine relevance and importance scores for all users in a particular organization of which the electronic mail system **250** is part. Further, the electronic mail system **250** has a mail presentation module **260**, which is generally used to provide a user his/her electronic mail (along with features therewith) in accordance with the relevance and importance scores determined by the relevance analysis engine **252**. Moreover, the electronic mail system **250** has a search module **262**,

which is generally used to provide a user with search capabilities in accordance with the relevance and importance scores determined by the relevance analysis engine **252**.

[0051] Still referring to Figure 2, the electronic mail system **250**, or any combination of the modules shown therein, may run on any type of computer system (e.g., a client-side device, a server-side device, a single processor computer, a multiprocessor computer). Further, the electronic mail system **250**, or any combination of the modules shown therein, may be contained on a computer-readable medium (e.g., memory, a hard drive, a compact disc (CD), a digital video disc (DVD), a magnetic tape drive, an optical drive, a universal serial bus (USB) drive, a flash drive). As contained on such a computer-readable medium, the functionality of the electronic mail system **250**, or of any combination of the modules shown therein, may be partly or wholly carried out in response to one or more processors executing software instructions stored on the computer-readable medium.

[0052] Further, based on the description provided below, it will be apparent which of the modules shown in Figure 2 may be applicable to the various techniques, uses, and functions described herein. For example, the monitoring module **252** may be used for performing any of the techniques described herein relating to tracking the actions of users in a particular environment or associated with a particular system. The relevance analysis engine **254** may be used for determining or calculating any of the relevance related values described herein. More particularly, for example, the relevance determination module **256** may be used for determining or calculating any one or more of the various types of relevance scores described herein, and the importance determination

module **258** may be used for determining or calculating any one or more of the various types of importance scores described herein. Further, the mail presentation module **260** may be used for generating any one or more of the various types of web pages described herein (it is noted that in one or more embodiments, non-web user interfaces may also or instead be presented by the mail presentation module **260**). Moreover, for example, the search module **262** may be used for performing any of the search techniques described herein.

[0053] An electronic mail application in accordance with one or more embodiments may help users in performing one or more of various different types of actions. For example, in a browsing mode, a user may browse through a folder of content, such as looking through newly received electronic mail objects in their inbox. In this case, the electronic mail objects may be sorted or ranked according to the relevance and importance of each. In, for example, a searching mode, a user may actively look for a specific piece of information or for information of a given type based on given search criteria (e.g., keywords, dates, specific senders). In this case, for example, a piece of content may be deemed to be more relevant to a search if two of the words in the subject are keyword search terms as opposed to if there was only one matched word. Moreover, relevance scores may be displayed to the user in response to performing a particular electronic mail action.

[0054] In regard to what types of information may be extracted from an electronic mail system in accordance with one or more embodiments, mail data, or data in an electronic mail itself (e.g., keywords, creation dates, recipient information, job codes)

may be recovered. Further, mail action data, or information regarding users' actions on electronic mail content (e.g., the number of times an electronic mail has been read, the number of times an electronic mail has been forwarded), may be recovered. Further, users may themselves rate various aspects and entities in the system to improve an accuracy of the relevance analysis engine. For example, a user may provide feedback regarding the relevance of any given piece of content, and/or the user may provide feedback regarding the importance of any given entity. Moreover, a user may also change various weighting factors used to determine settings of the relevance analysis engine, where such settings may otherwise be self-adjusting.

[0055] A user may interact with an electronic mail system in accordance with one or more embodiments in one or more of various ways. For example, in the case of reading new electronic mail, when a user receives a new electronic mail object, the electronic mail system may indicate how useful to the user that content is likely to be. In another example, in the case of reading existing mail, the electronic mail system may indicate how important to the user each electronic mail object and submessage is likely to be. In yet another example, in the case of searching through electronic mail, the electronic mail system may help rank search results and indicate to the user how likely each piece of content is to be of interest.

[0056] An overview of information flow, including various elements and relationships described above, in an electronic mail system according to one or more embodiments is shown in Figure 6. In one aspect, what is shown in Figure 6 may be thought of as representing the information flow in a relevance analysis engine according

to one or more embodiments. As used in Figure 6, "single" aspects refer to individual units (e.g., content, users), and "collaborative"/"communal" aspects refer to collections of units (e.g., content, users).

[0057] Turning now more specifically to Figure 6, mail data **80** represents properties of each electronic mail, including the current submessage (single), as well as both history and future submessages (collaborative). Mail data **80** serves as an input to importance **82** in that it determines the entities listed in the mail environment and where they appear. Further, mail data **80** acts as an input to relevance **84** as the entities involved in the mail help to determine its initial and ongoing relevance values.

[0058] Users' actions **86** represent the ways users interact with the electronic mail system, e.g. reading mails, replying to mails, changing settings. Users' actions **86** encompass an individual user's own actions and other users' actions (the communal elements of the system). These actions may help to determine relevance, as relevance is calculated based on the way users interact with content. Users' actions **86** may also help determine importance, as the importance of entities reflects the significance with which users treat them.

[0059] Importance **82** of a given entity may be stated with respect to a local importance to a single entity (e.g., "how important is John to Doug?") or a set of entities (e.g., "how important is John to the Medix drug development team?"). Importance **82** may draw input from mail data **80**, which determines the "fields of engagement" for entities, and from users' actions **86**, which reflect how significant a given entity is to a given user or set of users. Importance **82** may also act as an input to relevance **84**, as it

helps to determine the relative weightings given to different actors in the electronic mail system. It is noted that importance values for different entities may be shown in a user's graphical interface.

[0060] Still referring to Figure 6, relevance **84** of a given piece of content may be stated with respect to a single entity (e.g., "how relevant is the electronic mail object to John?") or with respect to a set of entities (e.g., "how relevant is the electronic mail object to the Medix drug development team?"). Relevance **84** may have as inputs the mail data **80** (properties of the content in question), users' actions **86** (what users have done with the content), and importance **82** (the significance users attribute to the entities involved in the content). Various sub-aspects of relevance **84** may be identified: "local relevance" – the relevance value presented to the user as they browse through their mail folders; "incoming relevance" – estimated relevance to a recipient before an electronic mail object is opened; "post-reading relevance" – relevance assigned after an electronic mail object is opened and dealt with by a recipient; and "search relevance" – the relevance assigned by the system to items identified in a specific search by a user.

[0061] User interaction **88** is present so that when users are searching for something, they have some idea of what they are looking for and can guide the system to that content by stating criteria and changing their search settings.

[0062] Browsing **90** is an activity a user spends most of their time doing in an electronic mail system – navigating through mail folders and reading content. Local relevance acts as an input to browsing **90**, helping to guide the user to the content they might wish to do deal with first. In one or more embodiments, there may be only one

local relevance value for each piece of content with respect to a given user. Further, it is noted that local relevance may initially be the incoming relevance value (described above), and once the user in question has read the electronic mail object, the local relevance value may be the post-reading relevance value (described above).

[0063] Continuing with reference to Figure 6, searching **92** is an activity that a user may use to search electronic mail amongst potentially thousands of electronic mails the user has sent and received. It is noted that the search relevance value (described above) for content relating to a given search may vary depending on the search criteria settings and/or the user's settings.

[0064] In addition to the components of the information flow shown in Figure 6, an electronic mail system in accordance with one or more embodiments may have access to relatively independent data points by, for example, indirect "watching" of users' actions and mail data and direct "questioning" (e.g., asking users whether the results are accurate). However, such data points may vary according to users' opinions of the significance of entities and content presented to them by the electronic mail system.

[0065] As described above, mail data represents the properties of a particular electronic mail object. Such properties may include, but are not limited to, the following: keywords in subject; keywords in message body; keywords in other fields; keywords in attachments (both names and content); a signature assigned to the message; a date the content was created; the recipients of the content; whether the electronic mail object was sent to:/cc:/bcc:; the sender of the content; the body size; the attachment size; the attachment type; and flags for follow up (e.g., the sender indicates the electronic mail

requires action by the recipient). Thus, for example, in one or more embodiments, a user may build a search query looking for different submessages that make up a particular electronic mail object and may specify different creation dates, senders, recipients, and/or other details for each submessage they are looking for that make up the electronic mail object.

[0066] Other criteria may also be used to leverage the relational architecture of an electronic mail system according to one or more embodiments: job codes – a mechanism for linking an entire submessage to a specific business process; context – a mechanism for linking text in a submessage to a specific process, either on a user or enterprise-wide basis; groups – sets of users across the system; domains – a naming directory across the system; priority – an urgency tag linked to notification processes; sensitivity – a confidentiality tag linked to readership and search rights; validity date – the period of time for which the content is viewable by the user; retention date – the period of time for which the content is retained by the corporation; allowed readers – the users who are allowed to read the mail; allowed geographies – the geographies to which the mail content is allowed to be sent; rule contravention – content sent that had weak conflicts (e.g., a history submessage could not be included for one of twenty recipients) or strong conflicts (e.g., a user tries to send an electronic mail object to another user with whom they are not allowed to communicate) with compliance rules, and was altered or managed appropriately by the application; other compliance interactions – various other interactions with the governance engine; other meta-style information stored relationally with respect to the mail; and other business processes that are triggered by the mail event.

[0067] Examples of scenarios in which mail data can affect relevance are now provided. In other words, the following description lists examples of ways in which mail data parameters, such as those described above, might affect relevance to a given user, either while browsing or as part of a search. In one case, if a user received an electronic mail object as one of fifty other recipients, the electronic mail object is less likely to be of relevance to the user than if the user received the electronic mail object as the sole recipient. In another situation, if an electronic mail object is flagged for follow up by the sender, it is more likely to be of significant relevance to the user. In another case, if the keyword a user is looking for is in the body, then the electronic mail object is likely to be relevant to the user. In yet another example, if a user is looking for a particular file, then electronic mails with attachments of the requested file type are relevant to the search. In another scenario, if a user is looking for a piece of content they received some time ago from a particular, important department, then criteria such as date period, priority, and department domain may be relevant to the search. In another case, if a user is looking for all content relating to a particular business transaction, and that content has been tagged with one or more job codes, then content with the appropriate job codes may be relevant to the search. In still another example, if a user is looking for content that may have caused compliance violations, then all content that caused weak or strong interactions with a given set of compliance rules may be relevant to the search. In another situation, for example, if a user is looking for all electronic mail objects that will trigger a financial purchase over the next 30 days, then the user can search for electronic mails that will cause a purchase over the defined date range.

[0068] Other scenarios may be used to leverage a relational structure of an electronic mail system in accordance with one or more embodiments to determine the "communication context" of electronic mail. Such scenarios may rely on history submessages (older submessages in the chain of the electronic mail object) and future submessages (submessages written later as part of other electronic mail objects that contain a present submessage). For example, if an electronic mail object is received by a user as a reply to an electronic mail object they themselves sent, then a submessage in the history will map to the "current submessage" – the one at the top of the chain – in an electronic mail object the user's "sent folder." In this case, the incoming electronic mail object is likely to be of significant relevance. In another example, if a particular submessage appears several times in a user's "inbox," then a new electronic mail object containing this submessage is likely to be of relevance to the user. In yet another example, if a new electronic mail object received by a user contains one or more history submessages written by the user, then it is likely that the new electronic mail object is of relevance to the user. In still another example, if historical submessages in a new electronic mail object were found to have high relevance ratings, then it is likely that the new submessage will have significant relevance to the user. In a further example, if a submessage, A does not contain a given keyword, but many submessages that were sent in replying to and forwarding submessage A do contain the keyword, then submessage A is likely to still be relevant in a search for the given keyword.

[0069] Now turning to users' actions, as described above, users' actions represent the ways users interact with an electronic mail system in accordance with one or more

embodiments. The electronic mail system monitors user actions inside the electronic mail system. It is noted that such actions may be deemed important as they reflect the users' opinions of the content and the value they place on it. Examples of users' actions include, but are not limited to, the following: a user composes an electronic mail object; a user reads an electronic mail object; a user forwards an electronic mail object; a user replies to an electronic mail object; a user deletes an electronic mail object; a user files an electronic mail object in a custom folder; a user prints an electronic mail object; a user saves a local copy of an electronic mail object; a user copies an electronic mail object to their clipboard; a user links an electronic mail object (or part of a submessage) to a context; a user finds an electronic mail object as part of a search; a user includes an electronic mail object in a visualization; a user opens or saves an attachment; and a user forwards an attachment.

[0070] The relevance of users' actions relies on the user's actions to help determine content relevance for content in the user's mail folders, both sent and received. The following examples illustrate how a user's own actions may be used as an input to a relevance analysis engine in accordance with one or more embodiments. In one scenario, if a user receives an electronic mail object from another user and immediately forwards it to three other users with comments, then the electronic mail object is likely to have had relevant content. In another case, if a user receives an electronic mail object and replies to it, then the content of the original electronic mail object is likely to be of significant relevance. In another situation, if a user receives an electronic mail object from another user and immediately deletes it, it is unlikely that they perceived the content to be of

significant relevance. In a further example, if a user reads an electronic mail object and then moves it to a custom folder out of the "inbox," then it is likely that the content had relevance to the user. In still another case, if a user has read a given electronic mail object many times, then it is likely to be of significant relevance to the user. In another scenario, the longer the user has an electronic mail object displayed on a screen, the more likely it is to be relevant to the user. In still another situation, if a user copies text from a given electronic mail object to their clipboard, then the electronic mail object will likely be more relevant than an electronic mail object from which the user did not copy text. In another example, if a user links a piece of content to a context, it is more likely to be of relevance to them. In yet another situation, if a user always deletes content from a specific user quickly, then it is unlikely that any content from this user is of significant relevance to them. Further, in another case, if a user always reads content from another specific user before reading content from other users in any given new set of new mail received in their inbox, then it is likely that content from this specific user is of significant relevance. In still another scenario, if a user typically only reads short electronic mail objects rather than long ones, then the chances are that short ones are more relevant to that user than longer ones. In a different case, if a user never reads messages sent to a particular group or domain, then content sent to that group or domain is likely to be of less relevance to the user. In yet another situation, if a user never reads electronic mail objects tagged with a particular job code, then this content is likely to be of little relevance to the user. In yet another example, if a user rarely reads electronic mail objects designated as lowest priority (or reads them last when browsing their

"inbox"), then that lowest priority content is likely to be of little relevance to the user. In another case, if a user sends an electronic mail object to a many users, it may potentially be more relevant than an electronic mail object sent to a single user. In still another situation, if a user sends an electronic mail object and then later forwards it to other users, it is likely to be of more relevance than an electronic mail object that was sent once but never resent. In a further example, if a user reads an electronic mail object from their "sent folder," then the content is likely to be of significant relevance. Continuing with another scenario, if a user saves an attachment to their local hard disk or network system, both the attachment and the accompanying electronic mail object are likely to be of significant relevance. Moreover, in another case, if a user prints an attachment or electronic mail object, it is likely to be of high relevance to the user. In yet another case, if a user sends an electronic mail object to their supervisor, the electronic mail is likely to be more relevant than an electronic mail object sent to a junior secretary. In still another situation, if a user sends an electronic mail object with extensive formatting, it is more likely to be relevant to the user than if they only wrote a few lines of plain text. In another example, if a user has not accessed the folder in which the content is stored for over a year, then it is likely that the content is not particularly relevant.

[0071] As described above, the actions of other users may also be used to determine the relevance of content in a given user's folders, both sent and received. Examples of such actions are now described. In one case, if a user sends an electronic mail object to someone, and that someone reads the electronic mail object, then that electronic mail object is more likely to be relevant (to both the receiver and sender) than if the recipient

did not read it. In another situation, if a user sends an electronic mail object to a group of users, and most of the recipients reply to the electronic mail object, forward the electronic mail object, and/or copy the electronic mail object to custom folders, the electronic mail object is more likely to be generally or globally relevant than if the recipients immediately deleted the electronic mail object. In another scenario, if most other recipients of an electronic mail object delete it as soon as it arrives in their "inbox," then the recipient in question will probably not find the content particularly relevant. In still another example, if most other recipients copy a given electronic mail object to a custom folder (or reply or forward it) once they have read it, then the recipient in question will probably find the content to be of significant relevance. In yet another case, if a given user is always copying another user on their mail content (e.g., an employee copying their supervisor), then the content is less likely to be of relevance to that other individual. In a further example, if a user sends a submessage to another user, but the recipient does not read the submessage, then this may have an affect on the relevance of the content. In another situation, if the vast majority of the users in a given department read a piece of content sent globally, then the content is likely to be relevant to another user in that department. In still another scenario, if the majority of other users tend to read content linked to a particular job code, then content attached to that job code is more likely to be relevant to the user in question. Further, for example, if the majority of other users tend to read content linked to a particular context, then content attached to that context is more likely to be relevant to the user in question. In another situation, if most users read content from a particular author (e.g., a supervisor), then content from this author is

typically of greater relevance than content from a different author that is not widely read. In yet another case, if most users reply to content from a particular author (e.g., a supervisor), then content from this author is typically of greater relevance than content from a different author not widely replied to.

[0072] In one or more embodiments, in addition to considering content and users' action with content, an electronic mail system may analyze the importance of particular authors and entities when determining relevance. Such entities include, but are not limited to, the following: users; groups; domains; folders; job codes; contexts (categories to which pieces of content can be assigned); signatures; other business process information; and back-end/automated actors (e.g., external computer systems). As an example of a use of relying on the importance of authors and entities when determining relevance, in an organization, every member or employee may have an importance rating for other members/employees. In such an environment, it is likely that a member/employee would find other individuals who report to them and those to whom they themselves report to be of greater importance than individuals working in a different department.

[0073] Importance, as determined by a relevance analysis engine in accordance with one or more embodiments, may be a relative value, respective to each user (or other entity) in the system. For example, users, groups, domains, and other entities may all have different relative importance values describing a particular job code. As such, it is possible to calculate an overall importance value for any given entity in the system, which describes that entity's importance to the entire organization, where "organization"

may be defined as all the users of all the entities. This calculated importance value may be referred to as the "global importance value" for that entity.

[0074] Further, in one or more embodiments, the importance for a given contact may be represented as a pair of different values. A first value may reflect the user's interest in content they received from the contact, and a second value may reflect the user's interest in content they sent to the contact. While in some cases, these two values may be similar, in other cases, they may not be similar. For example, if a particular supervisor is extremely talkative or verbose, he/she may have low importance when sending content to other users, but those other users may view the supervisor as being very important when they themselves are initiating the communication.

[0075] Moreover, it is noted that different users' actions have different effects on relevance, due to varying importance factors. Thus, for example, if a user, such as a supervisor, deletes a particular electronic mail object, then this will have a greater effect than if a junior secretary deletes the electronic mail object. Also, importance may have multiple feedback effects. For example, if user B receives an electronic mail object from user A, and then user B forwards the electronic mail object to user C who is of high importance, then this affects the importance of all three users.

[0076] Now turning to another example illustrating the use of importance values, in an organization with, for example, 100,000 employees, it is likely that any given user will communicate with new contacts each week. As described above, importance is a relative factor for each pair of users' (or entities') interactions and is stored with respect to each user. However, when a new contact communicates with a user for the first time, the

system may still make a "best guess" about that new contact's importance based on communal importance (described above with reference to Figure 6). This importance value may then change as the user interacts with this new contact. The first "best guess" value may be important nonetheless because in a system with a large population, the usability may be hampered if the importance value of every new contact started at a particular default value.

[0077] In one or more embodiments, an organization may decide to tie importance to roles instead of or in addition to specific users. This approach may be particularly likely in highly-regimented environments, such as the military, where roles are of greater significance than the actual human personality involved.

[0078] Now turning to some examples illustrating the use of roles for determining importance, if a user moves position within a company, their role will likely change, and, in turn, they will have greater relevance to some users and less relevance to others. Thus, an organization may choose to integrate their personnel system with a mail relevance analysis engine in accordance with one or more embodiments. When a person moves positions, the learning sensitivity of the engine with respect to that user may be increased, so that the relevance and importance values may reach their new equilibrium more quickly. Further, the increased sensitivity may be decreased according to an appropriate statistical function over time, until the sensitivity is at regular levels.

[0079] In another situation, if a user moves into a role that was previously filled, the relevance identity for the previous occupant of the role may also act as an input to the individual taking on the new role. This may be referred to as "role importance," which is

a factor of a role holder's local and global importance values. This further reflects the blend of role and the user's own identity in determining importance values and relevance of associated electronic mail content. Thus, for example, a person moving into a new role may still attach the same importance to content received from their spouse, but would likely move to a different importance value for their new supervisor.

[0080] Further, it is noted that may exist scenarios in which a single person has multiple roles, in which case, they may be able to send and receive electronic mail using different user accounts. Another possible scenario is when a single role is shared by a number of different individuals, all of whom write and receive content using a single account. In this case, an organization may choose to treat the role as a single entity in isolation (so there is no correlation with the users using the account). Alternatively, in one or more embodiments, an organization may specify a blend of the person's identity and the role being used, in which case performing an action while signed into one account affects the values in other accounts.

[0081] As described above, in a large system, a user may often be interacting with a new entity for the first time. One approach to setting an initial importance value in these cases is to recognize correlated or cognate entities. This may be done, for example, through the use of tree structures. For example, one human resources employee is generally likely to have similar importance as another one, especially if they are of similar position in a tree structure (e.g., they both work in employee appraisal form processing). Further, it is noted that in one or more embodiments, branched structures, such as knowledge taxonomies for contexts and directories for users and job code

hierarchies, may be used as a correlative basis for importance on an ongoing basis. Thus, for example, an importance change for one user may have a cascading "ripple effect" for similar users.

[0082] In regard to recognizing "cognate" entities are referred to above, an electronic mail system in accordance with one or more embodiments may identify a "cognate group" for a user, based on historical actions. More particularly, for each user, over time, the system may identify a cognate group of the *n* other users regarded as most similar. Such a group may be defined by similar past user actions and assessments, and/or a direct input by a user. Further, it is noted that it may not be inherent that the population of recipients falls into separate and mutually exclusive cognate groups. Rather, each recipient may have their own cognate group of those the system regards as most similar to that recipient. Thus, for example, in a certain case, user A may be in user B's cognate group because user A is part of a group of *n* users most similar to B, but at the same time user B may not be in A's cognate group because there are *n* other users more similar to user A than is user B. Accordingly, it may be assumed in certain instances that the actions of this particular group on an electronic mail object will give more indication of its relevance to the user than would a global reaction to the electronic mail.

[0083] In addition to the actions that users perform in an electronic mail system in accordance with one or more embodiments, the electronic mail system may take direct user input to help determine relevance as described above with reference to Figure 6. In this regard, users may be thought of as having an "importance slider" for other users in

the electronic mail system. For example, a user might make their boss "very important" and mark their junior, sales clerk "rarely important." Such weightings may be considered by the mail relevance analysis engine described above.

[0084] As to the "importance slider" referred to immediately above, a piece of mail content, or a submessage, may be presented to the user along with a slider that indicates its current local relevance score. The user may move this slider and apply the change in relevance to that specific piece of content, or the user may add that change back in to how relevance is calculated across the whole user's mailbox.

[0085] Further, to assist in providing tangible data points for the relevance analysis engine, one or more embodiments may also directly query the user on the relevance of content sampled from the system. It is noted that such a sampling mechanism may also be used as a way to corroborate data. For example, in one case, researchers may be asked to rate the usefulness of a piece of content to the project they are currently working on. In another case, rather than asking the user to rate the relevance of a piece of content, the electronic mail system may query the user to rate their estimation of whether the content is true. Moreover, sampling may be conducted in a random or structured manner.

[0086] In addition to an end user being able to affect relevance directly for their own "inbox," certain "super-user" roles may have the ability to affect relevance across the entire system, or at least a considerable portion thereof. In other words, for example, a "relevance administrator" may provide a boost to communications from a particular department. At least one reason for such manipulation is that relevance may not be a purely user-centric phenomenon – a separate question may relate to relevance for the

entire organization.

[0087] As described above, an electronic mail system in accordance with one or more embodiments may have a relational mail architecture, in which each piece of content written by a user is managed as a single instance object during its life in the mail system. The granularity of such an architecture may allow the electronic mail system to record all actions by users in the system, ensuring that the system has audited the entire lifecycle of the content. Further, relevance determination techniques in accordance with one or more embodiments may be applied to electronic mail at different levels. For example, relevance analysis may be applied to an electronic mail unit, where such a unit is considered as a whole regardless of what other electronic mails the unit contains. Such electronic mail unit analysis may be considered as a basic approach to determining relevance in electronic mail content. It allows relevance to be based on simple criteria such as, for example, the number of recipients, whether the user in question received the electronic mail object as a to:, cc: or bcc: recipient, the importance flag of the electronic mail object, and the time it was sent.

[0088] At another level, relevance analysis may be applied to relationally restructured electronic mail, where electronic mail is broken into its component messages, and the relational structure of the electronic mail system is mapped to the electronic mail store. In one or more embodiments, an electronic mail system may have the ability to analyze regular electronic mail content and convert it into a relational system. When the electronic mail system structures information once it has already been sent, the criteria for search may be, for example, keywords and other standard inputs,

user mail flow actions (e.g., compose, reply, and forward), and communal user mail flow actions (e.g., compose, reply, and forward).

[0089] When an electronic mail system in accordance with one or more embodiments is integrated with an electronic mail environment at the application level (as opposed to being a purely "downstream" process), then one or more of the following criteria may be used: user mail management actions (e.g., read, forward, reply, delete and file to custom folder); user attachment management actions (e.g., view, forward, save to disk); and mail content path (e.g., tying together multiple instances of the same electronic mail). By using a relational architecture for regular electronic mail, an electronic mail system in accordance with one or more embodiments allows for the searching of regular electronic mail, with a similar relevance system to that of a search engine used for an electronic mail system operating on electronic mail objects as described above with reference to Figure 1B.

[0090] Integration with a conventional email system may be achieved in various ways. In client-side integration, for example, a plug-in to the email client may be used to monitor end users' actions on the mail content. In server-side integration, a relational electronic mail server (e.g., one developed by BLUESPACE SOFTWARE) may be connected to the legacy email server and tag content with "flags" so that it can track what happens to the content. Examples of these two mechanisms are shown in Figure 7.

[0091] Turning to the processes shown in Figure 7, in step 1, new email received by a legacy email server (e.g., MICROSOFT EXCHANGE) **100** is retrieved by a relational knowledge server (e.g., one developed by BLUESPACE SOFTWARE) **102**. In step 2,

the knowledge server **102** adds meta-tags to each email in an effort to track that content through the lifecycle. In step 3, the tagged emails (represented as having flags) are then replaced on the email server **100**. In step 4, an email client (e.g., Microsoft Outlook) **104** downloads the tagged emails. In step 5, a plug-in to the email client **104** immediately retrieves the current relevance values for the content. Then, in step 6, the user carries out actions on the content in the email client **104**, and the plug-in records these actions and sends them back to the relational knowledge server **102**. In step 7, one of the actions carried out by the end user is to send emails in replies and forwards – these include the meta-tags positioned by the relational knowledge server **102**.

[0092] Still referring to Figure 7, it is noted that steps 5 and 6 may form a continuous feedback loop, where each time the user carries out an action, a corresponding relevance value changes. Further, plug-in to the email client **104** may also interact with search results to include and reflect the relevance values in the search listings.

[0093] Now turning to the use of one or more embodiments for research or forensics purposes, as the volume of email transacted has increased, so has the challenge for productively searching through vast quantities of email. A relevance analysis engine in accordance with one or more embodiments may be applied to regular electronic mail (as described above with reference to Figure 7), with one or more important benefits for forensics processes. In one aspect, the relevance analysis engine may combine a granular analysis process with a particular relevance algorithm (detailed examples of which are further described below). Content that meets certain criteria including relevance

parameters may be flagged for human inspection, and the most important individuals in a community may be inferred automatically. In another aspect, the relevance analysis engine may allow for the determination of similar or "cognate" entities and users by automatic analysis of users' actions (as opposed to attempting taxonomic structuring based on automated content analysis). This is based on "fingerprinting" different users (and entities) for similar direct and passive behavior. In still another aspect, high-level representations of user communities and sub-communities may be created based on a granular understanding of that community's interactions. This may reflect the importance of different individuals, based on the same flow analysis mechanisms for relevance described herein. In yet another aspect, increased granular capabilities in an electronic mail system in accordance with one or more embodiments allows for advanced visualization of communication communities. In another aspect, the approaches described immediately above may be combined for improved content sampling, through either or both of fully automated sampling and interactive human sampling. Further, it is noted that one or more of these approaches may be essentially language neutral.

[0094] Turning now to a description of user interface aspects relating to an electronic mail system in accordance with one or more embodiments, there may be several pages of an electronic mail application that include one or more of the various relevance and importance aspects described above. Such pages may include, for example, a mail view (when a user is viewing a single electronic mail object), a folder view (a folder listing showing many electronic mail objects), an address book (a list of contacts, groups, and domains), and preferences (a list of information). Each of these is

described in detail below.

[0095] In a mail view, a user is looking at a single electronic mail object and is able to see its entire content. An extension of this view is a reply view, as shown in 8. There are various relevance sliders in Figure 8, one for each submessage in the mail view and one for the entire electronic mail object, taking into account the current submessage and all history submessages. Relevance sliders **118** for history submessages may only be shown when the submessage is expanded, as in the second history submessage in Figure 8.

[0096] When a user hovers their mouse over a relevance slider, a pop-up window may appear, indicating why a particular relevance value is as it is, thereby potentially helping the user understand how the relevance value was determined. A similar pop-up window may appear for importance sliders, as further described below.

[0097] The user may provide feedback to the relevance analysis engine by changing the level of any given slider. When a user clicks and drags the level, the application may display a pop-up window above the slider indicating the value to which the user has moved the slider. The respective values may, for example, be: "irrelevant" – 0% to 10%; "slightly relevant" – 11% to 30%; "relevant" – 31% to 70%; "highly relevant" – 71% to 90%; and "critical" – 91% to 100%.

[0098] Further, the user interface may indicate to the user which submessages they have already seen before (e.g., in the case of a reply history) and which are new (e.g., in the case of a forward history). This is shown in Figure 8 via the state of the illustrated

"eyes" **110**, **112**, **114** in each submessage. If the eye is open (e.g., eyes **110**, **114**), the user has seen the submessage, and if the eye is closed (e.g., eye **112**), the user has not seen the submessage. Thus, as shown in Figure 8, if the active user (Wilson Ron) opened the submessage from Smith John, the respective eye icon for that submessage would become open.

[0099] Figure 8 shows further icons. Icon **120** allows the user to see a custom search folder, showing every electronic mail object that contains the given submessage. Another searching icon **122** allows the user to search for similar content to the given submessage, as determined by a search algorithm. Two more icons **124**, **126** allow the user to view the lifecycle of any given submessage, spawning pop-up windows (not shown in Figure 8) that visualize its path through the system and all users that have carried out actions on it.

[00100] In addition to relevance feedback, the mail view may present information on importance. Job code importance and other custom created variables may be given importance sliders **116**, as shown in Figure 8. A user may alter an importance slider, and either lock it at the new value or allow it to recalculate over time.

[00101] Further, in one or more embodiments, in addition to or instead of being able to provide feedback to a relevance analysis engine by moving relevance sliders, a user may also be requested to provide direct feedback on a sampled basis. In this case, for example, the user may be asked to directly rate the relevance of a given electronic mail objects, thus presenting the relevance analysis engine with a direct relevance data point, as opposed to correlating relevance with the user's actions. Such direct user questioning

is illustrated in box **129** shown in Figure 9.

[00102] As depicted in Figure 9, a user may not be forced to set the relevance value before being able to leave the mail view; however, in one or more embodiments, the interface may indicate that it is important for the user to do so. In one or more embodiments, it may be important that a high response rate be achieved for such relevance rating questions, as the effectiveness of the relevance analysis engine may correlate to the lack of bias in the survey sampling.

[00103] The number of times a user is asked to rank content may be determined based on their usage frequency (e.g., every 100 mail views) or absolute time (e.g., 5 times per week). A system administrator may determine this value, with the option to allow users to customize it through their preferences. Further, in one or more embodiments, a system administrator may also make direct user interaction mandatory for scenarios in which user feedback is required.

[00104] In one or more embodiments, direct feedback values submitted when users assign relevance to content through the survey function may be particularly helpful in the early stages of a given system. Such values help provide reliable relevance values to stabilize the weightings given to each type of user action. A system administrator may thus be able select a time-based factor to decrease the frequency of surveying users. Alternatively, a system administrator may base decreases in sampling frequency on statistical results derived from, for example, the stability of the global weighting factors.

[00105] Still referring to Figure 9, in regard to users' actions and users' importance

(as well as that of groups and domains), rather than clutter the interface with all this information by default, the application may allows the user to click details icon **128**, which, when selected, displays a new box (not shown in Figure 9) that presents this information, as shown in Figure 10. The box **130** shows, for example, the sender and each recipient of the electronic mail object, their actions on the content (forwarding, replying), and the importance of each recipient.

[00106] Further, it is noted that importance may change over time. For example, when a user is interacting strongly with another user on a given project, then that other user may be of greater importance at that point than a year after the time when the project has finished. In another example, if a user is working on a deal linked to a specific job code, that job code may be of greater relevance during the deal than six months after the deal has finished. The chart icons **132** in Figures 8 and 10 may be used to generate a chart showing the evolution of importance over time (for a definable period set by the user), as shown in the pop-up window **134** of Figure 11.

[00107] Now turning to a folder view as described above, reference is made to Figure 12. As shown in Figure 12, there is a relevance column **137**, which shows a relevance slider (shown, but not labeled) for every electronic mail object in the folder listing. As in the mail view described above, a user may move a relevance slider to provide feedback to the electronic mail application. In addition, the user may click the details icon **136** which, when selected, presents a similar box as that shown in Figure 10, indicating users' actions and importance indicators.

[00108] Now turning to an address book as described above, reference is made to

Figure 13. In Figure 13, importance slider **139** is provided for each user, group, or domain, and allows the user to directly interact with the importance values given to their contacts. The user may either lock the value or provide it as a new start value from which the electronic mail system equilibrates. Further, it is noted, for example, that when user A adds user B to their "favorites," the importance of user B with respect to user A may be given a boost.

[00109] Now turning to preference settings as described above, reference is made to Figure 14. Using preferences, the user can affect the different inputs into the relevance analysis engine. This may allow the user to tune the relevance analysis engine to provide the content that the user is interested in. For example, the user may affect the relative strength of the different actions by users, as well as affecting the relative strength of the different entities with importance (e.g., job codes, users).

[00110] In addition to varying the relative strengths of various inputs to the relevance analysis engine, the user may also alter the importance of the different entities in the system, using an interface similar to that depicted in Figure 13, as linked from Figure 14. The user may either lock the input at that value or let it equilibrate over time (the slider for the impact of linking content to a context is locked **138** in Figure 14).

[00111] Further, the actions shown in Figure 14 may be customized. Such customizations may be done on implementation or by a system administrator. Moreover, it is noted that the mechanism of surveying users described above may provide the raw data that usually determines the weighting values. A given user's input may partially determine both their own weighting values and, possibly to a much smaller extent, the

overall default weightings across the entire system. At least one reason for taking a content surveying approach to determining weightings is that weightings may be regarded as an esoteric factor, and it is very difficult for a typical user to determine these values directly.

[00112] Moreover, in one or more embodiments, determining weightings values may be done via a consulting exercise by an integrator entity implementing the system. The electronic mail application may have the ability to serve automated surveys to users via the weightings interface, in which users are "walked through" various scenarios and their responses determine the weighting values. This may be provided via the "Wizard" button **140** shown in Figure 14. Additionally, the user may restore the sliders back to their default values.

[00113] Another option for searching in one or more embodiments is to display a search capability on every page of the application, allowing the user a familiar search function from wherever they are in the application. This choice may appear under the "My settings" tab **142** in Figure 14.

[00114] Now turning to search interface aspects relating an electronic mail system in accordance with one or more embodiments, reference is first made to Figure 15. A search tool may be divided into two areas – a search form and the search results. The tab **144** shown to be selected in Figure 15 is the simple search option, in which a user just enters any keywords they are interested in. A search algorithm then searches for this text in all possible criteria (e.g., users, groups, job codes, message subjects, message bodies). If the user clicks the "options" roll down **146** in the search form interface as shown in

Figure 15, the application may then provides some additional search options, such as shown in Figure 16. It is noted that in one or more embodiments, the application may remember when the user clicks the options roll down **146** and serves the search interface to them with the options already showing the next time the user goes to the search page. The options roll down **146** allows the user to specify, for example, which mail folders to look in, people involved, key words to search for, the date range, and/or whether the electronic mail object should have attachments.

[00115] Further, the search interface shown in Figure 16 provides the user with two pairs of options relating to what happens to search results. Each of these options may have descriptive help in the form of a pop-up window that appears when the user moves their mouse over the option. A user may save any search they are doing as a "smart folder" that appears in a custom folder in their folder tree. A static "smart folder" **150** keeps the historical results of the search, so the user sees the same result set each time they click the "smart folder." A dynamic "smart folder" **152** may carry out the search again each time the user clicks the "smart folder." The user manages "smart folders" through a standard folder management interface. To edit any given "smart folder," the user may click the folder, and the application presents the folder's respective search query through the appropriate search interface. A final decision for the user to make is whether they would like the folder view or extended view for their search results.

[00116] Referring again to Figure 15, the "more options" tab **148** provides another screen similar to that shown in Figure 16, but with more options for the user to choose through, for example, checkboxes and text fields.

[00117] Now turning to a description of advanced search aspects related to an electronic mail system in accordance with one or more embodiments, reference is made to Figure 17. With the interface shown in Figure 17, a user may build search sentences and then use the sentences to do searches. A description of Figure 17 is further detailed below in regard to sentence types, electronic mail object actions, search request boxes, phrase palettes, and search options.

[00118] In one or more embodiments, there may be a plurality of different sentence types in an advanced search interface: people **154** – user, groups, and domains that the mail is from and to; keywords **156** – keywords in any field in the system (e.g., the subject or message body); details **158** – details and properties of the mail (e.g., whether it has attachments or has been assigned a job code); actions **160** – user's actions on mail content, such as reading or replying; and custom **162** – access to a library of other sentences. Such sentence types may be presented to the user through a tabbed interface as shown in Figure 17, allowing the user to switch between the different types as they build up their search request.

[00119] Mail actions **164** represent a large area of the user interface shown in Figure 17, because it is that sentence type that is shown in the interface. There may be a plurality of subsections within this part of the view: phrase **166** – the area in which the user selects the type of action they wish to use in their search request; sentence **168** – the area in which the user builds up the sentence they wish to add to their search request; and sub-phrase **170** – the area from which the user can introduce modifications to their basic phrase.

[00120] In Figure 17, an action selected is that of a person forwarding a mail to other people (be they users, groups, or domains). This phrase appears in the sentence box **168**, allowing the user to "fill in the blanks." In Figure 17, the user has already selected the first blank – that they themselves forwarded the mail. The user is in the process of filling in the second blank. By clicking the "people" field **172** with their mouse, they have caused the application to spawn a pop-up window **174**. This window **174** allows the user to select the most requested choice or go to a modified view of their address book to select the people they wish.

[00121] The user may customize the phrase in the sentence box by adding in sub-phrases from the sub-phrase screen area **170** . Each of the different sub-phrase types may have a color-coordinated arrow next to a given sub-phrase and a matching slot in the sentence box. The user can drag and drop sub-phrases into the sentence box **168** and then customize them as they wish. Figure 17 shows two sub-phrases have been added in – the "all of" sub-phrase **176** and the "from [date] to [date]" sub-phrase **178**. The date sub-phrase may require further user interaction to specify the dates, which the user has already carried out using the associated calendar buttons **180**. Each calendar button **180** spawns a calendar that allows the user to specify a date. It is noted that date formatting may be a user preference determined globally across the application via the "personal settings" tab.

[00122] Once a user has built a search sentence that they wish to add to their search request, they may click the "move to search request" button **182**, which will add the sentence to the bottom of the list in the search request box. In addition, the user may use

an "anchor" system. Any hard anchor, depicted in Figure 17, may be dragged to an anchor slot, subject to certain rules regarding how phrases can be combined. The anchor system may be common to the sentence area, the search request box, and the phrase palette. It is further noted that in order to assist the user, there may be various user pop-up help boxes that appear when the user moves their mouse over the different parts of the interface.

[00123] A search request box **184** is the screen area in which the user builds up their search criteria, drawing on sentences from each of the different sentence types. Each sentence is shown, along with anchors allowing each sentence to be extended with other sub-phrases. For example, the sentence, "have history submessages from Cattlin, Simon or Ash, Duncan" may be extended at the "from" point to "have history submessages (from Cattlin, Simon or Ash, Duncan) or contain the keywords 'sales proposition.'" Each sentence (in both the search request box **184** and the sentence box **168**) may be changed between positive and negative logic. The user may change the "have" dropdown box to "have not." There may be weightings for each sentence and each of its sentence criteria. These are shown by sliders in Figure 17, and the user can move the positions of these sliders to affect the relative importance of the respective inputs in the search request. In the example shown in Figure 17, electronic mail objects having history submessages from Simon Cattlin would be given higher ranking than those from Duncan Ash in the search results listing. Further, the first sentence regarding history submessage authorship would be given slightly greater weighting than the requirement to have at least one attachment.

[00124] Another part of a search request box shown in Figure 17 is an "actions" box **186**. This part may contain icons that allow the user to manage the sentences that make up their query. The top row from left to right allows the user to copy a sentence (e.g., to search for two very similar sentences), move a sentence up the list, move a sentence down the list, and delete a sentence.

[00125] Still referring to Figure 17, an icon may allow the user to toggle the weighting sliders on and off. When the user turns the weighting sliders off, the ordering of the list is particularly important, as the system automatically attaches greater weighting to sentences further up the list. It is noted that when a weighting slider is off, such may be indicated in the user interface. If the weighting sliders are hidden, an icon may be displayed to show the slider to indicate that it would toggle the weighting sliders back on again.

[00126] Another icon in the actions box **186** is a "?" symbol that provides the user with help on using the advanced search interface. When clicked, this icon launches a debugging interface in a pop-up window that walks the user through debugging their sentence to ensure that it is consistent. Such interactive debugging may be useful in helping to ensure that the interface is usable by most individuals, as well as being both sophisticated and powerful.

[00127] Further, it is noted that if a user starts designing a search in the regular search interface shown in Figure 16 and then moves to the advanced search interface as described above with reference to Figure 17, the user's initial input may be represented as a sentence in the advanced search interface. The user may then further customize and

extend the search from there. This provides a convenient first step for a user attempting to move from the regular search interface to the advanced search interface.

[00128] Still referring to Figure 17, a phrase palette **188** may be provided to offer a "dumping ground" for a user to move phrase and sub-phrases about. It is shown rolled up in Figure 17, but the user can expand it by clicking the appropriate icon/tab. The phrase palette may be particularly useful to the user when transferring phrases from their custom sentence type to mix with other sentence types.

[00129] Many of the elements in the search options box **190** may be similar to those provided in other web pages of an electronic mail system in accordance with one or more embodiments. For example, the user may save a search as a "smart folder" that appears in their "custom smart folder" set. The "smart folder" may either be static, in which case it keeps the same result listing as when the search was first run, or dynamic, in which case the application reruns the query each time the user views the smart folder.

[00130] Search results may either be displayed in folder view or in extended view, as described above. An additional option for the user is whether to include the historical submessages or not.

[00131] Still referring to Figure 17, a custom sentence type provides access to a library of phrases and sub-phrases. Such a library may be a centralized resource held server-side across an organization. The user may type in phrase fragments and search for phrases that match their desired input, thereby effectively bridging the gap between natural language searching and logic-based searching. The organization may also update

the library with new phrases and sub-phrases, mapping these to new logic elements at the search level. Examples of additional phrases include: "that I wrote that became part of a tree linked to job code X"; and "that I wrote and were read by Clarke, John." Other phrases may act to alter the search algorithm's action, such as: "that will help me find someone's name"; "that will help me find a historical document"; and "that will help me find out about a project." These phrases alter the search results in such a way as to assist the user in finding the type of information they are looking for.

[00132] Now referring to Figure 18, a user may interact with the settings for the search algorithm and results ranking. Using the interface shown in Figure 18, a user may change the relative strengths of the inputs to the search algorithm, in a similar manner as a user can adjust inputs to a relevance analysis engine as described above. The user may have the option of tying the weighting settings to be the same as their relevance settings. The user may have additional options they can change under "search other settings" **192** including, for example, the following: the number of results per page in both search result views; the information to include in the extended view; and the columns to include in the folder view (including local relevance and search relevance).

[00133] In addition, a user may add in new buttons that allow them to further alter the search results. These may be accessed via a library similar to that described above. Examples of different purposes of such custom buttons include: "limit it to mail I read"; "show me mail read by more people"; "show me more recent content"; and "let me change the weightings" (presents the user with a pop-up window to change the weightings of their search criteria).

[00134] Further, a user may enforce any changes on either a session basis (during a particular log in session of searches) or for all future searches. This option allows the user to increase the weighting towards specific criteria for a given search in which those criteria are important.

[00135] Now turning to a description of search results aspects of an electronic mail system in accordance with one or more embodiments, the electronic mail system may have a plurality of views for presenting search results. For example, a "folder view" represents a way a user may be accustomed to viewing search results. In another view, an "extended view" more closely resembles the way users are accustomed to viewing web search results. It is noted that providing a plurality of such views and then allowing the user to choose which they prefer as they conduct their search promotes user satisfaction. Further, the electronic mail application may "remember" the user's previous choice and defaults to this view when the user conducts a new search.

[00136] Figure 19 shows an example of folder view search results. In one aspect, this view may be regarded as being intuitive. The table indicates the author **194** and recipients **196** of each mail, the subject, any designated job code, key dates (delivered and read), and the mailbox the electronic mail object resides in. In addition, the table may include an information icon that allows access to a similar pop-up window as that shown in Figure 10. Further, the table may display relevance sliders for each electronic mail object returned by the search.

[00137] Still referring to Figure 19, the row above the column headings includes some potentially useful capabilities for the user. The user may refine the search and may

also search within the results they have already been presented with. "Back" and "forward" buttons **198** allow the user to move up and down their search history with a single click. The user may also sort each column ascending or descending to further customize the results view.

[00138] Now turning to a description of search results presented in an extended view, reference is made to Figure 20. As in the folder view for search results described above with reference to Figure 19, the extended view interface shown in Figure 20 includes options to refine the search **200**, search within results **202**, and move backward and forward **204** through search history.

[00139] Still referring to Figure 20, a "news" content set **206** is the results of searching an organization's knowledge base and intranet site, not electronic mail objects. In the example shown, the first two items are links to external content that a search spider found on the intranet site, while the second two items have been written by other employees as part of the organizational knowledge base. In is noted that the search results presented in accordance with one or more embodiments are highly customizable, and additional content feeds may be placed into these top items. The organization may assign content feeds to each user through the administrative interface or may allow the user themselves to choose different content feeds.

[00140] The next three items shows in Figure 20 are not expanded, but the user may select one of the icons to display each set of search hits. The content makes up these categories includes, for example, the following: contexts (e.g., development, sales) **208**; job codes **210**; and people (e.g., users, domains, groups) **212**. Each of these sets may

constitute a set of entities and the occurrences of the search term within them. The number shown in brackets denotes the number of search hits in each category.

[00141] Still referring to Figure 20, below the categories described above (e.g., news **206**, contexts **208**, job codes **210**, people **212**), the search results list hits within electronic mail content. Each search hit highlights the content that caused the item to appear in the search results, with, for example, the following display schema: the subject of the submessage; a magnifying glass containing a question mark; the relevance of the submessage; the body text showing the instances of the search term; the people involved (senders and recipients); a set of property icons similar to those described above with reference to Figure 8; a job code or attachment name if they are part of the search hit; the date the mail was received and read; and the mail folder in which the respective electronic mail object appears.

[00142] In one or more embodiments, clicking a user (or group or domain) name causes the application to display the respective "smart folder" of the searcher's interactions with that population entity. Clicking a job code, context, or other entity may have a similar effect.

[00143] Clicking the magnifying glass **214** with a question mark inside may spawn a pop-up window indicating to the user how the result matches their query. This allows the user a quick way to understand why the result would be of relevance to their query and how the application arrived at that decision. It also allows the user to provide feedback to the electronic mail application by indicating whether the result was what they were looking for.

[00144] It is noted that the relevance slider **216** for the submessage with the subject "don't bother reading – just a test message" **218** is quite low given that this is the first page of the search results. This highlights that the relevance sliders may not be respective to the search being conducted – they are local relevance values based on the output of the relevance algorithm. Relevance in search is used to determine the ranking for the search results (though, this may also be displayed as a user preference). The test submessage was deemed to have low relevance as few people read the submessage, given that the submessage was a test and not real content. However, because the body of the submessage includes many instances of the word "test," this helped drive the submessage up the ranking of the search results given the query was for the word "test."

[00145] As shown in Figure 20, the search result "don't bother reading – just a test message" **218** does not show the instances of the search term in the message body. This is because the submessage is unread, and in one or more embodiments, an electronic mail system may not allow partial reading of a piece of content in a user's incoming electronic mail. In this case, the user may click the submessage to read it if they wish to view the content.

[00146] In one or more embodiments, search terms displayed in search results may be highlighted in one or more of a plurality of ways. For example, as shown in Figure 20, search terms are displayed in bold to help the user identify the terms. Another option, used at the bottom of the interface shown in Figure 20, is to highlight the search terms **220**. Such a feature may be particularly useful for searches for multiple terms, as it may highlight each search term in a different color.

[00147] As described above, an electronic mail system in accordance with one or more embodiments may use a relevance analysis engine for one or more of various purposes (e.g., managing and presenting electronic mail, forensics, searching). Now will be provided a more detailed description of how such a relevance analysis engine may be implemented in either software and/or hardware. "Importance" and "relevance" are at least two components or concepts associated with the relevance analysis engine. A "relevance" score attaches to electronic mail and indicates to the recipient the priority they should give to dealing with this piece of electronic mail. Because a given piece of electronic mail may actually have more than one "sender" (e.g., an individual and a job code, a forwarded user), the relevance of the piece of electronic mail may relate to a set {s} of senders. An "importance" score attaches to an individual user (though this could sometimes be, for example, a job code rather than an individual). Importance is a metric used to help to determine the relevance of electronic mail objects by reference to who sent or forwarded them.

[00148] In order to describe one or more embodiments of a relevance analysis engine, the following components will be used:

1. Mail Data Component: a mail data component is a feature of or attachment to a particular electronic mail object (e.g., date or flag for follow up). [symbol: ${}^{x}MDC$ for mail data component of electronic mail object x; this is calculated from the various mail data factors ${}^{x}MDF_i$]
2. Individual Relevance Component: an IRC value is a relevance score for a piece of content respective to a given individual user, and is based on the actions of that user. [symbol: ${}^{x}IRC_r$ for recipient r's actions on electronic mail object x from senders {s}]
3. Individual Importance Component: an IIC value is an importance

score for an entity (e.g., a sender or job code) s respective to a given individual user, and is based on the actions of that user. [symbol: IIC_{sr} for recipient r's actions on past mail from sender s]

4. Global Importance: a global importance value is an importance score for an entity respective to the entire system, and is a weighted average of all IIC values for that entity. [symbol: I_s^g for sender-entity s]

5. Global Relevance: a global relevance value is a relevance score for a piece of content respective to the entire system, and is a weighted average of all IRC values for that piece of content. The weights used relate to the global importance values of those to whom the IRCs relate. [symbol: ${}^xR^g$ for a piece of content x from senders {s}]

6. Cognate Group: for each user, over time, the system identifies a "cognate group" of the n other users regarded as most similar. This group is defined by similar past user actions and assessments, and/or a direct input by a user. It is noted that it is not assumed that the population of recipients falls into separate and mutually exclusive cognate groups. Rather, each recipient may have their own cognate group of those the system regards as most similar to that recipient. Thus, in some cases, for example, user A may be in user B's cognate group (because A is in the n users most similar to B), but at the same time, B may not be in A's cognate group (because there are n other users more similar to A than B is).

7. Cognate Relevance Value: a cognate relevance value is a relevance score for a piece of content respective to a given recipient based on a weighted average of the IRC's of the n users considered to be the cognate group for that recipient. [symbol: ${}^xR_r^c$ for the relevance input from the cognate group c of recipient r, on electronic mail object x sent by senders {s} to recipient r]

8. Cognate Importance Value: a cognate importance value is an importance score for an entity respective to a given user, based on a weighted average of the IIC's of the n users considered to be the cognate group for that recipient. [symbol: I_{sr}^c for the sender s to recipient r in cognate group c]

9. Local Importance: a local importance value is the importance value for an entity (e.g., a sender, a job code) to a given user which is shown to that user through the interface. It is a function of the user's own IIC for that sender, together with global and cognate importance values for the sender. [symbol: I_{sr}^l for local importance of sender s to recipient r]

10. Combined Importance: local importance values are stored locally

for each recipient for each sender or entity with which that recipient has interacted. There may be a set {s} of senders associated with a given incoming electronic mail object, and the local importance factor germane to that electronic mail object is then not a single sender local importance but a value arising from the combination of local importance values for that set of senders. Whereas local, cognate and global importance values may each relate only to individual senders, the combined importance is a combination of importance values that relate to the particular set associated with a particular electronic mail object. [symbol: ${}^{x}CI_{\{s\}r}^{l}$ for local importance to recipient r of the particular set of senders {s} associated with electronic mail object x]

11. Local Relevance: a local relevance value is the relevance value for a piece of content to a given user that is shown to this user through the interface; it is a function of the user's own IRC, the global and cognate relevance scores, and the combined importance relating to that mail. [symbol: ${}^{x}R_{\{s\}r}^{l}$ for content x sent by senders {s} to recipient r]

[00149] Further, Figure 21 shows an example chronological scenario for the relevance analysis engine. Figure 21 shows at a high level how the terms described above relate to each other and the process by which they are calculated in the functioning of relevance analysis engine. Below is provided a more in-depth description of relevance techniques and computational implementation. It is noted that the scenario in Figure 21 depicts a final relevance mechanism (described below), which includes cognate relevance.

[00150] The steps exemplified in Figure 21 are as follows. In step 1, John sends an electronic mail object to Sue, Harry, and Chris. In step 2, a "knowledge server" (or "object-based electronic mail server") (e.g., TransAct developed by BLUESPACE SOFTWARE) in accordance with one or more embodiments carries out the following processes: sends the first electronic mail object; and calculates the combination

importance for each recipient. Then, in step 3, Sue and Harry are active or logged in. Subsequently, in step 4, the knowledge server initiates a batch cognate relevance recalculation for Sue, and so calculates the first cognate relevance value for this electronic mail object (at this point, the values is zero because there are no IRCs yet, so it has no input to the local relevance when evaluated). In step 5, Sue accesses her inbox. Then, in step 6, the knowledge server calculates Sue's local relevance value for the electronic mail object (in order to show the value in her folder listing). Thereafter, in step 7, Sue reads the electronic mail object. In step 8, the knowledge server carries out the following processes: marks the electronic mail object as having been read by Sue; calculates Sue's IRC value for the electronic mail object; recalculates Sue's IIC value with respect to John; and recalculates Sue's local relevance value (which changes because she has done something with it). Then, in step 9, the knowledge server calculates the first global relevance value based on Sue's IRC value. Subsequently, in step 10, the knowledge server initiates a batch cognate relevance recalculation for Harry, and so calculates Harry's first cognate relevance value for this electronic mail object, based on Sue's IRC value. In step 11, Sue replies to the electronic mail object. Then, in step 12, the knowledge server carries out the following processes: sends the replying electronic mail object; calculates the combination importance for each recipient for the replying electronic mail object (based on both Sue's importance and John's importance as the two authors) (it is noted that this replying electronic mail object may initiate a separate process thread); and updates Sue's IRC for the original electronic mail object and her IIC value for John. Thereafter, in step 13, the knowledge server recalculates

Harry's cognate relevance value for the original electronic mail object based on Sue's IRC value. In step 14, Harry accesses his inbox. Then, in step 15, the knowledge server calculates Harry's local relevance value (for his folder listing). In step 16, Harry reads the TransMail. Subsequently, in step 17, the knowledge server carries out the following processes: marks the electronic mail object as having been read by Harry; calculates Harry's IRC value for the electronic mail object; recalculates Harry's IIC value with respect to John; and recalculates Harry's local relevance value. Thereafter, in step 18, the knowledge server initiates a batch cognate relevance recalculation for Sue, and so recalculates Sue's cognate relevance value for the (original) electronic mail object, based on Harry's IRC value. In step 19, the knowledge server calculates the second global relevance value, based on the IRCs from Sue and Harry. Then, in step 20, Chris logs in or becomes active. Subsequently, in step 21, the knowledge server calculates Chris' cognate relevance value for the original electronic mail object for Chris based on the IRC value from Sue (Harry is not in Chris's cognate group). In step 22, the knowledge server continues to recalculate cognate relevance values and the global relevance value at specified, but, for example, increasing intervals. Then, in step 23, the knowledge server, during a period of low activity (e.g., in the evening), recalculates global importance, cognate importance, local importance, and combination importance values. A detailed visualization of how these different variables interact is exemplified in Figure 21.

[00151] As apparent from the description above in regard to a relevance analysis engine in accordance with one or more embodiments, the relevance analysis engine may continually update a set of inter-related values, based on user action, so that everyone's

actions affect everyone else. Thus, in view of this, each user/post may have a different individual set of importance and relevance criteria. Further, an electronic mail system of this kind in accordance with one or more embodiments may also require the input of user-set values for the relevance and importance criteria, which both correct the values for the user concerned and may also offer a check on the effectiveness of particular relevance mechanisms being used. Moreover, it is noted that for organizations with many posts attached to user teams, an electronic mail system of this kind may allow users and posts to be treated separately with varying levels of inter-correlation.

[00152] Now turning to a description of a core of a relevance analysis engine in accordance with one or more embodiments, it is assumed that (i) there exists a set of people a, b, c, . . . d involved in sending and receiving electronic mail objects; and (ii) senders {s} (plural because more than one may have been involved in forwarding it) send a piece of electronic mail x to the recipient r. Accordingly, an aim of the relevance analysis engine is to produce for recipient r a "good" value for the local relevance ${}^{x}R_{r}^{l}$ for that piece of electronic mail. Thus, such an aim may be expressed as follows: "if the recipient r were to open a set of incoming mails and consciously assign relevance scores and rankings, these would correlate highly with the ${}^{x}R_{r}^{l}$ values calculated by the system."

[00153] It may also be possible to define the relevance ranking of an incoming electronic mail x to recipient r before it is opened (its "before-reading" local relevance ${}^{x}R_{r}^{lb}$) and after it is opened and read (its "post-reading" local relevance ${}^{x}R_{r}^{lp}$). This may not be exactly the same as relevance of incoming and stored electronic mail object,

because some electronic mail objects may be stored without being opened. The icon in the user's view indicates that the electronic mail object has been read by an open letter symbol, so that it is clear which is being given as the local relevance value for a given content. Opening may be only one of the activities comprising the IRC for that recipient and electronic mail object (it could, for example, be deleted without opening it), but could be a key one in terms of ensuring that the actions done are based on full knowledge of the nature of the electronic mail object.

[00154] Further, there may also be a concept of search relevance ${}^{x}R_r^{S}$, which is a number assigned to the results of a particular electronic mail object search. These values may be computed using the local relevance values together with the search criteria and weights for them of the weighting slider icon (if utilized) given by the user to change the level of impact each action or entity has on the search request s.

[00155] For the Sth search, if ${}^{x}C_i^{s}$ is a dichotomous variable indicating the presence or absence of that search criterion in the electronic mail object (i.e., will have a value 0 or 1), and W_i^{s} is the sliding weight given it by the user for that search, then the search relevance may be calculated by:

$$ {}^{x}R_r^{S} = f\left({}^{x}R_r^{l}; \sum_{i=1}^{m} {}^{x}C_i^{s} W_i^{s}\right). $$

[00156] Search relevance values may be stored locally while the search results are being examined (unless explicitly stored by the user). It is noted an objective is for these to be allocated by the system to closely resemble the ranking which the user would assign were they to be able to look at the overall view and contents of their electronic mail.

[00157] Now turning to a description of factors used by a relevance analysis engine in accordance with one or more embodiments, a plurality of factors may make up the IRC. The first are indicator functions, ${}^{x}IRCD_{ri}$, which are dichotomous (i.e., are simply "on" or "off") depending on whether the action has happened or not. This includes attributes such as factors that are available only for past-reading local relevance: "the user opens the electronic mail object" (if so, then D=1, if not, then D=0); "the user clicks on a hyperlink in the electronic mail object" (if so, then D=1, if not, then D=0); and "the user replies to the electronic mail object" (if so, then D=1, if not, then D=0).

[00158] A second set of variables ${}^{x}IRCC_{rj}$ are values on cardinal scale (or ordinal scale transformed for the model to a cardinal scale): "how many extra times the user opens the electronic mail object"; and "to how many people the user forwards the electronic mail object." These may be placed into suitable functions $f({}^{x}C_{rj})$ that are between 0 and 1. Some of these factors may also "decay" over time. While it may be difficult to determine this without some indication from the user, there clearly may be some electronic mail objects connected with projects that, once completed, leave the user unlikely to wish to consult the electronic mail objects again. Some feedback may be useful connected with a particular job or project code. Further, there may be occasions when the relevance of an incoming electronic mail object is highly affected by the life cycle of a project. This means that it may be useful to include a time related factor of importance that has a discontinuity break when the project is completed. Moreover, this would mean that the relevance could be multiplied by a time-related factor ${}^{x}G_{r}(t)$. For example, ${}^{x}G_{r}(t)$ may reflect relevance of a project up to a deadline time, t_0, beyond which

the relevance begins to decay:

$$^{x}G_r(t) = \begin{cases} 1 & t \le t_0 \\ e^{-\lambda(t-t_0)} & t > t_0 \end{cases},$$

where λ is a decay constant (high for rapid decay, low for slow decay). In one or more embodiments, most stored electronic mail objects may have a time-decay-based element of relevance.

[00159] As described above, various mail data features may be considered that may affect relevance and that are inherent to the electronic mail object itself. For example, the electronic mail object may contain a flag for action from the sender. In another example, the electronic mail object may itself be a reply to an earlier electronic mail object. These factors $^{x}MDF_i$ enter the relevance calculations for the inherent or mail data electronic mail object component ^{x}MDC.

[00160] Another factor in relevance calculations is the importance of the initial and/or a subsequent sender (e.g., a forwarder), and other entities involved in the communication, e.g., job codes. In one or more embodiments, importance may be on a scale of numbers, and may be put into a function (e.g., $f(I_{sr}^{l})$, which is between 0 and 1). There may be an absolute global importance and also an importance to the particular recipient. For example, a recipient's manager may have low importance in the great scheme of things, but high importance to the recipient him/herself. Still further, as described above, the importance of the sender to the cognate or similar group within which a recipient is located may be considered. Accordingly, this means that there may be at least three factors for consideration as to importance:

1. I_s^g is the "global importance" of the sender s. This number may be the same for all recipients of all electronic mails from that sender at around that point in time, though clearly can change over time as the commercial scene changes. It may often depend, in particular, on the importance of the post as located in the organization chart of an organization.
2. I_{sr}^l is the "local importance" of the sender s to that individual recipient r. This may depend, for example, on where within the organizational chart the sender is relative to the receiver, or whether the receiver is working with the sender on some particular inter-company project.
3. I_{sr}^c is the "cognate importance" of the sender to a particular recipient, which is the sender's importance to the cognate or similar group within which the recipient is located. This may be thought of as a more focused element of communal importance than the global importance, which is universal.

[00161] There are two elements of both global and local importance: the "role importance" attaching to the post(s) held, and the importance that may attach to an individual per se. For example, suppose that the assistant to a director in one department moves to become the assistant a director in another department. Though the assistant may retain her electronic mail identity as an individual, much of both the global and local importance values may often attach to the posts held rather than the individual.

[00162] Further, there may also be importance factors involved in cases where the immediate sender *s* is forwarding mail from one or more previous senders, who also have their own importance values. This set is represented by {s}. This is why the input into the local relevance is actually a weighted average of the local importance values of all the senders, the combined importance $^xCI_{\{s\}r}^l$, which relates specifically to the combination of senders in this electronic mail object x as received by r. The local, global and cognate importance values, however, may relate to individual senders or entities in the system,

instead of to combinations.

[00163] Now turning to factors for consideration in importance determination, the factors labeled ${}^{x}D_{ri}$ and ${}^{x}C_{rj}$ are things that either happen or do not happen (e.g., "does the recipient reply?"). Their values may not be readjusted in the light of the success or otherwise of the prediction of relevance values for incoming electronic mail objects – only the weightings given to them are adjusted (using a model based on linear regression as described below).

[00164] The importance factors I_{sr}^{c} may be different, because they are not ascertainable variables but are (estimated) parameters constructed within the model itself. In this case, the values of these factors are affected by the electronic mail traffic that has occurred. For example, the importance of a user may be affected by the receipt of electronic mail from other users – his/her importance perturbs or is shared by the importance of the other users. This interactive process may be taken into account before the importance is fed into an importance determination mechanism.

[00165] If recipient i receives electronic mail from sender s, there may be some true local importance value ${}^{\Omega}I_{sr}{}^{l}$, which would be the importance r would give s with perfect information. In practice, this may not be static (importance values may constantly changing), but it may serve as a useful idealization. Initially, a starting value of ${}^{0}I_{sr}{}^{l}$ may be set globally, based on judgments of levels of responsibility within the organization. An example objective is that starting from ${}^{0}I_{sr}{}^{l}$, the system will converge over time to a value near to ${}^{\Omega}I_{sr}{}^{l}$.

[00166] The global importance value of a sender ${}^{\Omega}I_s{}^{g}$ is in some ways more simple because each sender has only one value, but it is also more complex insofar as it is arrived at as part of a large interconnecting matrix of importance values (so that, for example, a sender is more important if fewer but more important other users so regard him or her). The global importance may be an average of all the local importance values, but nonetheless may be a weighted average.

[00167] The cognate importance value of a sender ${}^{\Omega}I_{sr}{}^{c}$ to a recipient r is computed in run time, based on the individual importance components of those in the cognate group to this particular user. It may also be necessary to add in a factor that enables importance to be carried over from posts. For example, if user A moved from post S to take over post T, then user A's global importance will be a function of their previous importance while in post S and the importance the previous holder of post T had. Likewise, local importance may also be a function of post as well as individual.

[00168] Further, importance attaches to the sender, and may also attach to the job code and/or context, which are also (passive) entities within the system. The job code type carries an importance value alongside the sender. This added complexity means that the relevance determination mechanism may contain weightings of both types of importance – sender and job code type, among other factors.

[00169] In regard to computing the individual importance component, significant local action may be assessed. These are actions based on past receipts. These are weighted, based on a time-dependent function as described above, so that more recent

actions have greater weighting. For example, taken into consideration may be the percentage of times previous "unsolicited" electronic mail objects have been deleted or stored without opening. The following may be used to compute such a percentage, where n_s is the number received:

$$\frac{1}{n_s}\sum_{x=1}^{n_s} f^t(del;no).$$

Here, f^t is a time-based function for electronic mail object x, which is a function of whether the electronic mail object is deleted or stored unopened.

[00170] Further with respect to past received electronic mail, taken into consideration may be the "attention ratio" of t_x / L_x for electronic mail objects opened, which is the time spent reading the mail in relation to its length. The following may be used to compute an attention ration, where m_s is the number that have been opened:

$$\frac{1}{m_S}\sum_{x=1}^{m_s} f^t(t_x/L_x).$$

[00171] In addition, taken into consideration may be any past relevance scores given by the recipient to electronic mail from this sender (a weighted average if more than one has been graded). If n relevance scores have been given for past electronic mail, the significance of awarded rankings will decay over time according to the following:

$${}^{x}R_{rs}^{assigned} = \sum_{x=1}^{v} f^t(R(T_x)).$$

This may also be extended to user-assigned relevance values for particular job codes.

[00172] In regard to user-assigned importance scores, when an electronic mail object is opened, there is a slider showing the local importance of the sender as calculated by

the system (as described above with reference to Figure 8). The recipient is able to adjust this if they feel that it is inappropriately low or high. If this is the first time an electronic mail object has been received from this sender, then the recipient can select a setting that will flash to remind him or her that they have not previously received any electronic mail object from this source, and that readjustment of the calculated importance of the sender might therefore be particularly useful.

[00173] The effect of this action by the user may be very great when it is first set. Users would find it odd were they, for example, give a ranking of 5 and then immediately find that system assigns a local relevance of 1 based on their past actions. They may be in a better position to know the import of their past actions (if any) than the system.

[00174] Over time, however, if, for example, all the electronic mail from a sender assigned a high importance were in fact deleted without opening, the system would be expected to begin to adjust the local importance. In this case, an appropriate function, however, is $f^t(I_{rs}^{assigned})$. The assignment of importance may either be "locked" by the user (so that it does not change) or can be left for the effect of the assigned score to decay over time as just described above.

[00175] The individual importance component (IIC) may therefore be a weighted function of all these various elements of the individual recipient's actions towards sender s, as follows:

$$IIC_{sr} = w_1 \frac{1}{n_s} \sum_{x=1}^{n_s} f^t(del;no) + w_2 \frac{1}{m_S} \sum_{x=1}^{m_s} f^t(t_{x1}/L_{x1})$$

$$+ \quad w_3 \sum_{x=1}^{v} f^t(R(T_x) \quad + \quad w_4 f^t(I_{rs}^{assigned}) .$$

[00176] A factor that affects local relevance is a weighted combination of the individual local importance values of the set $\{s^x\}$ of all the senders (and other entities, such as job codes) involved in sending and transmitting the electronic mail object x (which may be effectively a kind of mail thread with those past senders recorded in the mail data).

[00177] Now turning to a description of computing global importance, global importance may not be updated at run time, but on a periodic review system. It may be a function of the elements now immediately described below. One element may relate to the weighted sum of individual importance components. The global importance value will be a weighted aggregate of all the individual importance components (IICs). It may be weighted because if, for example, 10 important people believe that John Doe is important, this is more indicative than if 100 unimportant ones find John Doe to be important. It is noted that this may imply a degree of circularity in the process, because global importance values are therefore mutually dependent. As described below, this requires a system of equations solved iteratively in what is, for example, essentially a Gauss-Seidel algorithm procedure until convergence to a sensible precision. In essence, however, this key element of global importance is:

$$f\left(\sum_{r=1}^{n} f\{IIC_{SR}, {}^{(k-1)}I_r^g\} \right),$$

where k-1 represents that global importance values in the previous period in this

calculation are being used).

[00178] Another element for computing global importance may relate to electronic mail objects received. Receipt of an electronic mail object from a more important sender increases the importance of a given recipient, dependent on how many others receive the same electronic mail object. "Shared importance" may therefore depend on I_s^g/N_x, where N_x is the number of recipients of the xth electronic mail object. The results may depend both on the number of electronic mail objects and the average of the shared importance values. The factor affecting the global importance at time k of a user may be given by:

$$f\left(\sum_{j=1}^{m}{}^{(k-1)}I_s^g/N_j;\frac{1}{m}\sum_{j=1}^{m}{}^{(k-1)}I_s^g/N_j\right),$$

where this represents a summation across the m senders of m electronic mail objects in time k-1 to k.

[00179] Another element for computing global importance may relate to directly assigned values. For example, in some cases within an organization, the human resources section may wish to directly intervene in the importance rankings. For convenience, this may be termed $f\,({}^{o}HR^{t}_{s})$, where O represents the HR section of organization O, s represents the user (potential sender) concerned, and t represents the time at which the intervention was made (this may be locked).

[00180] Still another element for computing global importance may relate to transferred post importance. This may, for example, relate to human resources, and may involve any change in importance due to assuming another or upgraded post or role in the organization: $f({}^{o}HR^{P}_{s})$.

[00181] With respect to calculating global importance in accordance with one or more embodiments, a global importance determination mechanism may readjust the global importance in time k from that of k-1 using various functions, where value α indicates how fast the global importance should change over time:

$$
\begin{aligned}
{}^{k}I_s^g = \boldsymbol{\alpha}^{(k-1)} I_s^g + (1-\boldsymbol{\alpha}) \Big[& w_1 f\left(\sum_{r=1}^{n} f\{IIC_{sr}; {}^{(k-1)}I_r^g\} \right) \\
& + w_2 f\left(\sum_{j=1}^{m} {}^{(k-1)}I_s / N_j ; \frac{1}{m} \sum_{j=1}^{m} {}^{(k-1)}I_s / N_j \right) \\
& + w_3 f({}^{o}HR_s^t) \\
& + w_3 f({}^{o}HR_s^P) \Big].
\end{aligned}
$$

[00182] These global importance values may be recomputed by the system every period at intervals as determined by a system operator/administrator. They may be heavily inter-related and so this system of equations may be solved iteratively in what is essentially a Gauss-Seidel algorithm procedure until convergence to a sensible precision. Further, it is noted that as a given implementation is developed and customized, more powerful iterative systems may be applied.

[00183] Now turning to a description regarding computing cognate importance in accordance with one or more embodiments, an electronic mail system may store for each user, a set of *n* other users considered to be most similar to the user concerned – this is that user's "cognate group." The composition of this cognate group may be altered by a periodic review (a rapid change not being vital). The cognate importance for a particular sender is computed periodically using the individual importance components of each of

the cognate group. In some cases, there may be no individual importance components for a particular member of the cognate group. In this case, the value may be computed from those for which components do exist, with a decreased weighting in the final formula for local importance:

$$I_{sr}{}^{c} = \frac{1}{n_C}\sum_{i=1}^{n_c} f(IIC_i),$$

where n_c is the number in the cognate group with non zero IICs.

[00184] In regard to computing local importance in accordance with one or more embodiments, the local importance represents the local importance to the recipient of the sender of an incoming piece of mail. It may be a function of, for example, the following: the individual importance component (IIC) (based on the user's own actions); the cognate importance component (based on the IICs of the users in the group currently identified as cognate to the present recipient); and the global importance value (based on the global values assigned at the last periodic review made by the system). Thus, this may be expressed as:

$$^{k}I_{sr}{}^{l} = f(IIC_{sr}\,;\, ^{(k-1)}I_{sr}{}^{c}\,;\, ^{(k-1)}I_{s}{}^{g}).$$

[00185] If this is the first time the recipient has received any communication from s, then there may be no value for $IIC_{sr}{}^{k}$ at all, in which case $^{0}I_{s}{}^{1}$ will be based purely on $I_{sr}{}^{c}$ (if this exists) and $^{(k-1)}I_{s}{}^{g}$. If the sender is a completely new user unknown to the system and having no importance value at all (not even one assigned by a human resources department), then the system may not have an obvious way to decide how important is

the sender, and so how relevant the mail. Because the default importance in such a case will be low, "newcomers" may be advised to encourage their human resources departments and/or potential recipients to set importance values for them, and in the case of, for example, a sales team, contacting a new customer they need to set the importance values for that customer accordingly. Each customer may form their own policy on how to assign initial global importance values, with correlation to rank or remuneration being most likely.

[00186] A well-established ${}^{k}IIC_{sr}$ may have a greater weighting over the ${}^{k}I_{s}^{g}$ value, and, as already described, a specific $f^{t}(I_{rs}^{assigned})$ gives high weighting within the ${}^{k}IIC_{sr}$ to any specific assignation of importance by the user. Further, if the individual user has specifically said that a particular sender (e.g., a spouse) is important, their lack of global importance or cognate will not be expected to alter this.

[00187] In regard to computing combined importance in accordance with one or more embodiments, it is noted that the combined importance may be a function of the local importance values of the set {s} of senders (and other entities, such as job codes) involved in electronic mail object x:

$$ {}^{x}CI_{\{s\}r}{}^{l} = f\left[\sum_{\{s\}} w_s I_{sr}{}^{l}\right]. $$

[00188] As described above, a sample may be collected by the system of the actual relevance scores that a user would have assigned, having looked at electronic mail, on a sliding scale that effectively works as a percentage. A "good" system would produce a

priori relevance weightings that turned out to be near to the ones assigned post-opening by the recipient. The formula described above is actually a linear combination of the various given values of the D, C and I's. It is therefore possible, given enough feedback data, to use linear regression to derive the weights which would minimize the error-squared terms in predicting the scores for relevance as given by the user. The database used may essentially be a very large matrix, with a large number of zeros and can often become ill-conditioned.

[00189] Such analysis enables a plurality of features. First, for example, different recipients who provide feedback may have their own individual weighting estimated. Further, if patterns of weightings emerge for different types or levels or importance scores of users, then global best estimates of weightings on a user category basis may be refined over time and these can be useful particularly in assigning initial weighting to a new user.

[00190] An example of interrelationships among at least some of the different types of values described above is shown in Figure 22. It is noted that the local importance and relevance values are what are shown to the user. Figure 22 also indicates the timings by which the different elements are calculated (further described in detail below).

[00191] Still referring to Figure 22, an individual user's actions on a particular piece of electronic mail may have at least two effects: (i) they contribute scores to that user's individual relevance components (IRCs) for that electronic mail object x; and (ii) they contribute scores to that user's individual importance components (IICs) for the sender(s) {s} of that electronic mail object. These may be stored locally to that user. In the

communal sphere: the user's IRCs then contribute to a global relevance score, and are available to contribute to any cognate relevance scores for other users; the user's IICs then contribute to a global importance score, and are available to contribute to any cognate importance for other users; global importance scores for senders are calculated periodically using the IIC's; and global relevance scores for each mail are calculated periodically using the IRCs (being revised more frequently in initial stages of an electronic mail). For each user, a cognate group of n users is identified, whose IICs correlate most closely with that user. As described above, this means that it is quite possible for John to be in Joan's cognate group but not vice versa. The identification of these n users does not have to change rapidly, and may be done on periodic review, with the identities of the n stored locally for user r. Further, for each user, a local importance value for each individual sender with whom a recipient r has interacted is calculated periodically and stored locally to r. This may utilize the recipients own IIC's for s, the global importance of s, and a value for cognate importance for s found from the combination of IICs amongst the current n identified as cognate to r. Still further, for each user, when an electronic mail object arrives, its mail data may include inputs from a set of senders {s} who have forwarded or contributed to the history of that mail. A combination importance for that electronic mail object x may be calculated when it arrives, using the mail data to identify senders {s} (as well as any other entities involved, such as job codes), and stored locally to r. If any key senders in {s} have not previously interacted with this recipient, there may be no stored value for local importance for them. In this case, a local importance figure is calculated for that sender when the mail arrives,

using just the global and cognate importance values because there are no IICs at that point.

[00192] In regard to finding local relevance in accordance with one or more embodiments, a local relevance value for recipient *r* for an electronic mail object *x* is calculated using the following, for example, as inputs: mail data (as incorporated in the electronic mail object itself); combination importance of senders {s} of electronic mail object x (as cached on arrival using periodically reviewed local importance values); global relevance of electronic mail object x (as stored on periodic review); recipient r's IRCs for electronic mail object x (i.e., any actions done on the mail which affect local relevance immediately); and cognate relevance for electronic mail object x (looking at the IRCs of the n in the cognate group).

[00193] Further, it is noted that the recipient r's IRCs may include actions done by him or her on this electronic mail object. The user's IICs may include both of these and of any past actions done on the electronic mail object from the same sender. These may be stored locally to the user.

[00194] It is additionally noted that the cognate group composition may be determined by correlating the IICs. The n other users with the most similar IICs to r will be the user's cognate group. This may be done on periodic review basis, as it does not need to change rapidly.

[00195] Moreover, it is noted that the sender s's global importance may be found from all the IICs (including that of r) on a periodic review basis. As described above,

because importance values are used to weight the IICs to determine importance values, this is circular, but it is solved to a given accuracy on this periodic review basis.

[00196] Further, it is noted that the cognate group may have at least two effects. One is through IICs of Cognate Group on these senders, to cognate importance values, to local importance values, to combined importance of senders, to local relevance, which reflects how the cognate group in general has rated the senders. Another effect is through the direct actions of the cognate group on this particular electronic mail object (e.g., IRCs of cognate group on this electronic mail, to cognate relevance, to local relevance).

[00197] Additionally, it is noted that an aim of one or more embodiments is to produce a "good" figure for local relevance, which is the relevance value seen by the recipient when they are browsing through the interface. The system offers a rapid effect of actions on the mail of (in, for example, decreasing order of importance and so rapidity of change): the actions of user r him/herself on electronic mail object x; the actions on the electronic mail object x of the group cognate to r; and the actions of all the other recipients on the electronic mail object.

[00198] Each of the different relevance and importance values may have different characteristics that help to determine how frequently the values must be updated and the mechanisms for doing so. One aim may be to provide as "up-to-date" a value as possible for local relevance scores for a given piece of electronic mail in a given user's area. To do this, there may be different frequencies at which changes have to be made to the values used. Local relevance may be based on, for example, the following: mail data factors; individual importance components; individual relevance components; global

importance values; global relevance values; cognate importance values; and cognate relevance values. One or more of these may be available in a browse mode. Whereas older historical mail may not need the local relevance to update so rapidly, it may be particularly important that the most recent relevance figure be available for incoming or recent electronic mail objects structured as described above with reference to Figure 1B. Further, it may be particularly useful to a given user to know what, for example, cognate users, have done with the electronic mail object soon after it arrives, as this will help him/her to assess priorities.

[00199] As to slowly changing values, the identification of a cognate group may be done on a periodic review basis. If a user is new to the system, it may be desirable to revise their cognate group more rapidly in the early stages, as opposed to for a well-established user. However, it is noted that may not be necessary that the identities of the cognate group be very recent for a relevance mechanism to be useful.

[00200] Further, global importance values may not need to change rapidly. They represent the overall importance of entities in the system. It is note that these values may need to change more rapidly for newer entities than for existing ones that already have the history. Moreover, updates to the global importance may be carried out on either a queue or periodic review basis.

[00201] As to moderately changing values, cognate importance may change faster than global importance, but perhaps not as fast as local importance. Cognate importance is based on purposefully biased samples, but still involves the views of a significant number of individuals rather than just one user.

[00202] In regard to local importance, the elements of local importance that relate to actions in the IIC may be used at run-time to update the value of local importance. An important attribute of local importance is that once a value has been assigned based on a reasonable data set, the local importance value may not need to change rapidly. New entities interacting with a given user have the ability for their importance value to change more quickly, as the data set is smaller. This is a property of increasing sample size on average in any case. Effects on local importance of the wider issues of global changes may be done by periodic review, as are those of cognate importance.

[00203] Still referring to moderately changing values, if global relevance values are the basis for the communal component of local relevance, then they may have to be kept up-to-date. If not, the value of the communal component in local relevance may be greatly reduced. For example, the impact of five other recipients deleting the electronic mail object before the user in question reads it will not affect his/her local relevance value until after he/she has dealt with the electronic mail object. As with local relevance, the rate of change in a global relevance value for a given piece of content should be greater the newer the content is (as the first users interact with it). In practice, it may not be dealt with in run-time, and thus may be queued and periodically reviewed. It is noted that the rapid reassessment in the early stages of the electronic mail object's history, may, however, be important.

[00204] As to rapidly changing values, cognate relevance values may be similar to global relevance values, in that they change rapidly when content is new, reflecting the first real user validation of the relevance predictors for the content. It is noted that the

number of calculations is potentially large. Examples of such scenarios are now described below based on the following assumptions: an electronic mail is sent to 200 people; and n=20 is the number in a cognate group. To find the global relevance requires data from 200 sources and gives one figure for global relevance of the mail. Each of the 200, however, may have a cognate group, and the calculation of cognate relevance scores for each one will require up to 200x20=4000 figures to be combined into 200 individual cognate relevance scores. Thus, the cognate relevance needs to be updated faster than the global, because it is what offers the best guide to the user as to whether to prioritize looking at the current electronic mail object.

[00205] In one scenario for calculating cognate relevance, cognate relevance may be recalculated whenever the user acted on electronic mail or refreshed the browse mode. This means that if, for example, there are 150 active mail items in the box, then each time the user acts or refreshes, there needs to be 150x20 = 3000 pieces of information combined into 150 separate figures. Refreshing many of these maybe unnecessary, if they have been around a long time. Changes are more likely in the early stages of their life cycle.

[00206] In another scenario for calculating cognate relevance, cognate relevance may be recalculated on a periodic basis, with a recycle period which starts short and gets longer during the mail's life cycle. This means that each time the review appears, if 200 received the mail, then there has to be up to 200x20 =4000 retrievals and 200 calculations. This may happen even if the user is not accessing or reading mail (e.g., if the user is on vacation). At least one advantage is that this will be frequent only in the

early stages of the mail's lifetime.

[00207] In still another scenario for calculating cognate relevance, cognate relevance may be generally calculated on a periodic review basis, with a shorter cycle in the earlier stages of the mail's life cycle. The actual period in these early stages, however, may depend on whether the user was active on his or her mail. When the mail arrives, it has some local time-flag indicating periodicity of review of cognate relevance, and this decreases if the user refreshed the browse-mail or did a mail action. This contrasts with the scenario described above in which there is only a single time-flag for any given electronic mail, and a single mathematical series for review of all recipient's cognate relevance for that mail simultaneously. Mathematically, it may involve a function with, for example, an exponentially increasing review period, but with its parameters also dependent on mail-action frequency.

[00208] In yet another scenario for calculating cognate relevance, a cognate relevance calculation may be triggered by recipient refresh or mail-action in its early stages, but thereafter, on a periodic review basis, with an increasing cycle length. Within a certain time of arrival (e.g., the first two hours), which has been termed "phase one," any refresh or action by the user on other mail would trigger the recalculation of the cognate relevance for this particular electronic mail object. After that cut-off period, i.e., during "phase two," recalculation may be done on a periodic review basis. This scenario may have the advantage of focusing a greater frequency of recalculation where it is most needed. It makes it also possible to build in a "minimum recalculation cycle" during phase one, so that it is not literally every mail-action that causes recalculation, but that

minimum-period reviews take place on this electronic mail object during times the mail browse is "active." In the critical early stages of the electronic mail object, review occurs when relevant (e.g., when the user is looking), but later in its life, its cognate relevance may not be triggered unnecessarily often.

[00209] In another scenario for calculating cognate relevance, cognate relevance calculation may take place on a periodic review system during the time that the user is active. Each review recalculates cognate relevance for any mail received during the last P minutes, where P depends on the time since that active user-session began. The effect of this is to review more recent mail more often, but only during an active session. This system may be combined with a general recalculation during the system's "quiet time," which adjusts older mail relevance scores.

[00210] In one or more embodiments, when the user logs in or becomes active, a time-line begins, with recalculations every T minutes. T may be a function of the global importance of the user r, and the time t since login or recommencement of activity:

$$T = f(\boldsymbol{I}_r^{g}; t).$$

[00211] The system may allow more important users a more rapid period of recalculation, but the users themselves may not determine their own recalculation period pattern because everyone would otherwise opt for the shortest periods possible. It is noted that potentially an organization may choose to vary their recalculation frequency factors correlated to user type, e.g., seniority. The time P may itself be a function of t so that the cognate relevance of older mail is recomputed less often:

$$P = f(t).$$

When the user logs out or becomes inactive, the process is paused, waiting to recommence once the user logs back in or becomes active again.

[00212] The periodicity of recalculation and time periods involved may be varied on a per implementation basis depending on the computational resources and priorities of a given system. It may also be automated to modify itself automatically if the computational pool recalculating cognate relevance values reaches full capacity and begins building an extended queue. Such a system may have one or more advantages. For example, recalculation is done only when either the user is active or during computational quiet times. It may not be necessary to have up-to-date relevance values (requiring repeated recalculation) if the user is not looking at them anyway. As to another advantage, cognate relevance values of most recently arrived mail are recalculated more frequently, up-to-date information on these being the most useful to the user. As to a further advantage, the application may not need to consider each cognate relevance value to see whether it must be recalculated. Instead, it may run a single process, starting at the top of each user's inbox at the time correlated to their activity, and recalculate each cognate relevance value until it reaches mail that is older than the period currently being recalculated. At this point, it may halt, and does not need to evaluate any further electronic mail objects for cognate relevance recalculation. Another advantage is that the system may allow separate recalculation of, for example, the last month's worth of, cognate relevance values for each user on each weekday night ("quiet time"), and recalculation of the last years (or more) worth of cognate relevance values for each user

over the weekend (quiet time). This may ensure that the older mail values do not become anachronistic, by using space computational capacity in off-peak periods.

[00213] Accordingly, in one or more embodiments, an electronic mail system may ensure that users have up-to-date (including, for example, as many actions of other users as possible) cognate relevance values for recent content, while acting to minimize overall computational load, but still ensuring that cognate relevance values for historical content (which should change much more slowly) are recalculated periodically in computational "quiet time." As an example of this, a scenario is now described. A user becomes inactive or logs out at time t = m. Then, the user becomes active or logs in at time t = n. If n - m > 20, the system calculates initial cognate relevance values for electronic mail objects received between t = m and t = n. If n - m < 20, the system recalculates cognate relevance values for electronic mail objects received between t = n-20 and t = n. If the user is still active or has not logged out at time t = n+5, the system recalculates cognate relevance values for electronic mail objects received between t = n-10 and t = n+5. If the user is still active or has not logged out at time t = n+10, the system recalculates cognate relevance values for electronic mail objects received between t = n-5 and t = n+10. If the user is still active or has not logged out at time t = n+15, the system recalculates cognate relevance values for electronic mail objects received between t = n-45 and t = n+15. If the user is still active or has not logged out at time t = n+20, the system recalculates cognate relevance values for electronic mail objects received between t = n+5 and t = n+20. If the user is still active or has not logged out at time t = n+25, the system recalculates cognate relevance values for electronic mail objects received between t =

n+10 and t = n+25. If the user is still active or has not logged out at time t = n+30, the system recalculates cognate relevance values for electronic mail objects received between t = n-150 and t = n+30. If the user is still active or has not logged out at time t = n+35, the system recalculates cognate relevance values for electronic mail objects received between t = n+20 and t = n+35. If the user is still active or has not logged out at time t = n+40, the system recalculates cognate relevance values for electronic mail objects received between t = n+25 and t = n+40. If the user is still active or has not logged out at time t = n+45, the system recalculates cognate relevance values for electronic mail objects received between t = n-15 and t = n+45. If the user is still active or has not logged out at time t = n+50, the system recalculates cognate relevance values for electronic mail objects received between t = n+35 and t = n+50. If the user is still active or has not logged out at time t = n+55, the system recalculates cognate relevance values for electronic mail objects received between t = n+40 and t = n+55. If the user is still active or has not logged out at time t = n+60, the system recalculates cognate relevance values for electronic mail objects received between t = n-300 and t = n+60. If the user is still active or has not logged out at time t = n+65, the system recalculates cognate relevance values for electronic mail objects received between t = n+50 and t = n+65. If the user is still active or has not logged out at time t = n+70, the system recalculates cognate relevance values for electronic mail objects received between t = n+55 and t = n+70. If the user is still active or has not logged out at time t = n+75, the system recalculates cognate relevance values for electronic mail objects received between t = n+15 and t = n+75. If the user is still active or has not logged out at time t = n+80, the

system recalculates cognate relevance values for electronic mail objects received between t = n+65 and t = n+80. If the user is still active or has not logged out at time t = n+85, the system recalculates cognate relevance values for electronic mail objects received between t = n+70 and t = n+85. If the user is still active or has not logged out at time t = n+90, the system recalculates cognate relevance values for electronic mail objects received between t = n-90 and t = n+90. If the user is still active or has not logged out at time t = n+95, the system recalculates cognate relevance values for electronic mail objects received between t = n+80 and t = n+95. If the user is still active or has not logged out at time t = n+100, the system recalculates cognate relevance values for e-mal objects received between t = n+85 and t = n+100. If the user is still active or has not logged out at time t = n+105, the system recalculates cognate relevance values for electronic mail objects received between t = n+45 and t = n+105. If the user is still active or has not logged out at time t = n+110, the system recalculates cognate relevance values for electronic mail objects received between t = n+95 and t = n+110. If the user is still active or has not logged out at time t = n+115, the system recalculates cognate relevance values for electronic mail objects received between t = n+100 and t = n+115. If the user is still active or has not logged out at time t = n+120, the system recalculates cognate relevance values for electronic mail objects received between t = n+240 and t = n+120. Thus, at 5, 10, 20, 25, 35, 40, 50 and 55 minutes, the system recalculates the last 15 minutes worth of electronic mail cognate relevance values. At 15 and 45 minutes, the system recalculates the last 30 minutes worth of electronic mail cognate relevance values. At 30 minutes, the system recalculates the last 3 hours of electronic mail cognate relevance values. At 60 minutes,

the system recalculates the last 6 hours of electronic mail cognate relevance values.

[00214] When the user logs out or becomes inactive, the process is paused, waiting to recommence once the user logs back in or becomes active again. It is noted that the times described above are just examples. In practice, the periodicity of recalculation and time periods involved may depend on the global importance of the user, and may be also varied on a per implementation basis depending on the computational resources and priorities of a given system.

[00215] For organizations involved, for example, in forensic analysis, investigation, or intelligence, a greater level of content analysis may be appropriate. A relevance analysis engine in accordance with one or more embodiments may have the capability to go to greater levels of granularity than submessages, for example, by resolving paragraphs of text or individual words. This may involve the system itself using spiders to identify terms within the content of any electronic mail object that statistically increase the likelihood that it is relevant to the recipient or person performing discovery.

[00216] Keywords of high relevance in a community of users may be automatically discovered based on statistical techniques. The frequencies of words in a set of mail content may be compared to regular (average) values, to discover the words that are used more than average in the volume of electronic mail. The more frequent words may then be correlated to content relevance, to highlight more frequent words (compared to average values) that are used in the more relevant content.

[00217] Further, such techniques may allow for "dashboarding" and visualization of

sets of content (corresponding to, for example, communities or sub-communities), to provide insight into the nature of the community for an analyst or investigator. This may be particularly useful when very large amounts of mail communications are being analyzed, and content must be semi-automatically sampled for further investigation.

[00218] In certain cases, users may seek to find ways to increase their own global importance. Thus, for example, two friends on similar grades may arrange to send each other large numbers of electronic mail objects that they would leave open for a while – each therefore increasing the other's importance. Each can monitor the other's actions over time, and large numbers of read electronic mail objects may increase both their global importance values. At least one advantage of this to them is that both may then get their electronic mail objects read faster when sending to others. In another scenario, when the content analysis is advanced, senders may deliberately introduce phrases and words to get their electronic mail objects read. Accordingly, one or more embodiments may incorporate detection devices identifying such behavior, and to penalize, mitigate, and/or halt the behavior where and as appropriate. In addition, this type of behavior may be flagged to administrators inside an organization.

[00219] To further describe an electronic mail system in accordance with one or more embodiments, a series of stages are now described to illustrate a software implementation. In a first stage, local relevance of content may be based on some factors of the mail data (e.g., features inherent in or attached to an electronic mail object) and each individual user's actions on content in their own folders (effectively recording what each user does with each mail, and treating every mail separately). Analysis may be

based on the current submessage, and not retrospectively applied to historical submessages. Incoming mail relevance may be based on mail data contained within the electronic mail object itself (e.g., a flag for follow up or urgency flag).

[00220] In a next stage, in addition to local relevance based on IRC on this electronic mail object, each entity that a user deals with is given a local importance value derived from that user's actions – i.e., using the IIC. The input to local relevance is through a combined local importance value, considering all the senders involved in this particular electronic mail object. At this stage, there may be no communal aspects that would estimate importance using other user actions; the importance values of senders to a recipient are based on the past actions of that recipient. Analysis may still be based on entities in each electronic mail object's respective current submessage, and not changing importance values of entities in the history portion.

[00221] In a third stage, user opinions are sampled, to confirm that the relevance and importance values assigned are accurate reflections of each user's real opinions. Also, there may be a facility for optional feedback via the sliders into the algorithm (as opposed to the sliders being for information presentation only).

[00222] In a next stage, relevance values of historical, and implicitly "future," submessages are changed based on each user's actions on a current submessage. Then, in a next stage, changes are made to the local importance values of entities involved in historical (and future) submessages based on each user's actions on a current submessage. Subsequently, relevance and importance are integrated with search, each again on a single-recipient basis (not taking any communal actions or values of other

users into the implementation other than those directly implied in the mail data itself). In a next stage, relevance and importance values are decayed effectively, so that older content loses its weight over time unless it is re-accessed by the user. Thereafter, users are allowed to vary their own settings, including the impact different actions have on relevance and importance.

[00223] In a next stage, global importance values are calculated and periodically updated for each entity, based mainly on a combination of each user's actions (IIC), which indicate importance of the sender. This may be the first step in communal importance. There may be human resources input also to the global importance values. Incoming mail relevance may now be based also on global importance values to give meaningful global (or communal) relevance.

[00224] In a next stage, global relevance values are calculated for each piece of content, based on the overall result of each user's local relevance values. Initial relevance for content can then include input from the global values (as well as, for example, local importance), and global relevance can still form a component of each user's local relevance even after they have read the content. This may be the first step in communal relevance. Incoming mail relevance can now be based also on global relevance scores, which are derived from a weighted aggregate of local relevance scores.

[00225] In a next stage, collaborative search takes place, allowing users to search other users' content, according to the rules set by a compliance policy. In a subsequent stage, there may be integration with human resources to assist in assigning similar users, to allow an indication of when a user changes roles. Also, in another stage, correlation of

groups, domains, and other aliases with users is done, so that the system can imply a degree of change on users' other identities when an action causes one of their identities to change. In a next stage, there is integration with enterprise naming systems, such as job codes, so that similar types and categories can be used in calculating importance values for related entities.

[00226] In a next stage, the system identifies a cognate group of n other users most similar to a given user/recipient. This is done, directly or indirectly, based on past actions. Calculation of each user's local importance values for other entities is based on their own actions, and also actions of those in their cognate group. An individual importance component value (IIC) or individual value inputting into the cognate importance component (I_{sr}^{c}) may be assigned when a user has interacted with an entity.

[00227] In a next stage, a cognate relevance value is now input into the final figure for local relevance. This is based on the actions of those in the group cognate to the recipient on the particular electronic mail object. In one or more embodiments, this may operate at run-time, but, in practice, it may have to be on a periodic review that initially has a more rapid review period.

[00228] In a next stage, communal weighting adjustment (of the impacts of different actions as described above) is allowed in a feedback loop, inside, for example, a single organization. Further, in another stage, smaller fragments are analyzed, such as paragraphs and words, as subsets of the overall electronic mail, to determine importance for these entities as well. In a next stage, federation of action weighting and other system constants is facilitated, based on other organizations' data (possibly anonymously) to

improve the accuracy of one or more of the mechanisms described herein.

[00229] As described above, a cognate group may be an important aspect of determining mail relevance in accordance with one or more embodiments. In some cases, a single global value for relevance of a given piece of content may not be the most appropriate measure of its relevance to a given community. For example, an electronic mail object regarding human resource policies might be most relevant to the human resources department, and of less relevance to other users. To improve the accuracy of the algorithm, it may be better to compare a user to a set of similar or cognate users rather than to the global population. In regard to how such similarity or identification of the cognate group is determined, one or more of the following approaches may be used: role or group based; collaborative filtering; fingerprinting; cluster models; search-based methods; frequency analysis; and importance value correlation. Each of these different approaches, which may be used singly or in any combination, will now be described in turn.

[00230] In regard to the role or group based approach, determining similarity may be dealt with in part by the role-based component in the effective local relevance value for each end user. Components may also be included for groups, domains, and other sets of users (or other entities). Users in the same groups or roles may be assumed to be similar. For example, if a number of users in a role receive an electronic mail object, and other users in that role have already viewed it and deleted it, then it is likely to be of low relevance to another user in that role. Further, users within the same group or role may have their historical behaviors correlated, to see which users within each group do, in

fact, behave in a similar manner. This method compares historical relevance values for content that users in the group have received, with those of the user in question (for the same content). This technique for extending the basic mechanism is essentially communal filtering (further described below). It is important to properly recognize the group (and other) local relevance and importance values used for determining relevance for content sent to those identities, with their role in ascribing similarity between users. A resulting statement of similarity from role or group methods might be, for example, "other users who also work in the development team found this mail to be of high relevance."

[00231] As to communal filtering, this mechanism may use the behavior of other recipients of a given piece of content to determine the content's relevance to the user in question – communal relevance. However, rather than just using the global relevance value for the content, the application may correlate the past behavior of other users to the user in question to improve the relevance predictor. A resulting statement of similarity from communal filtering might be, for example, "other users whose relevance values typically correlate to the user in question (based on comparison for historical content) found this mail to be of high relevance."

[00232] In regard to fingerprinting, it adopts a historical comparison method similar to that of communal filtering, but reduces the sparseness of the matrices by focusing on a sample of content with large numbers of recipients. Fingerprinting may rely on analyzing users' relevance values for a set of data that has been sent to very large numbers of users. This then provides relevance values for a common set of data,

allowing a "fingerprint" of each user's relevance values to be calculated, and then compared to that of other users. User fingerprint correlation may then be used as a basis for predicting a given user's relevance values for a new piece of content, based on other similar user's relevance values for the same piece of content. Thus, as more users read a new piece of mail, the initial relevance values for that content for the remaining recipients should be predicted more accurately. The correlation between users' behavior for mails with few recipients and large numbers of recipients may vary from one community to another. However, the feedback mechanisms in the relevance algorithm allow the strength of this correlation to be tested, and the suitability of fingerprinting for use in predicting behavior for all mail content to be evaluated. It is noted that for mail content sent with very large numbers of recipients, fingerprinting may be a reasonably accurate method for determining similarity. While this description has focused on users and content, other entities may also be subject to fingerprinting. The same matrix with content in one dimension may also be applied to groups and domains. A resulting statement of similarity from fingerprinting might be, for example, "other users whose relevance values typically correlate to the user in question (based on comparison for a non-representative sample of historical content) found this mail to be of high relevance."

[00233] As to cluster models, they work by segmenting users (and other entities) into categories based on their actions or attributes. In some ways, this approach may be similar to (and partly contains) role and group based similarity, but it may not assume separate relevance and importance values are stored for each cluster. Examples include, but are not limited to, the following: users who reply frequently, who forward frequently,

who never delete mails, who send mails to large groups, who receive certain mail volumes, who belong to particular groups, who work in particular departments, etc.; job codes, which have a particular job code parent or type, which are at a certain point in the domain hierarchy, which have a particular retention period, which result in particular external actions, which have certain allowed readers (access rights), etc.; and content, which has a large chain length (number of history submessages), which is of a certain length, which has attachments of certain size or type, which has a large number of recipients, which contains certain keywords, which is linked to a particular job code, which has a particular sensitivity rating, which interacts with certain external systems, etc. Further, it is noted that cluster models may potentially be applied to both entities and content in determining similarity and using it as a predictor for user relevance. A resulting statement of similarity from cluster models might be, for example, "other users who also forward mails regularly found this mail to be of high relevance."

[00234] In regard to search-based methods, they may operate with respect to content, or as a filter to then perform correlative analysis of users' historical relevance values. Search-based methods with respect to content may treat the similarity problem as a search for similar content in the user's historical content store. They then compare the relevance values from that user for similar historical content to predict the relevance of the new content. This essentially moves the similarity question to being a function of third party applications, but still relies on the user's own actions to determine relevance. In another approach, using search-based methods as a filter for comparing other users' historical behavior, the application finds similar content to that in question and then

inspects the relevance values given to it by other users. The approach then uses these relevance values from other users for similar content as an input to determine the initial relevance of the given content for the user in question.

[00235] As to frequency analysis, this may be another mechanism for reducing data size by communal filtering determining cognate groups. In this method, each user may have a set of other users associated with them based not on similarity, but on frequency. The method may work by finding the users who have the most content in common with the user in question (e.g., the most frequent to:, cc: and potentially bcc: recipients). Further, this method may not imply similarity among, for example, 100 users, who have most content in common with the user in question. Rather, it may provide a rich (not sparse) matrix of content on which to search for correlations (and anti-correlations) in user relevance values. On a periodic basis, such as each month, the ten most infrequent users in common with a given user may be removed from the matrix, and the ten most frequent users in common (not already identified in the matrix) may be added. This acts to keep the matrix rich with common values. From an importance perspective, this mechanism may be used to compare a user's own importance values given to entities (e.g., job codes) to those of other users. These correlations may then be used to provide more accurate estimates of initial importance for when a user interacts with that job code for the first time (as opposed to being forced to rely on the global value).

[00236] To further illustrate frequency analysis, consider, for example, a user, A, and four other users, B, C, D and E, with whom A has a significant volume of content in common. The table below lists the relevance values ascribed to historical content,

electronic mail objects A through T, that user A has in common with most or all of the other users.

	User				
Content	*User A*	*User B*	*User C*	*User D*	*User E*
TransMail T	5	9	9	4	1
TransMail S	6	2	3	7	4
TransMail R	2	5	7		5
TransMail Q	8		4	6	6
TransMail P	4	7		2	
TransMail O	2	6	9	4	
TransMail N	6	9	1	5	8
TransMail M	9	6	3	8	2
TransMail L	6	3	2	5	5
TransMail K	4	1	4	1	6
TransMail J	7	5		5	6
TransMail I	2	2	3	2	
TransMail H	4	8	6	3	9
TransMail G	3		7		8
TransMail F	1	4	6	3	4
TransMail E	8	7	8	7	2
TransMail D	5	8		7	
TransMail C	7	2	6	6	5
TransMail B	3	9	5	4	8
TransMail A	2	6	8		4

The list of relevance values for historical content allows each user's historical behavior to be correlated with each other user. The resulting linear correlation values are presented in the table below.

Correlations					
	User A	*User B*	*User C*	*User D*	*User E*
User A	1	0.042675	-0.35891	0.754378	-0.28248
User B	0.042675	1	0.301382	0.154122	0.122467
User C	-0.35891	0.301382	1	-0.0957	-0.31789
User D	0.754378	0.154122	-0.0957	1	-0.43545
User E	-0.28248	0.122467	-0.31789	-0.43545	1

Correlation may be a statistical function that indicates the extent of the similarity

between two data sets. From the table above, it can be seen that user A behaves very differently to user B, often has opposite relevance values to user B, and ascribes very similar relevance values to the same content as user D. Thus, the behavior (and resulting relevance value) of user D on a new piece of content sent to both users A and D would be a good indicator of the likely relevance of that content to user A.

[00237] User behavior may evolve over time, so the correlation formula may be weighted towards the newer content (a non-linear distribution of the square of the error terms in the regression). In addition, there is no requirement that the relationship be linear; thus non-linear correlation techniques may also be appropriate. In a system of very high computational power, the regression formula for each similar pairing could be used to "forecast" a suitable value for the communal component.

[00238] Another way to determine cognate users relates to correlating importance values. An identification may be to find the n users whose matrices of importance values correlate most highly with those of the given user r. This approach assumes that the effects of other factors (e.g., user actions) have fed through into the importance figures. In one or more embodiments, a single mechanism for determining the identities in cognate groups may rely primarily on the correlation of importance values, and these may be reconsidered during computational quiet periods as described above.

[00239] As described above, approaches for determining cognate groups may be combined. For example, cluster models may be added as another dimension to frequency analysis. Further, it is noted that for one or more of the cognate group determination approaches described above, the contribution of more recent content to the similarity

function may be greater than that of older content, reflecting the changing nature of the user population as the organization and their roles within it evolve. As described above, an exponential function may be applied to these contributions to decay them over time. Further, older data may be removed as new data is added, to ensure that the similarity structures may be maintained at a given size. The frequency of updating various relevance and importance values may also be determined on a per organization basis, depending on, for example, the business requirements and computational power available.

[00240] As apparent from the description herein of an electronic mail system in accordance with one or more embodiments, various advantages may be achieved with the electronic mail system. In one aspect, the electronic mail system "watches" how users interact with content, as well as implicitly other users and entities in the system, and takes each user's behavior as the primary input into a relevance analysis engine of the electronic mail system. An attraction of this approach might be that it collects vast amounts of statistical input data to the engine, but does not require additional user time. Further, the electronic mail system may be applied at a single user level only, based on a user's own actions. It may also be applied across an entire organization in a manner looking communally at other users' actions, which greatly increases its accuracy by providing more contextual data. In another aspect, the electronic mail system may be integrated with traditional email systems. Further, the electronic mail relevance system may also be applied to historical content, providing "dashboard" summaries and visual representations of communication communities. In a further aspect, the electronic mail

relevance system may enhance an investigative process, improving sampling, and ease of analysis.

[00241] While the invention has been described with respect to a limited number of embodiments, those skilled in the art, having benefit of the above description, will appreciate that other embodiments may be devised which do not depart from the scope of the present invention as described herein. Accordingly, the scope of the present invention should be limited only by the appended claims.

CLAIMS

What is claimed is:

1. A system, comprising:

a monitoring module configured to track actions by a plurality of users associated with an application for managing electronic content;

a relevance analysis module coupled to the monitoring module and configured to determine a relevance to a user of a piece of electronic content sent from a sender to the user, the relevance being determined dependent on:

a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content, and

a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content.

2. The system of claim 1, wherein the relevance is further determined dependent on a measure of local importance describing an importance of the sender to the user.

3. The system of claim 1, wherein the relevance is further determined dependent on a measure of global importance describing an importance of the sender within an organization with which the plurality of users are associated.

4. The system of claim 1, wherein the relevance analysis module is further configured to identify at least one other of the plurality of users that is a recipient of the

piece of electronic content and is cognate to the user and wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content dependent upon one or more actions of the at least one cognate recipient with respect to the piece of electronic content.

5. The system of claim 1, wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content responsive to contents of the piece of electronic content.

6. The system of claim 1, wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content responsive to one or more measures from the set consisting of:

a local importance describing an importance of the sender to the user;

a global importance describing an importance of the sender within an organization with which the plurality of users are associated;

a content relevance determined responsive to contents of the piece of electronic content; and

one or more actions with respect to the piece of electronic content of one or more of the plurality of users determined to be cognate to the user.

7. The system of claim 1, further comprising:

a content presentation module configured to present the piece of electronic content to the user in accordance with the determined relevance and to present to the user a value describing the determined relevance.

8. The system of claim 1, further comprising:

a search module configured to search the electronic content in accordance with the determined relevance.

9. The system of claim 1, wherein the piece of electronic content comprises an electronic mail object in a relational mail system.

10. The system of claim 1, wherein the electronic content comprises at least one of an email, a document, an instant message, and an SMS message.

11. A computer program product having a computer-readable medium having computer program instructions encoded thereon comprising:

a monitoring module configured to track actions by a plurality of users associated with an application for managing electronic content;

a relevance analysis module coupled to the monitoring module and configured to determine a relevance to a user of a piece of electronic content sent from a sender to the user, the relevance being determined dependent on:

a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content, and

a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content.

12. The computer program product of claim 11, wherein the relevance is further determined dependent on a measure of local importance describing an importance of the sender to the user.

13. The computer program product of claim 11, wherein the relevance is further determined dependent on a measure of global importance describing an importance of the sender within an organization with which the plurality of users are associated.

14. The computer program product of claim 11, wherein the relevance analysis module is further configured to identify at least one other of the plurality of users that is a recipient of the piece of electronic content and is cognate to the user and wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content dependent upon one or more actions of the at least one cognate recipient with respect to the piece of electronic content.

15. The computer program product of claim 11, wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content responsive to contents of the piece of electronic content.

16. The computer program product of claim 11, wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content responsive to one or more measures from the set consisting of:

a local importance describing an importance of the sender to the user;

a global importance describing an importance of the sender within an organization with which the plurality of users are associated;

a content relevance determined responsive to contents of the piece of electronic content; and

one or more actions with respect to the piece of electronic content of one or more of the plurality of users determined to be cognate to the user.

17. A method, comprising:

tracking actions by a plurality of users associated with an application for managing electronic content;

determining a relevance to a user of a piece of electronic content sent from a sender to the user, the relevance being determined dependent on:

a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content, and

a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content.

18. The method of claim 17, wherein the relevance is further determined dependent on a measure of local importance describing an importance of the sender to the user.

19. The method of claim 17, wherein the relevance is further determined dependent on a measure of global importance describing an importance of the sender within an organization with which the plurality of users are associated.

20. The method of claim 17, wherein determining the relevance to the user further comprises:

identifying at least one other of the plurality of users that is a recipient of the piece of electronic content and is cognate to the user; and

determining the relevance to the user of the piece of electronic content dependent upon one or more actions of the at least one cognate recipient with respect to the piece of electronic content.

21. The method of claim 17, wherein the relevance is further dependent on contents of the piece of electronic content.

22. The method of claim 17, wherein the relevance is further dependent on one or more measures from the set consisting of:

a local importance describing an importance of the sender to the user;

a global importance describing an importance of the sender within an organization with which the plurality of users are associated;

a content relevance determined responsive to contents of the piece of electronic content; and

one or more actions with respect to the piece of electronic content of one or more of the plurality of users determined to be cognate to the user.

DETERMINING RELEVANCE OF ELECTRONIC CONTENT

ABSTRACT

In an electronic mail system, a user's actions with an application for managing electronic content are monitored. Based on the user's actions and/or on the actions of other users in the system, a relevance analysis module determines a relevance to the user of a piece of electronic content based on a likelihood of interest of the electronic content piece to the user and/or a measure of a level of significance the user has to another entity associated with the electronic content piece.

Electronic Patent Application Fee Transmittal	
Application Number:	
Filing Date:	
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Filer:	Brian M. Hoffman
Attorney Docket Number:	23682-12381
Filed as Small Entity	

Utility Filing Fees

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Utility filing Fee (Electronic filing)	4011	1	75	75
Utility Search Fee	2111	1	250	250
Utility Examination Fee	2311	1	100	100
Pages:				
Claims:				
Claims in excess of 20	2202	2	25	50
Independent claims in excess of 3	2201	3	100	300
Miscellaneous-Filing:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				
Miscellaneous:				
		Total in USD ($)		**775**

Electronic Acknowledgement Receipt	
EFS ID:	1442182
Application Number:	11623653
International Application Number:	
Confirmation Number:	5243
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Customer Number:	00758
Filer:	Brian M. Hoffman
Filer Authorized By:	
Attorney Docket Number:	23682-12381
Receipt Date:	16-JAN-2007
Filing Date:	
Time Stamp:	18:22:44
Application Type:	Utility

Payment information:

Submitted with Payment	yes
Payment was successfully received in RAM	$775
RAM confirmation Number	948
Deposit Account	192555
The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows: Charge any Additional Fees required under 37 C.F.R. Section 1.16 and 1.17	

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)	Multi Part /.zip	Pages (if appl.)
1	Drawings	12381_Drawings.pdf	3458830	no	22
Warnings:					
Information:					
2	Application Data Sheet	12381_ADS_Form_SB_14.pdf	46151	no	4
Warnings:					
Information:					
This is not an USPTO supplied ADS fillable form					
3		12381_Mail_Relevance_spec.pdf	419462	yes	106

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Specification	1	99
	Claims	100	105
	Abstract	106	106

Warnings:					
Information:					
4	Fee Worksheet (PTO-06)	fee-info.pdf	8546	no	2
Warnings:					
Information:					
			Total Files Size (in bytes):	3932989	

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.

Electronic Acknowledgement Receipt	
EFS ID:	1442182
Application Number:	11623653
International Application Number:	
Confirmation Number:	5243
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Customer Number:	00758
Filer:	Brian M. Hoffman
Filer Authorized By:	
Attorney Docket Number:	23682-12381
Receipt Date:	16-JAN-2007
Filing Date:	
Time Stamp:	18:22:44
Application Type:	Utility

Payment information:

Submitted with Payment	yes
Payment was successfully received in RAM	$775
RAM confirmation Number	948
Deposit Account	192555
The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows: Charge any Additional Fees required under 37 C.F.R. Section 1.16 and 1.17	

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)	Multi Part /.zip	Pages (if appl.)
1	Drawings	12381_Drawings.pdf	3458830	no	22
Warnings:					
Information:					
2	Application Data Sheet	12381_ADS_Form_SB_14.pdf	46151	no	4
Warnings:					
Information:					
This is not an USPTO supplied ADS fillable form					
3		12381_Mail_Relevance_spec.pdf	419462	yes	106

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Specification	1	99
	Claims	100	105
	Abstract	106	106

Warnings:					
Information:					
4	Fee Worksheet (PTO-06)	fee-info.pdf	8546	no	2
Warnings:					
Information:					
			Total Files Size (in bytes):	3932989	

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.

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	DRAWINGS	TOT CLAIMS	IND CLAIMS
11/623,653	01/16/2007	2161	475	23682-12381	22	22	3

CONFIRMATION NO. 5243

758
FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA94041

FILING RECEIPT

Date Mailed: 02/13/2007

Receipt is acknowledged of this regular Patent Application. It will be considered in its order and you will be notified as to the results of the examination. Be sure to provide the U.S. APPLICATION NUMBER, FILING DATE, NAME OF APPLICANT, and TITLE OF INVENTION when inquiring about this application. 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 mail to the Commissioner for Patents P.O. Box 1450 Alexandria Va 22313-1450. Please provide a copy of this Filing Receipt with the changes noted thereon. If you received a "Notice to File Missing Parts" for this application, please submit any corrections 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 (if appropriate).**

Applicant(s)

Justin Marston, Austin, TX;
Paul Marston, Lancashire, UNITED KINGDOM;

Assignment For Published Patent Application

BlueSpace Software Corporation, Austin, TX

Power of Attorney: None

Domestic Priority data as claimed by applicant

This appln claims benefit of 60/758,828 01/13/2006

Foreign Applications

If Required, Foreign Filing License Granted: 02/12/2007

The country code and number of your priority application, to be used for filing abroad under the Paris Convention, is **US11/623,653**

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

Non-Publication Request: No

Early Publication Request: No

** SMALL ENTITY **

Title

Preliminary Class

707

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 (1-866-999-4158).

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).

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	ATTORNEY DOCKET NUMBER
11/623,653	01/16/2007	Justin Marston	23682-12381

CONFIRMATION NO. 5243

758
FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041

FORMALITIES LETTER

Date Mailed: 02/13/2007

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 and pay any fees required 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 oath or declaration is missing. *A properly signed oath or declaration in compliance with 37 CFR 1.63, identifying the application by the above Application Number and Filing Date, is required.*
  Note: If a petition under 37 CFR 1.47 is being filed, an oath or declaration in compliance with 37 CFR 1.63 signed by all available joint inventors, or if no inventor is available by a party with sufficient proprietary interest, is required.

The applicant needs to satisfy supplemental fees problems indicated below.

The required item(s) identified below must be timely submitted to avoid abandonment:

• To avoid abandonment, a surcharge (for late submission of filing fee, search fee, examination fee or oath or declaration) as set forth in 37 CFR 1.16(f) of $65 for a small entity in compliance with 37 CFR 1.27, must be submitted with the missing items identified in this letter.

SUMMARY OF FEES DUE:

Total additional fee(s) required for this application is **$65** for a small entity

- **$65** Surcharge.

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.
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.

Office of Initial Patent Examination (571) 272-4000, or 1-800-PTO-9199, or 1-800-972-6382
PART 3 - OFFICE COPY

Electronic Patent Application Fee Transmittal					
Application Number:					
Filing Date:					
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT				
First Named Inventor/Applicant Name:	Justin Marston				
Filer:	Brian M. Hoffman				
Attorney Docket Number:	23682-12381				
Filed as Small Entity					
Utility Filing Fees					
Description		**Fee Code**	**Quantity**	**Amount**	**Sub-Total in USD($)**
Basic Filing:					
Utility filing Fee (Electronic filing)		4011	1	75	75
Utility Search Fee		2111	1	250	250
Utility Examination Fee		2311	1	100	100
Pages:					
Claims:					
Claims in excess of 20		2202	2	25	50
Independent claims in excess of 3		2201	3	100	300
Miscellaneous-Filing:					

Adjustment date: 02/09/2007 MNGUYEN
01/17/2007 INTEFSW 00000948 192555 11623653
05 FC:2201 300.00 CR

PTO/SB/06 (12-04)
Approved for use through 7/31/2006. 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.

PATENT APPLICATION FEE DETERMINATION RECORD
Substitute for Form PTO-875

11/623,653

APPLICATION AS FILED – PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA	SMALL ENTITY RATE ($)	SMALL ENTITY FEE ($)		OTHER THAN SMALL ENTITY RATE ($)	OTHER THAN SMALL ENTITY FEE ($)
BASIC FEE (37 CFR 1.16(a), (b), or (c))				75			
SEARCH FEE (37 CFR 1.16(k), (i), or (m))				250			
EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))				100			
TOTAL CLAIMS (37 CFR 1.16(i))	22 minus 20 =	2	X$25=	50	OR	X$50	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	3 minus 3 =		X$100=			X$200=	
APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $250 ($125 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR						
MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))			N/A			N/A	
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL	475		TOTAL	

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 ($)	SMALL ENTITY ADDITIONAL FEE ($)		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))					N/A		OR	N/A	
						TOTAL ADD'T FEE		OR	TOTAL ADD'T FEE	

		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)		RATE ($)	ADDITIONAL FEE ($)
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))					N/A		OR	N/A	
						TOTAL ADD'T FEE		OR	TOTAL ADD'T 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 number found in the appropriate box in column 1.

This collection of information is required by 37 CFR 1.16. 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 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.

IN THE UNITED STATES

PATENT AND TRADEMARK OFFICE

APPLICANT(S):	Justin Marston et al.
APPLICATION NO.:	11/623,653
FILING DATE:	January 16, 2007
TITLE:	Determining Relevance of Electronic Content
EXAMINER:	Not Yet Known
GROUP ART UNIT:	2161
ATTY. DKT. NO.:	23682-12381

MAIL STOP MISSING PARTS
COMMISSIONER FOR PATENTS
P.O. BOX 1450
ALEXANDRIA, VA 22313-1450

RESPONSE TO NOTICE TO FILE MISSING PARTS OF APPLICATION

SIR:

Responsive to the Notice to File Missing Parts dated February 13, 2007, received in the above-identified patent application,

☒ Enclosed are:

☒ a copy of the Notice to File Missing Parts; and

☒ an original, signed Declaration;

☒ Power of Attorney to Prosecute Before the USPTO;

☒ Statement Under 3.73(b); and

☒ payment in the amount of $65.00 for the

☒ missing parts surcharge.

☒ Applicant claims small entity status under 37 C.F.R. § 1.27.

Respectfully submitted,
JUSTIN MARSTON ET AL.

Dated: April 11, 2007

By:/Brian Hoffman/

Brian Hoffman, Reg. No.: 39,713
Fenwick & West LLP
Silicon Valley Center
801 California Street
Mountain View, CA 94041
Tel.: (415) 875-2484
Fax.: (415) 281-1350

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	ATTORNEY DOCKET NUMBER
11/623,653	01/16/2007	Justin Marston	23682-12381

CONFIRMATION NO. 5243

758
FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041

FORMALITIES LETTER

Date Mailed: 02/13/2007

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 and pay any fees required 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 oath or declaration is missing. *A properly signed oath or declaration in compliance with 37 CFR 1.63, identifying the application by the above Application Number and Filing Date, is required.*
  Note: If a petition under 37 CFR 1.47 is being filed, an oath or declaration in compliance with 37 CFR 1.63 signed by all available joint inventors, or if no inventor is available by a party with sufficient proprietary interest, is required.

The applicant needs to satisfy supplemental fees problems indicated below.

The required item(s) identified below must be timely submitted to avoid abandonment:

• To avoid abandonment, a surcharge (for late submission of filing fee, search fee, examination fee or oath or declaration) as set forth in 37 CFR 1.16(f) of $65 for a small entity in compliance with 37 CFR 1.27, must be submitted with the missing items identified in this letter.

SUMMARY OF FEES DUE:

Total additional fee(s) required for this application is **$65** for a small entity

- **$65** Surcharge.

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.
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.

Office of Initial Patent Examination (571) 272-4000, or 1-800-PTO-9199, or 1-800-972-6382
PART 1 - ATTORNEY/APPLICANT COPY

DECLARATION FOR UTILITY OR DESIGN PATENT APPLICATION (37 CFR 1.63)		
☐ Declaration Submitted With Initial Filing OR ☒ Declaration Submitted after Initial Filing (surcharge (37 CFR 1.16 (e)) required)	Attorney Docket Number	23682-12381
	First Named Inventor	Justin Marston
	COMPLETE IF KNOWN	
	Application Number	11/623,653
	Filing Date	January 16, 2007
	Art Unit	2161
	Examiner Name	Not yet known

I hereby declare that:

Each inventor's residence, mailing address, and citizenship are as stated below next to their name.

I believe the inventor(s) named below to be the original and first inventor(s) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

DETERMINING RELEVANCE OF ELECTRONIC CONTENT

(Title of the Application)

the specification of which

☐ is attached hereto

OR

☒ was filed on (MM/DD/YYYY) 01/16/2007 as United States Application Number or PCT International Application Number 11/623,653 and was amended on (MM/DD/YYYY) _____ (if applicable).

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment specifically referred to above.

I acknowledge the duty to disclose information which is material to patentability as defined in 37 CFR 1.56, including for continuation-in-part applications, material information which became available between the filing date of the prior application and the national or PCT international filing date of the continuation-in-part application.

I hereby claim foreign priority benefits under 35 U.S.C. 119(a)-(d) or (f), or 365(b) of any foreign application(s) for patent, inventor's or plant breeder's rights certificate(s), or 365(a) of any PCT international application which designated at least one country other than the United States of America, listed below and have also identified below, by checking the box, any foreign application for patent, inventor's or plant breeder's rights certificate(s), or any PCT international application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application Number(s)	Country	Foreign Filing Date (MM/DD/YYYY)	Priority Not Claimed	Certified Copy Attached? Yes	No
			☐	☐	☐
			☐	☐	☐
			☐	☐	☐
			☐	☐	☐

☐ Additional foreign application numbers are listed on a supplemental priority data sheet PTO/SB/02B attached hereto.

DECLARATION -- Utility or Design Patent Application

Direct all correspondence to: ☒ Customer Number: **758** OR ☐ Correspondence address below			
Name			
Address			
City	State	ZIP	
Country	Telephone	Fax	

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 18 U.S.C. 1001 and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

NAME OF SOLE OR FIRST INVENTOR:	☐ A petition has been filed for this unsigned inventor
Given Name (first and middle [if any]): **Justin**	Family Name or Surname: **Marston**
Inventor's Signature: *Justin Marston*	Date: 5 April 07

Residence: City	State	Country	Citizenship
Austin	**TX**	**US**	**UK**

Mailing Address: **P.O. Box 30570**			
City: **Austin**	State: **TX**	ZIP: **78755**	Country: **US**

NAME OF SECOND INVENTOR:	☐ A petition has been filed for this unsigned inventor
Given Name (first and middle [if any]): **Paul**	Family Name or Surname: **Marston**
Inventor's Signature: *Paul Marston*	Date: 4 April 07

Residence: City	State	Country	Citizenship
Lancashire		**UK**	**UK**

Mailing Address: **143 Longmeanygate, Midge Hall**			
City: **Lancashire**	State:	ZIP: **PR26 6TD**	Country: **UK**

☐ Additional inventors or a legal representative are being named on additional sheet(s)

23682/12381/SF/5193175.1

POWER OF ATTORNEY TO PROSECUTE APPLICATIONS BEFORE THE USPTO

I hereby revoke all previous powers of attorney given in the application identified in the attached statement under 37 CFR 3.73(b)

I hereby appoint:

☒ Practitioners associated with the Customer Number 758

OR

☐ Practitioner(s) named below (if more than ten patent practitioners are to be named, then a customer number must be used):

Name	Registration Number

as attorney(s) or agent(s) to represent the undersigned before the United States Patent and Trademark Office (USPTO) in connection with any and all patent applications assigned only to the undersigned according to the USPTO assignment records or assignment documents attached to this form in accordance with 37 CFR 3.73(b).

Assignee Name and Address:

BlueSpace Software Corporation
P.O. Box 30570
Austin, TX 78755

A copy of this form, together with a statement under 37 CFR 3.73(b) (Form PTO/SB/96 or equivalent) is required to be filed in each application in which this form is used. The statement under 37 CFR 3.73(b) may be completed by one of the practitioners appointed in this form if the appointed practitioner is authorized to act on behalf of the assignee, and must identify the application in which this Power of Attorney is to be filed.

SIGNATURE of Assignee of Record
The individual whose signature and title is supplied below is authorized to act on behalf of the assignee

Name	Justin Marston		
Signature	Justin Marston	Date	9/9/6
Title	President and CEO	Telephone	(512) 386-3940

STATEMENT UNDER 37 CFR 3.73(b)

Applicant/Patent Owner: Justin Marston et al.

Application No./Patent No.: 11/623,653 Filed/Issue Date: January 16, 2007

Entitled: Determining Relevance of Electronic Content

BlueSpace Software Corporation, a Corporation
(Name of Assignee) (Type of Assignee, e.g., corporation, partnership, university, government agency, etc.)

states that it is:

1. ☒ the assignee of the entire right, title, and interest; or

2. ☐ an assignee of less than the entire right, title and interest.
The extent (by percentage) of its ownership interest is _____%
in the patent application/patent identified above by virtue of either:

A. [X] An assignment from the inventor(s) of the patent application/patent identified above. The assignment was recorded in the United States Patent and Trademark Office at Reel 019142, Frame 0251, or for which a copy thereof is attached and is concurrently submitted for recordation.

OR

B. [ ] A chain of title from the inventor(s), of the patent application/patent identified above, to the current assignee as shown below:

1. From _____ To: _____
The document was recorded in the United States Patent and Trademark Office at Reel _____, Frame _____, or for which a copy thereof is attached.

2. From _____ To: _____
The document was recorded in the United States Patent and Trademark Office at Reel _____, Frame _____, or for which a copy thereof is attached.

3. From _____ To: _____
The document was recorded in the United States Patent and Trademark Office at Reel _____, Frame _____, or for which a copy thereof is attached.

[ ] Additional documents in the chain of title are listed on a supplemental sheet.

[] Copies of assignments or other documents in the chain of title are attached.
[NOTE: A separate copy (*i.e.*, a true copy of the original document) must be submitted to Assignment Division in accordance with 37 CFR Part 3, if the assignment is to be recorded in the records of the USPTO. See MPEP 302.08]

The undersigned (whose title is supplied below) is authorized to act on behalf of the assignee.

April 11, 2007 — Date
Brian M. Hoffman — Typed or printed name

(415) 875-2484 — Telephone number
/Brian Hoffman/ — Signature

Attorney for Assignee — Title

Electronic Patent Application Fee Transmittal

Application Number:	11623653
Filing Date:	16-Jan-2007
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Filer:	Brian M. Hoffman
Attorney Docket Number:	23682-12381

Filed as Small Entity

Utility Filing Fees

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Late filing fee for oath or declaration	2051	1	65	65
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Miscellaneous:				
			Total in USD ($)	**65**

Electronic Acknowledgement Receipt	
EFS ID:	1672228
Application Number:	11623653
International Application Number:	
Confirmation Number:	5243
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Customer Number:	758
Filer:	Brian M. Hoffman
Filer Authorized By:	
Attorney Docket Number:	23682-12381
Receipt Date:	11-APR-2007
Filing Date:	16-JAN-2007
Time Stamp:	16:53:12
Application Type:	Utility

Payment information:

Submitted with Payment	yes
Payment was successfully received in RAM	$65
RAM confirmation Number	627
Deposit Account	192555
The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows: Charge any Additional Fees required under 37 C.F.R. Section 1.16 and 1.17	

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)	Multi Part /.zip	Pages (if appl.)
1	Applicant Response to Pre-Exam Formalities Notice	12381_Response_to_Notice_of_Missing_Parts.pdf	95528	no	4
Warnings:					
Information:					
2	Oath or Declaration filed	12381_Declaration.pdf	87228	no	2
Warnings:					
Information:					
3		12381_POA.pdf	78866	yes	2
	Multipart Description/PDF files in .zip description				
	Document Description		**Start**	**End**	
	Power of Attorney		1	1	
	Assignee showing of ownership per 37 CFR 3.73(b).		2	2	
Warnings:					
Information:					
4	Fee Worksheet (PTO-06)	fee-info.pdf	8151	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	269773		

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
11/623,653	01/16/2007	2161	540	23682-12381	22	3

CONFIRMATION NO. 5243

UPDATED FILING RECEIPT

758
FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA94041

Date Mailed: 04/24/2007

Receipt is acknowledged of this regular Patent Application. It will be considered in its order and you will be notified as to the results of the examination. Be sure to provide the U.S. APPLICATION NUMBER, FILING DATE, NAME OF APPLICANT, and TITLE OF INVENTION when inquiring about this application. 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 mail to the Commissioner for Patents P.O. Box 1450 Alexandria Va 22313-1450. Please provide a copy of this Filing Receipt with the changes noted thereon. If you received a "Notice to File Missing Parts" for this application, please submit any corrections 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 (if appropriate).**

Applicant(s)

Justin Marston, Austin, TX;
Paul Marston, Lancashire, UNITED KINGDOM;

Assignment For Published Patent Application

BlueSpace Software Corporation, Austin, TX

Power of Attorney: The patent practitioners associated with Customer Number 758

Domestic Priority data as claimed by applicant

This appln claims benefit of 60/758,828 01/13/2006

Foreign Applications

If Required, Foreign Filing License Granted: 02/12/2007

The country code and number of your priority application, to be used for filing abroad under the Paris Convention, is **US11/623,653**

Projected Publication Date: 08/02/2007

Non-Publication Request: No

Early Publication Request: No

** SMALL ENTITY **

Title

Preliminary Class

707

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 (1-866-999-4158).

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).

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
11/623,653	01/16/2007	Justin Marston	23682-12381

CONFIRMATION NO. 5243

758
FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA94041

Title: DETERMINING RELEVANCE OF ELECTRONIC CONTENT

Publication No. US-2007-0179945-A1
Publication Date: 08/02/2007

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 Office of Public Records. The Office of Public Records can be reached by telephone at (703) 308-9726 or (800) 972-6382, by facsimile at (703) 305-8759, by mail addressed to the United States Patent and Trademark Office, Office of Public Records, 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 http://pair.uspto.gov/. 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.

Pre-Grant Publication Division, 703-605-4283

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICANTS:	Justin P. Marston et al.
APPLICATION NO.:	11/623,653
FILING DATE:	January 16, 2007
TITLE:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
EXAMINER:	Jeffrey A. Gaffin
GROUP ART UNIT:	2165
ATTY. DKT. NO.:	23682-12381

CERTIFICATE OF ELECTRONIC (EFS-WEB) TRANSMISSION

I hereby certify that this correspondence is being transmitted via the Office electronic filing system in accordance with 37 C.F.R. § 1.8(a)(i)(C) from the Pacific Time Zone of the United States on the local date shown below.

Dated: _____ By: _____

Brian M. Hoffman, Reg. No. 39,713

INFORMATION DISCLOSURE STATEMENT
Under 37 CFR §§ 1.56 and 1.97-98

SIR:

Pursuant to the provisions of 37 CFR §§ 1.56 and 1.97-98, enclosed herewith is modified form PTO/SB/08A listing references for consideration by the Examiner.

The filing of this Information Disclosure Statement shall not be construed as a representation regarding the completeness of the list of references, or that inclusion of a reference in this list is an admission that it is prior art or is pertinent to this application, or that a search has been made, or as an admission that the information listed is, or may be considered to be, material to patentability, or that no other material information exists, and shall not be construed as an admission against interest in any manner.

This Information Disclosure Statement is being filed:

- ☒ within three months of the filing date of the application, or date of entry into the national stage of an international application, or before the mailing date of a first office action on the merits, whichever event last occurred;
- ☐ before the mailing of a first official action after the filing of a request for continued examination (RCE) under 37 CFR § 1.114;
- ☐ after three months of the filing date of this national application or the date of entry of the national stage in an international application, or after the

mailing date of the first official action on the merits, whichever event last occurred, but before the mailing date of the first to occur of either: (1) a final action under 37 CFR §1.113; or (2) an action that otherwise closes prosecution in the application, and:

- ☐ attached hereto is the fee set forth under 37 CFR §1.17(p) for submission of this Information Disclosure Statement under 37 CFR.§ 1.97(c); OR
- ☐ Applicant certifies pursuant to 37 CFR § 1.97(e) that:
  - ☐ each item of information contained in this Information Disclosure Statement was first cited in a communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of this Statement; OR
  - ☐ no item of information contained in this 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 this certification after making reasonable inquiry, no item of information contained in this Statement was known to any individual designated under 37 CFR § 1.56(c) more than three months prior to the filing of this Statement;

☐ on or before the payment of the issue fee but after the mailing date of the first to occur of either: (1) a final action under 37 CFR § 1.113; (2) a notice of allowance under 37 CFR § 1.311; or (3) an action that otherwise closes prosecution in the application, and:

- ☐ Applicant certifies pursuant to 37 CFR. § 1.97(e) that:
  - ☐ each item of information contained in this Information Disclosure Statement was cited in a communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of this Statement;
  - ☐ no item of information contained in this 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 this certification after making reasonable inquiry, no item of

information contained in this Statement was known to any individual designated under 37 CFR § 1.56(c) more than three months prior to the filing of this Statement; AND

☐ attached hereto is the fee set forth under 37 CFR §1.17(p) for submission of this Information Disclosure Statement under 37 CFR.§ 1.97(d); OR

☐ after the payment of the issue fee. Applicant requests that the information contained in this Information Disclosure Statement be placed in the file according to 37 CFR § 1.97(i), although the information may not be considered by the USPTO.

☐ Enclosed is a copy of each listed reference that may be material to the examination of this application, and for which there may be a duty to disclose.

☐ This application relies, under 35 U.S.C. § 120, on the earlier filing date of prior application No. _____ , filed on _____ , and the references cited therein are hereby referenced, but are not required to be provided in this application under 37 CFR § 1.98(d).

☒ Copies of any foreign patent documents and non-patent literature cited herein are enclosed.

☐ Each item of information contained in this Information Disclosure Statement was cited in a communication from a foreign patent office in a counterpart application, and the communication was not received by any individual designated in 37 CFR § 1.56(c) more than thirty days prior to the filing of this Information Disclosure Statement. 37 CFR § 1.704(d).

☒ Applicant submits that no fee is required for the consideration of this Information Disclosure Statement.

Consideration of the listed references and favorable action are solicited.

Respectfully submitted,
JUSTIN P. MARSTON ET AL.

Dated: November 29, 2007

By: /Brian Hoffman/
Brian M. Hoffman, Reg. No. 39,713
Attorney for Applicant
Fenwick & West LLP
801 California Street
Mountain View, CA 94041
Tel.: (415) 875-2484
Fax: (415) 281-1350

Substitute for form 1449A/PTO				Complete if Known	
INFORMATION DISCLOSURE STATEMENT BY APPLICANT				**Application No.**	11/623,653
				Filing Date	January 16, 2007
				First Named Inventor	Justin Marston
				Art Unit	2165
				Examiner Name	Jeffrey A. Gaffin
Sheet	1	of	1	Attorney Docket Number	23682-12381

U.S. PATENT DOCUMENTS				
		Document No.		
Examiner Initials*	Cite No.[1]	Number – Kind Code[2] (if known)	Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document
	A1	US-2004/0133645 A1	07-08-2004	Massanelli et al.
	A2	US-2003/0191806 A1	10-09-2003	Osterberg, Jr., et al.
	A3	US-2002/0136279 A1	09-26-2002	Al-Kazily et al.
	A4	US-2002/0122543 A1	09-05-2002	Rowen

FOREIGN PATENT DOCUMENTS					
		Foreign Patent Document			
Examiner Initials*	Cite No.[1]	Country Code[3] – Number[4] Kind Code[5] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	T[6]
	B1	EP 1 085 444 A2	03-21-2001	Microsoft Corporation	
	B2	EP 1 085 444 A3	04-09-2003	Microsoft Corporation	

OTHER REFERENCES – NON-PATENT LITERATURE DOCUMENTS			
Examiner Initials*	Cite No.[1]	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, page(s), volume-issue number(s), publisher, city and/or country where published	T[6]
	C1	International Search Report and Written Opinion, PCT/US06/24859, April 26, 2007, 7 pages.	
	C2	EDWARDS, K., "The Design and Implementation of the MONTAGE Multimedia Mail System," Communications for Distributed Applications and Systems, Chapel Hill, April 18, 1991, pp. 47-57, Proceedings of the Conference on Communications Software, New York, U.S.	

Examiner Signature		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.

[1]Applicant's unique citation designation number (optional). [2]See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3]Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4]For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5]Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST. 16 if possible. [6]Applicant is to place a check mark here if English language Translation is attached.

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 085 444 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
21.03.2001 Bulletin 2001/12

(51) Int. Cl.[7]: **G06F 17/60**

(21) Application number: **00308239.3**

(22) Date of filing: **20.09.2000**

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **20.09.1999 US 399417**

(71) Applicant: **MICROSOFT CORPORATION**
Redmond, Washington 98052-6399 (US)

(72) Inventors:
- **Ahmed, Muhammad A**
  NO 202 Seattle, WA 98105 (US)
- **Alam, Muhammad Shabbir**
  NO 2801 Bellevue, Washington 98007 (US)

(74) Representative:
Orr, William McLean
URQUHART-DYKES & LORD
5th Floor, Tower House
Merrion Way
Leeds West Yorkshire, LS2 8PA (GB)

(54) Thread based e-mail

(57) Systems and methods for providing electronic messaging services to multiple users by storing a single copy of an electronic message at a central location and notifying recipients of the stored single copy. An electronic message includes a distribution list and a message content. A distribution list identifying multiple recipients causes prior art systems to duplicate the entire message for each recipient, placing potentially large demands on both processing power and storage space. In contrast, the systems and methods disclosed herein store a single copy or a limited number of copies of an electronic message addressed to multiple recipients and provide each recipient with a relatively small notification. In addition to providing information regarding content and origin, the notification also provides access to the stored message. Furthermore, the methods and systems also aid in organizing replies to electronic messages. Replies are associated with an initial message through a message identifier. The association helps to organize electronic messages by subject and provides context without requiring an author to duplicate the content of the initial message with the reply.

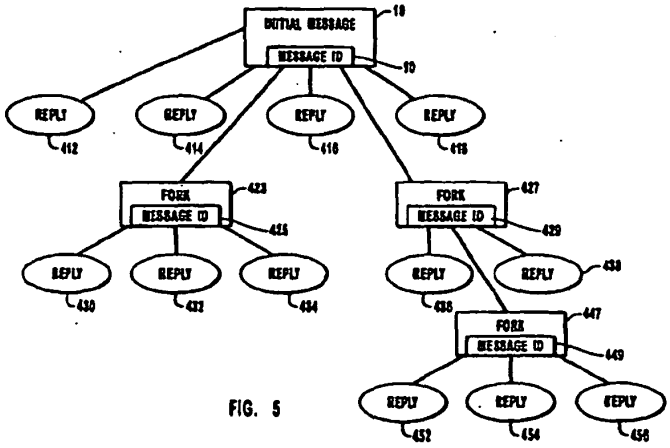


FIG. 5

EP 1 085 444 A2

Description

BACKGROUND OF THE INVENTION

1. The Field of the Invention

[0001] The present invention relates to methods and systems for delivering and organizing electronic messages. More specifically, the present invention relates to methods and systems that may be used to reduce the computing resources necessary for delivering a single message to multiple recipients and to improve the organization of responses to the message.

2. The Prior State of the Art

[0002] Traditional electronic messaging systems are a huge improvement over the paper analogues they augment or replace. By reducing the human and material resources necessary to transcribe, photocopy, and distribute paper communications, electronic messaging allows for resources to be focused on core business objectives rather than wasted on overhead. However, the contribution of prior art electronic messaging systems can be characterized as replacing physical inefficiencies with electronic ones.

[0003] Specifically, prior art electronic messaging systems distribute a single message to multiple recipients by duplicating the message and delivering a copy to each recipient. The two major inefficiencies that result include both processing time to duplicate the message and storage space to hold the copies. For example, consider what occurs when prior art systems deliver a three kilobyte message to all three thousand employees of a particular business. The prior art electronic messaging systems must dedicate processing resources to duplicate the three kilobyte message three thousand times, producing some nine megabytes of data that must be stored in the computer systems of the business.

[0004] Admittedly, with today's computer systems, duplicating a file three thousand times and storing nine megabytes of data are minor concerns, when viewed in isolation. However, several other factors render this drain on computer resources more serious. For example, electronic messaging systems operate on many individual messages. Duplicating a single message three thousand times is the equivalent of receiving three thousand messages at once. When combined with the normal operating conditions of the electronic messaging system, a sudden requirement to process three thousand messages can dramatically reduce system performance.

[0005] This problem is further complicated when recipients begin replying to widely distributed electronic messages. Because replies may be difficult to understand in isolation, the original message is often copied with the reply to avoid misunderstandings. Furthermore, the replies are often distributed to all original recipients, multiplying the initial nine megabytes of data by a factor of two and adding another three thousand messages the size of the reply.

[0006] However, copying the original message with a reply provides only minimal context. By its very nature, recipients will read and respond to a widely distributed electronic message over a substantial period of time. Some recipients may be out of the office, for hours at a meeting or days on vacation, when certain electronic messages are delivered. Some recipients will be involved in other work when the message is delivered and postpone reading and/or response to a more convenient time. Others, however, will read and respond immediately. This response time delay leads prior art electronic messaging system to intersperse responses with unrelated messages received in the interim. Moreover, if a specific electronic message generates a large amount of discussion, it may be advantageous to divide an initial topic into several subtopics. However, prior art electronic messaging systems provide only limited support for organizing electronic messages by topic, such as rules for sorting messages by sender.

[0007] Furthermore, electronic messaging typically provides only limited formatting capabilities. Many messages, however, require fairly sophisticated presentation. Enhanced formatting may be important as a means of efficiently communicating information. Perhaps a table of data quickly conveys the message that would otherwise require significant amounts of text. Alternatively, higher degrees of formatting generally connote an increased level of importance or authority. The effect of requiring a high degree of formatting often leads to electronic messages containing attachments of word processing documents that provide the increased formatting capabilities. Unfortunately, the formatting information of a document may contribute as much to a document's size as the text of the document itself and introduces a certain amount of overhead. Thus, three kilobytes is an extremely conservative size for an email message with an attached word processing document. It is not uncommon for even simple word processing documents to be in the range of ten to twenty kilobytes. A twenty kilobyte attachment sent to three thousand recipients with a single reply that includes the original message represents 120 megabytes of data - for a single communication. The problem may be at its worst when a fifty kilobyte image or 100 kilobyte program is involved.

[0008] Because the copies are electronic, little consideration is given to the resources that are required for delivering electronic messages. It is a simple matter for the average person to recognize the effort in physically duplicating something three thousand times. Few if any would even consider making a short (or even humorous) comment in a paper system that requires three thousand copies of the comment, three thousand duplicates of the original message, time to attach the com-

ment to the duplicates, and manual distribution to three thousand people. However, in an electronic messaging context, senders give little consideration to the computing resources necessary for processing widely distributed messages and the possible impact of replies.

[0009] Nevertheless, some prior art electronic messaging systems provide ill-suited solutions to the processing and storage problems introduced by multiple recipient messages. For example, public folders and newsgroups both provide a mechanism for placing information in a common area where it can be accessed by multiple requestors without being duplicated for each. While newsgroups and public folders provide a topical arrangement of information, they lack the ability to direct messages to an easily modifiable group of recipients. Specifically, newsgroups and public folders require a user to post information and then depend on those who are interested in the topic to periodically check for new messages. This operation renders newsgroups and public folders virtually useless for ad hoc electronic discussions.

[0010] Moreover, newsgroups and public folders provide limited capabilities for controlling access to posted information. For example, access to newsgroups and public folders can be controlled by users, but access cannot be tailored on a message by message basis. This limitation imposes a need for newsgroups and public folders to be created based on both topic and anticipated distribution, leading to a proliferation of newsgroups and public folders that must be checked periodically for new messages.

[0011] Finally, growing enthusiasm for the Internet further contributes to the problems in the prior art because an increasing number of people have access to electronic messaging services. This increased access enhances the ability to interact with one another and increases the likelihood that a single message will be meaningful to multiple recipients, further aggravating processing and storage problems. The prior art demands either sacrificing performance, limiting distribution, or designing systems to support the inefficiencies at their peak demand.

SUMMARY AND OBJECTS OF THE INVENTION

[0012] The problems in the prior state of the art have been successfully overcome by the present invention which is directed to methods and systems for improved distribution of electronic messages directed to multiple recipients and improved organization of replies to an electronic message. The improvement eliminates inefficiencies in processing power and storage space caused by duplicating electronic messages and provides enhanced organization by grouping replies with the original message.

[0013] More specifically, in one embodiment of the present invention, a host system providing electronic messaging services efficiently distributes electronic messages to multiple recipients. In contrast to the prior art's need to duplicate a message for each recipient, the host system stores as few as one copy of the electronic message, notifies each recipient of the stored message, and includes with the notification a link for accessing the stored message. Therefore, the host system saves both the processing time that would otherwise be used for duplicating and the corresponding storage space required to hold each of the duplicated messages.

[0014] Alternatively, the invention can be practiced using multiple host systems, each providing electronic messaging services to a particular set of clients. With multiple host systems, the host system originating a message creates a list identifying host systems providing electronic messaging services for each recipient and a list of recipients for each of the identified host systems. Then, for each host system, the originating system sends a host-specific list of recipients and a copy of the electronic message. However, in a multiple host system environment, an originating host system may determine that another host system providing electronic messaging service to certain recipients is unsupported in that it does not provide the full functionality of the invention disclosed herein. Upon identifying an unsupported host system, the originating host system sends the electronic message using another protocol that is compatible with the unsupported host system.

[0015] Yet another embodiment of the present invention provides for organizing related electronic messages. Replies may be associated with an earlier message by assigning an identifier to the initial electronic message. This association enhances a client's ability to establish the context of a reply. For example, assigning an identifier allows for the automated creation of a message-specific storage mechanism to hold the original message as well as all subsequent replies to that initial message. As electronic messages with identifiers are distributed, at least one storage mechanism for each initial message is created. Electronic messages that are replies are associated with their corresponding initial message by being placed in the storage mechanism previously created for each particular initial message.

[0016] The storage mechanism for initial messages may also comprise means for indicating the presence of electronic messages that have not been opened by the client. When a client selects a particular storage mechanism, only electronic messages associated with the message identifier will be presented. Within each storage mechanism, replies may be organized in a tree, hierarchical (parent-child), or other suitable arrangement; replies may be further ordered by time of receipt. Associating replies in this manner also allows for the elimination of common text among messages. There is no motivation for a reply to duplicate an earlier message within the reply itself because it is a relatively simple matter to access the earlier message.

[0017] Although presented separately, the aspects of the present invention summarized above should not

be considered as mutually exclusive. Rather, when combined, these aspects cooperate to greatly enhance the art of electronic messaging. Furthermore, some attributes are common among diverse embodiments of the present invention. For example, the ability to alter a distribution list for granting or denying access to one or more recipients, message specific rules that determine the lifetime of a particular message or control access to the message, and host-specific rules that provide a similar function all may enhance the embodiments described above.

[0018] It is, therefore, an object of the present invention to provide methods and systems for reducing both the processing and storage space resources associated with electronically sending a message to multiple recipients.

[0019] It is a further object of the present invention to provide methods and systems for tracking and dynamically controlling access to an electronic message.

[0020] It is yet another object of the present invention to provide methods and systems for organizing related electronic messages.

[0021] It is also an object of the present invention to provide methods and systems capable of interacting with host systems that do not operate according to the present invention.

[0022] Additional objects and advantages of the invention will be set forth in the description which follows, and in part will be obvious from the description, or may be learned by practicing the invention. The objects and advantages of the invention may be realized and obtained by means of the instruments and combinations particularly pointed out in the appended claims. These and other objects and features of the present invention will become more fully apparent from the following description and appended claims, or may be learned by the practice of the invention as set forth hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0023] In order that the manner in which the above-recited and other advantages and objects of the invention are obtained, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments thereof which are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

Figure 1 is an example system that provides a suitable operating environment for the present invention;

Figure 2 illustrates the basic elements of one embodiment of an electronic message;

Figure 3 illustrates the basic components of a host system that provides electronic messaging services;

Figure 4 illustrates how a preferred embodiment of the present invention processes electronic messages within a single host system environment;

Figure 5 illustrates how replies are associated with original messages; and

Figure 6 illustrates how a preferred embodiment of the present invention processes electronic messages within a multiple host system environment where one host system providing electronic messaging services does not operate according to the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0024] The present invention addresses the problems that prior art electronic messaging services face when distributing a single message to multiple recipients. Rather than expending processing power to duplicate the electronic message for each recipient and rather than dedicating storage resources to retain each of the newly made copies, the present invention stores a single copy of the electronic message and sends a relatively small notification to each recipient. The notification provides access to the stored copy of the electronic message and may also include certain information regarding its content and origin.

[0025] The present invention also aids in organizing replies to an electronic message. By assigning a message identifier, each reply can be associated with the message that prompted the response. The association is helpful to a recipient because it provides context for the reply without requiring its author to duplicate the initial message with the response. Furthermore, the context established through association is applicable to each of possibly many replies, whereas, without association, a recipient would need to separately establish context for each reply received.

[0026] The invention is described below by using diagrams to illustrate either the structure or processing of embodiments used to implement the methods and systems of the present invention. Using the diagrams in this manner to present the invention should not be construed as limiting its scope. The present invention contemplates both methods and systems for processing electronic messages. The embodiments of the present invention may comprise a special purpose or general purpose computer comprising various computer hardware.

[0027] Embodiments within the scope of the present invention also include computer-readable media having computer-executable instructions or data structures stored thereon. Such computer-readable media can be any available media which can be accessed by a general purpose or special purpose

computer. By way of example, and not limitation, such computer-readable media can comprise RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired executable instructions or data structures and which can be accessed by a general purpose or special purpose computer. When information is transferred or provided over a network or other communications connection to a computer, the computer properly views the connection as a computer-readable medium. Thus, such a connection is also properly termed a computer-readable medium. Combinations of the above should also be included within the scope of computer-readable media. Computer-executable instructions comprise, for example, instructions and data which cause a general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions. The computer-executable instructions and associated data structures represent an example of program code means for executing the steps of the invention disclosed herein.

[0028] Figure 1 and the following discussion are intended to provide a brief, general description of a suitable computing environment in which the invention may be implemented. Although not required, the invention will be described in the general context of computer-executable instructions, such as program modules, being executed by a personal computer. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Moreover, those skilled in the art will appreciate that the invention may be practiced with other computer system configurations, including hand-held devices, multi-processor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe computers, and the like. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

[0029] With reference to Figure 1, an exemplary system for implementing the invention includes a general purpose computing device in the form of a conventional computer 20, including a processing unit 21, a system memory 22, and a system bus 23 that couples various system components including the system memory to the processing unit. The system bus 23 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. The system memory includes read only memory (ROM) 24 and random access memory (RAM) 25. A basic input/output system (BIOS) 26, containing the basic routines that help to transfer information between elements within the computer 20, such as during start-up, may be stored in ROM 24. The computer 20 may also include a magnetic hard disk drive 27 for reading from and writing to a magnetic hard disk, not shown, a magnetic disk drive 28 for reading from or writing to a removable magnetic disk 29, and an optical disk drive 30 for reading from or writing to removable optical disk 31 such as a CD-ROM or other optical media. The magnetic hard disk drive 27, magnetic disk drive 28, and optical disk drive 30 are connected to the system bus 23 by a hard disk drive interface 32, a magnetic disk drive-interface 33, and an optical drive interface 34, respectively. The drives and their associated computer-readable media provide nonvolatile storage of computer readable instructions, data structures, program modules and other data for the computer 20. Although the exemplary environment described herein employs a magnetic hard disk 27, a removable magnetic disk 29 and a removable optical disk 31, it should be appreciated by those skilled in the art that other types of computer readable media which can store data that is accessible by a computer, such as magnetic cassettes, flash memory cards, digital video disks, Bernoulli cartridges, random access memories (RAMs), read only memories (ROM), and the like, may also be used in the exemplary operating environment.

[0030] A number of program modules may be stored on the hard disk, magnetic disk 29, optical disk 31, ROM 24 or RAM 25, including an operating system 35, one or more application programs 36, other program modules 37, and program data 38. A user may enter commands and information into the computer 20 through input devices such as a keyboard 40 and pointing device 42. Other input devices (not shown) may include a microphone, joy stick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 21 through a serial port interface 46 that is coupled to system bus 23, but may be connected by other interfaces, such as a parallel port, game port or a universal serial bus (USB). A monitor 47 or other type of display device is also connected to system bus 23 via an interface, such as video adapter 48. In addition to the monitor, personal computers typically include other peripheral output devices (not shown), such as speakers and printers.

[0031] The computer 20 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 49. Remote computer 49 may be another personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the computer 20, although only a memory storage device 50 has been illustrated in Figure 1. The logical connections depicted in Figure 1 include a local area network (LAN) 51 and a wide area network (WAN) 52 that are presented here by way of example and not limitation. Such networking environments are common-

place in offices enterprise-wide computer networks, intranets and the Internet.

[0032] When used in a LAN networking environment, the computer 20 is connected to the local network 51 through a network interface or adapter 53. When used in a WAN networking environment, the computer 20 typically includes a modem 54 or other means for establishing communications over the wide area network 52, such as the Internet. The modem 54, which may be internal or external, is connected to the system bus 23 via the serial port interface 46. In a networked environment, program modules depicted relative to the computer 20, or portions thereof, may be stored in the remote memory storage device. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

[0033] Figure 2 illustrates the basic elements of electronic message 10. Although the invention will be described below in reference to electronic message 10 of Figure 2, the invention can alternatively be practiced with other electronic messages having more or fewer elements. Specifically, distribution list 60 describes all intended recipients of electronic message 10. Although not shown separately, distribution list 60 includes direct recipients as well as those receiving normal or blind copies. Including blind copy recipients in distribution list 60 does not necessarily alter the rules of who will receive copies of a subsequent reply. That is, appearing on distribution list 60 does not mean that a client receiving a blind copy of an electronic message will be identifiable to other recipients and therefore potentially receive any later replies to the electronic message. Timestamp 70 is a general indication of a message's temporal context. It may represent the actual creation time, the time a message is received by an electronic messaging service for delivery, the actual delivery time, etc., or any combination thereof. Subject 72 is a field where the sender can summarize a message's content. Subject 72 aids recipients in evaluating the relative importance of any particular message without having to review the entire message. Finally, message content 80 represents a message's substantive content and may include minimally formatted text, binary data such as sound, images, executable files, and/or word processing documents, etc., that may be included with a message as an attachment or as the message itself.

[0034] Figure 3 illustrates one example of the host systems that can be used to practice the present invention. In general, the host systems of the invention can be any system that provides electronic messaging services to users. Host system 100, which should not be seen as imposing any structural or architectural limitations on the invention, includes a message server 112 and clients 114. Message server 112 can be any network server or other computing device capable of managing electronic messaging services for clients 114. In this example, message server 112 is associated with a message store 113, which can be any suitable data storage system capable of storing message content 80 of an electronic message 10, as shown in Figure 2. The process of storing message content 80 is further described below in reference to Figures 4-6. Clients 14, which are individually numbered client 1, client 2, client 3, client 4, . . . client n, can be as few as one client or as many clients as host system 100 and message server 112 are capable of serving.

[0035] In the example of Figure 3, message server 112 and clients 114 are connected by means of network 116, which may be a local area network, or a wide area network. In many cases, network 116 can be an organization-wide LAN configured to provide messaging services to users in the organization. Accordingly, clients 114 can be conventional client systems used in LANs or WANs, or any other suitable computing systems. For instance, clients 114 can be the computer 20 of Figure 1. As shown in Figure 3, host system 100 can be linked to other host systems. Host system 100 and, optionally, the other host systems designated by Figure 3 represent one example of a messaging system.

[0036] Figure 4 combines the basic elements of an electronic message 10, as illustrated in Figure 2, and the basic elements of a host system 100 providing electronic messaging services, as illustrated in Figure 3, to demonstrate the operation of a preferred embodiment of the present invention. The electronic messaging services of host system 100 function as follows. Limited copies of the message content 80 of Figure 2 are stored as message 180 in, for example, message store 113 of Figure 3. Generally, message 180 refers to a single electronic message. However, at times message 180 may also be viewed appropriately as an initial electronic message with associated replies and/or a forked electronic message as described below in connection with Figure 5.

[0037] The limited copies are in contrast to traditional electronic messaging services that store a separate copy for each client who is a recipient of an electronic message. At a minimum, the preferred embodiment of Figure 4 stores one copy of message 180. However, clients of host system 100 may be organized into diverse groups, with each group defining, for example, resources accessible by its members (among other things). Where clients are organized by group, host system 100 may store one copy of message 180 for each group having a member as an intended recipient of message 180.

[0038] Host system 100 also stores distribution list 60 of Figure 2 as distribution list 160. Next, the host system 100 identifies which of clients 114 are recipients as listed in distribution list 160 and notifies each recipient of message 180 through notification 110. In this example of Figure 4, clients 1, 3, and n are identified as recipients. Clients 1, 3, and n may be or may not be members of a common group.

[0039] Distribution list 160 may serve several pur-

poses. In addition to containing the intended recipients, distribution list 160 may also provides means for tracking various actions that recipients perform with regard to message 180, such as recording recipients who access the message. However, this tracking also may be performed by message 180 itself, it may be part of the rules included with either the electronic message or host system providing electronic messaging services, or tracking may be implemented separately or as a subpart of some collection of functionality.

[0040] Notification 110 may include various data such as timestamp 170 and subject 172 that correspond to the timestamp 70 and subject 72 of Figure 2. More importantly, notification 110 may include message link 182 for enabling the client to access the stored message 180. Timestamp 170, subject 172, and message link 182 are one example of a preferred embodiment's notification means for making message 180 available to recipients. A notification means may include other information such as a few lines of text, the message sender, etc. The present invention contemplates a variety of information as part of notification 110, including the above mentioned options being present either singly or in combination. It is only required that notification 110 provide some means for indicating the existence of message 180. This notification is in contrast to prior art systems that would duplicate electronic message 10, of Figure 2, and store a copy for each recipient.

[0041] Another aspect of the embodiment of Figure 4 is the ability to track access to message 180. Tracking may be useful as an auditing feature that allows the electronic messaging services of host system 100 to determine and report who has accessed message 180. Moreover, tracking access may be useful when considering rules for governing access and rules for determining lifetime as discussed below. Furthermore, one embodiment of the present invention notifies recipients when the distribution list changes and provides tools for examining the historical content of the distribution list, including comments to describe why specific changes were made.

[0042] The electronic messaging services also comprise rules for governing access to message 180. These rules may be defined as default rules to be used by the electronic messaging services or they may be defined on a message by message basis. Regardless, the rules may control access by determining, for example, if message 180 can be forwarded to new recipients, if message 180 can be downloaded from the message server to a local storage device, if distribution list 160 can be altered, and who is authorized to perform such actions. Altering distribution list 160 to add a new recipient grants access to message 180 through notification 110 as if the new recipient had been an original recipient. Altering distribution list 160 to remove a recipient revokes access to message 180 as if access had never been granted; future updates will not be available to the removed recipient. However, a recipient on distribution list 160 may have a local copy of message 180. By maintaining a local copy, a recipient removed from distribution list 160 will retain access to message 180 as it existed at the time of removal from the distribution list 160. Revocation may be accompanied by a message indicating that message 180 is no longer available to the client.

[0043] The electronic messaging services may also include rules for determining the lifetime of message 180. By way of example and not limitation, message 180 may only be useful through a given date, until everyone has read it, or until recipients no longer access it for a particular time. The rules may choose one of the above mentioned criteria or a combination thereof, or some other suitable parameter or combination for determining the lifetime of message 180. Rules may also indicate if the lifetime of message 180 can be extended and who is authorized to make an extension.

[0044] Figure 5 illustrates the organizational component of one embodiment of the present invention. In order to associate replies with the message that prompted them, each initial message 10 is assigned a message identifier 90. Message identifier 90 allows electronic messaging services of the host system to group replies by subject matter. As practiced in the prior art, conventional electronic messaging services intersperse replies with other electronic messages making it difficult for a recipient to follow a particular topic. Furthermore, as practiced in conventional systems, in an effort to aid other users in following the discussion, responses may duplicate the content of initial message 10 in the response, forcing the user to wade through the content of message 10 to reach the reply's substance.

[0045] In contrast, the embodiment of the present invention illustrated in Figure 5 associates replies to an initial message 10 using message identifier 90. Figure 5 shows a hierarchical tree model by way of example rather than limitation; any hierarchical organization is within the scope of the present invention. Reply 412, reply 414, reply 416, and reply 418 all represent reply messages associated with initial message 10 through message identifier 90. This organization allows, for example, a recipient to view a display with initial message 10 for context along with the text of replies 412, 414, 416, and 418 without having the content of message 10 duplicated in each reply.

[0046] Figure 5 also shows fork 423 and fork 427. In prior art systems, unintentional forks occur when individual topics begin to be discussed separately. Forking is a natural process that occurs based on the individual interests of each particular recipient. Participating in a forked discussion is like trying to listen to more than one person at the same time and makes establishing context a difficult task. Moreover, forking is further complicated when the original set of recipients splinters into subgroups discussing individual topics in detail. Forking is beneficial because it allows those with specific interests to focus their expertise without being subjected to

larger discussions that may be irrelevant to them. Unintentional forking becomes especially problematic in prior art systems when recipients are added to a distribution list, because the added recipients must depend on others to provide them with the discussion's history. In conventional systems, it becomes difficult to separate the history from the current discussion due to the volume of information duplicated in the name of context. Moreover, context must be established on a message by message basis by examining timestamps and sorting through text copied from earlier messages.

[0047] In contrast to forking experienced in the prior art, fork 423 and fork 427 represent intentional forks according to an embodiment of the present invention. These intentional forks are based on recipients determining that a particular topic is more suited to a small group. Fork 423 generates a new message identifier 425 and includes associated replies 430, 432, and 434. Fork 427, with new message identifier 429, includes associated replies 436 and 438 as well as fork 447 having message identifier 449 and associated replies 452, 454, and 456. The advantage of the embodiment exemplified in Figure 5 is that organization is a function of the electronic messaging services rather than the whim of an individual recipient saving messages on a given topic. Furthermore, adding a recipient to the appropriate distribution list provides the new recipient with the topic history as well, because the new recipients access is identical to that of an original recipient. Moreover, the hierarchy provides context that carries from message for message. For example, replies 430, 432, and 434 all relate to the topic division made in fork 423.

[0048] Although the discussion associated with Figures 2 and 4 spoke in terms of message 180, whether a message is an initial message, a reply, or a response to a reply is immaterial to how the message is delivered to clients who are recipients. Therefore, Figures 2 and 4 could have been described accurately in terms of message/reply rather than simply message. With the preceding discussion introducing replies and their corresponding organization, Figure 6 will describe messages and replies identically with regard to how they are distributed as any distinction is arbitrary. Nevertheless, the terms "message" and "reply" may be employed to indicate the relationship between several messages. Any descriptions of relationship with regard to distribution are by way of example and not intended as a limitation.

[0049] Figure 6 illustrates the operation of an embodiment of the present invention within a multiple host system environment where at least one of the host systems providing electronic messaging services does not operate according to the present invention. Originating host system 100 notifies its clients who are recipients of an electronic message as described in relation to Figure 4. However, message 180, shown in Figure 4, may be either an initial message or reply and is therefore designated message/reply 180 in Figure 6.

[0050] As the originating host system of message/reply 180, host system 100 is responsible for notifying other host systems serving clients who are recipients of message/reply 180. To this end, host system 100 creates host list 164 identifying each host system (e.g., host systems 100, 200, and 300) providing electronic messaging services to the recipients of message/reply 180. Host system 100 also creates a recipient by host list 162. List 162 organizes the recipients according to the host system that provides electronic messaging services to the recipients.

[0051] Host system 200 is identified by host system 100 as a host system that operates according to the current invention. Host system 100 transmits message/reply 180 and the portion of list 162 necessary for host system 200 to notify clients of host system 200 that are recipients of message/reply. Host system 200 stores the transmitted copy of message/reply 180 as message/reply 280 and the transmitted portion of list 162 as distribution list 260. In this example, host system 100 must also indicate how host system 200 should send replies to message/reply 180. This may be accomplished by identifying host system 100 as the originator of message/reply 180, meaning that all replies should be directed to host system 100 for distribution. Alternatively, host system 100 may transmit host list 164 and recipient by host list 162 to host system 200 so that host system 200 may distribute replies independent of host system 100. After receiving message/reply 180, stored as message/reply 280, and distribution list 160, stored as distribution list 260, host system 200 distributes message/reply 280 according to the description of host system 100 in Figure 4 and organizes message/reply 280 according to the description of Figure 5.

[0052] Host system 300 is identified as a host system that does not operate according to the present invention. For example, host system 300 may be incapable of storing a single copy of message/reply 180 and notifying recipients of the stored single copy. In response, host system 100 sends message/reply 180 using a protocol, which is not the one described in reference to host system 200, but is instead compatible with host system 100 and host system 300. One skilled in the art will recognize that several suitable protocols for use with host system 300 are available. In one example of using a protocol compatible with host system 300, host system 100 duplicates electronic message 10 of Figure 2 for each recipient who is a client of host system 300. As shown in Figure 6, electronic message 310 illustrates each of the duplicated electronic messages, containing a full distribution list 360, timestamp 370, subject 372, and message/reply 380, corresponding to the similarly number elements of electronic message 10 of Figure 2.

[0053] The present invention may be embodied in other specific forms without departing from its spirit or essential characteristics. The described embodiments are to be considered in all respects only as illustrative

and not restrictive. The scope of the invention is, therefore, indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

Claims

1. In a messaging system that provides clients with electronic messaging services, a method for distributing an electronic message to each client who is a recipient of the electronic message, comprising the steps of:

   creating one or more distribution lists including each client of a host system who is a recipient of the electronic message;
   storing at least one copy of the electronic message at the host system; and
   using the one or more distribution lists to notify each client of the host system who is a recipient of the electronic message, such that the at least one copy of the electronic message is made available to each notified client.

2. A method as recited in claim 1, further comprising the step of tracking each client who accesses the at least one stored copy of the electronic message.

3. A method as recited in claim 1, further comprising the step of altering the one or more distribution lists.

4. A method as recited in claim 3 wherein the step of altering the one or more distribution lists comprises the step of removing a recipient, thereby denying the recipient subsequent access to the at least one copy of the electronic message.

5. A method as recited in claim 3 wherein a recipient maintains a local copy of the electronic message and the step of altering the one or more distribution lists comprises the step of removing a recipient, thereby denying the recipient subsequent access to future updates of the electronic message but leaving the local copy intact.

6. A method as recited in claim 3 wherein the step of altering the one or more distribution lists comprises the step of adding a recipient, thereby granting the recipient access to the at least one copy of the electronic message.

7. A method as recited in claim 1 wherein the electronic message comprises rules that govern access to the at least one copy of the electronic message.

8. A method as recited in claim 7, further comprising the step of using the rules to prevent a client who is a recipient from distributing the electronic message to another client who is not a recipient.

9. A method as recited in claim 7, further comprising the step of using the rules to delete the at least one copy of the electronic message at a time designated by the rules.

10. A method as recited in claim 1 wherein a single copy of the electronic message is stored at the host.

11. A method as recited in claim 1 wherein clients are organized into groups and one copy of the electronic message is stored for each group having a client who is a recipient of the electronic message.

12. A method as recited in claim 1 wherein the notification to each recipient comprises the electronic message's subject and a link to access the electronic message.

13. In a messaging system that includes an originating host system and at least one other host system, each host system providing electronic messaging services to a particular set of clients, a method for distributing an electronic message comprising the steps of:

    the originating host system performing the steps of:

    creating one or more distribution lists including each recipient of the electronic message;
    identifying each host system providing electronic messaging services for each recipient of the electronic message;
    determining whether or not each identified host system is capable of using the one or more distribution lists to notify recipients of at least one copy of an electronic message stored at each identified host system;
    grouping each recipient of the electronic message according to the host system providing electronic messaging services in order to form one or more host-specific distribution lists for each identified host system; and
    transmitting at least one copy of the electronic message to each identified host system.

14. A method as recited in claim 13 wherein at least one host system performs the steps of:

    storing the at least one copy of the electronic message and the corresponding one or more

host-specific distribution lists; and

using the one or more host-specific distribution lists to notify each client of the at least one host system who is a recipient of the electronic message, such that the at least one copy of the electronic message is made available to each notified client.

15. A method as recited in claim 13 wherein at least one host system performs the steps of:

storing at least one copy of the electronic message; and

tracking each client who accesses the at least one copy of the electronic message stored at the at least one host system.

16. A method as recited in claim 13, further comprising the step of altering one or more distribution lists.

17. A method as recited in claim 13 wherein the electronic message comprises rules that govern access to the at least one copy of the electronic message.

18. A method as recited in claim 13 wherein at least one host system stores a single copy of the electronic message.

19. A method as recited in claim 13 wherein at least one host's clients are organized into groups and the at least one host stores one copy of the electronic message for each group having a client who is a recipient of the electronic message.

20. A method a recited in claim 13 wherein the originating host system further performs the steps of:

determining that at least one host system is unsupported in that it is not capable of using one or more distribution lists to notify recipients of at least one copy of an electronic message stored at the at least one host system; and

sending the electronic message to the unsupported host system using another electronic messaging protocol compatible with the unsupported host system.

21. In a messaging system that provides clients with electronic messaging services, a method for organizing one or more replies to an electronic message, comprising the steps of:

assigning a unique identifier to the electronic message; and

using the unique identifier to associate the one or more replies with the electronic message.

22. A method as recited in claim 21, further comprising the steps of:

assigning a reply identifier to a first reply created by a recipient of the electronic message; and

using the reply identifier to associate, with the first reply, each of one or more subsequent replies created in response to the first reply.

23. A method as recited in claim 21, further comprising the step of separately presenting the electronic message and the one or more replies to a recipient.

24. A method as recited in claim 22, further comprising the step of separately presenting the electronic message, the first reply and the one or more subsequent replies to a recipient.

25. A method as recited in claim 24 wherein the step of separately presenting the electronic message, the first reply, and the one or more subsequent replies comprises the step of presenting to the recipient a tree arrangement wherein the electronic message is a trunk, the first reply is a first-level branch depending from the trunk, and the one or more subsequent replies are second-level branches depending from the first-level branch.

26. A method as recited in claim 21 wherein the method includes steps for distributing the electronic message to clients of a host system, comprising the steps of:

creating one or more distribution lists including each client of the host system who is a recipient of the electronic message;

storing at least one copy of the electronic message at the host system; and

using the one or more distribution lists to notify each client of the host system who is a recipient of the electronic message, such that the at least one copy of the electronic message is made available to each notified client.

27. A method as recited in claim 26 wherein the method includes steps for distributing replies to the electronic message, comprising the steps of:

assigning a unique list identifier to the distribution list;

storing at least one copy of a reply created in response to the electronic message; and

notifying each client who is a recipient of the reply using the unique list identifier, such that the at least one copy of the reply is made available to each client who is a recipient of the reply.

28. A method as recited in claim 21, further comprising the steps of:

assigning a reply identifier to a first reply; and
using the reply identifier to associate, with the first reply, each of one or more subsequent replies created in response to the first reply.

29. A system for distributing electronic messages in a messaging environment that provides clients with electronic messaging services, the system comprising:

processor means for creating one or more distribution lists that include each client of a host system who is a recipient of the electronic message;
storage means for storing at least one copy of the electronic message at the host system; and
notification means for using the one or more distribution lists to notify each client of the host system who is a recipient of the electronic message, such that the at least one copy of the electronic message is made available to each notified client.

30. A system as recited in claim 29, further comprising means for tracking each client who accesses the at least one stored copy of the electronic message.

31. A system as recited in claim 29, further comprising means for altering the one or more distribution lists.

32. A system as recited in claim 29 wherein the electronic message comprises means for governing access to the at least one copy of the electronic message.

33. A system as recited in claim 29 wherein the storage means stores a single copy of the electronic message.

34. A system as recited in claim 29 wherein the host's clients are organized into groups and the storage means stores one copy of the electronic message for each group having a client who is a recipient of the electronic message.

35. A system for organizing replies to an electronic message, the system comprising:

processor means for uniquely identifying the electronic message; and
processor means for associating one or more replies with the uniquely identified electronic message.

36. A system recited in claim 35 wherein a recipient of the electronic message can optionally fork a first reply, the system further comprising:

processor means for uniquely identifying the first reply; and
processor means for associating, with the first reply, each of one or more subsequent replies created in response to the first reply.

37. A system as recited in claim 35, the system further comprising means for separately presenting the electronic message and the one or more replies to a recipient.

38. A system as recited in claim 36, the system further comprising means for separately presenting to a recipient the electronic message, the first reply and the one or more subsequent replies.

39. A system as recited in claim 35 wherein the system includes a host system and means for distributing the electronic message, the system further comprising:

processor means for creating one or more distribution lists including each client of the host system who is a recipient of the electronic message;
storage means for storing at least one copy of the electronic message at the host system; and
notification means for using the one or more distribution lists to notify each client of the host system who is a recipient of the electronic message, such that the at least one copy of the electronic message is made available to each notified client.

40. A system as recited in claim 39 wherein the system includes means for distributing replies to an electronic message, comprising:

processor means for assigning a unique list identifier to the distribution list;
storage means for storing a single copy of a reply created in response to the electronic message; and
notification means for notifying each client who is a recipient of the reply using the unique list identifier, such that the single copy of the reply is made available to each notified client.

41. A system as recited in claim 40 wherein the system includes means for a recipient of the electronic message to optionally fork a first reply, comprising:

processor means for assigning a reply identifier to the first reply; and
processor means for using the unique identifier

to associate, with the first reply, each of one or more subsequent replies created in response to the first reply.

42. A computer program product for implementing a method for distributing an electronic message wherein the method is capable of being implemented in a messaging system that includes a host system and provides clients with electronic messaging services, the computer program product comprising:

a computer-readable medium carrying computer-executable instructions for implementing the method wherein the computer-executable instructions comprise:

program code means for creating one or more distribution lists including each client of the host system who is a recipient of the electronic message;

program code means for storing at least one copy of the electronic message at the host system; and

program code means for using the one or more distribution lists to notify each client of the host system who is a recipient of the electronic message, such that the at least one copy of the electronic message is made available to each notified client.

43. A computer program product as recited in claim 42 wherein the computer-executable instructions further comprise program code means for altering the one or more distribution lists.

44. A computer program product as recited in claim 42 wherein the computer-executable instructions further comprise program code means for processing rules governing access to the at least one stored copy of the electronic message.

45. A computer program product for implementing a method for organizing replies to an electronic message wherein the method is capable of being implemented in a messaging system that includes a host system and provides clients with electronic messaging services, the computer program product comprising:

a computer-readable medium carrying computer-executable instructions for implementing the method wherein the computer-executable instructions comprise:

program code means for assigning a unique identifier to the electronic message; and

program code means for using the unique identifier to associate one or more replies with the electronic message.

46. A computer program product as recited in claim 45 wherein the computer-executable instructions further comprise program code means for distributing the electronic message, comprising:

program code means for creating one or more distribution lists including each client of the host system who is a recipient of the electronic message;

program code means for storing at least one copy of the electronic message at the host system; and

program code means for using the one or more distribution lists to notify each client of the host system who is a recipient of the electronic message, such that the at least one copy of the electronic message is made available to each notified client.

47. A computer program product as recited in claim 45 wherein the computer-executable instructions further comprise program code means for distributing one or more replies the electronic message, comprising:

program code means for assigning a unique list identifier to the distribution list;

program code means for storing at least one copy of a reply created in response to the electronic message; and

program code means for notifying each client who is a recipient of the reply using the unique list identifier, such that the at least one copy of the reply is made available to each client who is a recipient of the reply.

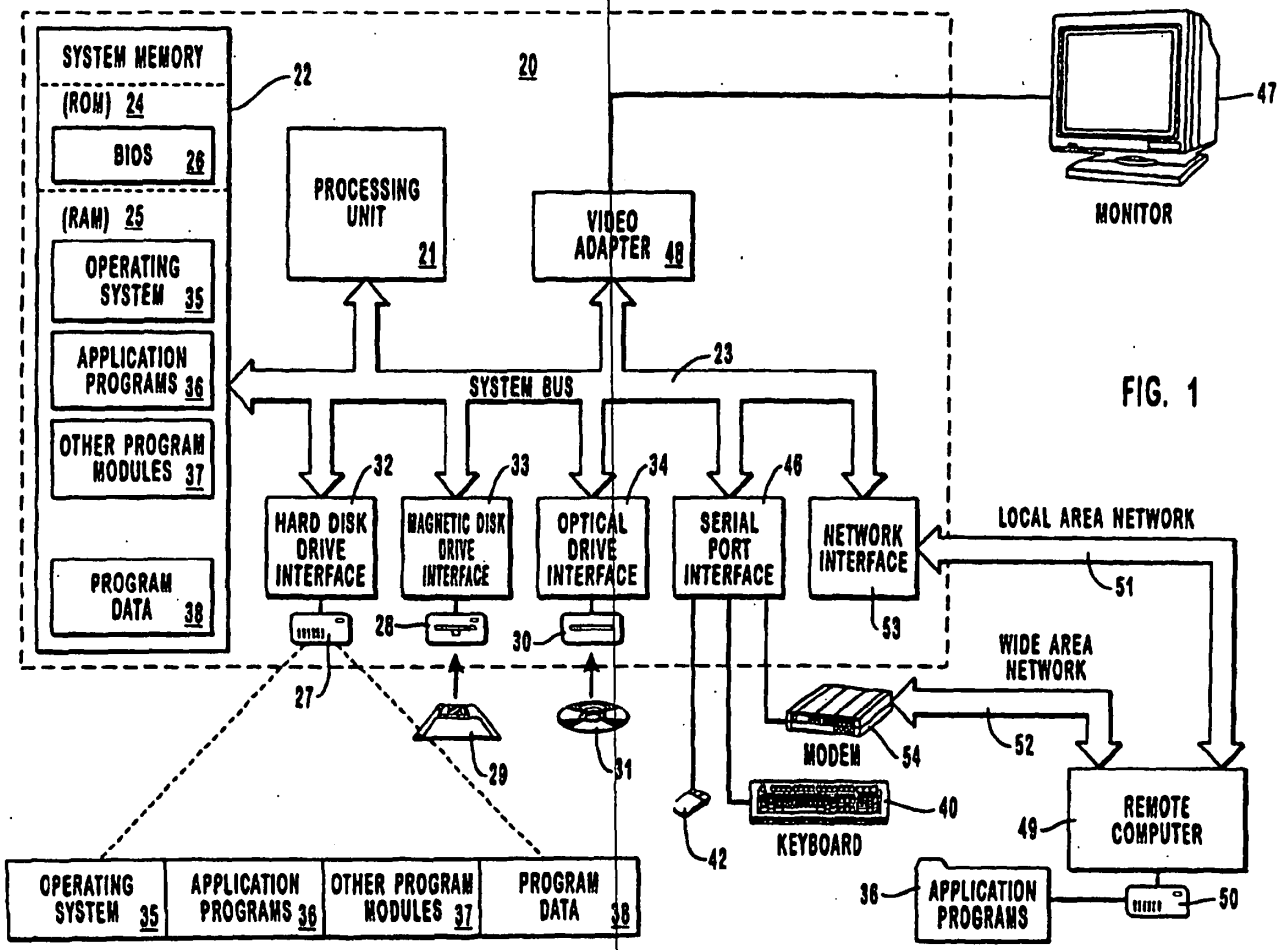

SYSTEM MEMORY
(ROM) 24
BIOS 26
(RAM) 25
OPERATING SYSTEM 35
APPLICATION PROGRAMS 36
OTHER PROGRAM MODULES 37
PROGRAM DATA 38
22
20
PROCESSING UNIT 21
VIDEO ADAPTER 48
47
MONITOR
23
SYSTEM BUS
32
33
34
46
HARD DISK DRIVE INTERFACE
MAGNETIC DISK DRIVE INTERFACE
OPTICAL DRIVE INTERFACE
SERIAL PORT INTERFACE
NETWORK INTERFACE
LOCAL AREA NETWORK
51
53
28
30
27
29
31
WIDE AREA NETWORK
MODEM
54
52
40
KEYBOARD
42
49
REMOTE COMPUTER
36
APPLICATION PROGRAMS
50
OPERATING SYSTEM 35
APPLICATION PROGRAMS 36
OTHER PROGRAM MODULES 37
PROGRAM DATA 38

FIG. 1

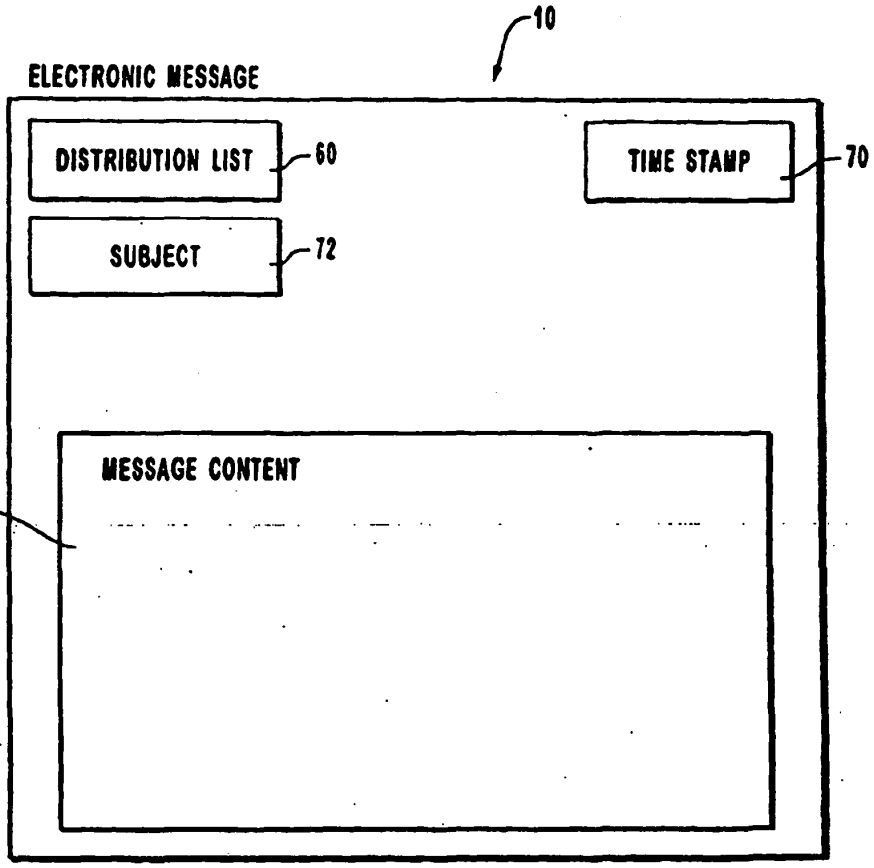

10
ELECTRONIC MESSAGE
DISTRIBUTION LIST
60
TIME STAMP
70
SUBJECT
72
MESSAGE CONTENT
80

FIG. 2

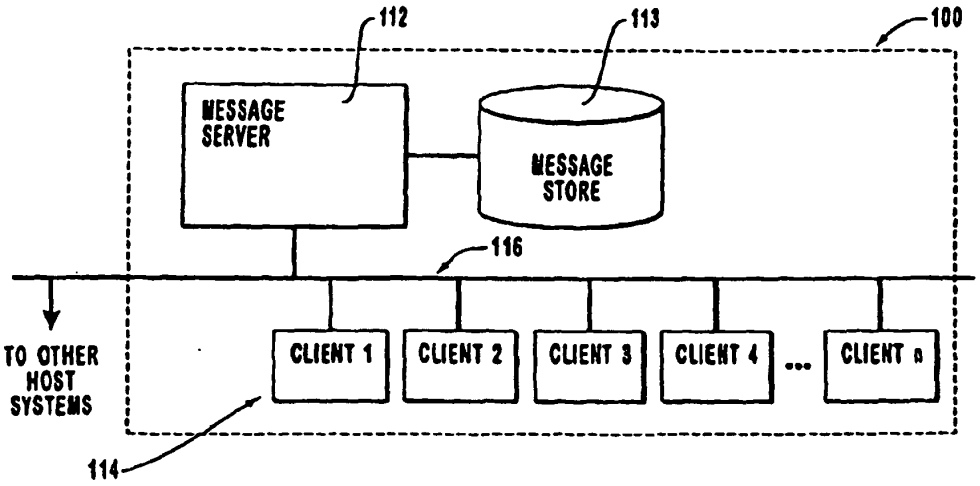

112
113
100
MESSAGE SERVER
MESSAGE STORE
116
TO OTHER HOST SYSTEMS
CLIENT 1
CLIENT 2
CLIENT 3
CLIENT 4
...
CLIENT n
114

FIG. 3

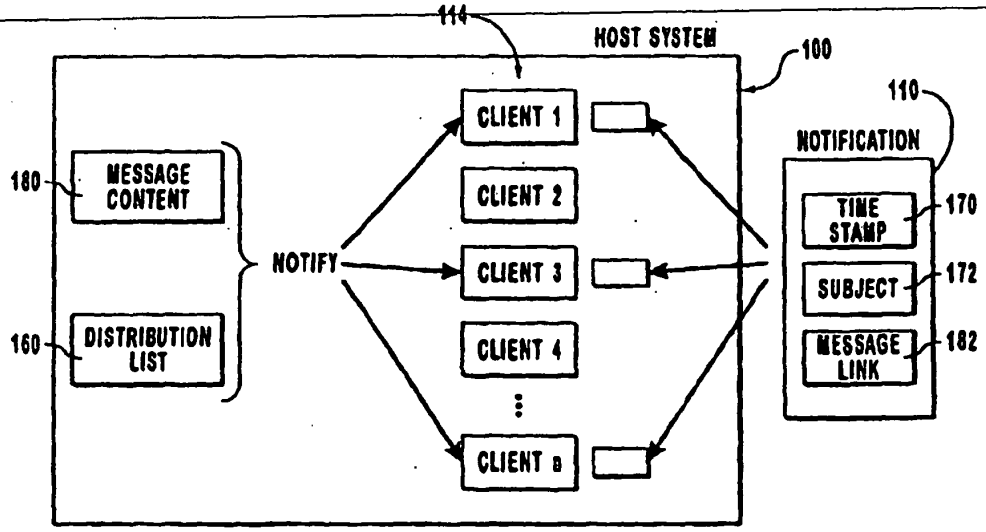

114
HOST SYSTEM
100
110
NOTIFICATION
CLIENT 1
CLIENT 2
CLIENT 3
CLIENT 4
CLIENT n
180
MESSAGE CONTENT
160
DISTRIBUTION LIST
NOTIFY
TIME STAMP
170
SUBJECT
172
MESSAGE LINK
182

FIG. 4

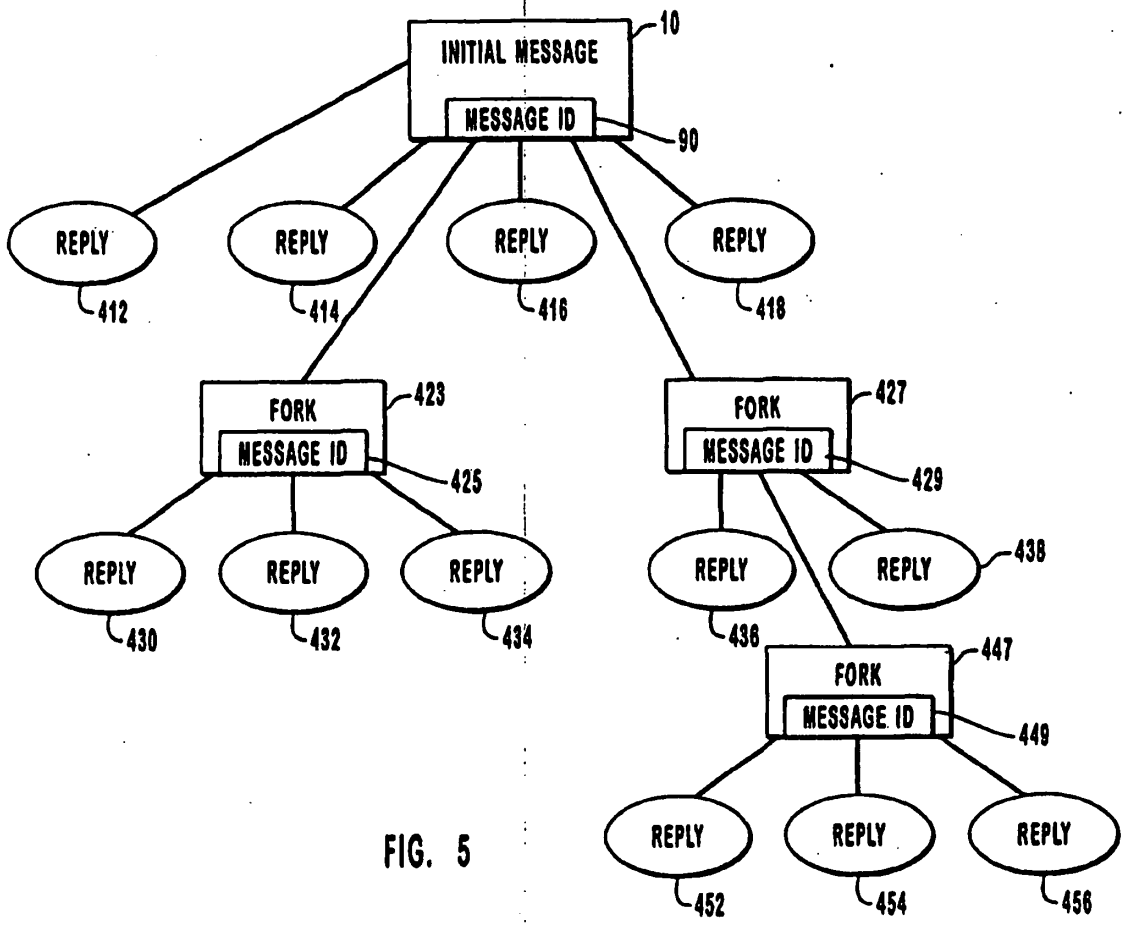


FIG. 5

EP 1 085 444 A2

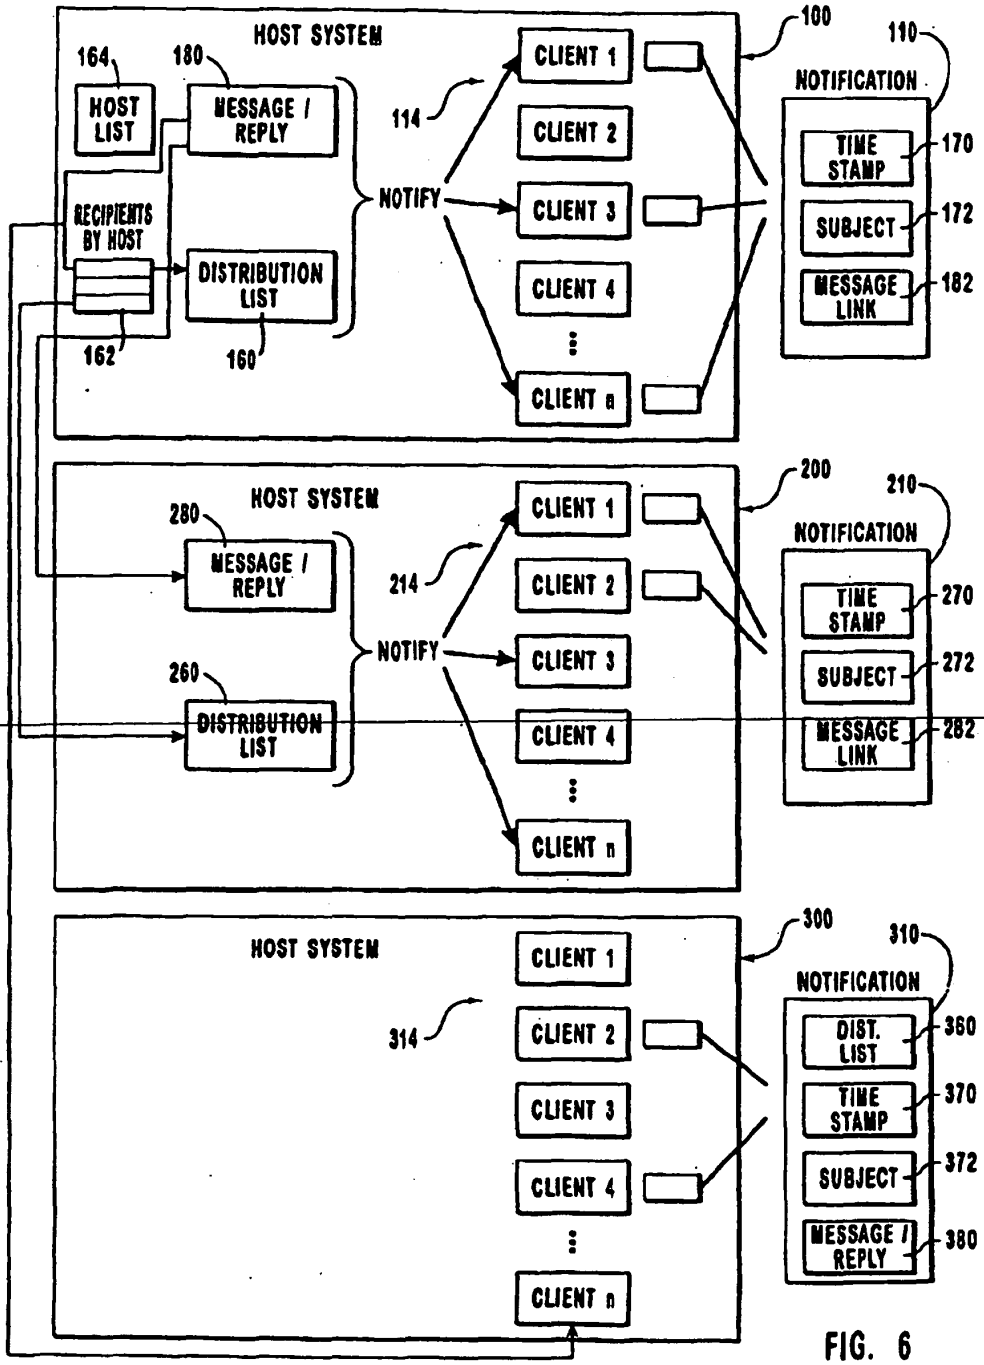

HOST SYSTEM
164
HOST LIST
180
MESSAGE / REPLY
RECIPIENTS BY HOST
DISTRIBUTION LIST
162
160
114
NOTIFY
CLIENT 1
CLIENT 2
CLIENT 3
CLIENT 4
CLIENT n
100
110
NOTIFICATION
TIME STAMP
170
SUBJECT
172
MESSAGE LINK
182
HOST SYSTEM
280
MESSAGE / REPLY
260
DISTRIBUTION LIST
214
NOTIFY
CLIENT 1
CLIENT 2
CLIENT 3
CLIENT 4
CLIENT n
200
210
NOTIFICATION
TIME STAMP
270
SUBJECT
272
MESSAGE LINK
282
HOST SYSTEM
314
CLIENT 1
CLIENT 2
CLIENT 3
CLIENT 4
CLIENT n
300
310
NOTIFICATION
DIST. LIST
360
TIME STAMP
370
SUBJECT
372
MESSAGE / REPLY
380

FIG. 6

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 085 444 A3

(12) EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
09.04.2003 Bulletin 2003/15

(51) Int Cl.[7]: G06F 17/60

(43) Date of publication A2:
21.03.2001 Bulletin 2001/12

(21) Application number: 00308239.3

(22) Date of filing: 20.09.2000

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 20.09.1999 US 399417

(71) Applicant: MICROSOFT CORPORATION
Redmond, Washington 98052-6399 (US)

(72) Inventors:
- Ahmed, Muhammad A
NO 202 Seattle, WA 98105 (US)
- Alam, Muhammad Shabbir
NO 2801 Bellevue, Washington 98007 (US)

(74) Representative: Orr, William McLean
URQUHART-DYKES & LORD
5th Floor, Tower House
Merrion Way
Leeds West Yorkshire, LS2 8PA (GB)

(54) Thread based e-mail

(57) Systems and methods for providing electronic messaging services to multiple users by storing a single copy of an electronic message at a central location and notifying recipients of the stored single copy. An electronic message includes a distribution list and a message content. A distribution list identifying multiple recipients causes prior art systems to duplicate the entire message for each recipient, placing potentially large demands on both processing power and storage space. In contrast, the systems and methods disclosed herein store a single copy or a limited number of copies of an electronic message addressed to multiple recipients and provide each recipient with a relatively small notification. In addition to providing information regarding content and origin, the notification also provides access to the stored message. Furthermore, the methods and systems also aid in organizing replies to electronic messages. Replies are associated with an initial message through a message identifier. The association helps to organize electronic messages by subject and provides context without requiring an author to duplicate the content of the initial message with the reply.

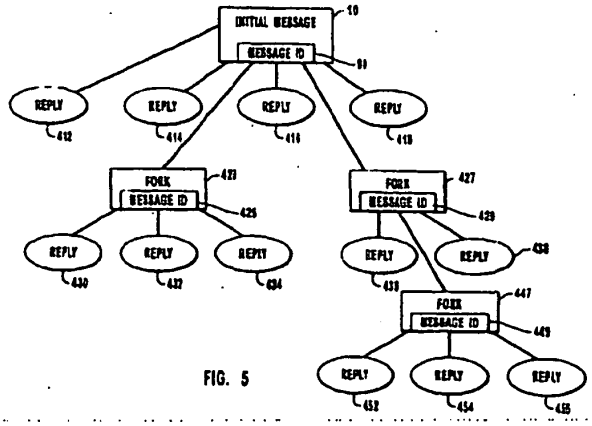


FIG. 5

EP 1 085 444 A3

European Patent Office

EUROPEAN SEARCH REPORT

Application Number
EP 00 30 8239

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	US 5 930 471 A (SMITH THOMAS M ET AL) 27 July 1999 (1999-07-27) * abstract; figures 2,3 * * column 3, line 40 - column 7, line 21 * Appendix 1	1-3,7, 10,12, 29-31, 42-44	G06F17/60
A	* the whole document *	4-6,8,9, 11, 13-20, 32-34	
X	EP 0 471 639 A (IBM) 19 February 1992 (1992-02-19) * the whole document *	1,3,6-8, 10,12, 29,31, 42-44	
X	EP 0 739 115 A (AT & T CORP) 23 October 1996 (1996-10-23) * the whole document *	1,2,7, 10-12, 28,29, 30,42-44	TECHNICAL FIELDS SEARCHED (Int.Cl.7) H04L G06F
X	'DISTRIBUTING A DOCUMENT USING A REGIONAL REFERENCE' IBM TECHNICAL DISCLOSURE BULLETIN, IBM CORP. NEW YORK, US, vol. 33, no. 3A, 1 August 1990 (1990-08-01), page 251 XP000123925 ISSN: 0018-8689 * page 1 - page 1 *	1,7,10	
A	EP 0 785 661 A (AT & T CORP) 23 July 1997 (1997-07-23) * abstract *	9	

~~The present search report has been drawn up for all claims~~

Place of search	Date of completion of the search	Examiner
MUNICH	3 December 2002	Huber, A

CATEGORY OF CITED DOCUMENTS

X : particularly relevant if taken alone
Y : particularly relevant if combined with another document of the same category
A : technological background
O : non-written disclosure
P : intermediate document

T : theory or principle underlying the invention
E : earlier patent document, but published on, or after the filing date
D : document cited in the application
L : document cited for other reasons

& : member of the same patent family, corresponding document

EPO FORM 1503 03.82 (P04C01)

European Patent Office

Application Number
EP 00 30 8239

CLAIMS INCURRING FEES

The present European patent application comprised at the time of filing more than ten claims.

☐ Only part of the claims have been paid within the prescribed time limit. The present European search report has been drawn up for the first ten claims and for those claims for which claims fees have been paid, namely claim(s):

☐ No claims fees have been paid within the prescribed time limit. The present European search report has been drawn up for the first ten claims.

LACK OF UNITY OF INVENTION

The Search Division considers that the present European patent application does not comply with the requirements of unity of invention and relates to several inventions or groups of inventions, namely:

see sheet B

☐ All further search fees have been paid within the fixed time limit. The present European search report has been drawn up for all claims.

☐ As all searchable claims could be searched without effort justifying an additional fee, the Search Division did not invite payment of any additional fee.

☐ Only part of the further search fees have been paid within the fixed time limit. The present European search report has been drawn up for those parts of the European patent application which relate to the inventions in respect of which search fees have been paid, namely claims:

☒ None of the further search fees have been paid within the fixed time limit. The present European search report has been drawn up for those parts of the European patent application which relate to the invention first mentioned in the claims, namely claims:

1-20, 29-34, 42-44

European Patent Office

LACK OF UNITY OF INVENTION
SHEET B

Application Number
EP 00 30 8239

The Search Division considers that the present European patent application does not comply with the requirements of unity of invention and relates to several inventions or groups of inventions, namely:

1. Claims: 1-20, 29-34, 42-44

Delivering a single electronic message to multiple recipients by storing a copy of said electronic message at a message host, and by sending notification messages including a reference to said copy of said electronic message from said message host to said recipients.

2. Claims: 21-28, 35-41, 45-47

Organising replies to electronic messages by assigning a unique identifier to an electronic message, and by using this unique identifier to associate the one or more replies with the electronic message.

ANNEX TO THE EUROPEAN SEARCH REPORT ON EUROPEAN PATENT APPLICATION NO. EP 00 30 8239

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

03-12-2002

Patent document cited in search report		Publication date		Patent family member(s)		Publication date
US 5930471	A	27-07-1999	WO	9829994	A1	09-07-1998
EP 0471639	A	19-02-1992	US	5144557	A	01-09-1992
			EP	0471639	A2	19-02-1992
			JP	1993558	C	22-11-1995
			JP	6067999	A	11-03-1994
			JP	7027499	B	29-03-1995
EP 0739115	A	23-10-1996	CA	2172746	A1	21-10-1996
			EP	0739115	A2	23-10-1996
			JP	8307452	A	22-11-1996
			US	5941946	A	24-08-1999
EP 0785661	A	23-07-1997	US	5781614	A	14-07-1998
			AU	1019197	A	24-07-1997
			CA	2189089	A1	20-07-1997
			EP	0785661	A2	23-07-1997
			JP	9233197	A	05-09-1997

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

PATENT COOPERATION TREATY

From the INTERNATIONAL SEARCHING AUTHORITY

To: Brian M. Hoffman
Fenwick & West LLP
Silicon Valley Center
801 California Street
Mountain View, CA 94041

PCT

NOTIFICATION OF TRANSMITTAL OF THE INTERNATIONAL SEARCH REPORT AND THE WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY, OR THE DECLARATION

(PCT Rule 44.1)

	Date of mailing *(day/month/year)*
Applicant's or agent's file reference 11745 PCT	**FOR FURTHER ACTION** See paragraphs 1 and 4 below
International application No. PCT/US 06/24859	International filing date *(day/month/year)* 26 June 2006 (26.06.2006)
Applicant Bluespace Group LTD	

1. ☒ The applicant is hereby notified that the international search report and the written opinion of the International Searching Authority have been established and are transmitted herewith.

   Filing of amendments and statement under Article 19:
   The applicant is entitled, if he so wishes, to amend the claims of the international application (see Rule 46):

   When? The time limit for filing such amendments is normally two months from the date of transmittal of the international search report.

   Where? Directly to the International Bureau of WIPO, 34 chemin des Colombettes 1211 Geneva 20, Switzerland, Facsimile No.: +41 22 740 14 35

   For more detailed instructions, see the notes on the accompanying sheet.

2. ☐ The applicant is hereby notified that no international search report will be established and that the declaration under Article 17(2)(a) to that effect and the written opinion of the International Searching Authority are transmitted herewith.

3. ☐ **With regard to the protest** against payment of (an) additional fee(s) under Rule 40.2, the applicant is notified that:

   ☐ the protest together with the decision thereon has been transmitted to the International Bureau together with the applicant's request to forward the texts of both the protest and the decision thereon to the designated Offices.

   ☐ no decision has been made yet on the protest; the applicant will be notified as soon as a decision is made.

4. **Reminders**

   Shortly after the expiration of **18 months** from the priority date, the international application will be published by the International Bureau. If the applicant wishes to avoid or postpone publication, a notice of withdrawal of the international application, or of the priority claim, must reach the International Bureau as provided in Rules 90*bis*.1 and 90*bis*.3, respectively, before the completion of the technical preparations for international publication.

   The applicant may submit comments on an informal basis on the written opinion of the International Searching Authority to the International Bureau. The International Bureau will send a copy of such comments to all designated Offices unless an international preliminary examination report has been or is to be established. These comments would also be made available to the public but not before the expiration of 30 months from the priority date.

   Within **19 months** from the priority date, but only in respect of some designated Offices, a demand for international preliminary examination must be filed if the applicant wishes to postpone the entry into the national phase **until 30 months** from the priority date (in some Offices even later); otherwise, the applicant must, **within 20 months** from the priority date, perform the prescribed acts for entry into the national phase before those designated Offices.

   In respect of other designated Offices, the time limit of **30 months** (or later) will apply even if no demand is filed within 19 months.

   See the Annex to Form PCT/IB/301 and, for details about the applicable time limits, Office by Office, see the *PCT Applicant's Guide*, Volume II, National Chapters and the WIPO Internet site.

Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201	Authorized officer: Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774

Form PCT/ISA/220 (January 2004) *(See notes on accompanying sheet)*

PATENT COOPERATION TREATY

From the INTERNATIONAL SEARCHING AUTHORITY

To: Brian M. Hoffman
Fenwick & West LLP
Silicon Valley Center
801 California Street
Mountain View, CA 94041

PCT

NOTIFICATION OF TRANSMITTAL OF THE INTERNATIONAL SEARCH REPORT AND THE WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY, OR THE DECLARATION

(PCT Rule 44.1)

	Date of mailing *(day/month/year)* 26 APR 2007
Applicant's or agent's file reference 11745 PCT	**FOR FURTHER ACTION** See paragraphs 1 and 4 below
International application No. PCT/US 06/24859	International filing date *(day/month/year)* 26 June 2006 (26.06.2006)
Applicant Bluespace Group LTD	

1. ☒ The applicant is hereby notified that the international search report and the written opinion of the International Searching Authority have been established and are transmitted herewith.

   Filing of amendments and statement under Article 19:
   The applicant is entitled, if he so wishes, to amend the claims of the international application (see Rule 46):

   When? The time limit for filing such amendments is normally two months from the date of transmittal of the international search report.

   Where? Directly to the International Bureau of WIPO, 34 chemin des Colombettes 1211 Geneva 20, Switzerland, Facsimile No.: +41 22 740 14 35

   For more detailed instructions, see the notes on the accompanying sheet.

2. ☐ The applicant is hereby notified that no international search report will be established and that the declaration under Article 17(2)(a) to that effect and the written opinion of the International Searching Authority are transmitted herewith.

3. ☐ **With regard to the protest** against payment of (an) additional fee(s) under Rule 40.2, the applicant is notified that:

   ☐ the protest together with the decision thereon has been transmitted to the International Bureau together with the applicant's request to forward the texts of both the protest and the decision thereon to the designated Offices.

   ☐ no decision has been made yet on the protest; the applicant will be notified as soon as a decision is made.

4. **Reminders**

   Shortly after the expiration of **18 months** from the priority date, the international application will be published by the International Bureau. If the applicant wishes to avoid or postpone publication, a notice of withdrawal of the international application, or of the priority claim, must reach the International Bureau as provided in Rules 90*bis*.1 and 90*bis*.3, respectively, before the completion of the technical preparations for international publication.

   The applicant may submit comments on an informal basis on the written opinion of the International Searching Authority to the International Bureau. The International Bureau will send a copy of such comments to all designated Offices unless an international preliminary examination report has been or is to be established. These comments would also be made available to the public but not before the expiration of 30 months from the priority date.

   Within **19 months** from the priority date, but only in respect of some designated Offices, a demand for international preliminary examination must be filed if the applicant wishes to postpone the entry into the national phase **until 30 months** from the priority date (in some Offices even later); otherwise, the applicant must, **within 20 months** from the priority date, perform the prescribed acts for entry into the national phase before those designated Offices.

   In respect of other designated Offices, the time limit of **30 months** (or later) will apply even if no demand is filed within 19 months.

   See the Annex to Form PCT/IB/301 and, for details about the applicable time limits, Office by Office, see the *PCT Applicant's Guide*, Volume II, National Chapters and the WIPO Internet site.

Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201	Authorized officer: Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774

Form PCT/ISA/220 (January 2004) *(See notes on accompanying sheet)*

PATENT COOPERATION TREATY

PCT

INTERNATIONAL SEARCH REPORT

(PCT Article 18 and Rules 43 and 44)

Applicant's or agent's file reference 11745 PCT	**FOR FURTHER ACTION**	see Form PCT/ISA/220 as well as, where applicable, item 5 below.
International application No. PCT/US 06/24859	International filing date *(day/month/year)* 26 June 2006 (26.06.2006)	(Earliest) Priority Date *(day/month/year)* 24 June 2005 (24.06.2005)
Applicant Bluespace Group LTD		

This international search report has been prepared by this International Searching Authority and is transmitted to the applicant according to Article 18. A copy is being transmitted to the International Bureau.

This international search report consists of a total of 9 sheets.

☒ It is also accompanied by a copy of each prior art document cited in this report.

1. **Basis of the report**
   a. With regard to the **language**, the international search was carried out on the basis of:
      ☒ the international application in the language in which it was filed
      ☐ a translation of the international application into _____, which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b))
   b. ☐ With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, see Box No. I.

2. ☐ **Certain claims were found unsearchable** (see Box No. II)

3. ☐ **Unity of invention is lacking** (see Box No. III)

4. With regard to the **title**,
   ☒ the text is approved as submitted by the applicant
   ☐ the text has been established by this Authority to read as follows:

5. With regard to the **abstract**,
   ☒ the text is approved as submitted by the applicant
   ☐ the text has been established, according to Rule 38.2(b), by this Authority as it appears in Box No. IV. The applicant may, within one month from the date of mailing of this international search report, submit comments to this Authority

6. With regard to the **drawings**,
   a. the figure of the **drawings** to be published with the abstract is Figure No. 1
      ☒ as suggested by the applicant
      ☐ as selected by this Authority, because the applicant failed to suggest a figure
      ☐ as selected by this Authority, because this figure better characterizes the invention
   b. ☐ none of the figures is to be published with the abstract

Form PCT/ISA/210 (first sheet) (April 2005)

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US 06/24859

A. CLASSIFICATION OF SUBJECT MATTER
IPC(8) - G06F 15/16 (2007.01)
USPC - 709/206
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC(8) - G06F 15/16 (2007.01)
USPC - 709/206

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
USPC: 709/204, 206, 232, 249

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

Electronic data base: USPTO WEST (PGPB, USPT, EPAB, JPAB); DIALOG PRO
Search Terms Used: e-mail or lectronic message, categorizing or assigning context or content, identifying or monitoring e-mails, retaining or storing or database or server emails, business or organization or enterprise, time or period etc.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2002/0136279 A1 (Al-Kazily et al.) 26 September 2002 (26.09.2002) (abstract, figs. 1-3, and para [0010]-[0015], [0042]-[0046], [0067]-[0075])	1-23
A	US 2004/0133645 A1 (Massanelli et al.) 08 July 2004 (08.07.2004)	1-23
A	US 2003/0191806 A1 (Osterberg, JR. et al.) 09 October 2003 (09.10.2003)	1-23
A	US 2002/0122543 A1 (Rowen) 05 September 2002 (05.09.2002)	1-23

☐ Further documents are listed in the continuation of Box C. ☐

* Special categories of cited documents:
"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
23 February 2007 (23.02.2007)	26 APR 2007
Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201	Authorized officer: Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774

Form PCT/ISA/210 (second sheet) (April 2005)

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

To: Brian M. Hoffman
Fenwick & West LLP
Silicon Valley Center
801 California Street
Mountain View, CA 94041

PCT

WRITTEN OPINION OF THE
INTERNATIONAL SEARCHING AUTHORITY

(PCT Rule 43*bis*.1)

Date of mailing *(day/month/year)*	26 APR 2007
Applicant's or agent's file reference 11745 PCT	**FOR FURTHER ACTION** See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US 06/24859	26 June 2006 (26.06.2006)	24 June 2005 (24.06.2005)

International Patent Classification (IPC) or both national classification and IPC
IPC(8) - G06F 15/16 (2007.01)
USPC - 709/206

Applicant Bluespace Group LTD

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☐ Box No. VII Certain defects in the international application
- ☐ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notified the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

3. For further details, see notes to Form PCT/ISA/220.

Name and mailing address of the ISA/US	Date of completion of this opinion	Authorized officer:
Mail Stop PCT, Attn: ISA/US Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201	23 February 2007 (23.02.2007)	Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774

Form PCT/ISA/237 (cover sheet) (April 2005)

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

International application No.
PCT/US 06/24859

Box No. I Basis of this opinion

1. With regard to the **language**, this opinion has been established on the basis of:

   ☒ the international application in the language in which it was filed

   ☐ a translation of the international application into _____ , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b)).

2. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application and necessary to the claimed invention, this opinion has been established on the basis of:

   a. type of material

   ☐ a sequence listing

   ☐ table(s) related to the sequence listing

   b. format of material

   ☐ on paper

   ☐ in electronic form

   c. time of filing/furnishing

   ☐ contained in the international application as filed

   ☐ filed together with the international application in electronic form

   ☐ furnished subsequently to this Authority for the purposes of search

3. ☐ In addition, in the case that more than one version or copy of a sequence listing and/or table(s) relating thereto has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

4. Additional comments:

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

International application No.
PCT/US 06/24859

Box No. V **Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement**

1. Statement

Novelty (N)	Claims	None	YES
	Claims	1-23	NO
Inventive step (IS)	Claims	None	YES
	Claims	1-23	NO
Industrial applicability (IA)	Claims	1-23	YES
	Claims	None	NO

2. Citations and explanations:

Claims 1-23 lack novelty under PCT Article 33(2) as being anticipated by US 2002/0136279 A1 to Al-Kazily et al. (hereinafter 'Al-Kazily'). Al-Kazily teaches a system/method/computer program product for providing an electronic messaging comprising:

1) message module (para [0039])
2) context module (para [0045])
3) definition module (para [0044])
4) association module (para [0043])
5) rights module (para [0073])
6) properties module (para [0045])
7) operations module (para [0070])

As to claims 1, 9, and 17, Al-Kazily teaches a system/method/computer program product for providing an electronic messaging comprising: a message module adapted to control a message database storing messages exchanged in the messaging system each message comprising one or more message components having content, and a context module adapted to control a context database storing a knowledge taxonomy having context categories and contexts associating content of the message components with the context categories. (abstract, figs. 1-3, and para [0010]-[0015], [0042]-[0046], [0067]-[0075]).

As to claims 2, 10, and 18, Al-Kazily teaches structuring context categories representing activities performed by the enterprise (para [0039]-[0045]).

As to claims 3-4, 11-12, and 19, Al-Kazily teaches monitoring messages exchange by the messaging system and creating contexts associating content of the messages with context categories (para [0042]-[0045]).

As to claims 5, 13, and 20, Al-Kazily specifying actions an end-user is allowed to perform with respect to a context category (para [0072]-[0074]).

As to claims 6-8, 14-16, and 21-22, Al-Kazily teaches specifying retention/validity period for a context category during which an association of content to the context category can be made, and performing operations in the messaging system (para [0070]-[0075]).

As to claim 23, Al-Kazily teaches allowing the user to create a context associating content of a message component with a context category, and/or view contexts associated with the context category (para [0070]-[[0075]).

Claims 1-23 have industrial applicability as defined by PCT Article 33(4) because the subject matter can be made or used in industry.

The Design and Implementation of the MONTAGE Multimedia Mail System

W. Keith Edwards

Georgia Institute of Technology
Software Engineering Research Center
Multimedia Computing Group
Atlanta, GA 30332-0280
keith@cc.gatech.edu

ABSTRACT

Electronic mail systems capable of transmitting compositions consisting of various unconventional media (such as voice, video, and images) have attracted a substantial amount of interest in recent years. It is important that mailers capable of delivering such documents present them in an organized fashion. We present the Montage multimedia electronic mail system, along with its model for multimedia documents. Montage makes use of a simpler format than more generalized hypermedia systems. It is our belief that the Montage model is more effective than general hypertext for the task of creating user-to-user messages. Furthermore, Montage is designed to be runtime extensible to new media types by its users. Thus the system does not have to know ahead of time all the possible media users may want to send. The system is built on top of existing mail transport protocols for flexibility and portability. We discuss the design of the mailer along with experiences gained from its implementation. The user interface to the system is also presented.

KEYWORDS

Multimedia mail, electronic mail, compound documents.

CH2955-3/91/0000-0047$01.00 ©1991 IEEE

INTRODUCTION

mon•tage män-'täzh, *n.*

1 *the production of a rapid succession of images in a motion picture to illustrate an association of ideas*

2a *a literary, musical, or artistic composite of juxtaposed more or less heterogeneous elements*

2b *a composite picture made by combining several separate pictures*

3 *a heterogeneous mixture*

Webster's Ninth New Collegiate Dictionary
Copyright © 1989

In the past decade, the proliferation of fast, inexpensive, networked computer workstations has produced an explosion in the use of electronic mail. Electronic mail systems have traditionally been limited to the transmission of pure textual information only.

As computer workstations increase in power, and as the use of windowing interfaces becomes more widespread and standardized, it becomes apparent that capabilities now exist for the transmission and reception of complex multimedia documents consisting of voice, images, video, and other media types in addition to plain text.

The chief advantage of such a system is the increased communication bandwidth between the users of the system. A multimedia mail system should allow people to communicate as freely and without restriction as conventional "paper" mail systems do. In current paper mail systems, users can seal most

BNSDOCID: <XP_____10039767A__I_>

anything within an envelope and expect prompt delivery. Electronic multimedia mail systems should allow similar flexibility.

Furthermore, to be useful for the widest range of people, attention should be paid to current standards for electronic mail transport. It should be possible to build multimedia mail systems on top of existing lower-level mail transport protocols and thereby use existing mail routing software for transport. Furthermore, the system should allow users to send plain "flat text" mail messages to users who do not have multimedia capabilities.

We present an experimental multimedia electronic mail system, called *Montage*, which has been developed at the Georgia Institute of Technology's Software Engineering Research Center. We believe that this system provides a flexible and convenient means to send and receive complex multimedia documents. This system is built strictly on top of existing lower-level mail routing protocols and thus should work on most systems which support the Unix operating system, X Windows, and Simple Mail Transport Protocol (SMTP) mail transfer systems.

WHY MULTIMEDIA MAIL?

Traditional electronic mail systems have been limited to sending only textual data. While text-based store and forward communication systems are useful, the old adage that "a picture is worth a thousand words" is often true. For example, it is nearly impossible to concisely describe a complex architectural drawing by writing a textual description of it. Many applications involve the use of graphical information and it is extremely inconvenient to transport this information using existing electronic mail systems.

Furthermore, the mode of information delivery most familiar to humans is direct voice communication. The work of Rohr [Rohr 86] indicates that some concepts are inherently graphical while others are inherently verbal. Other research indicates that in many applications, concepts are understood better when presented in a "mixed media" mode [Guastello 89].

While computer-based electronic mail systems enjoy many benefits over paper mail, they continue to lag behind in many respects. With paper mail systems, users can post basically any type of document which can be printed on paper. This includes images and formatted text, all possibly with annotations marked on the document.

With a more flexible electronic mail system it should be possible to combine the benefits of both systems: the convenience and speed of existing computer systems with the flexibility of paper mail systems. Furthermore, it should be possible to use the dynamic aspects of computing technology to open up mail systems to media which have previously gone unused in a store and forward environment, such as audio and video.

DESIGN

Our design for the Montage system was influenced by several desired goals. Foremost, we wanted to create a mail system which implements our model of multimedia messages. Secondly, to be widely used, the system had to be built on top of existing protocols and standards where ever possible. Finally, to meet unforeseen needs, the system had to be extensible by end users.

Requirements

Our design was constrained by several factors. First, we wanted to distinguish between multimedia messages and more complex general multimedia "documents." Montage is not a hypermedia system. Section 4, *A Model for Multimedia Messages* , ellaborates on this.

With this design decision, we were able to greatly simplify implementation. Montage messages proceed along one basic stream of control, unlike hypermedia messages which have a potentially large number of paths through the information.

Another design constraint was that the system should function cleanly over existing low-level transport mechanisms. We do not require that Montage must deliver its messages over TCP/IP connections. Since mail messages may be transferred between several intervening machines which may or may not run Montage, the messages should look for all intents like ordinary Simple Mail Transport Protocol (SMTP) [Crocker 82][Stallings 87] messages. Montage does not rely on any protocols higher than SMTP for transport.

Finally, we deemed that the system must be as extensible as possible. Since we cannot anticipate the needs of all users, we must provide a mechanism to dynamically add support for new media types, without the need to alter and recompile Montage.

Platform

We chose our platform based on availability and portability. While the Montage mail format may be implemented on a variety of platforms, including PCs and Macintoshes, our prototype implementation is for Unix workstations running the X Window

BNSDOCID: <XP_____10039767A__I_>

System. This platform ensures a reasonably high degree of portability for the system.

The interprocess communication abilities of Unix and the high level of flexibility in X also greatly simplify the implementation of some parts of Montage, as we shall see.

Specifically, our implementation of Montage was done on a Sun SPARCstation-1/GX machine running SunOS 4.0.3 and X11R4. The interface was built using Hewlett-Packard's widget set for X Windows. Audio support is native on the SPARCstation, video support is provided by a RasterOps video frame buffer board.

The hardware dependent parts of the Montage implementation (audio and video) are localized to simplify the porting process. It should be possible to run Montage on any Unix-based platform with X support. The audio and video functionality can be ported if necessary, but Montage will function perfectly well without audio and video support. In other words, Montage will function as a formatted text and image mailer if these facilities are not present.

Extensibility

It is impossible to know *a priori* all the things users may want to mail with a multimedia mail system. Almost by definition, multimedia mail systems should be able to send and receive a wide array of media.

Therefore, it was a primary goal that Montage be highly extensible by its end users. Montage allows users to modify the behavior of the program at runtime, without the need to recompile the system. Support may be added for new media types and new message transport techniques.

The facilities available for user extensibility will be addressed in following sections.

A MODEL FOR MULTIMEDIA MESSAGES

It is our belief that unlike more general purpose multi- or hypermedia documents, mail messages typically tend to be short and attempt to convey one central idea from the author. This belief has driven our design for the Montage message format.

Hypermedia documents tend to present the user with a large body of information in various media. The user then "navigates" through the document along a path that he or she chooses. Hypermedia systems generally tend to be quite interactive--the user is presented with almost all relevant information about a topic and then browses the document to find interesting information [Nielsen 90].

Contrast this model to a typical mail message in which the author is generally trying to convey some small number of central tenets. Unlike more general hypermedia documents, mail tends to be less interactive. The author has an idea to present and in some sense defines the reader's path through the message at composition time.

We believe that by limiting the generality of our message format we can achieve greater levels of usability, simplicity (for both the users and the implementors), and understandability of messages.

Still, however, since various media will be combined within a single message there must be some way to navigate through the packaged information. We have taken a simple approach to this task without having to resort to the less-constrained approach of a full-blown hypermedia system.

In Montage, all media are one of two classes: static or dynamic. The defining characteristic of a dynamic medium is that the information presented changes with time. Static media do not have this temporal component. Examples of dynamic media include video and audio clips. Examples of static media include simple text, formatted (rich) text, and still images.

Obviously it is of great interest to be able to combine various media freely within a single message. The differing characteristics of our two media classes make this packaging difficult. It is natural to think of a message window containing text and image data freely interspersed. Similarly it is natural to think of video and sound data being played at the same time. This mixing of media is common and is something users are used to experiencing on a day-by-day basis.

But does it make sense to combine static and dynamic media? Consider a message containing only text. When the message is opened a window appears that contains the text of the message. Now consider a message containing only voice. When the message is opened the voice clip is replayed. But what about a message containing both voice and text? Does one begin playing the voice as soon as the text window is opened? What about video and text? Does one open two windows, one for the text and one for the video? How does the user know which window to focus attention on? What if the moving video is placed within the text window? The video would become unviewable if the text were scrolled.

BNSDOCID: <XP_____10039767A_I_>

In both of these circumstances the mode of interaction would become one where the user's attention is focused on one media while the other is ignored. Some means must be provided to let users control the playback of the dynamic media so that they can decide when to focus attention on the various media. In addition, in our example above where video is presented in one window and text is presented in another, the user has no idea which contains the "central" part of the message. In other words, the user does not know which component of the message to focus on first.

Research supports the idea that users' interactions with multimedia systems tend to be largely "media-modal." That is, they segregate the information presented to them based on the medium of interaction [Laurel 90].

Primary Media Classes

To overcome these difficulties, Montage introduces the concept of a *primary media class*. Each message in Montage has a primary media class. This class is either static or dynamic. The primary media class is the type of the media presenting the "main thrust" of the communication from author to reader. Various media with the same media class (either static or dynamic) as the primary media class may be combined freely. For example, if a message's primary media class is dynamic, the author may freely intersperse audio and video in a message.

All of the components of the primary media class which compose the principal part of the message are collectively called the *primary component*.

The primary component is presented to the user in a single window. The window contains mechanisms for controlling the presentation of the primary component contained therein (in the case of dynamic media, the controls may be buttons for play, pause, fast forward, and reverse; in the case of static media the controls may be scrollbars to view various parts of the message).

We use the concept of the primary component to project a single "path" through the message, much as in existing mail systems. The primary media and primary component concepts also serves to give the author a means for expressing the central ideas in a message.

If Montage were restricted to using only media with types the same as the primary class, the system would not be very flexible. There is an obvious need for the incorporation of any type of media within a message, regardless of the primary media class.

As we shall see, the Montage model allows the use of any arbitrary media in the form of attachments.

Attachments

To augment the power of the system, Montage also makes use of *attachments*. Attachments may be placed on any message of a given primary media class. Attachments themselves are basically submessages and may be of any media type, regardless of primary media class.

Attachments may be considered to be "margin notes" that are not central to the message but still convey useful information. Since they are not presented to the user immediately when the window is opened they are, in some sense, secondary in importance to the information contained in the primary component.

Attachments give us a way to convey additional information while retaining our easy-to-play back, easy-to-understand main message thread. Without attachments, messages would degenerate to a conglomeration of mixed media within a single window which would in many cases be unmanageable.

In Montage, attachments are presented along side the primary component window in the form of icons. The image of the icon represents the type of media in the attachment. Furthermore, attachment icons are "connected" to a particular location in the primary component of the message. A use of attachments may be to provide annotation or supplementary information to the information contained in the primary component.

In the case of static primary media, attachments are connected to a certain physical point in the message. Thus, a voice attachment may be connected to a particular line number in a text message. In the case of dynamic media, attachments are connected to a certain time range in the primary message. Thus, as the dynamic message "plays back," the various attachments are presented to the user during the time they are relevant. These attachments may be selected as they appear to give additional information provided by the author.

Consider two examples of the use of primary media and annotations to construct a message. A geographically distributed group may be collaborating on a document. The primary component may be the document in question. Attached to this at various points may be audio annotations requesting changes, image data to be considered for review (image data may also appear within the document

BNSDOCID: <XP_____10039767A_!_>

itself), and video message clips to the various members of the group.

As another example, consider a researcher mailing a video of a presentation to a colleague. The primary media class here is dynamic, and the primary component of the message contains the video and associated audio of the talk. Attached to this at appropriate points are textual and audio annotations (the playback of the presentation may be stopped at any point to review the annotations), image data of the slides used in the presentation, text of papers, program source, and so forth.

Summary of Model

To summarize, the Montage model projects a single path of message traversal. This primary path is reflected in the primary component of the message. The primary component may contain mixed media, but the media it contains must be either all static or all dynamic to focus reader attention.

Additional information may be placed in a message in the form of attachments. An attachment is a piece of annotational information which is secondary to the primary flow through the message. Attachments are not restricted to being of the same class as the primary component.

IMPLEMENTATION NOTES

We have nearly finished a prototype implementation of Montage. This section covers the details of our implementation, including our assumptions and design decisions, the format Montage uses for message interchange, and the user interface to the system.

Overview

It was our desire that Montage be as flexible as possible in the types of media it can use. Therefore, we were determined to provide support even for media with very high bandwidth requirements, although the actual use of such media may be awkward on today's hardware because of pragmatic constraints.

Montage provides full support for interchange of text, still images, audio, and video. Local area networks commonly found today, such as Ethernet, have bandwidth in the 10 Mbps range. Even with Ethernet, applications will almost never see the full theoretical bandwidth available in the system.

The bandwidth provided by such a network is sufficient for text and small image interchange, workable for audio and large image interchange, and unwieldy for video interchange. Nevertheless, network speed and capacity will improve, so it is important to lay a software groundwork for applications which can make use of these faster networks.

Fortunately, the batch-processing style of mail makes full utilization of network resources less important. Whereas a real-time video conferencing system would require a guaranteed portion of the communication bandwidth to function, mail systems can be much less picky. As long as the message arrives in what the user perceives to be a "reasonable" amount of time the system is meeting its goals.

Montage relies on the underlying mail transport agents for actual mail delivery across the network. So as newer mail delivery agents that can efficiently transport large message across a network become available, Montage will be able to make use of these systems.

Interchange Format

While Montage mail headers conform to the SMTP standard, Montage makes use of a custom format for the bodies of electronic mail messages. This format was created because we felt that existing standards were either inflexible in the types of media they allow, or were too unconstrained to present media in an understandable format appropriate for mail messages.

The limitations of existing mail transport agents require that the contents of Montage mail message consist entirely of printable ASCII characters. Although there are some experimental binary mail transport agents, far and away the largest number of mailers available on Unix platforms only support ASCII transfer.

Because of this restriction, we are forced to encode the bodies of Montage messages into ASCII. Our reference implementation uses the standard Unix *uuencode* program to accomplish this. The ASCII-encoded file's size is expanded by 35% after this encoding process. To compensate for this, the message body is first compressed before it is mapped into ASCII.

But what is actually contained in this message body? In Montage, the message body actually consists of several discrete units, called *chapters*. When mail is received by Montage, the body is separated from the header, it is converted to binary from its ASCII format, uncompressed, and then broken down into its components. In our implementation, each of these chapters is stored as a separate file. There is one special chapter, called the *table of contents*, which contains information on the layout of the entire mes-

BNSDOCID: <XP_____10039767A__I_>

sage. The table of contents references and connects all other chapters in the message.

Our first implementation uses existing Unix tools in an effort to rapidly produce an operational version of Montage. We use the *tar* Unix archive program to combine the various chapter files upon message creation, and to break out the chapters upon message receipt.

Montage adds one line to the SMTP mail message header which specifies the version number for the encoding scheme. Since the above method is only one of many possible encodings, future Montage mailers will be able to determine the encoding method by the version number in the header. Note that this is the only addition we need to make to the header to be able to send multimedia messages.

X.400

In 1984, CCITT released a set of standards for electronic mail systems. These standards do not deal with the user interfaces of mail systems, but rather they specify the services available for sending messages across the network.

The X.400 model defines two agents that make up an electronic mail system: the User Agent (UA) and the Message Transfer Agent (MTA). The UA provides the user interface for the system and may interact with other UAs. UAs hand messages off to MTAs for transport. X.400 specifies the interactions between UAs and MTAs, but does not specify the interactions between UAs and the users [Cunningham 84][Cunningham 85].

Montage implements a subset of X.400. Montage provides most of the header fields which are used in communication between UAs and MTAs, but does not provide the inter-UA communication facilities.

It should again be noted that Montage is largely independent of any underlying mail transport agent. The system can be configured to use an administrator-defined mail agent. On our prototype system this is the Unix *sendmail* program. While we have not tried using a full X.400 mail transport agent with Montage we believe that it should function properly.

Chapters

As stated before, a message body consists of one or more message chapters along with a table of contents. The table of contents specifies the relations between the various chapters (relative placement within a message, chapter format, etc.)

Each chapter contains a "section" of the message. Each chapter is represented in one medium. The type of this medium must be specified in the table of contents so that the recipient mailer will know how to "play back" the chapter.

A chapter is basically a single file containing a single message component. Chapters may contain either part of the primary component, an attachment, or the table of contents.

Note that the principle component of the message may be composed of many chapters, as long as the media in those chapters is of the same class as the primary media class (either static or dynamic). Thus, it is possible to have text interspersed with images in the primary component of the message.

Montage treats each chapter as raw data. That is, it associates no real semantic information with each chapter. Instead, the chapter data is handed to a playback module, or handler, which may be internal or external to Montage. The handler which is invoked on a particular chapter depends on the media type of the chapter. These handlers are specified by the users of the system, and Montage provides a mechanism for users to specify external programs which they can use to play back and record message chapters.

Media types are identified to Montage by *tags* which are associated with each chapter via the table of contents. Montage does not predefine any tags and indeed does not even associate any meaning with tags. Tags are defined by the users of the system. When a chapter is encountered the catalog of tags is searched and the user-specified handler (which may be an external program) is invoked on the chapter. The mappings from tags to handlers may be completely specified by users in a per-user database (with "sensible" defaults provided in a system-wide database).

Thus Montage is completely runtime extensible by the user in the media domain.

This ability provides Montage with a great deal of its flexibility. Users can choose their favorite tool for "recording" text (i.e., they can use their choice of editors (or word processors or spreadsheets) to enter text into the mail system). Similarly, they can use their choice of tools for viewing received audio and video chapters. Montage provides some simple means for playing and recording certain simple media. But because the system is not limited to those media which have built-in handlers, users can automatically add support for new media by specifying external programs to be used for media playback and recording. Additionally, this mechanism allows

BNSDOCID: <XP_____10039767A_I_>

a great deal of user customizability and support for individual user preferences, by allowing users to use their choice of mechanisms for playback and recording of media.

Table of Contents

The table of contents (or *TOC*) is a special chapter that specifies the relations of the other chapters to one another. The format of the TOC is relatively simple, reflecting our simple model of mail usage.

We have chosen a line-oriented ASCII format for the table of contents. Since most TOCs will be relatively small, we felt that the space savings accomplished by encoding the information in a machine readable format would not be as important as the simplicity and ease of debugging provided by a human-readable format.

Each TOC must have at least 3 records. Each record gives some piece of information and each appears on a separate line. The required records are:

1. `TOCVersion` The version identifier for the TOC format.
2. `Class` The class of the primary message component, either static or dynamic.
3. `Primary` The name and type of the message chapter containing a part of the primary component of the message.

In addition to these required records, there are several optional records:

1. `Attachment` The name, tag, and position of an attachment. Position is given as a line number in the primary component if the message is static, or as a time offset in the primary component if the message is dynamic.
2. `Author` The mail address and (optional) name of the sender.
3. `CreationDate` The date the TOC was created.
4. `Subject` The subject of the message.
5. `ID` A unique identifier for the message on the machine it was created.
6. `Comment` Signifies that the rest of the line is to be treated as a comment (ignored). This is primarily used as a debugging tool.

Since it is possible to combine several chapters within the primary section of the message (as long as these media are either all static or all dynamic), there may be several `Primary` fields in the TOC. These chapters will be presented sequentially to the user.

Additionally, some media (for example, a proprietary format for a general purpose multimedia document) may mix static and dynamic media within a single window, even though this mixing is something that "native" Montage does not allow. In such a case, the media in question may reference files which it expects to contain the various media components it needs. These external files may be bundled together with the rest of a Montage message, and are referred to as "subchapters" since they are not used directly by Montage itself, but rather by one of the media that Montage is dealing with. The media which uses these external files must be responsible for managing them. Subchapters are not specified in the TOC since individual applications, rather than Montage itself, deal with them.

Here is a sample table of contents for a Montage message. This message consists of a simple text primary component, with three attachments, a μ-law sound clip (with tag "CODEC88"), another simple text segment (with tag "SimpText"), and an X Bitmap image (tag "XBM").

```
Author: Keith Edwards <keith@cc>
CreationDate: 18 Oct 90 10:18:52 PDT
Subject: Notes from the meeting
ID: 2848.AA08665
TOCVersion: 1.0
Class: static
Primary: txt.2848.txt SimpText
Attachment: snd.2848.snd CODEC88 27
Attachment: txt.2848.txt SimpText 30
Attachment: xbm.2848.xbm XBM 49
```

Interface

We built the interface for our implementation on top of the X Window System. We perceived two primary reasons for using X. Most importantly, X is accepted as a standard throughout the workstation environment we were targeting as our audience. A primary reason for this acceptance is the portability of X and applications developed for X. As a result of this portability, the interface portions of Montage should recompile cleanly on any workstation which supports X Windows.

Secondly, X provides certain general mechanisms for user customization at runtime. We were able to make use of these mechanisms to create a highly configurable mail system--not only is the interface of Montage configurable, but the actual types of media the system can handle can be changed at runtime. It would have been possible to get this degree of configurability without using X, but we were able to prototype the system much more rapidly by using the facilities already available to X applications.

BNSDOCID: <XP_____10039767A__1_>

Our front end is built using a set of user interface objects (*widgets* in the X parlance) provided by the Hewlett-Packard company. We choose the HP widget set based on its perceived completeness and orthogonality. A secondary consideration was that the HP widget set is freely available, unlike some widget libraries (such as Motif and OpenLook) which are available only to sites which pay a licensing fee.

Montage makes use of a multiple-window interface. The first thing a Montage user sees is the *Control Palette*. The Control Palette presents the user with a group of on-screen "buttons" representing the major options available in the system. The primary options available from the Control Palette are

- New Mail check in any newly received mail into the Montage system. New mail must be "checked in" before it can be read.
- **Info** presents the user with a "pop-up" window, giving information about the program's origin.
- Messages presents the user with a scrollable list of messages. Users may then view, forward, save, or delete these messages.
- Compose is an interface to message creation. Facilities are provided for composition in several media types. Support for new media types may be added dynamically.
- Iconify reduces all the Montage windows to a single small window. This is convenient for when the user is not working with the system at the moment.
- Preferences brings up a panel which allows the user to customize his or her default preferences.
- Quit quits the system.

We will now go into some detail on the mechanics of the Montage interface.

Checking In New Mail Clicking the New Mail button causes any mail in the user's mail to be "checked in" to Montage. Whenever new mail is checked in, the system reads in all new messages from the user's mail spool file. As the messages are read in, the headers are parsed and a mail cache file is built in the user's home directory. The cache file contains important header information such as the message sender and subject, and the number and format of any attachments. The system also calculates the byte offset of the start of the actual message body so that the message data can be separated from the header. After the cache file is updated, the individual messages are stored in a subdirectory of the user's directory, on a one-message-per-file basis. The system spool file is then removed to free up system disk resources.

Users are notified when new mail arrives by the playback of a sound file. Users may also enable an automatic mail check-in function via the preferences panel (see **Preferences**, below).

Basic Message Control The most commonly used Control Panel items are Messages and Compose. These two items give the interface to message creation, viewing, deletion, and most other common functions.

When the user clicks the Messages button to view the list of current messages, the system scans the mail cache file to retrieve the header information. This keeps Montage from having to open, scan, and parse all the message files individually. A new window is opened and the user is presented with a scrollable list of the current messages.

From the *Messages* window (see Figure 1) users have the option to view messages, save messages to a file, delete messages, or reply to or forward messages. All actions on messages are accomplished by highlighting the desired message or messages, and then clicking the button to perform some action on those messages. Since viewing is perhaps the most common operation, users may view messages by double clicking on the message.

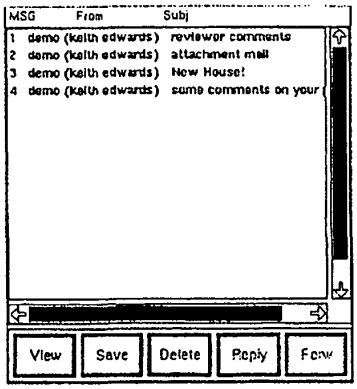


Figure 1: The Message Window

Saving a message brings up a window which prompts

BNSDOCID: <XP_____10039767A__I_>

for the filename in which to save the message. Currently, messages are saved in their native, bundled format. We will be adding support for saving individual message components.

Replying to a message puts the user into the *Compose* window (to be discussed shortly). When replying, the user is not restricted to using the same primary media class as the original message.

Forwarding simply allows the message to be resent to a new destination.

Message Viewing When a message is viewed, the *View* window (Figure 2) appears on the user's display. The message is automatically unpacked and the primary component of the message is presented in the main portion of the window. Any attachments are displayed along the side of the window and are represented by icons which depict the types of the attachments. There is a single scrollbar attached to the primary window and attachment list.

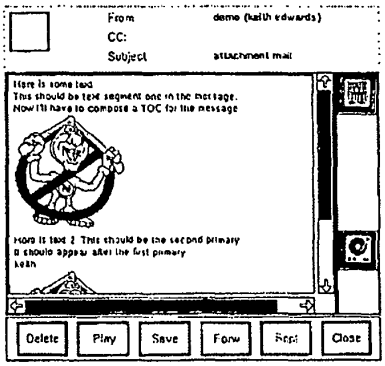


Figure 2: The View Window

What is actually displayed in the primary component window depends on the media themselves. Typically, static media will be presented in a scrolling window. The handling of dynamic media is more complex, however. Users have control over sound via a control panel much resembling a tapedeck. Video messages are controlled by via a window with a set of controls much like a video cassette recorder. As the dynamic media are played, any attachments scroll by the user to the side of the control panel. If desired, users may stop playback and view attachments relevant to the portion of the primary component they are viewing.

The *View* window also gives users options to delete, save, forward, and reply to messages. These functions work the same as those presented in the *Messages* window.

Composing New Messages Clicking the Compose button brings up the *Compose* window, which is an interface to message creation (see Figure 3). Users fill in the necessary components of the message header (recipient, subject line, and any carbon copies) by typing in the appropriate fields.

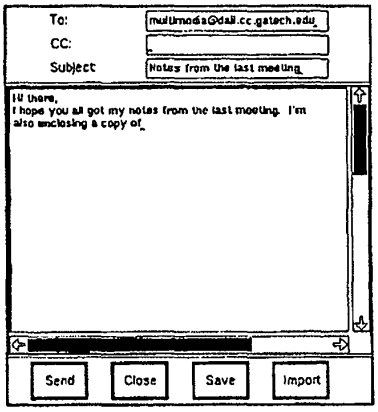


Figure 3: The Compose Window

The main portion of the *Compose* window is taken up by a region where the primary component is created. Users may type into this window directly, or they may import files of various types via the **Import** button.

The **Import** button brings up a window which allows users to insert external files as either attachments or as parts of the primary component. The user must specify the type of the inclusion so that Montage may build the table of contents for the message correctly.

Because of the ability to type directly into the *Compose* window's primary component area, the input of text is simple. However, it is desirable to be able to record other media as easily as text. It is often cumbersome to have to use tools external to Montage to record some message component into a file and then import the file into a message. Thus the *View*

BNSDOCID: <XP_____.10039767A__I_>

window provides a menu for Compose Tools. These tools allow users to incorporate media into messages almost as easily as text.

For example, Montage supports a sound recording tool from the Compose Tools menu. This tool (resembling a tapedeck, see Figure 4) allows users to record and edit sound clips. A similar video tool allows the recording and editing of video messages. Users may configure Montage to add new items and actions to the tools menu.

Users click the Send button to pass the message on to the mail transport agent. The *View* window contains buttons to close the window, and save the message to a file.

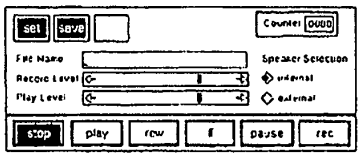


Figure 4: A Sound Tool

Hiding Montage Montage uses a multiple-window interface. Each window in the system may be shrunk or "iconified" individually.

It may be desirable at times to iconify all the windows of the application at once to effectively hide the application. This may be accomplished by the Iconify button on the Control Palette. This button reduces all the Montage windows to a single small window. Any windows which have been iconified individually are "joined" into the Montage icon. All windows will be returned to their original states when the icon is clicked.

User Preferences Montage supports user-configurable preferences via the Preferences button. This button brings up a control panel which allows users to select and modify a variety of system parameters, including changing the speakers to which audio output will be routed, setting audio play and record levels, setting whether or not message headers will be displayed, and so on.

PERFORMANCE

Overall mail system performance depends on many variables: the underlying network used for mail transport, the size of the messages, and even the habits of the users (some users may want to send only sound clips while others may be content to send simple text most of the time).

Designers of mail systems must be concerned with two aspects of performance: mail delivery latency and message size. Mail delivery latency is the actual time required for transmission of a typical message across the network. Message size is the amount of physical storage required to actually store the mail message.

The delivery latency is dependent on the speed of the underlying network which is in turn dependent on the mail delivery agent chosen for message transport. Since Montage is usable with any RFC-822-compliant mail transport agent, we are not directly concerned with network latency and performance may be characterized by one metric: message size. In most cases message size will be proportional to delivery latency for a given network and transport agent, so message size is an effective measure of performance.

Sending a message via Montage will add somewhat to the "native" size of the information being transmitted. Essentially the native size of a Montage message is the sum of the sizes of the attachments, along with the primary message component and the message header. When the message is sent, a table of contents is added. Furthermore, all of the message body undergoes a packaging process which adds some to the size of the message.

Our goal is to minimize the increase in size by (1) optimizing the packaging process for size, and (2) making the table of contents as small as possible. We will address each of these issues in turn.

The packaging process basically consists of "bundling" the various message components into one file. To minimize the size of the resulting file, it is compressed using adaptive Lempel-Ziv coding [Welch 84] as implemented by the Unix *compress* program). The amount of compression obtained by using this algorithm depends on the size of the input and the distribution of common substrings. Typically, for English text, the compression will be in the range of 50% to 60%. Some files, such as μ-law sound files, are already encoded and cannot be compressed further.

Since the result of this compression is a binary file, it must be mapped into ASCII characters for transmission via RFC-822 mailers. This mapping process expands the file's size by 35% (3 bytes become 4, plus some control information is added). Thus, for fairly long text messages, Montage can achieve message sizes smaller than those for messages containing

BNSDOCID: <XP_____10039767A__I_>

the same, but unencoded text (original message size × .55 × 1.35 = .74 of the original message size).

Of course added to any encoded message is a table of contents. This table of contents specifies the "layout" of the message. Our format for the table of contents is ASCII-based. We deliberately chose this format for simplicity in implementation and to ease debugging. Usually the table of contents for a given message will be very small.

The minimum size for a table of contents is variable. Since the table of contents contains entries for things such as sender name and subject there is no fixed minimum size. But in general, a simple table of contents will be on the order of 300 bytes. Each attachment will add something on the order of 60 bytes for information regarding attachment type and placement. So we see that the contribution of the table of contents to the size of the total message is negligible.

STATUS

The design and a first-pass implementation of Montage have been completed. Currently the system allows static primary media, and attachments of either text, audio, or several image formats. Facilities have been completed to allow easy composition of messages by typing or importing text, importing images, and recording sound via a simple sound editor mechanism.

Work is progressing on the addition of full dynamic primary media capabilities. We hope to soon have store-and-forward video transmission along with synchronized sound output.

Montage is currently running as a complete functional mail system which can serve as a replacement for existing Unix mail systems. That is, Montage may already be used to replace users' text mail composition and reading programs; the remaining work lies in the area of support for additional media.

CONCLUSIONS, CAVEATS, FUTURE DIRECTIONS

Our work on Montage was based on a set of assumptions which we believe greatly enhanced the usability of the system, as well as simplified the implementation. The initial response within the local research community at the Software Engineering Research Center and outside sponsors has been favorable.

In the near future, we plan to polish the implementation of the system to the point where we have a robust, workable mail system which may then be integrated into an actual work environment. Further work will be done on enhancing the video capabilities of the system and integrating the mailer with existing tools and media.

Much work still needs to be done on increasing the usability of the user interface. An online, context-sensitive help system is also needed.

We feel that Montage represents a flexible, workable system for the exchange of multimedia documents while presenting a friendly, intuitive interface to users.

ACKNOWLEDGEMENTS

The research and development of this system were supported by a grant from BellSouth Corporation.

REFERENCES

Crocker, David H. [1982] Standard for the Format of ARPA Internet Text Messages, Internet Request For Comment (RFC) 822, August 13, 1982.

Cunningham, I. [1984] Electronic Mail Standards to Get Rubber-Stamped and Go Worldwide. *Data Communications*, May 1984.

Cunningham, I., and I. Kerr. [1985] New Electronic Mail Standards. *Telecommunications*, July 1985.

Guastello, S., M. Traut, and G. Korienek. [1989] Verbal Versus Pictorial Representations of Objects in a Human-Computer Interface. In *International Journal of Man-Machine Studies*, July 1989, Vol. 31, No. 1, pp. 99-120.

Laurel, Brenda, Tim Oren, and Abbe Don. [1990] Issues in Multimedia Interface Design: Media Integration and Interface Agents. In *ACM SIGCHI Proceedings, 1990.* (Seattle, Washington April 1-5). pp. 133-139.

Nielsen, Jakob [1990] *Hypertext and Hypermedia*. Academic Press Inc., San Diego, CA, 1990.

Rohr, G. [1986] Using Visual Concepts. *Visual Languages*, S. Chang, T. Ichikawa, and P. Ligomenides, eds., Plenum Press, New York, 1986, pp. 325-348.

Stallings, W. [1987] *Handbook of Computer-Communications Standards, Volume 3: Department of Defense (DoD) Protocol Standards*. New York: Macmillan, 1987.

Welch, Terry A. [1984] A Technique for High Performance Data Compression. *IEEE Computer*, vol. 17, no. 6 (June 1984), pp. 8-19.

BNSDOCID: <XP_____10039767A__I_>

Electronic Acknowledgement Receipt

EFS ID:	2527196
Application Number:	11623653
International Application Number:	
Confirmation Number:	5243
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Customer Number:	758
Filer:	Brian M. Hoffman
Filer Authorized By:	
Attorney Docket Number:	23682-12381
Receipt Date:	29-NOV-2007
Filing Date:	16-JAN-2007
Time Stamp:	19:03:02
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 Letter	12381_US_IDS.pdf	122937 bb19316af3e939a54c7abd067df71e1cdb94dce8	no	3

Warnings:	
Information:	

2	Information Disclosure Statement (IDS) Filed	12381_US_Form_8A.pdf	111567 66dd8627dcb2e89e6db431b95a1acb62b41f765a	no	1
Warnings:					
Information:					
This is not an USPTO supplied IDS fillable form					
3	Foreign Reference	EP1085444A2.pdf	1107918 c9ec261836d84a9f42b9e014c7fe4db1b0f0dc95	no	17
Warnings:					
Information:					
4	Foreign Reference	EP1085444A3.pdf	163975 fc403db2539727b29d4af6f892f4ce396a31c139	no	5
Warnings:					
Information:					
5	NPL Documents	11745_PCT_ISR_WO.pdf	371998 4f86cc4918901badfac4d8d522f7b2baef797222	no	7
Warnings:					
Information:					
6	NPL Documents	Edwards.pdf	3494867 b4bb28f76032dd59e4d368cc5fbb625c28cc3090	no	11
Warnings:					
Information:					
			Total Files Size (in bytes):	5373262	

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

APPLICANTS:	Justin P. Marston et al.
APPLICATION NO.:	11/623,653
FILING DATE:	January 16, 2007
TITLE:	Determining Relevance of Electronic Content
EXAMINER:	Jeffrey A. Gaffin
GROUP ART UNIT:	2165
ATTY. DKT. NO.:	23682-12381

CERTIFICATE OF ELECTRONIC (EFS-WEB) TRANSMISSION

I hereby certify that this correspondence is being transmitted via the Office electronic filing system in accordance with 37 C.F.R. § 1.8(a)(i)(C) from the Pacific Time Zone of the United States on the local date shown below.

Dated: _____ By: _____

Briam M. Hoffman, Reg. No. 39,713

COMMISSIONER FOR PATENTS
P.O. BOX 1450
ALEXANDRIA, VA 22313-1450

INFORMATION DISCLOSURE STATEMENT
Under 37 CFR §§ 1.56 and 1.97-98

SIR:

Pursuant to the provisions of 37 CFR §§ 1.56 and 1.97-98, enclosed herewith is modified form PTO/SB/08A listing references for consideration by the Examiner.

The filing of this Information Disclosure Statement shall not be construed as a representation regarding the completeness of the list of references, or that inclusion of a reference in this list is an admission that it is prior art or is pertinent to this application, or that a search has been made, or as an admission that the information listed is, or may be considered to be, material to patentability, or that no other material information exists, and shall not be construed as an admission against interest in any manner.

This Information Disclosure Statement is being filed:

☒ within three months of the filing date of the application, or date of entry into the national stage of an international application, or before the mailing date of a first office action on the merits, whichever event last occurred;

- ☐ before the mailing of a first official action after the filing of a request for continued examination (RCE) under 37 CFR § 1.114;
- ☐ after three months of the filing date of this national application or the date of entry of the national stage in an international application, or after the mailing date of the first official action on the merits, whichever event last occurred, but before the mailing date of the first to occur of either: (1) a final action under 37 CFR §1.113; or (2) an action that otherwise closes prosecution in the application, and:
  - ☐ attached hereto is the fee set forth under 37 CFR §1.17(p) for submission of this Information Disclosure Statement under 37 CFR.§ 1.97(c); OR
  - ☐ Applicant certifies pursuant to 37 CFR § 1.97(e) that:
    - ☐ each item of information contained in this Information Disclosure Statement was first cited in a communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of this Statement; OR
    - ☐ no item of information contained in this 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 this certification after making reasonable inquiry, no item of information contained in this Statement was known to any individual designated under 37 CFR § 1.56(c) more than three months prior to the filing of this Statement;
- ☐ on or before the payment of the issue fee but after the mailing date of the first to occur of either: (1) a final action under 37 CFR § 1.113; (2) a notice of allowance under 37 CFR § 1.311; or (3) an action that otherwise closes prosecution in the application, and:
  - ☐ Applicant certifies pursuant to 37 CFR. § 1.97(e) that:
    - ☐ each item of information contained in this Information Disclosure Statement was cited in a communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of this Statement;

☐ no item of information contained in this 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 this certification after making reasonable inquiry, no item of information contained in this Statement was known to any individual designated under 37 CFR § 1.56(c) more than three months prior to the filing of this Statement; AND

☐ attached hereto is the fee set forth under 37 CFR §1.17(p) for submission of this Information Disclosure Statement under 37 CFR.§ 1.97(d); OR

☐ after the payment of the issue fee. Applicant requests that the information contained in this Information Disclosure Statement be placed in the file according to 37 CFR § 1.97(i), although the information may not be considered by the USPTO.

☐ This application relies, under 35 U.S.C. § 120, on the earlier filing date of prior application No. _____ , filed on _____ , and the references cited therein are hereby referenced, but are not required to be provided in this application under 37 CFR § 1.98(d).

☐ Each item of information contained in this Information Disclosure Statement was cited in a communication from a foreign patent office in a counterpart application, and the communication was not received by any individual designated in 37 CFR § 1.56(c) more than thirty days prior to the filing of this Information Disclosure Statement. 37 CFR § 1.704(d).

Respectfully submitted,
JUSTIN P. MARSTON ET AL.

Dated: June 16, 2008

By: /Brian Hoffman/
Brian M. Hoffman, Reg. No.: 39,713
Fenwick & West LLP
Silicon Valley Center
801 California Street
Mountain View, CA 94041
Tel.: (415) 875-2484
Fax.: (650) 938-5200

Substitute for form 1449A/PTO				Complete if Known	
INFORMATION DISCLOSURE STATEMENT BY APPLICANT				**Application No.**	11/623,653
				Filing Date	January 16, 2007
				First Named Inventor	Justin P. Marston
				Art Unit	2165
				Examiner Name	Jeffrey A. Gaffin
Sheet	1	of	1	Attorney Docket Number	23682-12381

U.S. PATENT DOCUMENTS				
		Document No.		
Examiner Initials*	Cite No.[1]	Number – Kind Code[2] (if known)	Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document
	A1	US-2004/0078447 A1	04-22-2004	Malik et al.
		US-		
		US-		

FOREIGN PATENT DOCUMENTS					
		Foreign Patent Document			
Examiner Initials*	Cite No.[1]	Country Code[3] – Number[4] Kind Code[5] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	T[6]

OTHER REFERENCES – NON-PATENT LITERATURE DOCUMENTS			
Examiner Initials*	Cite No.[1]	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, page(s), volume-issue number(s), publisher, city and/or country where published	T[6]
	C1	PCT International Search Report and Written Opinion, PCT/US07/60587, February 14, 2008, 8 pages	

Examiner Signature		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.

[1]Applicant's unique citation designation number (optional). [2]See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3]Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4]For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5]Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST. 16 if possible. [6]Applicant is to place a check mark here if English language Translation is attached.

PATENT COOPERATION TREATY

From the INTERNATIONAL SEARCHING AUTHORITY

To: Brian M. Hoffman
Fenwick & West LLP
Silicon Valley Center
801 California Street
Mountain View, CA 94041

Fenwick & West LLP
FEB 19 2008
RECEIVED

PCT

NOTIFICATION OF TRANSMITTAL OF THE INTERNATIONAL SEARCH REPORT AND THE WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY, OR THE DECLARATION

(PCT Rule 44.1)

Applicant's or agent's file reference 12381 PCT	Date of mailing *(day/month/year)* **FOR FURTHER ACTION** See paragraphs 1 and 4 below
International application No. PCT/US 07/60587	International filing date *(day/month/year)* 16 January 2007 (16.01.2007)
Applicant Bluespace Software Corp.	

23682-

1. ☒ The applicant is hereby notified that the international search report and the written opinion of the International Searching Authority have been established and are transmitted herewith.

   Filing of amendments and statement under Article 19:
   The applicant is entitled, if he so wishes, to amend the claims of the international application (see Rule 46):

   When? The time limit for filing such amendments is normally two months from the date of transmittal of the international search report.

   Where? Directly to the International Bureau of WIPO, 34 chemin des Colombettes 1211 Geneva 20, Switzerland, Facsimile No.: +41 22 740 14 35

   For more detailed instructions, see the notes on the accompanying sheet.

2. ☐ The applicant is hereby notified that no international search report will be established and that the declaration under Article 17(2)(a) to that effect and the written opinion of the International Searching Authority are transmitted herewith.

3. ☐ **With regard to the protest** against payment of (an) additional fee(s) under Rule 40.2, the applicant is notified that:

   ☐ the protest together with the decision thereon has been transmitted to the International Bureau together with the applicant's request to forward the texts of both the protest and the decision thereon to the designated Offices.

   ☐ no decision has been made yet on the protest; the applicant will be notified as soon as a decision is made.

4. **Reminders**

   Shortly after the expiration of **18 months** from the priority date, the international application will be published by the International Bureau. If the applicant wishes to avoid or postpone publication, a notice of withdrawal of the international application, or of the priority claim, must reach the International Bureau as provided in Rules 90*bis*.1 and 90*bis*.3, respectively, before the completion of the technical preparations for international publication.

   The applicant may submit comments on an informal basis on the written opinion of the International Searching Authority to the International Bureau. The International Bureau will send a copy of such comments to all designated Offices unless an international preliminary examination report has been or is to be established. These comments would also be made available to the public but not before the expiration of 30 months from the priority date.

   Within **19 months** from the priority date, but only in respect of some designated Offices, a demand for international preliminary examination must be filed if the applicant wishes to postpone the entry into the national phase **until 30 months** from the priority date (in some Offices even later); otherwise, the applicant must, **within 20 months** from the priority date, perform the prescribed acts for entry into the national phase before those designated Offices.

   In respect of other designated Offices, the time limit of **30 months** (or later) will apply even if no demand is filed within 19 months.

   See the Annex to Form PCT/IB/301 and, for details about the applicable time limits, Office by Office, see the *PCT Applicant's Guide*, Volume II, National Chapters and the WIPO Internet site.

Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201	Authorized officer: Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774

Form PCT/ISA/220 (January 2004)

(See notes on accompanying sheet)

PATENT COOPERATION TREATY

From the INTERNATIONAL SEARCHING AUTHORITY

To: Brian M. Hoffman
Fenwick & West LLP
Silicon Valley Center
801 California Street
Mountain View, CA 94041

PCT

NOTIFICATION OF TRANSMITTAL OF THE INTERNATIONAL SEARCH REPORT AND THE WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY, OR THE DECLARATION

(PCT Rule 44.1)

	Date of mailing *(day/month/year)* 14 FEB 2008
Applicant's or agent's file reference 12381 PCT	**FOR FURTHER ACTION** See paragraphs 1 and 4 below
International application No. PCT/US 07/60587	International filing date *(day/month/year)* 16 January 2007 (16.01.2007)
Applicant Bluespace Software Corp.	

1. ☒ The applicant is hereby notified that the international search report and the written opinion of the International Searching Authority have been established and are transmitted herewith.

   Filing of amendments and statement under Article 19:
   The applicant is entitled, if he so wishes, to amend the claims of the international application (see Rule 46):

   When? The time limit for filing such amendments is normally two months from the date of transmittal of the international search report.

   Where? Directly to the International Bureau of WIPO, 34 chemin des Colombettes 1211 Geneva 20, Switzerland, Facsimile No.: +41 22 740 14 35

   For more detailed instructions, see the notes on the accompanying sheet.

2. ☐ The applicant is hereby notified that no international search report will be established and that the declaration under Article 17(2)(a) to that effect and the written opinion of the International Searching Authority are transmitted herewith.

3. ☐ **With regard to the protest** against payment of (an) additional fee(s) under Rule 40.2, the applicant is notified that:

   ☐ the protest together with the decision thereon has been transmitted to the International Bureau together with the applicant's request to forward the texts of both the protest and the decision thereon to the designated Offices.

   ☐ no decision has been made yet on the protest; the applicant will be notified as soon as a decision is made.

4. **Reminders**

   Shortly after the expiration of **18 months** from the priority date, the international application will be published by the International Bureau. If the applicant wishes to avoid or postpone publication, a notice of withdrawal of the international application, or of the priority claim, must reach the International Bureau as provided in Rules 90*bis*.1 and 90*bis*.3, respectively, before the completion of the technical preparations for international publication.

   The applicant may submit comments on an informal basis on the written opinion of the International Searching Authority to the International Bureau. The International Bureau will send a copy of such comments to all designated Offices unless an international preliminary examination report has been or is to be established. These comments would also be made available to the public but not before the expiration of 30 months from the priority date.

   Within **19 months** from the priority date, but only in respect of some designated Offices, a demand for international preliminary examination must be filed if the applicant wishes to postpone the entry into the national phase **until 30 months** from the priority date (in some Offices even later); otherwise, the applicant must, **within 20 months** from the priority date, perform the prescribed acts for entry into the national phase before those designated Offices.

   In respect of other designated Offices, the time limit of **30 months** (or later) will apply even if no demand is filed within 19 months.

   See the Annex to Form PCT/IB/301 and, for details about the applicable time limits, Office by Office, see the *PCT Applicant's Guide*, Volume II, National Chapters and the WIPO Internet site.

Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201	Authorized officer: Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774

Form PCT/ISA/220 (January 2004) *(See notes on accompanying sheet)*

PATENT COOPERATION TREATY

PCT

INTERNATIONAL SEARCH REPORT

(PCT Article 18 and Rules 43 and 44)

Applicant's or agent's file reference 12381 PCT	**FOR FURTHER ACTION**	see Form PCT/ISA/220 as well as, where applicable, item 5 below.
International application No. PCT/US 07/60587	International filing date *(day/month/year)* 16 January 2007 (16.01.2007)	(Earliest) Priority Date *(day/month/year)* 13 January 2006 (13.01.2006)
Applicant Bluespace Software Corp.		

This international search report has been prepared by this International Searching Authority and is transmitted to the applicant according to Article 18. A copy is being transmitted to the International Bureau.

This international search report consists of a total of 2 sheets.

☐ It is also accompanied by a copy of each prior art document cited in this report.

1. **Basis of the report**
   a. With regard to the **language**, the international search was carried out on the basis of:
      ☒ the international application in the language in which it was filed.
      ☐ a translation of the international application into _____ which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b)).
   b. ☐ This international search report has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43.6*bis*(a)).
   c. ☐ With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, see Box No. I.

2. ☐ **Certain claims were found unsearchable** (see Box No. II).

3. ☐ **Unity of invention is lacking** (see Box No. III).

4. With regard to the **title**,
   ☒ the text is approved as submitted by the applicant.
   ☐ the text has been established by this Authority to read as follows:

5. With regard to the **abstract**,
   ☒ the text is approved as submitted by the applicant.
   ☐ the text has been established, according to Rule 38.2(b), by this Authority as it appears in Box No. IV. The applicant may, within one month from the date of mailing of this international search report, submit comments to this Authority.

6. With regard to the **drawings**,
   a. the figure of the **drawings** to be published with the abstract is Figure No. 2
      ☐ as suggested by the applicant.
      ☒ as selected by this Authority, because the applicant failed to suggest a figure.
      ☐ as selected by this Authority, because this figure better characterizes the invention.
   b. ☐ none of the figures is to be published with the abstract.

Form PCT/ISA/210 (first sheet) (April 2007)

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US 07/60587

A. CLASSIFICATION OF SUBJECT MATTER
IPC(8) - G06F 7/00 (2007.01)
USPC - 707/5
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
707/5

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
USPC: 707/1, 2, 3

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
PubWEST (PGPB,USPT,USOC,EPAB,JPAB), Google Scholar - terms: relev$, similar$, profile, preference, manag$, user$, electronic, mail, email, content

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2004/0078447 A1 (MALIK et al.) 22 April 2004 (22.04.2004), para [0035], [0045], [0013], [0014], [0040]-[0046], [0072-0074], [0076], [0079] and [0080].	1-22
A	US 6,134,582 A (KENNEDY) 17 October 2000 (17.10.2000), col 5-12.	1-22
A	US 2005/0038863 A1 (ONYON et al.) 17 February 2005 (17.02.2005), para [0020]-[0040].	1-22

☐ Further documents are listed in the continuation of Box C. ☐

* Special categories of cited documents:
"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
26 September 2007 (26.09.2007)	14 FEB 2008
Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201	Authorized officer: Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774

Form PCT/ISA/210 (second sheet) (April 2007)

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

To: Brian M. Hoffman
Fenwick & West LLP
Silicon Valley Center
801 California Street
Mountain View, CA 94041

PCT

WRITTEN OPINION OF THE
INTERNATIONAL SEARCHING AUTHORITY

(PCT Rule 43*bis*.1)

Date of mailing *(day/month/year)*	14 FEB 2008
Applicant's or agent's file reference 12381 PCT	**FOR FURTHER ACTION** See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US 07/60587	16 January 2007 (16.01.2007)	13 January 2006 (13.01.2006)

International Patent Classification (IPC) or both national classification and IPC
IPC(8) - G06F 7/00 (2007.01)
USPC - 707/5

Applicant Bluespace Software Corp.

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☐ Box No. VII Certain defects in the international application
- ☐ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notified the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

3. For further details, see notes to Form PCT/ISA/220.

Name and mailing address of the ISA/US	Date of completion of this opinion	Authorized officer:
Mail Stop PCT, Attn: ISA/US Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201	26 September 2007 (26.09.2007)	Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774

Form PCT/ISA/237 (cover sheet) (April 2007)

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

International application No.
PCT/US 07/60587

Box No. I Basis of this opinion

1. With regard to the **language**, this opinion has been established on the basis of:

   ☒ the international application in the language in which it was filed.

   ☐ a translation of the international application into _____ which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b)).

2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43*bis*.1(a))

3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of:

   a. type of material

      ☐ a sequence listing

      ☐ table(s) related to the sequence listing

   b. format of material

      ☐ on paper

      ☐ in electronic form

   c. time of filing/furnishing

      ☐ contained in the international application as filed

      ☐ filed together with the international application in electronic form

      ☐ furnished subsequently to this Authority for the purposes of search

4. ☐ In addition, in the case that more than one version or copy of a sequence listing and/or table(s) relating thereto has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

5. Additional comments:

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

International application No.
PCT/US 07/60587

Box No. V **Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement**

1. Statement

Novelty (N)	Claims	None	YES
	Claims	1-22	NO
Inventive step (IS)	Claims	None	YES
	Claims	1-22	NO
Industrial applicability (IA)	Claims	1-22	YES
	Claims	None	NO

2. Citations and explanations:

Claims 1-22 lack novelty under PCT Article 33(2) as being anticipated by US 2004/0078447 A1 to Malik et al. (hereinafter 'Malik').

As per claims 1 and 11, Malik discloses a system or a computer program product having a computer-readable medium having computer program instructions encoded thereon (see fig 1, 2A and 2B), comprising: a monitoring module configured to track actions by a plurality of users associated with an application for managing electronic content (see para [0035] and [0045]), a relevance analysis module coupled to the monitoring module and configured to determine a relevance to a user of a piece of electronic content sent from a sender to the user (user(s) profile, see abastract and para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]), the relevance being determined dependent on: a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]), and a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claim 17, Malik discloses a method, comprising: tracking actions by a plurality of users associated with an application for managing electronic content (see para [0035] and [0045]), determining a relevance to a user of a piece of electronic content sent from a sender to the user (user(s) profile, see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]), the relevance being determined dependent on: a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]), and a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claims 2, 12 and 18, Malik further discloses the system of claim 1, wherein the relevance is further determined dependent on a measure of local importance describing an importance of the sender to the user (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claims 3, 13 and 19, Malik further discloses the system of claim 1, wherein the relevance is further determined dependent on a measure of global importance describing an importance of the sender within an organization with which the plurality of users are associated (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claims 4, 14 and 20, Malik further discloses the system of claim 1, wherein the relevance analysis module is further configured to identify at least one other of the plurality of users that is a recipient of the piece of electronic content and is cognate to the user and wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content dependent upon one or more actions of the at least one cognate recipient with respect to the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claims 5, 15 and 21, Malik further discloses the system of claim 1, wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content responsive to contents of the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claims 6, 16 and 22, Malik further discloses the system of claim 1, wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content responsive to one or more measures from the set consisting of: a local importance describing an importance of the sender to the user (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]), a global importance describing an importance of the sender within an organization with which the plurality of users are associated (see para [0010]), a content relevance determined responsive to contents of the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]) and one or more actions with respect to the piece of electronic content of one or more of the plurality of users determined to be cognate to the user (see para [0010]).

---------- Continued in Supplemental Box ---------

Supplemental Box

In case **the space in any of the preceding boxes is not sufficient.**

Continuation of:
Box No. V (Citations and explanations):

As per claim 7, Malik further discloses the system of claim 1, further comprising: a content presentation module configured to present the piece of electronic content to the user in accordance with the determined relevance and to present to the user a value describing the determined relevance (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claim 8, Malik further discloses the system of claim 1, further comprising: a search module configured to search the electronic content in accordance with the determined relevance (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claim 9, Malik further discloses the system of claim 1, wherein the piece of electronic content comprises an electronic mail object in a relational mail system (see para [0072], [0074], [0076], [0079] and [0080]).

As per claim 10, Malik further discloses the system of claim 1, wherein the electronic content comprises at least one of an email, a document, an instant message, and an SMS message (at least one of an email, see para [0072], [0074], [0076], [0079] and [0080]).

Claims 1-22 have industrial applicability as defined by PCT Article 33(4) because the subject matter claimed can be made or used in industry.

Electronic Acknowledgement Receipt

EFS ID:	3468558
Application Number:	11623653
International Application Number:	
Confirmation Number:	5243
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Customer Number:	758
Filer:	Brian M. Hoffman
Filer Authorized By:	
Attorney Docket Number:	23682-12381
Receipt Date:	18-JUN-2008
Filing Date:	16-JAN-2007
Time Stamp:	18:21:36
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 Letter	12381IDS.pdf	122709 f58a87147f6100d21c6ed171ae1ce1051fc625fb	no	3

Warnings:	
Information:	

2	Information Disclosure Statement (IDS) Filed	12381Form8A.pdf	110031 fba190f0ffd517d7008e1ef6a3d68075f991589d	no	1
Warnings:					
Information:					
This is not an USPTO supplied IDS fillable form					
3	NPL Documents	12381PCT-PCT-US07-60587-ISR-WO.pdf	513104 63999379a37a7f9fcacaa2665e482bf91ccca635	no	8
Warnings:					
Information:					
		Total Files Size (in bytes):	745844		

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

APPLICANTS:	Justin P. Marston et al.
APPLICATION NO.:	11/623,653
FILING DATE:	January 16, 2007
TITLE:	Determining Relevance of Electronic Content
EXAMINER:	Jeffrey A. Gaffin
GROUP ART UNIT:	2165
ATTY. DKT. NO.:	23682-12381

CERTIFICATE OF ELECTRONIC (EFS-WEB) TRANSMISSION

I hereby certify that this correspondence is being transmitted via the Office electronic filing system in accordance with 37 C.F.R. § 1.8(a)(i)(C) from the Pacific Time Zone of the United States on the local date shown below.

Dated: _____ By: _____

Briam M. Hoffman, Reg. No. 39,713

COMMISSIONER FOR PATENTS
P.O. BOX 1450
ALEXANDRIA, VA 22313-1450

SUPPLEMENTAL INFORMATION DISCLOSURE STATEMENT
Under 37 CFR §§ 1.56 and 1.97-98

SIR:

Pursuant to the provisions of 37 CFR §§ 1.56 and 1.97-98, enclosed herewith is modified form PTO/SB/08A listing references for consideration by the Examiner.

The filing of this Information Disclosure Statement shall not be construed as a representation regarding the completeness of the list of references, or that inclusion of a reference in this list is an admission that it is prior art or is pertinent to this application, or that a search has been made, or as an admission that the information listed is, or may be considered to be, material to patentability, or that no other material information exists, and shall not be construed as an admission against interest in any manner.

This Information Disclosure Statement is being filed:

☒ within three months of the filing date of the application, or date of entry into the national stage of an international application, or before the mailing date of a first office action on the merits, whichever event last occurred;

- ☐ before the mailing of a first official action after the filing of a request for continued examination (RCE) under 37 CFR § 1.114;
- ☐ after three months of the filing date of this national application or the date of entry of the national stage in an international application, or after the mailing date of the first official action on the merits, whichever event last occurred, but before the mailing date of the first to occur of either: (1) a final action under 37 CFR §1.113; or (2) an action that otherwise closes prosecution in the application, and:
  - ☐ attached hereto is the fee set forth under 37 CFR §1.17(p) for submission of this Information Disclosure Statement under 37 CFR.§ 1.97(c); OR
  - ☐ Applicant certifies pursuant to 37 CFR § 1.97(e) that:
    - ☐ each item of information contained in this Information Disclosure Statement was first cited in a communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of this Statement; OR
    - ☐ no item of information contained in this 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 this certification after making reasonable inquiry, no item of information contained in this Statement was known to any individual designated under 37 CFR § 1.56(c) more than three months prior to the filing of this Statement;
- ☐ on or before the payment of the issue fee but after the mailing date of the first to occur of either: (1) a final action under 37 CFR § 1.113; (2) a notice of allowance under 37 CFR § 1.311; or (3) an action that otherwise closes prosecution in the application, and:
  - ☐ Applicant certifies pursuant to 37 CFR. § 1.97(e) that:
    - ☐ each item of information contained in this Information Disclosure Statement was cited in a communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of this Statement;

☐ no item of information contained in this 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 this certification after making reasonable inquiry, no item of information contained in this Statement was known to any individual designated under 37 CFR § 1.56(c) more than three months prior to the filing of this Statement; AND

☐ attached hereto is the fee set forth under 37 CFR §1.17(p) for submission of this Information Disclosure Statement under 37 CFR.§ 1.97(d); OR

☐ after the payment of the issue fee. Applicant requests that the information contained in this Information Disclosure Statement be placed in the file according to 37 CFR § 1.97(i), although the information may not be considered by the USPTO.

☐ This application relies, under 35 U.S.C. § 120, on the earlier filing date of prior application No. _____ , filed on _____ , and the references cited therein are hereby referenced, but are not required to be provided in this application under 37 CFR § 1.98(d).

☐ Each item of information contained in this Information Disclosure Statement was cited in a communication from a foreign patent office in a counterpart application, and the communication was not received by any individual designated in 37 CFR § 1.56(c) more than thirty days prior to the filing of this Information Disclosure Statement. 37 CFR § 1.704(d).

Respectfully submitted,
JUSTIN P. MARSTON ET AL.

Dated: July 22, 2008

By: /Brian Hoffman/
Brian M. Hoffman, Reg. No.: 39,713
Fenwick & West LLP
Silicon Valley Center
801 California Street
Mountain View, CA 94041
Tel.: (415) 875-2484
Fax.: (650) 938-5200

Substitute for form 1449A/PTO				Complete if Known	
INFORMATION DISCLOSURE STATEMENT BY APPLICANT				**Application No.**	11/623,653
				Filing Date	January 16, 2007
				First Named Inventor	Justin P. Marston
				Art Unit	2165
				Examiner Name	Jeffrey A. Gaffin
Sheet	1	of	1	Attorney Docket Number	23682-12381

U.S. PATENT DOCUMENTS				
		Document No.		
Examiner Initials*	Cite No.[1]	Number – Kind Code[2] (if known)	Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document
	A1	US-2007/0288575 A1	12-13-2007	Gillum et al.
	A2	US-2005/0203929 A1	09-15-2005	Hazarika et al.
	A3	US-2005/0204009 A1	09-15-2005	Hazarika et al.

FOREIGN PATENT DOCUMENTS					
		Foreign Patent Document			
Examiner Initials*	Cite No.[1]	Country Code[3] – Number[4] Kind Code[5] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	T[6]

OTHER REFERENCES – NON-PATENT LITERATURE DOCUMENTS			
Examiner Initials*	Cite No.[1]	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, page(s), volume-issue number(s), publisher, city and/or country where published	T[6]

Examiner Signature		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.

[1]Applicant's unique citation designation number (optional). [2]See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3]Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4]For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5]Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST. 16 if possible. [6]Applicant is to place a check mark here if English language Translation is attached.

Electronic Acknowledgement Receipt

EFS ID:	3659317
Application Number:	11623653
International Application Number:	
Confirmation Number:	5243
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Customer Number:	758
Filer:	Brian M. Hoffman/meera yang
Filer Authorized By:	Brian M. Hoffman
Attorney Docket Number:	23682-12381
Receipt Date:	22-JUL-2008
Filing Date:	16-JAN-2007
Time Stamp:	17:35:03
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		12381_US_SIDS.pdf	332928 ddaeb337ab584b942aadccb9871440be91ecf646	yes	4

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Information Disclosure Statement Letter	1	3
	Information Disclosure Statement (IDS) Filed	4	4
Warnings:			
Information:			
	Total Files Size (in bytes):	332928	

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 NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
11/623,653	01/16/2007	Justin Marston	23682-12381	5243

758 7590 11/14/2008

FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041

EXAMINER	
HOCKER, JOHN P	
ART UNIT	**PAPER NUMBER**
2169	

MAIL DATE	DELIVERY MODE
11/14/2008	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)	
	11/623,653	MARSTON ET AL.	
	Examiner	Art Unit	
	JOHN P. HOCKER	2169	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE *3* MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *16 January 2007*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-22* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-22* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

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)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☒ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date *11/29/2007, 06/18/2008 and 07/22/2008* .
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

1. Claims 1 – 22 have been examined and are pending.

2. Claims 1 – 22 have been rejected (Non-Final Rejection).

Information Disclosure Statement

The information disclosure statement (IDS) submitted on 11/29/2007, 06/18/2008 and 07/22/2008 were received by examiner. The submission is in compliance with the provisions of 37 CFR 1.97. Accordingly, the information disclosure statement is being considered by the examiner.

Related Cases

This case is related to copending application no. 11/426118.

This case is related to PCT Written Opinion and Search Reports No. PCT/US06/24859 and No. PCT/US06/60587.

Identical claims 1 – 22 were submitted in a PCT application. Examiner finds no reason to disagree with the written opinion in the PCT International Search Report. (PCT/US07/60587) (February 14, 2008, 8 pages).

The PCT International Search Report and Written Opinion found the claims lacked novelty as anticipated by the Malik Reference

Claim Rejections - 35 USC § 102 and § 35 USC 103

Examiner's Note: The requirements for patentability are that a claimed invention must at a minimum meet the requirements of utility (35 USC 101), novelty (35 USC 102) and non-obviousness (35 USC 103). The claims are rejected under 102 as they seemingly are anticipated by the prior art. Alternatively, the claims are rejected under 103 as they are obvious in view of the prior art of record and the common sense of one of ordinary skill.

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

> (b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> **(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title,** if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. **Patentability shall not be negatived by the manner in which the invention was made.**

Claims 1 - 22 rejected under 35 U.S.C. 102(b) as anticipated by or, in the alternative, under 35 U.S.C. 103(a) as obvious over Malik (US Publication No. 2004/0078447)

The specification describes the exemplary embodiment of the claimed system to be what is known as an email system (Specification, pg. 1). The claim's allegedly novel and non-obvious feature is the determination of user relevance for the email (i.e. electronic content). (Specification pgs 1 – 2) This argument is unpersuasive though as shown in Malik. In the same fashioon, the email system of Malik utilizes user profiles to determine user relevance of electronic content.

As per claims 1 and 11, Malik discloses a system and program (see fig 1,2A and 2B), comprising:

(i) a monitoring module configured to track actions by a plurality of users associated with an application for managing electronic content (see para [0035] and [0045]),

(ii) a relevance analysis module coupled to the monitoring module and configured to determine a relevance to a user of a received piece of electronic content (user(s) profile, see abstract and para [0013], [0014], [0040]-[0046], [0072-0074] and [0079 – 0080]),

(iii) the relevance being determined dependent on: a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079] - 0080),

and (iv) a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079 - 0080]).

As per claims 2, 12 and 18, Malik further discloses the system of claim 1, wherein the relevance is further determined dependent on a measure of local importance describing an importance of the sender to the user (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079])

As per claims 3, 13 and 19, Malik further discloses the system of claim 1, wherein the relevance is further determined dependent on a measure of global importance describing an

importance of the sender within an organization with which the plurality of users are associated (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claims 4, 14 and 20, Malik further discloses the system of claim 1, wherein the relevance analysis module is further configured to identify at least one other of the plurality of users that is a recipient of the piece of electronic content and is cognate to the user and wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content dependent upon one or more actions of the at least one cognate recipient with respect to the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claims 5, 15 and 21, Malik further discloses the system of claim 1, wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content responsive to contents of the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claims 6, 16 and 22, Malik further discloses the system of claim 1, wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content responsive to one or more measures from the set consisting of: a ocal importance describing an importance of the sender to the user (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]), a i g obal mportance describing an importance of the sender within an organization with which the plurality of users are associated (see para [0010]), a

content relevance determined responsive to contents of the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]) and one or more actions with respect to the piece of electronic content of one or more of the plurality of users determined to be cognate to the user (see para [0010]).

As per claim 7, Malik further discloses the system of claim 1, further comprising: a content presentation module configured to present the piece of electronic content to the user in accordance with the determined relevance and to present to the user a value describing the determined relevance (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).
As per claim 8, Malik further discloses the system of claim 1, further comprising: a search module configured to search the electronic content in accordance with the determined relevance (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claim 9, Malik further discloses the system of claim 1, wherein the piece of electronic content comprises an electronic mail object in a relational mail system (see para [0072], [0074], [0076], [0079] and [0080]).

As per claim 10, Malik further discloses the system of claim 1, wherein the electronic content comprises at least one of an email, a document, an instant message, and an SMS message (at least one of an email, see para [0072], [0074], [0076], [0079] and [0080]).

As per claim 17, Malik discloses a method, comprising: tracking actions by a plurality of users associated with an application for managing electronic content (see para [0035] and [0045]), determining a relevance to a user of a piece of electronic content sent from a sender to the user (user(s) profile, see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]), the relevance being determined dependent on: a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]), and a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

Examiner's note: Although seemingly every claim limitation is taught by the prior art of record, it is arguable that each and every word is not verbatim 102 anticipated by the prior art. Examiner asserts that nonetheless, the claimed invention is properly rejected under 103 if the differences are insubstantial and the concept as a whole would be obvious to one of ordinary skill in the art with common sense which includes well known techniques and known substitutes. Examiner asserts the claim mappings above show the state of the prior art.

Examiner's note: In evaluating applicant's claims to patentability, Examiner was not persuaded by applicant's description of the prior art. Specifically, applicant argues that the prior art teaches ""unintelligent mechanisms" for storage and sorting of email (Specification [0004]. Examiner could not disagree more, the "intelligent" features applicant claims are the determination of relevance for a user. In the same respect, the email system of Malik determines

relevance of the content based on a user profile (i.e. based on a user). Thus applicant's claimed Background and Summary is a mischaracterization of the state of the prior art and the argument for patentability relies on this mischaracterization. The fact is that the prior art Malik teaches this intelligent mechanism for email storage and sorting. Specifically, the intelligent mechanism is for presenting the email based on user relevance (i.e. only pertinent to certain user profiles). Applicant has not named a credible improvement over the prior art.

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to John P. Hocker whose telephone number is 5712705004. The examiner can normally be reached on Mon - Thu 11:00 - 4:00 EST.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Mohammad Ali can be reached on 5712724105. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/John P. Hocker/ Art Unit 2169

/Etienne P LeRoux/
Primary Examiner, Art Unit 2161

Notice of References Cited	Application/Control No. 11/623,653	Applicant(s)/Patent Under Reexamination MARSTON ET AL.	
	Examiner JOHN P. HOCKER	Art Unit 2169	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	Classification
*	A	US-2004/0078447 A1	04-2004	Malik et al.	709/206
	B	US-			
	C	US-			
	D	US-			
	E	US-			
	F	US-			
	G	US-			
	H	US-			
	I	US-			
	J	US-			
	K	US-			
	L	US-			
	M	US-			

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	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.

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	19391	(707/1-3).ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/11/09 15:27
L2	3609	1 and (email or (electronic adj3 mail))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/11/09 15:28
L3	1159	2 and (profile)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/11/09 15:28
L4	1151	3 and user	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/11/09 15:29
L5	757	4 and (user near5 profile)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/11/09 15:29
L6	656	5 and (search)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/11/09 15:29
L7	404	6 and (track or tracking)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/11/09 15:30

L8	8	("20020122543" | "20020136279" | "20030191806" | "20040133645").PN.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/11/09 15:31
S1	0	("20040078447").PN.	USPAT; USOCR	OR	OFF	2008/11/09 11:58
S2	2	"20040078447"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/11/09 11:58

11/9/2008 3:44:57 PM
C:\Documents and Settings\jhocker\My Documents\EAST\Workspaces\Hocker Workspaces\11623653.wsp

Substitute for form 1449A/PTO				Complete if Known	
INFORMATION DISCLOSURE STATEMENT BY APPLICANT				**Application No.**	11/623,653
				Filing Date	January 16, 2007
				First Named Inventor	Justin P. Marston
				Art Unit	2165
				Examiner Name	Jeffrey A. Gaffin
Sheet	1	of	1	Attorney Docket Number	23682-12381

U.S. PATENT DOCUMENTS				
		Document No.		
Examiner Initials*	Cite No.[1]	Number – Kind Code[2] (if known)	Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document
	A1	US-2004/0078447 A1	04-22-2004	Malik et al.
		US-		
		US-		

FOREIGN PATENT DOCUMENTS					
		Foreign Patent Document			
Examiner Initials*	Cite No.[1]	Country Code[3] – Number[4] Kind Code[5] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	T[6]

OTHER REFERENCES – NON-PATENT LITERATURE DOCUMENTS			
Examiner Initials*	Cite No.[1]	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, page(s), volume-issue number(s), publisher, city and/or country where published	T[6]
	C1	PCT International Search Report and Written Opinion, PCT/US07/60587, February 14, 2008, 8 pages	

Examiner Signature	/John Hocker/	Date Considered	11/09/2008

*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.

[1]Applicant's unique citation designation number (optional). [2]See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3]Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4]For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5]Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST. 16 if possible. [6]Applicant is to place a check mark here if English language Translation is attached.

23682/12381/DOCS/1925025.1

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /J.H./

Substitute for form 1449A/PTO				Complete if Known	
INFORMATION DISCLOSURE STATEMENT BY APPLICANT				**Application No.**	11/623,653
				Filing Date	January 16, 2007
				First Named Inventor	Justin P. Marston
				Art Unit	2165
				Examiner Name	Jeffrey A. Gaffin
Sheet	1	of	1	Attorney Docket Number	23682-12381

U.S. PATENT DOCUMENTS				
		Document No.		
Examiner Initials*	Cite No.[1]	Number – Kind Code[2] (if known)	Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document
	A1	US-2007/0288575 A1	12-13-2007	Gillum et al.
	A2	US-2005/0203929 A1	09-15-2005	Hazarika et al.
	A3	US-2005/0204009 A1	09-15-2005	Hazarika et al.

FOREIGN PATENT DOCUMENTS					
		Foreign Patent Document			
Examiner Initials*	Cite No.[1]	Country Code[3] – Number[4] Kind Code[5] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	T[6]

OTHER REFERENCES – NON-PATENT LITERATURE DOCUMENTS			
Examiner Initials*	Cite No.[1]	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, page(s), volume-issue number(s), publisher, city and/or country where published	T[6]

Examiner Signature	/John Hocker/	Date Considered	11/09/2008

*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.

[1]Applicant's unique citation designation number (optional). [2]See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3]Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4]For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5]Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST. 16 if possible. [6]Applicant is to place a check mark here if English language Translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /J.H./

Substitute for form 1449A/PTO				Complete if Known	
INFORMATION DISCLOSURE STATEMENT BY APPLICANT				**Application No.**	11/623,653
				Filing Date	January 16, 2007
				First Named Inventor	Justin Marston
				Art Unit	2165
				Examiner Name	Jeffrey A. Gaffin
Sheet	1	of	1	Attorney Docket Number	23682-12381

U.S. PATENT DOCUMENTS				
		Document No.		
Examiner Initials*	Cite No.[1]	Number – Kind Code[2] (if known)	Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document
	A1	US-2004/0133645 A1	07-08-2004	Massanelli et al.
	A2	US-2003/0191806 A1	10-09-2003	Osterberg, Jr., et al.
	A3	US-2002/0136279 A1	09-26-2002	Al-Kazily et al.
	A4	US-2002/0122543 A1	09-05-2002	Rowen

FOREIGN PATENT DOCUMENTS					
		Foreign Patent Document			
Examiner Initials*	Cite No.[1]	Country Code[3] – Number[4] Kind Code[5] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	T[6]
	B1	EP 1 085 444 A2	03-21-2001	Microsoft Corporation	
	B2	EP 1 085 444 A3	04-09-2003	Microsoft Corporation	

OTHER REFERENCES – NON-PATENT LITERATURE DOCUMENTS			
Examiner Initials*	Cite No.[1]	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, page(s), volume-issue number(s), publisher, city and/or country where published	T[6]
	C1	International Search Report and Written Opinion, PCT/US06/24859, April 26, 2007, 7 pages.	
	C2	EDWARDS, K., "The Design and Implementation of the MONTAGE Multimedia Mail System," Communications for Distributed Applications and Systems, Chapel Hill, April 18, 1991, pp. 47-57, Proceedings of the Conference on Communications Software, New York, U.S.	

Examiner Signature	/John Hocker/	Date Considered	11/09/2008

*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.

[1]Applicant's unique citation designation number (optional). [2]See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3]Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4]For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5]Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST. 16 if possible. [6]Applicant is to place a check mark here if English language Translation is attached.

23682/12381/DOCS/1829301.1

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /J.H./

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. 5243

SERIAL NUMBER	FILING or 371(c) DATE RULE	CLASS	GROUP ART UNIT	ATTORNEY DOCKET NO.
11/623,653	01/16/2007	707	2169	23682-12381

APPLICANTS

Justin Marston, Austin, TX;
Paul Marston, Lancashire, UNITED KINGDOM;

**** CONTINUING DATA ***************************

This appln claims benefit of 60/758,828 01/13/2006

**** FOREIGN APPLICATIONS ***************************

**** IF REQUIRED, FOREIGN FILING LICENSE GRANTED **** ** SMALL ENTITY **
02/12/2007

Foreign Priority claimed ☐ Yes ☑ No 35 USC 119(a-d) conditions met ☐ Yes ☑ No Verified and Acknowledged /JOHN P HOCKER/ Examiner's Signature	☐ Met after Allowance Initials	STATE OR COUNTRY	SHEETS DRAWINGS	TOTAL CLAIMS	INDEPENDENT CLAIMS
		TX	22	22	3

ADDRESS

FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041
UNITED STATES

TITLE

DETERMINING RELEVANCE OF ELECTRONIC CONTENT

FILING FEE RECEIVED		
540	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

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	11623653	MARSTON ET AL.
	Examiner	Art Unit
	JOHN P HOCKER	2169

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	11/09/2008								
	1	✓								
	2	✓								
	3	✓								
	4	✓								
	5	✓								
	6	✓								
	7	✓								
	8	✓								
	9	✓								
	10	✓								
	11	✓								
	12	✓								
	13	✓								
	14	✓								
	15	✓								
	16	✓								
	17	✓								
	18	✓								
	19	✓								
	20	✓								
	21	✓								
	22	✓								

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	11623653	MARSTON ET AL.
	Examiner	Art Unit
	JOHN P HOCKER	2169

SEARCHED			
Class	Subclass	Date	Examiner
707	1-3	11/09/2008	JPH

SEARCH NOTES		
Search Notes	Date	Examiner
Inventive concept: " User Profiles for electronic content viewing (e.g. email)"	11/09/2008	JPH
1: IDS Reviewed for Prior Art.	11/09/2008	JPH
2: Examiner performed inventor name search.	11/09/2008	JPH
3: Examiner searched 'Computer Technology' NPL (e.g. IEEE (Institute of Electrical and Electronics Engineers), Association for Computing Machinery, Wikipedia and/or Google Scholar)	11/09/2008	JPH
4: Examiner searched International Technical & Patent Databases(US-PGPUB, USPAT, USOCR, FPRS, EPO, JPO, DERWENT and IBM_TDB).	11/09/2008	JPH
During Patent Searches, Examiner searched class/subclasses (as listed above) combined with text of the inventive concept. Please see East Search History for text searched.	11/09/2008	JPH

INTERFERENCE SEARCH			
Class	Subclass	Date	Examiner

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.
11/623,653	01/16/2007	Justin Marston	23682-12381	5243

758 7590 12/15/2008

FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041

EXAMINER	
HOCKER, JOHN P	
ART UNIT	PAPER NUMBER
2169	
MAIL DATE	DELIVERY MODE
12/15/2008	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Interview Summary	Application No.	Applicant(s)	
	11/623,653	MARSTON ET AL.	
	Examiner	Art Unit	
	JOHN P. HOCKER	2169	

All participants (applicant, applicant's representative, PTO personnel):

(1) *JOHN P. HOCKER*. (3)_____.

(2) *Chris King*. (4)_____.

Date of Interview: *10 December 2008*.

Type: a)☒ Telephonic b)☐ Video Conference
c)☐ Personal [copy given to: 1)☐ applicant 2)☐ applicant's representative]

Exhibit shown or demonstration conducted: d)☐ Yes e)☐ No.
If Yes, brief description: _____.

Claim(s) discussed: *1*.

Identification of prior art discussed: *Malik (See Non-Final Rejection mailed 11/14/2008)*.

Agreement with respect to the claims f)☐ was reached. g)☒ was not reached. h)☐ N/A.

Substance of Interview including description of the general nature of what was agreed to if an agreement was reached, or any other comments: *Applicant attempted to distinguish claimed invention over prior art of record. Arguments hinged around Claim 1, Limitation (iii) and (iv). Basically, Applicant argued that Malik did not determine relevance based on user actions "on a piece of electronic content". No agreement was reached as to the allowability of the claims.*.

(A fuller description, if necessary, and a copy of the amendments which the examiner agreed would render the claims allowable, if available, must be attached. Also, where no copy of the amendments that would render the claims allowable is available, a summary thereof must be attached.)

THE FORMAL WRITTEN REPLY TO THE LAST OFFICE ACTION MUST INCLUDE THE SUBSTANCE OF THE INTERVIEW. (See MPEP Section 713.04). If a reply to the last Office action has already been filed, APPLICANT IS GIVEN A NON-EXTENDABLE PERIOD OF THE LONGER OF ONE MONTH OR THIRTY DAYS FROM THIS INTERVIEW DATE, OR THE MAILING DATE OF THIS INTERVIEW SUMMARY FORM, WHICHEVER IS LATER, TO FILE A STATEMENT OF THE SUBSTANCE OF THE INTERVIEW. See Summary of Record of Interview requirements on reverse side or on attached sheet.

J.P.H	/Mohammad Ali/ Supervisory Patent Examiner, Art Unit 2169

IN THE UNITED STATES

PATENT AND TRADEMARK OFFICE

APPLICANT:	Justin Marston et al
APPLICATION NO.:	11/623,653
FILING DATE:	January 16, 2007
TITLE:	Determining Relevance of Electronic Content
EXAMINER:	John P. Hocker
GROUP ART UNIT:	2169
ATTY. DKT. NO.:	20423-12381

CERTIFICATE OF ELECTRONIC (EFS-WEB) TRANSMISSION

I hereby certify that this correspondence is being transmitted via the Office electronic filing system in accordance with 37 C.F.R. § 1.8(a)(i)(C) from the **Pacific Time Zone** of the United States on the local date shown below.

Dated: January 5, 2009 By: /Christopher King/
Christopher P. King, Reg. No. 60,985

RESPONSE A

Sir:

In response to the Non-Final Office Action mailed November 14, 2008, please consider the following remarks.

IN THE CLAIMS

No claims are amended, added, or canceled. All pending claims are produced below. In addition, the status of each is also indicated below and appropriately noted as "Original", "Currently Amended", "Canceled", "New", "Withdrawn", "Previously Presented", and "Not Entered" as requested by the Office.

1. (Original) A system, comprising:
   a monitoring module configured to track actions by a plurality of users associated with an application for managing electronic content;
   a relevance analysis module coupled to the monitoring module and configured to determine a relevance to a user of a piece of electronic content sent from a sender to the user, the relevance being determined dependent on:
      a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content, and
      a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content.

2. (Original) The system of claim 1, wherein the relevance is further determined dependent on a measure of local importance describing an importance of the sender to the user.

3. (Original) The system of claim 1, wherein the relevance is further determined dependent on a measure of global importance describing an importance of the sender within an organization with which the plurality of users are associated.

4. (Original) The system of claim 1, wherein the relevance analysis module is further configured to identify at least one other of the plurality of users that is a recipient of the piece of electronic content and is cognate to the user and wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content

dependent upon one or more actions of the at least one cognate recipient with respect to the piece of electronic content.

5. (Original) The system of claim 1, wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content responsive to contents of the piece of electronic content.

6. (Original) The system of claim 1, wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content responsive to one or more measures from the set consisting of:

a local importance describing an importance of the sender to the user;

a global importance describing an importance of the sender within an organization with which the plurality of users are associated;

a content relevance determined responsive to contents of the piece of electronic content; and

one or more actions with respect to the piece of electronic content of one or more of the plurality of users determined to be cognate to the user.

7. (Original) The system of claim 1, further comprising:

a content presentation module configured to present the piece of electronic content to the user in accordance with the determined relevance and to present to the user a value describing the determined relevance.

8. (Original) The system of claim 1, further comprising:

a search module configured to search the electronic content in accordance with the determined relevance.

9. (Original) The system of claim 1, wherein the piece of electronic content comprises an electronic mail object in a relational mail system.

10. (Original) The system of claim 1, wherein the electronic content comprises at least one of an email, a document, an instant message, and an SMS message.

11. (Original) A computer program product having a computer-readable medium having computer program instructions encoded thereon comprising:

a monitoring module configured to track actions by a plurality of users associated with an application for managing electronic content;

a relevance analysis module coupled to the monitoring module and configured to determine a relevance to a user of a piece of electronic content sent from a sender to the user, the relevance being determined dependent on:

a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content, and

a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content.

12. (Original) The computer program product of claim 11, wherein the relevance is further determined dependent on a measure of local importance describing an importance of the sender to the user.

13. (Original) The computer program product of claim 11, wherein the relevance is further determined dependent on a measure of global importance describing an importance of the sender within an organization with which the plurality of users are associated.

14. (Original) The computer program product of claim 11, wherein the relevance analysis module is further configured to identify at least one other of the plurality of users that is a recipient of the piece of electronic content and is cognate to the user and wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content dependent upon one or more actions of the at least one cognate recipient with respect to the piece of electronic content.

15. (Original) The computer program product of claim 11, wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content responsive to contents of the piece of electronic content.

16. (Original) The computer program product of claim 11, wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content responsive to one or more measures from the set consisting of:

a local importance describing an importance of the sender to the user;

a global importance describing an importance of the sender within an organization with which the plurality of users are associated;

a content relevance determined responsive to contents of the piece of electronic content; and

one or more actions with respect to the piece of electronic content of one or more of the plurality of users determined to be cognate to the user.

17. (Original) A method, comprising:

tracking actions by a plurality of users associated with an application for managing electronic content;

determining a relevance to a user of a piece of electronic content sent from a sender to the user, the relevance being determined dependent on:

a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content, and

a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content.

18. (Original) The method of claim 17, wherein the relevance is further determined dependent on a measure of local importance describing an importance of the sender to the user.

19. (Original) The method of claim 17, wherein the relevance is further determined dependent on a measure of global importance describing an importance of the sender within an organization with which the plurality of users are associated.

20. (Original) The method of claim 17, wherein determining the relevance to the user further comprises:

identifying at least one other of the plurality of users that is a recipient of the piece of electronic content and is cognate to the user; and

determining the relevance to the user of the piece of electronic content dependent upon one or more actions of the at least one cognate recipient with respect to the piece of electronic content.

21. (Original) The method of claim 17, wherein the relevance is further dependent on contents of the piece of electronic content.

22. (Original) The method of claim 17, wherein the relevance is further dependent on one or more measures from the set consisting of:

a local importance describing an importance of the sender to the user;

a global importance describing an importance of the sender within an organization with which the plurality of users are associated;

a content relevance determined responsive to contents of the piece of electronic content; and

one or more actions with respect to the piece of electronic content of one or more of the plurality of users determined to be cognate to the user.

REMARKS

Claims 1-22 are pending. In an Office Action dated November 14, 2008, claims 1-22 were rejected. Applicants thank the Examiner for examination of the claims pending in this application and address the Examiner's comments below.

Applicants' representative conducted a telephonic interview with Examiner Hocker on December 10, 2008, during which the rejection of claim 1 was discussed. No agreement was reached. A summary of the interview is incorporated into these Remarks.

In view of the following Remarks, Applicants respectfully request that the Examiner reconsider all outstanding rejections, and withdraw them.

Response to Rejections Under §§ 102 and 103

Claims 1-22 stand rejected under 35 U.S.C. § 102(b) or, in the alternative, under 35 U.S.C. § 103(a), as allegedly being unpatentable over Malik, U.S. Patent Application No. 2004/0078447. These rejections are traversed.

Independent claim 1 specifically recites:

> a monitoring module configured to track actions by a plurality of users associated with an application for managing electronic content;
> a relevance analysis module coupled to the monitoring module and configured to determine a relevance to a user of a piece of electronic content sent from a sender to the user, the relevance being determined dependent on:
> a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user **on the piece of electronic content**, and
> a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content.

Thus, the claimed invention recites determining relevance to a user of a piece of electronic content (e.g., an email). The relevance is determined dependent on a measure of a relevance of

the piece of electronic content as determined dependent on one or more **actions of the user** (and of at least one other of the plurality of users that is a recipient of the piece of electronic content) **on the piece of electronic content**. Examples of actions of the user and others on the piece of electronic content include forwarding, deleting, or printing an electronic email object. (Specification, paragraph 0069).

Malik does not disclose or suggest such features. Malik generally discloses a mailbox system, which differentiates between user and group mailboxes. A user profile can be associated with a user email folder and an instant messaging (IM) account, and this association facilitates defining properties of the user email folder and the user IM account using the user profile. (Malik, Abstract). The user profiles are a collection of settings, such as a listing of a user's email mailboxes and folders, IM accounts, usernames and passwords, address books, and the like. (Malik, paragraphs 0041-0044).

The Examiner cites Malik paragraphs 0013, 0014, 0040-0046, 0072-0074, and 0079-0080 as allegedly disclosing "a measure of relevance of the piece of electronic content as determined dependent one or more actions [of the user, and of at least one other of the plurality of users that is a recipient of the piece of electronic content] on the piece of electronic content." However, all of these portions fail to disclose or suggest that the measure of relevance is determined dependent on actions on **user actions on the piece of electronic content**, as claimed. Cited paragraphs 0013 and 0014 merely disclose that a user profile may be used and that such a profile has a collection of properties related to a user email folder and IM account. Cited paragraphs 0040-0046 disclose what the profile contains, such as usernames and passwords. Cited paragraphs 0072-0074 disclose that a single database used for both email and IM information can facilitate integration of email and IM systems. Cited paragraph 0079 again discloses that email

and IM information can be stored in a single database. Cited paragraph 0080 discloses email options that may be set using a GUI. However, none of these cited paragraphs show determining a measure of relevance of a piece of electronic content dependent on user actions on the piece of electronic content, and so are not relevant to the claimed elements in question.

In the interview of December 10, 2008, the Examiner noted, with respect to the first element ("a monitoring module configured to track actions by a plurality of users associated with an application for managing electronic content"), that Malik paragraph 0035 states that in some server embodiments, server processes are configured to track email access by individual users. However, such a general statement in no way discloses or suggests the specific elements recited by the claim—that is, the mere statement that a server may track email access does not show **determining a measure of relevance of a piece of electronic content** based on user actions on the piece of electronic content. Thus, Malik neither anticipates the claimed invention nor renders it obvious.

Independent claims 11 and 17 recite, for example, "a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content" and are allowable for at least the same reasons discussed above with respect to independent claim 1.

The remaining claims depend, directly or indirectly, from one of independent claims 1, 11, 17, and recite additional patentably distinguishable features and limitations. Thus, they are allowable for at least the same reasons discussed above with regard to their respective independent claims.

Based on the above amendment and the remarks, Applicants respectfully submit that for at least these reasons claims 1-22 are patentably distinguishable over the applied reference.

Therefore, Applicants respectfully request that the Examiner reconsider the rejections, and withdraw them.

Note that Applicants have not substantively amended independent claims 1, 11, or 17. Thus, any subsequent Office Action rejecting claims 1, 11, or 17 based on a new ground of rejection may not be made final. MPEP 706.07(a).

Applicants respectfully invite the Examiner to contact Applicants' representative at the number provided below if he believes it will help expedite furtherance of this application.

Respectfully submitted,

Dated: January 5, 2009

By: /Christopher King/

Christopher P. King, Reg. No. 60,985
Attorney for Applicant
Fenwick & West LLP
801 California Street
Mountain View, CA 94041
Tel.: (650) 335-7633
Fax: (650) 938-5200

Electronic Acknowledgement Receipt	
EFS ID:	4553063
Application Number:	11623653
International Application Number:	
Confirmation Number:	5243
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Customer Number:	00758
Filer:	Christopher-Charles P. King
Filer Authorized By:	
Attorney Docket Number:	23682-12381
Receipt Date:	05-JAN-2009
Filing Date:	16-JAN-2007
Time Stamp:	17:08:22
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		12381_Response_A.pdf	91082 526c4a8b7136b0076753e228d68b2676cded86bd	yes	10

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Amendment/Req. Reconsideration-After Non-Final Reject	1	1
	Claims	2	6
	Applicant Arguments/Remarks Made in an Amendment	7	10
Warnings:			
Information:			
	Total Files Size (in bytes):	91082	

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/06 (07-06)
Approved for use through 1/31/2007. 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.

PATENT APPLICATION FEE DETERMINATION RECORD

Substitute for Form PTO-875

Application or Docket Number	Filing Date	☐ To be Mailed
11/623,653	01/16/2007	

APPLICATION AS FILED – PART I

	(Column 1)	(Column 2)	SMALL ENTITY ☒		OR	OTHER THAN SMALL ENTITY	
FOR	NUMBER FILED	NUMBER EXTRA	RATE ($)	FEE ($)		RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A			N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A			N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A			N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*	X $ =		OR	X $ =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*	X $ =			X $ =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $250 ($125 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).						
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))							
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL			TOTAL	

APPLICATION AS AMENDED – PART II

		(Column 1)		(Column 2)	(Column 3)	SMALL ENTITY		OR	OTHER THAN SMALL ENTITY	
AMENDMENT	01/05/2009	CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)		RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	* 22	Minus	** 22	= 0	X $26 =	0	OR	X $ =	
	Independent (37 CFR 1.16(h))	* 3	Minus	***3	= 0	X $110 =	0	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	0	OR	TOTAL ADD'L FEE	

		(Column 1)		(Column 2)	(Column 3)					
AMENDMENT		CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	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 number found in the appropriate box in column 1.

Legal Instrument Examiner:
/VANESSA BARBER/

This collection of information is required by 37 CFR 1.16. 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 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.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICANTS:	Justin P. Marston et al.
APPLICATION NO.:	11/623,653
FILING DATE:	January 16, 2007
TITLE:	Determining Relevance of Electronic Content
EXAMINER:	John P. Hocker
GROUP ART UNIT:	2169
ATTY. DKT. NO.:	23682-12381

SUPPLEMENTAL INFORMATION DISCLOSURE STATEMENT
Under 37 CFR §§ 1.56 and 1.97-98

SIR:

Pursuant to the provisions of 37 CFR §§ 1.56 and 1.97-98, enclosed herewith is modified form PTO/SB/08A listing references for consideration by the Examiner.

The filing of this Information Disclosure Statement shall not be construed as a representation regarding the completeness of the list of references, or that inclusion of a reference in this list is an admission that it is prior art or is pertinent to this application, or that a search has been made, or as an admission that the information listed is, or may be considered to be, material to patentability, or that no other material information exists, and shall not be construed as an admission against interest in any manner.

This Information Disclosure Statement is being filed:

- ☐ within three months of the filing date of the application, or date of entry into the national stage of an international application, or before the mailing date of a first office action on the merits, whichever event last occurred;
- ☐ before the mailing of a first official action after the filing of a request for continued examination (RCE) under 37 CFR § 1.114;
- ☒ after three months of the filing date of this national application or the date of entry of the national stage in an international application, or after the mailing date of the first official action on the merits, whichever event last occurred, but before the mailing date of the first to occur of either: (1) a final action under 37 CFR §1.113; or (2) an action that otherwise closes prosecution in the application, and:

☒ attached hereto is the fee set forth under 37 CFR §1.17(p) for submission of this Information Disclosure Statement under 37 CFR.§ 1.97(c); OR

☐ Applicant certifies pursuant to 37 CFR § 1.97(e) that:

- ☐ each item of information contained in this Information Disclosure Statement was first cited in a communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of this Statement; OR
- ☐ no item of information contained in this 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 this certification after making reasonable inquiry, no item of information contained in this Statement was known to any individual designated under 37 CFR § 1.56(c) more than three months prior to the filing of this Statement;

☐ on or before the payment of the issue fee but after the mailing date of the first to occur of either: (1) a final action under 37 CFR § 1.113; (2) a notice of allowance under 37 CFR § 1.311; or (3) an action that otherwise closes prosecution in the application, and:

☐ Applicant certifies pursuant to 37 CFR. § 1.97(e) that:

- ☐ each item of information contained in this Information Disclosure Statement was cited in a communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of this Statement;
- ☐ no item of information contained in this 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 this certification after making reasonable inquiry, no item of information contained in this Statement was known to any individual designated under 37 CFR § 1.56(c) more than three months prior to the filing of this Statement; AND

☐ attached hereto is the fee set forth under 37 CFR §1.17(p) for submission of this Information Disclosure Statement under 37 CFR.§ 1.97(d); OR

☐ after the payment of the issue fee. Applicant requests that the information contained in this Information Disclosure Statement be placed in the file according to 37 CFR § 1.97(i), although the information may not be considered by the USPTO.

☐ This application relies, under 35 U.S.C. § 120, on the earlier filing date of prior application No. _____ , filed on _____ , and the references cited therein are hereby referenced, but are not required to be provided in this application under 37 CFR § 1.98(d).

☐ Each item of information contained in this Information Disclosure Statement was cited in a communication from a foreign patent office in a counterpart application, and the communication was not received by any individual designated in 37 CFR § 1.56(c) more than thirty days prior to the filing of this Information Disclosure Statement. 37 CFR § 1.704(d).

Respectfully submitted,
JUSTIN P. MARSTON ET AL.

Dated: March 12, 2009

By: /Christopher King/
Christopher P. King, Reg. No. 60,985
Attorney for Applicant
Fenwick & West LLP
801 California Street
Mountain View, CA 94041
Tel.: (650) 335-7633
Fax: (650) 938-5200

Substitute for form 1449A/PTO				Complete if Known	
INFORMATION DISCLOSURE STATEMENT BY APPLICANT				**Application No.**	11/623,653
				Filing Date	January 16, 2007
				First Named Inventor	Justin P. Marston
				Art Unit	2169
				Examiner Name	John P. Hocker
Sheet	1	of	1	Attorney Docket Number	23682-12381

U.S. PATENT DOCUMENTS				
		Document No.		
Examiner Initials*	Cite No.[1]	Number – Kind Code[2] (if known)	Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document
	A1	US-2004/0054744 A1	03-18-2004	Karamchedu et al.
	A2	US-6,249,807 B1	06-19-2001	Shaw et al.

FOREIGN PATENT DOCUMENTS					
		Foreign Patent Document			
Examiner Initials*	Cite No.[1]	Country Code[3] – Number[4] Kind Code[5] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	T[6]

OTHER REFERENCES – NON-PATENT LITERATURE DOCUMENTS			
Examiner Initials*	Cite No.[1]	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, page(s), volume-issue number(s), publisher, city and/or country where published	T[6]

Examiner Signature		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.

[1]Applicant's unique citation designation number (optional). [2]See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3]Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4]For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5]Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST. 16 if possible. [6]Applicant is to place a check mark here if English language Translation is attached.

Electronic Patent Application Fee Transmittal	
Application Number:	11623653
Filing Date:	16-Jan-2007
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Filer:	Christopher-Charles P. King/meera yang
Attorney Docket Number:	23682-12381
Filed as Large Entity	

Utility under 35 USC 111(a) Filing Fees

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Miscellaneous:				
Submission- Information Disclosure Stmt	1806	1	180	180
		Total in USD ($)		**180**

Electronic Acknowledgement Receipt	
EFS ID:	4956440
Application Number:	11623653
International Application Number:	
Confirmation Number:	5243
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Customer Number:	00758
Filer:	Christopher-Charles P. King
Filer Authorized By:	
Attorney Docket Number:	23682-12381
Receipt Date:	12-MAR-2009
Filing Date:	16-JAN-2007
Time Stamp:	16:00:25
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	Deposit Account
Payment was successfully received in RAM	$180
RAM confirmation Number	1458
Deposit Account	192555
Authorized User	
The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows: Charge any Additional Fees required under 37 C.F.R. Section 1.16 (National application filing, search, and examination fees) Charge any Additional Fees required under 37 C.F.R. Section 1.17 (Patent application and reexamination processing fees)	

Charge any Additional Fees required under 37 C.F.R. Section 1.19 (Document supply fees)

Charge any Additional Fees required under 37 C.F.R. Section 1.20 (Post Issuance fees)

Charge any Additional Fees required under 37 C.F.R. Section 1.21 (Miscellaneous fees and charges)

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1		12381_US_SIDS.pdf	176976 781bba98594c3e881a1d90054f2574059ef768d8	yes	4

	Multipart Description/PDF files in .zip description		
	Document Description	Start	End
	Information Disclosure Statement Letter	1	3
	Information Disclosure Statement (IDS) Filed (SB/08)	4	4

Warnings:

Information:

2	Fee Worksheet (PTO-06)	fee-info.pdf	30248 214483669a667e09346a2ea52ac48de92679f7aa	no	2

Warnings:

Information:

Total Files Size (in bytes):	207224

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 NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
11/623,653	01/16/2007	Justin Marston	23682-12381	5243

758 7590 05/12/2009

FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041

EXAMINER	
HOCKER, JOHN P	
ART UNIT	PAPER NUMBER
2158	

MAIL DATE	DELIVERY MODE
05/12/2009	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)	
	11/623,653	MARSTON ET AL.	
	Examiner	Art Unit	
	JOHN P. HOCKER	2158	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE *3* MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *05 January 2009*.

2a)☒ This action is **FINAL**. 2b)☐ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☐ Claim(s) _____ is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-22* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

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)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☒ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date *03/12/2009*.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

1. Claims 1 – 22 have been examined.

2. Claims 1 – 22 have been rejected (**Final Rejection**).

Examiner Interview

After further search and consideration, and after speaking with Mohammad Ali, the invention as claimed does not overcome 102/103 tests for novelty/non-obviousness. Examiner thanks applicant for interview to hone in precise arguments about allegedly novel technical features. Examiner also asked if there were any new arguments for allowability. At the present time, no new arguments were raised. The interview occurred on May 7th, 2009 with Chris King.

Information Disclosure Statement

The information disclosure statement (IDS) submitted 03/12/2009 has been considered by examiner. The submission is in compliance with the provisions of 37 CFR 1.97. Accordingly, the information disclosure statement is being considered by the examiner.

Related Cases

This case is related to copending application no. 11/426118.

This case is related to PCT Written Opinion and Search Reports No. PCT/US06/24859 and No. PCT/US06/60587.

Identical claims 1 – 22 were submitted in a PCT application. Examiner finds no reason to disagree with the written opinion in the PCT International Search Report. (PCT/US07/60587) (February 14, 2008, 8 pages).

The PCT International Search Report and Written Opinion found the claims lacked novelty as anticipated by the Malik Reference

Response to Arguments

By Interview and Written Correspondence, Applicant has raised a primary argument that the Malik Reference does not teach that "relevance" is based on a user's actions (Applicant's Remarks, filed: 01/05/2009, pgs. 1 - 4). This is language from Claim 1, Limitation iv.

Examiner reminds applicant that a reference is good for all it teaches, and as shown in the Non-Final rejection the Malik reference cited paragraphs (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079 - 0080]) for Claim 1, Limitation iv.

Clearly, [0044] which mentions the profiles of Malik may be based on "email rules by which (email messages) are sorted, forwarded, deleted, etc." would in fact meet this language. The language is stating relevance is based on user actions. The relevance of Malik is based on the profile, which in 0044 is based on user actions (e.g.sorting, forwarding, deleting). Thus the argument that Malik does not teach relevance based on user actions is unpersuasive. If applicant has any questions or concerns please do not hesitate to phone examiner at contact info below.

Claim Rejections - 35 USC § 102 and § 35 USC 103

Examiner's Note: The requirements for patentability are that a claimed invention must at a minimum meet the requirements of utility (35 USC 101), novelty (35 USC 102) and non-obviousness (35 USC 103). The claims are rejected under 102 as they seemingly are anticipated by the prior art. Alternatively, the claims are rejected under 103 as they are obvious in view of the prior art of record and the common sense of one of ordinary skill.

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

> (b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> **(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title,** if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. **Patentability shall not be negatived by the manner in which the invention was made.**

Claims 1 - 22 rejected under 35 U.S.C. 102(b) as anticipated by or, in the alternative, under 35 U.S.C. 103(a) as obvious over Malik (US Publication No. 2004/0078447)

The specification describes the exemplary embodiment of the claimed system to be what is known as an email system (Specification, pg. 1). The claim's allegedly novel and non-obvious feature is the determination of user relevance for the email (i.e. electronic content). (Specification pgs 1 – 2) This argument is unpersuasive though as shown in Malik. In the same fashion, the email system of Malik utilizes user profiles to determine user relevance of electronic content.

As per claims 1 and 11, Malik discloses a system and program (see fig 1,2A and 2B), comprising:

(i) a monitoring module configured to track actions by a plurality of users associated with an application for managing electronic content (see para [0035] and [0045]),

(ii) a relevance analysis module coupled to the monitoring module and configured to determine a relevance to a user of a received piece of electronic content (user(s) profile, see abstract and para [0013], [0014], [0040]-[0046], [0072-0074] and [0079 – 0080]),

(iii) the relevance being determined dependent on: a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079] - 0080),

and (iv) a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079 - 0080]).

As per claims 2, 12 and 18, Malik further discloses the system of claim 1, wherein the relevance is further determined dependent on a measure of local importance describing an importance of the sender to the user (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079])

As per claims 3, 13 and 19, Malik further discloses the system of claim 1, wherein the relevance is further determined dependent on a measure of global importance describing an

importance of the sender within an organization with which the plurality of users are associated (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claims 4, 14 and 20, Malik further discloses the system of claim 1, wherein the relevance analysis module is further configured to identify at least one other of the plurality of users that is a recipient of the piece of electronic content and is cognate to the user and wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content dependent upon one or more actions of the at least one cognate recipient with respect to the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claims 5, 15 and 21, Malik further discloses the system of claim 1, wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content responsive to contents of the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claims 6, 16 and 22, Malik further discloses the system of claim 1, wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content responsive to one or more measures from the set consisting of: a ocal importance describing an importance of the sender to the user (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]), a i g obal mportance describing an importance of the sender within an organization with which the plurality of users are associated (see para [0010]), a

content relevance determined responsive to contents of the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]) and one or more actions with respect to the piece of electronic content of one or more of the plurality of users determined to be cognate to the user (see para [0010]).

As per claim 7, Malik further discloses the system of claim 1, further comprising: a content presentation module configured to present the piece of electronic content to the user in accordance with the determined relevance and to present to the user a value describing the determined relevance (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claim 8, Malik further discloses the system of claim 1, further comprising: a search module configured to search the electronic content in accordance with the determined relevance (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

As per claim 9, Malik further discloses the system of claim 1, wherein the piece of electronic content comprises an electronic mail object in a relational mail system (see para [0072], [0074], [0076], [0079] and [0080]).

As per claim 10, Malik further discloses the system of claim 1, wherein the electronic content comprises at least one of an email, a document, an instant message, and an SMS message (at least one of an email, see para [0072], [0074], [0076], [0079] and [0080]).

As per claim 17, Malik discloses a method, comprising: tracking actions by a plurality of users associated with an application for managing electronic content (see para [0035] and [0045]), determining a relevance to a user of a piece of electronic content sent from a sender to the user (user(s) profile, see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]), the relevance being determined dependent on: a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]), and a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079]).

Examiner's note: Although seemingly every claim limitation is taught by the prior art of record, it is arguable that each and every word is not verbatim 102 anticipated by the prior art. Examiner asserts that nonetheless, the claimed invention is properly rejected under 103 if the differences are insubstantial and the concept as a whole would be obvious to one of ordinary skill in the art with common sense which includes well known techniques and known substitutes. Examiner asserts the claim mappings above show the state of the prior art.

Examiner's note: In evaluating applicant's claims to patentability, Examiner was not persuaded by applicant's description of the prior art. Specifically, applicant argues that the prior art teaches ""unintelligent mechanisms" for storage and sorting of email (Specification [0004]. Examiner could not disagree more, the "intelligent" features applicant claims are the **determination of relevance for a user.** In the same respect, the email system of Malik

determines relevance of the content based on a user profile (i.e. based on a user). Thus applicant's claimed Background and Summary is a mischaracterization of the state of the prior art and the argument for patentability relies on this mischaracterization. The fact is that the prior art Malik teaches this intelligent mechanism for email storage and sorting. Specifically, the intelligent mechanism is for presenting the email based on user relevance (i.e. only pertinent to certain user profiles). Applicant has not named a credible improvement over the prior art.

Conclusion

THIS ACTION IS MADE FINAL. Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to John P. Hocker whose telephone number is 5712705004. The examiner can normally be reached on Mon - Thu 11:00 - 4:00 EST.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Mohammad Ali can be reached on 5712724105. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/John P. Hocker/ Art Unit 2158

/Mohammad Ali/
Supervisory Patent Examiner, Art Unit 2158

Notice of References Cited	Application/Control No. 11/623,653	Applicant(s)/Patent Under Reexamination MARSTON ET AL.	
	Examiner JOHN P. HOCKER	Art Unit 2158	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	Classification
*	A	US-2004/0078447 A1	04-2004	Malik et al.	709/206
	B	US-			
	C	US-			
	D	US-			
	E	US-			
	F	US-			
	G	US-			
	H	US-			
	I	US-			
	J	US-			
	K	US-			
	L	US-			
	M	US-			

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	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
	11623653	MARSTON ET AL.
	Examiner	Art Unit
	JOHN P HOCKER	2169

SEARCHED			
Class	Subclass	Date	Examiner
707	1-3	11/09/2008	JPH

SEARCH NOTES		
Search Notes	Date	Examiner
Inventive concept: " User Profiles for electronic content viewing (e.g. email)"	11/09/2008	JPH
1: IDS Reviewed for Prior Art.	11/09/2008	JPH
2: Examiner performed inventor name search.	11/09/2008	JPH
3: Examiner searched 'Computer Technology' NPL (e.g. IEEE (Institute of Electrical and Electronics Engineers), Association for Computing Machinery, Wikipedia and/or Google Scholar)	11/09/2008	JPH
4: Examiner searched International Technical & Patent Databases(US-PGPUB, USPAT, USOCR, FPRS, EPO, JPO, DERWENT and IBM_TDB).	11/09/2008	JPH
During Patent Searches, Examiner searched class/subclasses (as listed above) combined with text of the inventive concept. Please see East Search History for text searched.	11/09/2008	JPH

INTERFERENCE SEARCH			
Class	Subclass	Date	Examiner

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	11623653	MARSTON ET AL.
	Examiner	Art Unit
	JOHN P HOCKER	2169

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	11/09/2008	05/07/2009							
	1	✓	✓							
	2	✓	✓							
	3	✓	✓							
	4	✓	✓							
	5	✓	✓							
	6	✓	✓							
	7	✓	✓							
	8	✓	✓							
	9	✓	✓							
	10	✓	✓							
	11	✓	✓							
	12	✓	✓							
	13	✓	✓							
	14	✓	✓							
	15	✓	✓							
	16	✓	✓							
	17	✓	✓							
	18	✓	✓							
	19	✓	✓							
	20	✓	✓							
	21	✓	✓							
	22	✓	✓							

Substitute for form 1449A/PTO				Complete if Known	
INFORMATION DISCLOSURE STATEMENT BY APPLICANT				**Application No.**	11/623,653
				Filing Date	January 16, 2007
				First Named Inventor	Justin P. Marston
				Art Unit	2169
				Examiner Name	John P. Hocker
Sheet	1	of	1	Attorney Docket Number	23682-12381

U.S. PATENT DOCUMENTS				
		Document No.		
Examiner Initials*	Cite No.[1]	Number – Kind Code[2] (if known)	Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document
	A1	US-2004/0054744 A1	03-18-2004	Karamchedu et al.
	A2	US-6,249,807 B1	06-19-2001	Shaw et al.

FOREIGN PATENT DOCUMENTS					
		Foreign Patent Document			
Examiner Initials*	Cite No.[1]	Country Code[3] – Number[4] Kind Code[5] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	T[6]

OTHER REFERENCES – NON-PATENT LITERATURE DOCUMENTS			
Examiner Initials*	Cite No.[1]	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, page(s), volume-issue number(s), publisher, city and/or country where published	T[6]

Examiner Signature	/John Hocker/	Date Considered	04/15/2009

*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.

[1]Applicant's unique citation designation number (optional). [2]See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3]Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4]For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5]Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST. 16 if possible. [6]Applicant is to place a check mark here if English language Translation is attached.

23682/12381/DOCS/2041672.1

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /J.H./

IN THE

UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICANT:	Justin Marston et al
SERIAL NO.:	11/623,653
FILING DATE:	January 16, 2007
TITLE:	Determining Relevance of Electronic Content
EXAMINER:	John P. Hocker
GROUP ART UNIT:	2169
ATTY. DKT. NO.:	23682-12381

CERTIFICATE OF ELECTRONIC (EFS-WEB) TRANSMISSION

I hereby certify that this correspondence is being transmitted via the Office electronic filing system in accordance with 37 C.F.R. § 1.8(a)(i)(C) from the **Pacific Time Zone** of the United States on the local date shown below.

Dated: August 11, 2009 By: /Christopher King/
Christopher P. King, Reg. No. 60,985

COMMISSIONER FOR PATENTS
P.O. BOX 1450
ALEXANDRIA, VA 22313-1450

STATEMENT OF SUBSTANCE OF INTERVIEW

Sir:

Applicants thank Examiners Hocker and Leroux for their time in conducting a telephone interview on August 10, 2009. During the telephone interview, Applicants' representative Chris King and the Examiners discussed the rejection of claims 1 and 7 in light of the Malik reference (U.S. Patent Publication No. 2004/0078447). Examiner Hocker agreed that Malik does not disclose or suggest that the measure of relevance of the piece of electronic content is determined dependent on one or more actions on the piece of electronic content of at least one other of the plurality of users that is a recipient of the piece of

electronic content, and that if claim 1 were amended to include this limitation, then claim 1 would be distinguished from Malik.

Respectfully Submitted,
Justin Marston et al

Date: August 11, 2009

By: /Christopher King/

Christopher P. King, Reg. No.: 60,985
Fenwick & West LLP
Silicon Valley Center
801 California Street
Mountain View, CA 94041
Tel.: (650) 335-7633
Fax.: (650) 938-5200

Electronic Acknowledgement Receipt	
EFS ID:	5870138
Application Number:	11623653
International Application Number:	
Confirmation Number:	5243
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Customer Number:	00758
Filer:	Christopher-Charles P. King
Filer Authorized By:	
Attorney Docket Number:	23682-12381
Receipt Date:	11-AUG-2009
Filing Date:	16-JAN-2007
Time Stamp:	19:39:24
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	12381_interview_summary.pdf	66294 84d1ad5aeaada5d7cb9306196114072d04a0e736	no	2

Warnings:
Information:

Total Files Size (in bytes):	66294

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.

PRE-APPEAL BRIEF REQUEST FOR REVIEW

Docket Number
23682-12381

Pursuant to 240 OG 45 and the *Legal Framework For EFS-Web*, I hereby certify that this follow-on correspondence is being officially submitted through the USPTO EFS-Web system from the Pacific Time Zone of the United States on the local date shown below.

on August 12, 2009

Signature /Christopher King/

Typed or printed name Christopher-Charles P. King

Application Number	Filed
11/623,653	January 16, 2007

First Named Inventor
Justin Marston

Art Unit	Examiner
2158	John P. Hocker

This request is being filed with a notice of appeal.

I am the

☐ applicant/inventor.

☐ assignee of record of the entire interest.
See 37 CFR 3.71. Statement under 37 CFR 3.73(b) is enclosed.

☒ attorney or agent of record.
Registration number 60,985.

☐ attorney or agent acting under 37 CFR 1.34.
Registration number if acting under 37 CFR 1.34 _____

/Christopher King/
Signature

Christopher-Charles P. King
Typed or printed name

(650) 335-7633
Telephone number

August 12, 2009
Date

NOTE: Signatures of all the inventors or assignees of record of the entire interest or their representative(s) are required. Submit multiple forms if more than one signature is required, see below*.

☒ *Total of 1 of 1 form is submitted.

ATTACHMENT TO THE PRE-APPEAL BRIEF REQUEST FOR REVIEW

Pre-appeal review is requested because the rejections in the Final Office Action of May 12, 2009, are improper and without factual or legal basis. Applicants respectfully request that the Panel indicate that at least claims 4, 6, 7, 14, 16, 20, 22 recite allowable subject matter.

I. Status of the Claims

Claims 1-22 are pending in this application and stand rejected under § 102—or alternatively, § 103—based on Malik, U.S. Patent Publication No. 2004/0078447.

II. Rejections of claims 1-22 as unpatentable over Malik

To render a claim unpatentable under 35 U.S.C. § 103, the prior art must disclose or suggest each and every limitation in the claim. *In re Fine*, 837 F.2d 1071 (Fed. Cir. 1988) (reversing § 103 rejection because examiner ignored material claimed limitation that was absent from reference); *In re Royka*, 490 F.2d 981 (CCPA 1974).

Independent claim 1, from which claim 4 depends, recites a method comprising:

> a monitoring module configured to track actions by a plurality of users associated with an application for managing electronic content;
> **a relevance analysis module** coupled to the monitoring module and configured to determine a relevance to a user of a piece of electronic content sent from a sender to the user, the relevance being determined dependent on:
> a measure of a relevance of the piece of electronic content as determined dependent on **one or more actions of the user** on the piece of electronic content, and
> a measure of a relevance of the piece of electronic content as determined dependent **on one or more actions of at least one other of the plurality of users that is a recipient** of the piece of electronic content.

Thus, claim 1 recites that the determined relevance to the user of the piece of electronic content is determined dependent on one or more actions of the user on the piece of electronic content,

and on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content. As noted to the Examiner during a telephone interview on August 10, 2009, examples of actions of the user and other recipients on the piece of electronic content include forwarding, deleting, or printing an electronic email object. (*See, e.g.,* Figure 6, item 86; paragraph 0069). The determined relevance can be displayed graphically to the user as a value associated with the a piece of electronic content such as an email, e.g., as in item 118 of Figure 8, which shows a relevance slider control graphically expressing the relevance value of an email submessage.

Claim 4 depends from claim 1 and further recites:

> wherein the relevance analysis module is further configured to identify at least one other of the plurality of users that is a recipient of the piece of electronic content and is cognate to the user and wherein the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content dependent upon one or more actions of the at least one cognate recipient with respect to the piece of electronic content.

Thus, claim 4 further determines the relevance to the user of the piece of electronic content dependent upon one or more actions of at least one cognate recipient with respect to the piece of electronic content, where the cognate recipient is a recipient of the piece of electronic content as in the last element of claim 1. Thus, the relevance is determined in part *dependent on one or more actions* ***with respect to the piece of electronic content*** *of at least one other of the plurality of users that is a recipient*, i.e., the cognate recipient.

Malik does not disclose or suggest the features of claim 4. Malik generally discloses a mailbox system that differentiates between user and group mailboxes. A user profile can be associated with a user email folder and an instant messaging (IM) account, and this association facilitates defining properties of the user email folder and the user IM account using the user profile. (Malik, Abstract). The user profiles are a collection of settings, such as a listing of a

user's email mailboxes and folders, IM accounts, usernames and passwords, address books, and the like. (Malik, paragraphs 0041-0044).

During the interview, the Examiner agreed that Malik does not disclose or suggest that the measure of relevance of the piece of electronic content is determined dependent on one or more actions ***on the piece of electronic content*** of at least one other of the plurality of users that is a recipient of the piece of electronic content. The Examiner further agreed that if claim 1 were amended to include this feature, then claim 1 would be distinguished from Malik. Applicants note that, as described above, claim 4 as filed **does** in fact include the feature in question, since it recites that the relevance analysis module is further configured to identify at least one other of the plurality of users that is a recipient of the piece of electronic content and is cognate to the user—i.e., a cognate recipient is one of the plurality of users that is a recipient of the piece of electronic content, as recited in claim 1—and further recites that the relevance to the user of the piece of electronic content is determined dependent upon one or more actions of this cognate recipient **with respect to the piece of electronic content**. Thus, claim 4 is patentable over Malik for at least the same reasons agreed to during the interview.

Indeed, the portions of Malik cited with respect to claim 4 do not disclose or suggest this claimed feature, i.e., that "the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content dependent upon one or more actions of the at least one cognate recipient with respect to the piece of electronic content." As with the rejections of the other claims, the Examiner cites Malik's paragraphs 0013-0014, 0040-0046, 0072-0074, and 0079-0080. Paragraphs 0013-0014 mention a user profile that has a collection of properties related to both email and instant messaging. Paragraphs 0040-0046 further describe the user profiles, listing examples of what a profile includes, such as a user's email mailboxes

and folders, IM accounts, usernames and passwords, address books, and the like. Paragraphs 0072-0074 and 0079-0080 mention various user interfaces for storing profile information related to instant messaging and to email, respectively. None of these paragraphs, or any other portion of Malik not cited, disclose or suggest that the determination of relevance of a piece of electronic content to a particular user includes actions of **other** recipients (such as cognate recipients) **with respect to the piece of electronic content**.

Claim 6 recites that the relevance analysis module is further configured to determine the relevance to the user of the piece of electronic content responsive to one or more measures from the set consisting of "… one or more actions **with respect to the piece of electronic content** of one or more of the plurality of users determined to be cognate to the user." Claims 14 and 16, and 20 and 22, recite limitations similar to those of claims 4 and 6. Thus, claims 6, 14, 16, 20, and 22 are patentably distinguishable from Malik for at least the same reasons discussed above with respect to claim 4.

Claim 7 recites a content presentation module configured to present the piece of electronic content to the user in accordance with the determined relevance and to present to the user **a value describing the determined relevance**. For example, as noted during the interview, relevance slider 118 of Figure 8 graphically depicts a relevance value of a particular email submessage. As with the other claims, the Examiner cites paragraphs 0013-0014, 0040-0046, 0072-0074, and 0079, but none of these paragraphs disclose or suggest the presentation of a value describing the determined relevance.

III. Summary

Based on the foregoing, Applicant respectfully submits that the pending rejections suffer

from clear deficiencies. Accordingly, Applicant requests that the rejections of independent claims 4, 6, 7, 14, 16, 20, and 22 be withdrawn.

Respectfully Submitted,

Date: August 12, 2009

By: /Christopher King/

Christopher P. King, Reg. No. 60,985
Attorney for Applicants
FENWICK & WEST LLP
801 California Street
Mountain View, CA 94041
Phone: (650) 335-7633
Fax: (650) 938-5200

Electronic Patent Application Fee Transmittal

Application Number:	11623653
Filing Date:	16-Jan-2007
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Filer:	Christopher-Charles P. King/Larisa Burshteyn
Attorney Docket Number:	23682-12381

Filed as Small Entity

Utility under 35 USC 111(a) Filing Fees

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Notice of appeal	2401	1	270	270
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Miscellaneous:				
			Total in USD ($)	**270**

Electronic Acknowledgement Receipt	
EFS ID:	5875235
Application Number:	11623653
International Application Number:	
Confirmation Number:	5243
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Customer Number:	00758
Filer:	Christopher-Charles P. King
Filer Authorized By:	
Attorney Docket Number:	23682-12381
Receipt Date:	12-AUG-2009
Filing Date:	16-JAN-2007
Time Stamp:	15:56:15
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	Deposit Account
Payment was successfully received in RAM	$270
RAM confirmation Number	1785
Deposit Account	192555
Authorized User	
The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows: Charge any Additional Fees required under 37 C.F.R. Section 1.16 (National application filing, search, and examination fees) Charge any Additional Fees required under 37 C.F.R. Section 1.17 (Patent application and reexamination processing fees)	

Charge any Additional Fees required under 37 C.F.R. Section 1.19 (Document supply fees)

Charge any Additional Fees required under 37 C.F.R. Section 1.20 (Post Issuance fees)

Charge any Additional Fees required under 37 C.F.R. Section 1.21 (Miscellaneous fees and charges)

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Notice of Appeal Filed	12381_Notice_of_Appeal.pdf	70009 ac3a51eb1e32b7cb6f37b6ed55e27e4ed5081808	no	1
Warnings:					
Information:					
2	Pre-Brief Conference request	12381_Pre-Appeal_Brief.pdf	155033 112a56fcde6758ec2a15d3ccedc2bb5c4585d3d7	no	6
Warnings:					
Information:					
3	Fee Worksheet (PTO-875)	fee-info.pdf	30300 8b3bc9be7d4d8253ad75879f85cb79676fb235c7	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	255342		

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

APPLICANT:	Justin Marston et al
SERIAL NO.:	11/623,653
FILING DATE:	January 16, 2007
TITLE:	Determining Relevance of Electronic Content
EXAMINER:	John P. Hocker
GROUP ART UNIT:	2169
ATTY. DKT. NO.:	23682-12381

CERTIFICATE OF ELECTRONIC (EFS-WEB) TRANSMISSION

I hereby certify that this correspondence is being transmitted via the Office electronic filing system in accordance with 37 C.F.R. § 1.8(a)(i)(C) from the Pacific Time Zone of the United States on the local date shown below.

Dated: August 12, 2009

By: /Christopher King/
Christopher P. King, Reg. No.: 60,985

NOTICE OF APPEAL

SIR:

Applicant hereby appeals to the Board of Patent Appeals and Interferences from the decision of the Examiner in the Final Office Action of May 12, 2009, finally rejecting claims 1-22.

The Commissioner is hereby authorized to charge to Deposit Account No. 19-2555 the required fee under 37 CFR § 41.20(b)(1 and any additional charges or fees as applicable for the payment of the filing of the Notice of Appeal, and the Pre-Appeal Brief Request for Review.

Dated: August 12, 2009

By: /Christopher King/
Christopher P. King, Reg. No.: 60,985
Fenwick & West LLP
Silicon Valley Center
801 California Street
Mountain View, CA 94041
Tel.: (650) 335-7633
Fax: (650) 938-5200

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.
11/623,653	01/16/2007	Justin Marston	23682-12381	5243

758 7590 08/25/2009

FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041

EXAMINER	
HOCKER, JOHN P	
ART UNIT	**PAPER NUMBER**
2158	
MAIL DATE	**DELIVERY MODE**
08/25/2009	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Interview Summary	Application No.	Applicant(s)	
	11/623,653	MARSTON ET AL.	
	Examiner	Art Unit	
	JOHN P. HOCKER	2158	

All participants (applicant, applicant's representative, PTO personnel):

(1) *JOHN P. HOCKER*. (3)*Etienne Leroux*.

(2) *Chris King*. (4)_____.

Date of Interview: *11 August 2009*.

Type: a)☒ Telephonic b)☐ Video Conference
c)☐ Personal [copy given to: 1)☐ applicant 2)☐ applicant's representative]

Exhibit shown or demonstration conducted: d)☐ Yes e)☐ No.
If Yes, brief description: _____.

Claim(s) discussed: *1-3 and 7*.

Identification of prior art discussed: *Malik*.

Agreement with respect to the claims f)☐ was reached. g)☒ was not reached. h)☐ N/A.

Substance of Interview including description of the general nature of what was agreed to if an agreement was reached, or any other comments: *Applicant and Examiner primarily discussed the differences between the disclosed invention (See Figs. 6 & 8, including "Relevance" bar) and the prior art Malik (an email storage system utilizing user profiles, and group profiles). The arguably novel feature is that the claimed invention bases the 'Relevance' on the actions of another user "ON The Piece of Content". In other words, not just any general actions of another (i.e. group administrator adjusting values for group (properties/preferences). Examiner's interpretation was that the claim 1, limitation read on these general actions. Applicant seemed to indicate an amendment further clarifying the Actions on the content in limitation (iv) was forth coming. Examiner stated further search would be needed but that the interview was helpful in clarifying the distinguishment over Malik. Examiner stated Malik alone does not seem to teach the Fig. 6 and Fig. 8 inventions. Applicant seemed to agree to amend to further specify the 'actions' in the last limitation were based on 'others' (i.e. another person forwards it, another person deletes it). No agreement was reached as to allowability as an Updated Search and Further Consideration in view of the prior art is necessary.*.

(A fuller description, if necessary, and a copy of the amendments which the examiner agreed would render the claims allowable, if available, must be attached. Also, where no copy of the amendments that would render the claims allowable is available, a summary thereof must be attached.)

THE FORMAL WRITTEN REPLY TO THE LAST OFFICE ACTION MUST INCLUDE THE SUBSTANCE OF THE INTERVIEW. (See MPEP Section 713.04). If a reply to the last Office action has already been filed, APPLICANT IS GIVEN A NON-EXTENDABLE PERIOD OF THE LONGER OF ONE MONTH OR THIRTY DAYS FROM THIS INTERVIEW DATE, OR THE MAILING DATE OF THIS INTERVIEW SUMMARY FORM, WHICHEVER IS LATER, TO FILE A STATEMENT OF THE SUBSTANCE OF THE INTERVIEW. See Summary of Record of Interview requirements on reverse side or on attached sheet.

/Etienne P LeRoux/ Primary Examiner, Art Unit 2161	

Summary of Record of Interview Requirements

Manual of Patent Examining Procedure (MPEP), Section 713.04, Substance of Interview Must be Made of Record
A complete written statement as to the substance of any face-to-face, video conference, or telephone interview with regard to an application must be made of record in the application whether or not an agreement with the examiner was reached at the interview.

Title 37 Code of Federal Regulations (CFR) § 1.133 Interviews
Paragraph (b)

In every instance where reconsideration is requested in view of an interview with an examiner, a complete written statement of the reasons presented at the interview as warranting favorable action must be filed by the applicant. An interview does not remove the necessity for reply to Office action as specified in §§ 1.111, 1.135. (35 U.S.C. 132)

37 CFR §1.2 Business to be transacted in writing.
All business with the Patent or Trademark Office should be transacted in writing. The personal attendance of applicants or their attorneys or agents at the Patent and Trademark Office is unnecessary. The action of the Patent and Trademark Office will be based exclusively on the written record in the Office. No attention will be paid to any alleged oral promise, stipulation, or understanding in relation to which there is disagreement or doubt.

The action of the Patent and Trademark Office cannot be based exclusively on the written record in the Office if that record is itself incomplete through the failure to record the substance of interviews.

It is the responsibility of the applicant or the attorney or agent to make the substance of an interview of record in the application file, unless the examiner indicates he or she will do so. It is the examiner's responsibility to see that such a record is made and to correct material inaccuracies which bear directly on the question of patentability.

Examiners must complete an Interview Summary Form for each interview held where a matter of substance has been discussed during the interview by checking the appropriate boxes and filling in the blanks. Discussions regarding only procedural matters, directed solely to restriction requirements for which interview recordation is otherwise provided for in Section 812.01 of the Manual of Patent Examining Procedure, or pointing out typographical errors or unreadable script in Office actions or the like, are excluded from the interview recordation procedures below. Where the substance of an interview is completely recorded in an Examiners Amendment, no separate Interview Summary Record is required.

The Interview Summary Form shall be given an appropriate Paper No., placed in the right hand portion of the file, and listed on the "Contents" section of the file wrapper. In a personal interview, a duplicate of the Form is given to the applicant (or attorney or agent) at the conclusion of the interview. In the case of a telephone or video-conference interview, the copy is mailed to the applicant's correspondence address either with or prior to the next official communication. If additional correspondence from the examiner is not likely before an allowance or if other circumstances dictate, the Form should be mailed promptly after the interview rather than with the next official communication.

The Form provides for recordation of the following information:
- Application Number (Series Code and Serial Number)
- Name of applicant
- Name of examiner
- Date of interview
- Type of interview (telephonic, video-conference, or personal)
- Name of participant(s) (applicant, attorney or agent, examiner, other PTO personnel, etc.)
- An indication whether or not an exhibit was shown or a demonstration conducted
- An identification of the specific prior art discussed
- An indication whether an agreement was reached and if so, a description of the general nature of the agreement (may be by attachment of a copy of amendments or claims agreed as being allowable). Note: Agreement as to allowability is tentative and does not restrict further action by the examiner to the contrary.
- The signature of the examiner who conducted the interview (if Form is not an attachment to a signed Office action)

It is desirable that the examiner orally remind the applicant of his or her obligation to record the substance of the interview of each case. It should be noted, however, that the Interview Summary Form will not normally be considered a complete and proper recordation of the interview unless it includes, or is supplemented by the applicant or the examiner to include, all of the applicable items required below concerning the substance of the interview.

A complete and proper recordation of the substance of any interview should include at least the following applicable items:
1) A brief description of the nature of any exhibit shown or any demonstration conducted,
2) an identification of the claims discussed,
3) an identification of the specific prior art discussed,
4) an identification of the principal proposed amendments of a substantive nature discussed, unless these are already described on the Interview Summary Form completed by the Examiner,
5) a brief identification of the general thrust of the principal arguments presented to the examiner,
   (The identification of arguments need not be lengthy or elaborate. A verbatim or highly detailed description of the arguments is not required. The identification of the arguments is sufficient if the general nature or thrust of the principal arguments made to the examiner can be understood in the context of the application file. Of course, the applicant may desire to emphasize and fully describe those arguments which he or she feels were or might be persuasive to the examiner.)
6) a general indication of any other pertinent matters discussed, and
7) if appropriate, the general results or outcome of the interview unless already described in the Interview Summary Form completed by the examiner.

Examiners are expected to carefully review the applicant's record of the substance of an interview. If the record is not complete and accurate, the examiner will give the applicant an extendable one month time period to correct the record.

Examiner to Check for Accuracy

If the claims are allowable for other reasons of record, the examiner should send a letter setting forth the examiner's version of the statement attributed to him or her. If the record is complete and accurate, the examiner should place the indication, "Interview Record OK" on the paper recording the substance of the interview along with the date and the examiner's initials.

.

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.
11/623,653	01/16/2007	Justin Marston	23682-12381	5243

758 7590 12/15/2009

FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041

EXAMINER	
HOCKER, JOHN P	
ART UNIT	**PAPER NUMBER**
2158	

MAIL DATE	DELIVERY MODE
12/15/2009	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of Panel Decision from Pre-Appeal Brief Review	Application/Control No.	Applicant(s)/Patent under Reexamination
	11/623,653	MARSTON ET AL.
	MOHAMMAD ALI	Art Unit: 2158

This is in response to the Pre-Appeal Brief Request for Review filed 12 August 2009.

1. ☐ **Improper Request** – The Request is improper and a conference will not be held for the following reason(s):

☐ The Notice of Appeal has not been filed concurrent with the Pre-Appeal Brief Request.
☐ The request does not include reasons why a review is appropriate.
☐ A proposed amendment is included with the Pre-Appeal Brief request.
☐ Other: .

The time period for filing a response continues to run from the receipt date of the Notice of Appeal or from the mail date of the last Office communication, if no Notice of Appeal has been received.

2. ☒ **Proceed to Board of Patent Appeals and Interferences** – A Pre-Appeal Brief conference has been held. The application remains under appeal because there is at least one actual issue for appeal. Applicant is required to submit an appeal brief in accordance with 37 CFR 41.37. The time period for filing an appeal brief will be reset to be one month from mailing this decision, or the balance of the two-month time period running from the receipt of the notice of appeal, whichever is greater. Further, the time period for filing of the appeal brief is extendible under 37 CFR 1.136 based upon the mail date of this decision or the receipt date of the notice of appeal, as applicable.

☐ The panel has determined the status of the claim(s) is as follows:
Claim(s) allowed: _____.
Claim(s) objected to: _____.
Claim(s) rejected: _____.
Claim(s) withdrawn from consideration: _____.

3. ☐ **Allowable application** – A conference has been held. The rejection is withdrawn and a Notice of Allowance will be mailed. Prosecution on the merits remains closed. No further action is required by applicant at this time.

4. ☐ **Reopen Prosecution** – A conference has been held. The rejection is withdrawn and a new Office action will be mailed. No further action is required by applicant at this time.

All participants:

(1) *MOHAMMAD ALI*.

(2) *Hosain Alam*.

(3) *John Hocker*.

(4)_____.

/Mohammad Ali/
Supervisory Patent Examiner, Art Unit 2158

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.
11/623,653	01/16/2007	Justin Marston	23682-12381	5243

758 7590 01/11/2010

FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041

EXAMINER	
HOCKER, JOHN P	
ART UNIT	PAPER NUMBER
2158	

MAIL DATE	DELIVERY MODE
01/11/2010	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Interview Summary	Application No.	Applicant(s)	
	11/623,653	MARSTON ET AL.	
	Examiner	Art Unit	
	JOHN P. HOCKER	2158	

All participants (applicant, applicant's representative, PTO personnel):

(1) *JOHN P. HOCKER*. (3)_____.

(2) *Chris King*. (4)_____.

Date of Interview: *04 January 2009*.

Type: a)☒ Telephonic b)☐ Video Conference
c)☐ Personal [copy given to: 1)☐ applicant 2)☐ applicant's representative]

Exhibit shown or demonstration conducted: d)☐ Yes e)☐ No.
If Yes, brief description: _____.

Claim(s) discussed: *1-22*.

Identification of prior art discussed: *Malik*.

Agreement with respect to the claims f)☐ was reached. g)☒ was not reached. h)☐ N/A.

Substance of Interview including description of the general nature of what was agreed to if an agreement was reached, or any other comments: *See Continuation Sheet*.

(A fuller description, if necessary, and a copy of the amendments which the examiner agreed would render the claims allowable, if available, must be attached. Also, where no copy of the amendments that would render the claims allowable is available, a summary thereof must be attached.)

THE FORMAL WRITTEN REPLY TO THE LAST OFFICE ACTION MUST INCLUDE THE SUBSTANCE OF THE INTERVIEW. (See MPEP Section 713.04). If a reply to the last Office action has already been filed, APPLICANT IS GIVEN A NON-EXTENDABLE PERIOD OF THE LONGER OF ONE MONTH OR THIRTY DAYS FROM THIS INTERVIEW DATE, OR THE MAILING DATE OF THIS INTERVIEW SUMMARY FORM, WHICHEVER IS LATER, TO FILE A STATEMENT OF THE SUBSTANCE OF THE INTERVIEW. See Summary of Record of Interview requirements on reverse side or on attached sheet.

/JPH/ (Art Unit: 2158)	/Mohammad Ali/ Supervisory Patent Examiner, Art Unit 2158

Summary of Record of Interview Requirements

Manual of Patent Examining Procedure (MPEP), Section 713.04, Substance of Interview Must be Made of Record
A complete written statement as to the substance of any face-to-face, video conference, or telephone interview with regard to an application must be made of record in the application whether or not an agreement with the examiner was reached at the interview.

Title 37 Code of Federal Regulations (CFR) § 1.133 Interviews
Paragraph (b)

In every instance where reconsideration is requested in view of an interview with an examiner, a complete written statement of the reasons presented at the interview as warranting favorable action must be filed by the applicant. An interview does not remove the necessity for reply to Office action as specified in §§ 1.111, 1.135. (35 U.S.C. 132)

37 CFR §1.2 Business to be transacted in writing.
All business with the Patent or Trademark Office should be transacted in writing. The personal attendance of applicants or their attorneys or agents at the Patent and Trademark Office is unnecessary. The action of the Patent and Trademark Office will be based exclusively on the written record in the Office. No attention will be paid to any alleged oral promise, stipulation, or understanding in relation to which there is disagreement or doubt.

The action of the Patent and Trademark Office cannot be based exclusively on the written record in the Office if that record is itself incomplete through the failure to record the substance of interviews.

It is the responsibility of the applicant or the attorney or agent to make the substance of an interview of record in the application file, unless the examiner indicates he or she will do so. It is the examiner's responsibility to see that such a record is made and to correct material inaccuracies which bear directly on the question of patentability.

Examiners must complete an Interview Summary Form for each interview held where a matter of substance has been discussed during the interview by checking the appropriate boxes and filling in the blanks. Discussions regarding only procedural matters, directed solely to restriction requirements for which interview recordation is otherwise provided for in Section 812.01 of the Manual of Patent Examining Procedure, or pointing out typographical errors or unreadable script in Office actions or the like, are excluded from the interview recordation procedures below. Where the substance of an interview is completely recorded in an Examiners Amendment, no separate Interview Summary Record is required.

The Interview Summary Form shall be given an appropriate Paper No., placed in the right hand portion of the file, and listed on the "Contents" section of the file wrapper. In a personal interview, a duplicate of the Form is given to the applicant (or attorney or agent) at the conclusion of the interview. In the case of a telephone or video-conference interview, the copy is mailed to the applicant's correspondence address either with or prior to the next official communication. If additional correspondence from the examiner is not likely before an allowance or if other circumstances dictate, the Form should be mailed promptly after the interview rather than with the next official communication.

The Form provides for recordation of the following information:
- Application Number (Series Code and Serial Number)
- Name of applicant
- Name of examiner
- Date of interview
- Type of interview (telephonic, video-conference, or personal)
- Name of participant(s) (applicant, attorney or agent, examiner, other PTO personnel, etc.)
- An indication whether or not an exhibit was shown or a demonstration conducted
- An identification of the specific prior art discussed
- An indication whether an agreement was reached and if so, a description of the general nature of the agreement (may be by attachment of a copy of amendments or claims agreed as being allowable). Note: Agreement as to allowability is tentative and does not restrict further action by the examiner to the contrary.
- The signature of the examiner who conducted the interview (if Form is not an attachment to a signed Office action)

It is desirable that the examiner orally remind the applicant of his or her obligation to record the substance of the interview of each case. It should be noted, however, that the Interview Summary Form will not normally be considered a complete and proper recordation of the interview unless it includes, or is supplemented by the applicant or the examiner to include, all of the applicable items required below concerning the substance of the interview.

A complete and proper recordation of the substance of any interview should include at least the following applicable items:
1) A brief description of the nature of any exhibit shown or any demonstration conducted,
2) an identification of the claims discussed,
3) an identification of the specific prior art discussed,
4) an identification of the principal proposed amendments of a substantive nature discussed, unless these are already described on the Interview Summary Form completed by the Examiner,
5) a brief identification of the general thrust of the principal arguments presented to the examiner,
   (The identification of arguments need not be lengthy or elaborate. A verbatim or highly detailed description of the arguments is not required. The identification of the arguments is sufficient if the general nature or thrust of the principal arguments made to the examiner can be understood in the context of the application file. Of course, the applicant may desire to emphasize and fully describe those arguments which he or she feels were or might be persuasive to the examiner.)
6) a general indication of any other pertinent matters discussed, and
7) if appropriate, the general results or outcome of the interview unless already described in the Interview Summary Form completed by the examiner.

Examiners are expected to carefully review the applicant's record of the substance of an interview. If the record is not complete and accurate, the examiner will give the applicant an extendable one month time period to correct the record.

Examiner to Check for Accuracy

If the claims are allowable for other reasons of record, the examiner should send a letter setting forth the examiner's version of the statement attributed to him or her. If the record is complete and accurate, the examiner should place the indication, "Interview Record OK" on the paper recording the substance of the interview along with the date and the examiner's initials.

Continuation of Substance of Interview including description of the general nature of what was agreed to if an agreement was reached, or any other comments: Applicant called to ask about the PreAppeal Decision to send to the Board of Patent Appeals and Interferences. Examiner cited paragraphs dealing with a user that is able to affect others in the group (i.e. actions of another on a piece of electronic content). Examiner spoke with Supervisory Patent Examiner Mohammad Ali about suggestions for allowance. The only suggestion was to expand the independent claims to further distinguish over the prior art. No agreement was reached as to the allowability of the claims. .

.

PTO/SB/30EFS (12-08)
Approved for use through 01/31/2009. 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.

REQUEST FOR CONTINUED EXAMINATION(RCE)TRANSMITTAL
(Submitted Only via EFS-Web)

Application Number	11/623,653	Filing Date	2007-01-16	Docket Number (if applicable)	23682-12381	Art Unit	2158
First Named Inventor	Justin Marston			Examiner Name	John P. Hocker		

This is a Request for Continued Examination (RCE) under 37 CFR 1.114 of the above-identified application.
Request for Continued Examination (RCE) practice under 37 CFR 1.114 does not apply to any utility or plant application filed prior to June 8, 1995, or to any design application. The Instruction Sheet for this form is located at WWW.USPTO.GOV

SUBMISSION REQUIRED UNDER 37 CFR 1.114

Note: If the RCE is proper, any previously filed unentered amendments and amendments enclosed with the RCE will be entered in the order in which they were filed unless applicant instructs otherwise. If applicant does not wish to have any previously filed unentered amendment(s) entered, applicant must request non-entry of such amendment(s).

☐ Previously submitted. If a final Office action is outstanding, any amendments filed after the final Office action may be considered as a submission even if this box is not checked.

- ☐ Consider the arguments in the Appeal Brief or Reply Brief previously filed on _____
- ☐ Other _____

☒ Enclosed

- ☒ Amendment/Reply
- ☒ Information Disclosure Statement (IDS)
- ☐ Affidavit(s)/ Declaration(s)
- ☐ Other _____

MISCELLANEOUS

☐ Suspension of action on the above-identified application is requested under 37 CFR 1.103(c) for a period of months _____ (Period of suspension shall not exceed 3 months; Fee under 37 CFR 1.17(i) required)

☐ Other _____

FEES

☒ **The RCE fee under 37 CFR 1.17(e) is required by 37 CFR 1.114 when the RCE is filed.**
The Director is hereby authorized to charge any underpayment of fees, or credit any overpayments, to Deposit Account No 192555

SIGNATURE OF APPLICANT, ATTORNEY, OR AGENT REQUIRED

☒ Patent Practitioner Signature

☐ Applicant Signature

PTO/SB/30EFS (12-08)
Approved for use through 01/31/2009. 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.

Signature of Registered U.S. Patent Practitioner			
Signature	/Christopher King/	Date (YYYY-MM-DD)	2010-01-15
Name	Christopher-Charles P. King	Registration Number	60985

This collection of information is required by 37 CFR 1.114. 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.

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 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.

IN THE UNITED STATES

PATENT AND TRADEMARK OFFICE

APPLICANT:	Justin Marston et al
APPLICATION NO.:	11/623,653
FILING DATE:	January 16, 2007
TITLE:	Determining Relevance of Electronic Content
EXAMINER:	John P. Hocker
GROUP ART UNIT:	2169
ATTY. DKT. NO.:	20423-12381

CERTIFICATE OF ELECTRONIC (EFS-WEB) TRANSMISSION

I hereby certify that this correspondence is being transmitted via the Office electronic filing system in accordance with 37 C.F.R. § 1.8(a)(i)(C) from the **Pacific Time Zone** of the United States on the local date shown below.

Dated: January 15, 2010 By: /Christopher King/
Christopher P. King, Reg. No. 60,985

AMENDMENT C AND RESPONSE

Sir:

This amendment is in response to the Final Office Action mailed May 12, 2009, and is timely filed with respect to the Notice of Appeal filed on August 12, 2009.

IN THE CLAIMS

Claims 1, 7, 8, 11, and 17 are amended herein. Claims 2-6, 12-16, and 18-22 are canceled, and claims 23-37 are added. All pending claims are produced below. In addition, the status of each is also indicated below and appropriately noted as "Original", "Currently Amended", "Canceled", "New", "Withdrawn", "Previously Presented", and "Not Entered" as requested by the Office.

1. (Currently Amended) A computer-implemented system for determining a relevance score of a piece of electronic content sent from a sender to a user, the relevance score determined with respect to the user, the computer-implemented system comprising:

a computer processor; and

a computer-readable storage medium storing computer program modules configured to execute on the computer processor, the computer program modules comprising:

a monitoring module configured to track actions by a plurality of users associated with an application for managing electronic content;

~~a relevance analysis module coupled to the monitoring module and configured to determine a relevance to a user of a piece of electronic content sent from a sender to the user, the relevance being determined dependent on:~~

a first relevance measurement module coupled to the monitoring module and configured to determine a first measure of [[a]] relevance of the piece of electronic content based at least in part ~~as determined dependent~~ on one or more actions of the user on the piece of electronic content;~~, and~~

a second relevance measurement module coupled to the monitoring module and configured to:

identify at least one other of the plurality of users that is a recipient of the piece of electronic content and that is cognate to the user, and

determine a second measure of [[a]] relevance of the piece of electronic content ~~as determined dependent~~ based at least in part on one or more measures derived from the group consisting of:

a local importance describing an importance of the sender to the user,

a global importance describing an importance of the sender within an organization with which the plurality of users are associated,

a content relevance determined responsive to content of the piece of electronic content, and

one or more actions on the piece of electronic content of the at least one other of the plurality of users that is a recipient of the piece of electronic content and that is cognate to the user; and

a relevance analysis module configured to determine the relevance score of the piece of electronic content sent from the sender to the user, the relevance score being determined based at least in part on the first determined measure of relevance and on the second determined measure of relevance.

2-6. (Canceled)

7. (Currently Amended) The system of claim 1, further comprising:

a content presentation module configured to present the piece of electronic content to the user in accordance with the determined relevance score and to present to the user a value describing the determined relevance score.

8. (Currently Amended) The system of claim 1, further comprising:

a search module configured to search the electronic content in accordance with the determined relevance score.

9. (Original) The system of claim 1, wherein the piece of electronic content comprises an electronic mail object in a relational mail system.

10. (Original) The system of claim 1, wherein the electronic content comprises at least one of an email, a document, an instant message, and an SMS message.

11. (Currently Amended) A computer program product for determining a relevance score of a piece of electronic content sent from a sender to a user, the relevance score determined with respect to the user, the computer-program product having a computer-readable storage medium having computer program instructions encoded thereon comprising:

a monitoring module configured to track actions by a plurality of users associated with an application for managing electronic content;

~~a relevance analysis module coupled to the monitoring module and configured to determine a relevance to a user of a piece of electronic content sent from a sender to the user, the relevance being determined dependent on:~~

a first relevance measurement module coupled to the monitoring module and configured to determine a first measure of [[a]] relevance of the piece of electronic content based at least in part ~~as determined dependent~~ on one or more actions of the user on the piece of electronic content;~~, and~~

a second relevance measurement module coupled to the monitoring module and configured to:

identify at least one other of the plurality of users that is a recipient of the piece of electronic content and that is cognate to the user, and

determine a second measure of [[a]] relevance of the piece of electronic content based at least in part ~~as determined dependent~~ on one or more measures derived from the group consisting of:

a local importance describing an importance of the sender to the user,

a global importance describing an importance of the sender within an organization with which the plurality of users are associated,

a content relevance determined responsive to content of the piece of electronic content, and

one or more actions on the piece of electronic content of the at least one other of the plurality of users that is a recipient of the piece of electronic content and that is cognate to the user; and

a relevance analysis module configured to determine the relevance score of the piece of electronic content sent from the sender to the user, the relevance score being determined based at least in part on the first determined measure of relevance and on the second determined measure of relevance.

12-16. (Canceled)

17. (Currently Amended) A computer-implemented method for determining a relevance score of a piece of electronic content sent from a sender to a user, the relevance score determined with respect to the user, the computer-implemented method[[,]] comprising:

tracking actions by a plurality of users associated with an application for managing electronic content;

~~determining a relevance to a user of a piece of electronic content sent from a sender to the user, the relevance being determined dependent on:~~

identifying at least one other of the plurality of users that is a recipient of the piece of electronic content and that is cognate to the user,

determining a first measure of [[a]] relevance of the piece of electronic content based at least in part ~~as determined dependent~~ on one or more actions of the user on the piece of electronic content;~~, and~~

determining a second measure of [[a]] relevance of the piece of electronic content ~~as determined dependent~~ based at least in part on one or more measures derived from the group consisting of:

a local importance describing an importance of the sender to the user,

a global importance describing an importance of the sender within an organization with which the plurality of users are associated,

a content relevance determined responsive to content of the piece of electronic content, and

one or more actions on the piece of electronic content of the at least one other of the plurality of users that is a recipient of the piece of electronic content and that is cognate to the user; and

determining the relevance score of the piece of electronic content sent from the sender to the user, the relevance score being determined based at least in part on the first determined measure of relevance and on the second determined measure of relevance.

18-22. (Canceled)

23. (New) The system of claim 1, wherein identifying at least one other of the plurality of users that is cognate to the user comprises analyzing historical actions of the at least one other of the plurality of users with respect to historical actions of the user.

24. (New) The system of claim 23, wherein the at least one other of the plurality of users is cognate to the user, and the user is not cognate to the at least one other of the plurality of users.

25. (New) The system of claim 9, wherein the electronic mail object includes one or more submessages.

26. (New) The system of claim 25, wherein the submessages are represented in the electronic mail object as references to submessage content stored in a data store of the relational mail system.

27. (New) The system of claim 9, wherein the electronic mail object has an associated job code representing a business process with which the electronic mail object is associated.

28. (New) The computer program product of claim 11, further comprising:

a content presentation module configured to present the piece of electronic content to the user in accordance with the determined relevance score and to present to the user a value describing the determined relevance score.

29. (New) The computer program product of claim 11, further comprising:

a search module configured to search the electronic content in accordance with the determined relevance score.

30. (New) The computer program product of claim 11, wherein the electronic content comprises at least one of an email, a document, an instant message, and an SMS message.

31. (New) The computer program product of claim 11, wherein the piece of electronic content comprises an electronic mail object in a relational mail system.

32. (New) The computer program product of claim 31, wherein the electronic mail object includes one or more submessages that are represented in the electronic mail object as references to submessage content stored in a data store of the relational mail system.

33. (New) The method of claim 17, further comprising:

presenting the piece of electronic content to the user in accordance with the determined relevance score and presenting to the user a value describing the determined relevance score.

34. (New) The method of claim 17, further comprising:

searching the electronic content in accordance with the determined relevance score.

35. (New) The method of claim 17, wherein the piece of electronic content comprises an electronic mail object in a relational mail system.

36. (New) The method of claim 35, wherein the electronic mail object includes one or more submessages that are represented in the electronic mail object as references to submessage content stored in a data store of the relational mail system.

37. (New) The method of claim 17, wherein identifying at least one other of the plurality of users that is cognate to the user comprises analyzing historical actions of the one other of the plurality of users with respect to historical actions of the user.

REMARKS

Claims 1-22 were finally rejected in a Final Office Action mailed May 12, 2009. A Notice of Appeal and Pre-Appeal Request for Review was filed on August 12, 2009. A Notice of Panel Decision from Pre-Appeal Brief Review was mailed on December 15, 2009. This amendment is filed with a Request for Continued Examination and thus reopens prosecution under 37 C.F.R. §1.114(d).

Applicants are amending claims 1, 7, 8, 11, and 17. Claims 2-6, 12-16, and 18-22 (15 claims total) are canceled, and claims 23-37 (also 15 claims) are added. In making these amendments, Applicants do not concede that the subject matter of such claims was in fact disclosed or taught by the cited prior art. Rather, Applicants reserve the right to pursue such protection at a later point in time and merely seek to pursue protection for the subject matter presented in this submission.

Applicants' representative spoke with Examiner Hocker on January 14 and January 15, 2010, in order to expedite prosecution of the case. Examiner Hocker agreed that claim 1, as amended herein, contains features not found in the prior art and is allowable. Accordingly, although not agreeing that the claims in their prior state were unpatentable over the cited prior art, Applicants are amending the claims herein in order to expedite allowance.

Applicants request the Examiner to contact Applicants' representative at the number provided below if any further issues arise prior to allowance.

Respectfully submitted,

Dated: January 15, 2010

By: /Christopher King/

Christopher P. King, Reg. No. 60,985
Attorney for Applicant
Fenwick & West LLP
801 California Street
Mountain View, CA 94041
Tel.: (650) 335-7633
Fax: (650) 938-5200

Electronic Patent Application Fee Transmittal	
Application Number:	11623653
Filing Date:	16-Jan-2007
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Filer:	Christopher-Charles P. King
Attorney Docket Number:	23682-12381
Filed as Large Entity	

Utility under 35 USC 111(a) Filing Fees

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Miscellaneous:				
Request for continued examination	1801	1	810	810
		Total in USD ($)		**810**

Electronic Acknowledgement Receipt	
EFS ID:	6825236
Application Number:	11623653
International Application Number:	
Confirmation Number:	5243
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Customer Number:	00758
Filer:	Christopher-Charles P. King
Filer Authorized By:	
Attorney Docket Number:	23682-12381
Receipt Date:	15-JAN-2010
Filing Date:	16-JAN-2007
Time Stamp:	20:37:17
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	Deposit Account
Payment was successfully received in RAM	$810
RAM confirmation Number	6204
Deposit Account	192555
Authorized User	
The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows: Charge any Additional Fees required under 37 C.F.R. Section 1.16 (National application filing, search, and examination fees) Charge any Additional Fees required under 37 C.F.R. Section 1.17 (Patent application and reexamination processing fees)	

Charge any Additional Fees required under 37 C.F.R. Section 1.19 (Document supply fees)	
Charge any Additional Fees required under 37 C.F.R. Section 1.20 (Post Issuance fees)	
Charge any Additional Fees required under 37 C.F.R. Section 1.21 (Miscellaneous fees and charges)	

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Transmittal Letter	12381_US_SIDS.pdf	120875 5495a67da980afe3a25a81006cf9791167b78b6b	no	3
Warnings:					
Information:					
2	Information Disclosure Statement (IDS) Filed (SB/08)	12381_US_Form_8A.pdf	120241 bca4d8eeef4986f1296dd0c59687de304c60c057	no	4
Warnings:					
Information:					
This is not an USPTO supplied IDS fillable form					
3	Foreign Reference	EP0739115A2.pdf	555991 1012edfab308d2481bef5317eaf883f4b924ffad	no	11
Warnings:					
Information:					
4	Foreign Reference	WO03036492.pdf	970290 6102512a795f0cb7f6ab38e41d3dccb1f29d5983	no	23
Warnings:					
Information:					
5	Foreign Reference	WO03073711.pdf	3493479 20600ade67f102900b1212aedc9f660df4767b4e	no	83
Warnings:					
Information:					
6	Foreign Reference	WO04012415.pdf	1251182 cfd34db6792e60f90a4ec1c70f81be3a8e398933	no	38
Warnings:					
Information:					
7	NPL Documents	BRUSSEE.pdf	4542578 c557a47b79808b024f41d85c331711bf57299662	no	95
Warnings:					
Information:					

8	NPL Documents	DAVISON.pdf	1086831 897dc23b35b9318115fd6c730049ece2933c24c6	no	16
Warnings:					
Information:					
9	NPL Documents	Gertzman.pdf	374704 aa93cdd9f2f08c77621cd90b373c4550cc2fb963	no	5
Warnings:					
Information:					
10	NPL Documents	Gupta.pdf	1153687 5f4eb3d5ab5247dcadf13a384e044ae84a40354d	no	11
Warnings:					
Information:					
11	NPL Documents	Mambo.pdf	949113 08df5666d79890f9aae3c7ee641e6331233d613a	no	10
Warnings:					
Information:					
12	NPL Documents	STANFORD.pdf	902308 9d3f774175c63f4436280f3e810321575bb8797a	no	12
Warnings:					
Information:					
13	NPL Documents	VAN.pdf	712776 2caef54496d07146d1fe10db86476d6c425bb57a	no	13
Warnings:					
Information:					
14	NPL Documents	Faden_Multilevel.pdf	240650 4e524223efb2138d6d4cca37349f2c746da5df17	no	6
Warnings:					
Information:					
15	NPL Documents	Faden_Solaris.pdf	810277 56c918c30aba0ba2b0f7e58e691cd0689bc0e5be	no	11
Warnings:					
Information:					
16	NPL Documents	PCT001845IPRP.pdf	217844 f6c9bab0964f6ea030b1b23e8d4a07fa496a0832	no	6
Warnings:					
Information:					

17	NPL Documents	PCT001845ISRWO.pdf	280295 9a8a637d096800f68f0d01c8eebd67d0a0554d8e	no	8
Warnings:					
Information:					
18	NPL Documents	PCT004609.pdf	345969 26bd088fc6fa392c5cd61f5021b6b4228f7c397a	no	9
Warnings:					
Information:					
19	NPL Documents	PCT000805.pdf	323336 d541a4c516657c044aa2a9d67619b1f072fbf9a5	no	9
Warnings:					
Information:					
20	NPL Documents	PCT001868.pdf	331961 a0fc52ebbc65a4b375d62154602d9017d88df781	no	8
Warnings:					
Information:					
21	NPL Documents	GOLDSZMIDT.pdf	757578 22e6840517320265b21806033eaa2f73e5d6534c	no	10
Warnings:					
Information:					
22	Request for Continued Examination (RCE)	12381_RCE.pdf	697436 d9fc8dce5aac195e2fce37b3e400e64bf999aed6	no	3
Warnings:					
Information:					
23		12381_AmendmentC.pdf	86957 38af5aae82aae6d6530942cd7325776b96adefbc	yes	11

Multipart Description/PDF files in .zip description

Document Description	Start	End
Amendment Submitted/Entered with Filing of CPA/RCE	1	1
Claims	2	9
Applicant Arguments/Remarks Made in an Amendment	10	11

Warnings:	
Information:	

24	Fee Worksheet (PTO-875)	fee-info.pdf	30199 8921a11391001abc7fd60ec80f0c219a1b27f782	no	2
Warnings:					
Information:					
			Total Files Size (in bytes):	20356557	

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

APPLICANT:	Justin Marston
APPLICATION NO.:	11/623,653
FILING DATE:	January 16, 2007
TITLE:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
EXAMINER:	John P. Hocker
GROUP ART UNIT:	2158
ATTY. DKT. NO.:	23682-12381

CERTIFICATE OF ELECTRONIC (EFS-WEB) TRANSMISSION

I hereby certify that this correspondence is being transmitted via the Office electronic filing system in accordance with 37 C.F.R. § 1.8(a)(i)(C) from the Pacific Time Zone of the United States on the local date shown below.

Dated: __January 15, 2010_____ By: _____/Christopher King/_____

Christopher P. King, Reg. No. 60,985

INFORMATION DISCLOSURE STATEMENT
Under 37 CFR §§ 1.56 and 1.97-98

SIR:

Pursuant to the provisions of 37 CFR §§ 1.56 and 1.97-98, enclosed herewith is modified form PTO/SB/08A listing references for consideration by the Examiner.

The filing of this Information Disclosure Statement shall not be construed as a representation regarding the completeness of the list of references, or that inclusion of a reference in this list is an admission that it is prior art or is pertinent to this application, or that a search has been made, or as an admission that the information listed is, or may be considered to be, material to patentability, or that no other material information exists, and shall not be construed as an admission against interest in any manner.

This Information Disclosure Statement is being filed:

- ☐ within three months of the filing date of the application, or date of entry into the national stage of an international application, or before the mailing date of a first office action on the merits, whichever event last occurred;
- ☒ before the mailing of a first official action after the filing of a request for continued examination (RCE) under 37 CFR § 1.114;

☐ after three months of the filing date of this national application or the date of entry of the national stage in an international application, or after the mailing date of the first official action on the merits, whichever event last occurred, but before the mailing date of the first to occur of either: (1) a final action under 37 CFR §1.113; or (2) an action that otherwise closes prosecution in the application, and:

- ☐ attached hereto is the fee set forth under 37 CFR §1.17(p) for submission of this Information Disclosure Statement under 37 CFR.§ 1.97(c); OR
- ☐ Applicant certifies pursuant to 37 CFR § 1.97(e) that:
  - ☐ each item of information contained in this Information Disclosure Statement was first cited in a communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of this Statement; OR
  - ☐ no item of information contained in this 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 this certification after making reasonable inquiry, no item of information contained in this Statement was known to any individual designated under 37 CFR § 1.56(c) more than three months prior to the filing of this Statement;

☐ on or before the payment of the issue fee but after the mailing date of the first to occur of either: (1) a final action under 37 CFR § 1.113; (2) a notice of allowance under 37 CFR § 1.311; or (3) an action that otherwise closes prosecution in the application, and:

- ☐ Applicant certifies pursuant to 37 CFR. § 1.97(e) that:
  - ☐ each item of information contained in this Information Disclosure Statement was cited in a communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of this Statement;
  - ☐ no item of information contained in this 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 this certification after making reasonable inquiry, no item of information contained in this Statement was known to any individual designated under 37 CFR § 1.56(c) more than three months prior to the filing of this Statement; AND

☐ attached hereto is the fee set forth under 37 CFR §1.17(p) for submission of this Information Disclosure Statement under 37 CFR.§ 1.97(d); OR

☐ after the payment of the issue fee. Applicant requests that the information contained in this Information Disclosure Statement be placed in the file according to 37 CFR § 1.97(i), although the information may not be considered by the USPTO.

☐ This application relies, under 35 U.S.C. § 120, on the earlier filing date of prior application No. _____ , filed on _____ , and the references cited therein are hereby referenced, but are not required to be provided in this application under 37 CFR § 1.98(d).

☐ Each item of information contained in this Information Disclosure Statement was cited in a communication from a foreign patent office in a counterpart application, and the communication was not received by any individual designated in 37 CFR § 1.56(c) more than thirty days prior to the filing of this Information Disclosure Statement. 37 CFR § 1.704(d).

Respectfully submitted,

Dated: January 15, 2010

/Christopher King/

Christopher P. King, Reg. No. 60,985
Attorney for Applicant
Fenwick & West LLP
801 California Street
Mountain View, CA 94041
Tel.: (650) 335-7633
Fax: (650) 938-5200

Substitute for form 1449A/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

Complete if Known	
Application No.	11/623,653
Filing Date	January 16, 2009
First Named Inventor	Justin Marston
Art Unit	2158
Examiner Name	John P. Hocker
Attorney Docket Number	23682-12381

Sheet 1 of 4

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Document No. Number – Kind Code[2] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document
	A1	US-2007/0038942 A1	02-15-2007	Chen et al.
	A2	US-2006/0168029 A1	07-27-2006	Fitzpatrick et al.
	A3	US-2005/0149498 A1	07-07-2005	Lawrence et al.
	A4	US-2005/0086340 A1	04-21-2005	Kang et al.
	A5	US-2005/0054334 A1	03-10-2005	Brown et al.
	A6	US-2004/0260761 A1	12-23-2004	Leaute et al.
	A7	US-2004/0177271 A1	09-09-2004	Arnold et al.
	A8	US-2004/0153515 A1	08-05-2004	Touboul et al.
	A9	US-2004/0064515 A1	04-01-2004	Hockey
	A10	US-2004/0054498 A1	03-18-2004	Shipp
	A11	US-2003/0140223 A1	07-24-2003	Desideri
	A12	US-2003/0055907 A1	03-20-2003	Stiers
	A13	US-2003/0009595 A1	01-09-2003	Collins
	A14	US-2002/0144154 A1	10-03-2002	Tomkow
	A15	US-2002/0007453 A1	01-17-2002	Nemovicher
	A16	US-7,412,437 B2	08-12-2008	Moody et al.
	A17	US-7,233,954	06-19-2007	Horvitz
	A18	US-7,103,794	09-05-2006	Malcolm
	A19	US-7,058,687	06-06-2006	Kucherawy
	A20	US-6,970,913	11-29-2005	Albert et al.
	A21	US-6,959,382	10-25-2005	Kinnis et al.
	A22	US-6,850,968	02-01-2005	Pfeffer et al.
	A23	US-6,789,107 B1	09-07-2004	Bates et al.
	A24	US-6,704,772	03-09-2004	Ahmed et al.
	A25	US-6,704,768	03-09-2004	Zombek
	A26	US-6,694,358 B1	02-17-2004	Swildens et al.
	A27	US-6,581,090 B1	06-17-2003	Lindbo et al.
	A28	US-6,578,076 B1	06-10-2003	Putzolu

Examiner Signature		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.

[1]Applicant's unique citation designation number (optional). [2]See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3]Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4]For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5]Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST. 16 if possible. [6]Applicant is to place a check mark here if English language Translation is attached.

Substitute for form 1449A/PTO				Complete if Known	
INFORMATION DISCLOSURE STATEMENT BY APPLICANT				**Application No.**	11/623,653
				Filing Date	January 16, 2009
				First Named Inventor	Justin Marston
				Art Unit	2158
				Examiner Name	John P. Hocker
Sheet	2	of	4	Attorney Docket Number	23682-12381

U.S. PATENT DOCUMENTS				
		Document No.		
Examiner Initials*	Cite No.[1]	Number – Kind Code[2] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document
	A29	US-6,553,413 B1	04-22-2003	Leighton et al.
	A30	US-6,502,125 B1	12-31-2002	Kenner et al.
	A31	US-6,484,143 B1	11-19-2002	Swildens et al.
	A32	US-6,480,893 B2	11-12-2002	Kriegsman
	A33	US-6,314,565 B1	11-06-2001	Kenner et al.
	A34	US-6,282,569 B1	08-28-2001	Wallis et al.
	A35	US-6,421,726 B1	07-16-2002	Kenner et al.
	A36	US-6,415,280 B1	07-02-2002	Farber et al.
	A37	US-6,405,252 B1	06-11-2002	Gupta et al.
	A38	US-6,370,571 B1	04-09-2002	Medin, Jr.
	A39	US-6,185,598 B1	02-06-2001	Farber et al.
	A40	US-6,185,551	02-06-2001	Birell et al.
	A41	US-6,182,136	01-30-2001	Ramanathan et al.
	A42	US-6,181,867 B1	01-30-2001	Kenner et al.
	A43	US-6,178,160 B1	01-23-2001	Bolton et al.
	A44	US-6,154,744	11-28-2000	Kenner et al.
	A45	US-6,134,582 A	10-17-2000	Kennedy
	A46	US-6,134,598 A	10-17-2000	Raman
	A47	US-6,122,632 A	09-19-2000	Botts et al.
	A48	US-6,112,239	08-29-2000	Kenner et al.
	A49	US-6,035,327	03-07-2000	Buckley et al.
	A50	US-6,029,175	02-22-2000	Chow et al.
	A51	US-6,009,462	12-28-1999	Birrell et al.
	A52	US-5,938,735	08-17-1999	Malik
	A53	US-5,938,732	08-17-1999	Lim et al.
	A54	US-6,003,030	12-14-1999	Kenner et al.

Examiner Signature		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.

[1]Applicant's unique citation designation number (optional). [2]See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3]Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4]For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5]Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST. 16 if possible. [6]Applicant is to place a check mark here if English language Translation is attached.

Substitute for form 1449A/PTO				Complete if Known	
INFORMATION DISCLOSURE STATEMENT BY APPLICANT				**Application No.**	11/623,653
				Filing Date	January 16, 2009
				First Named Inventor	Justin Marston
				Art Unit	2158
				Examiner Name	John P. Hocker
Sheet	3	of	4	Attorney Docket Number	23682-12381

FOREIGN PATENT DOCUMENTS					
		Foreign Patent Document			
Examiner Initials*	Cite No.[1]	Country Code[3] – Number[4] Kind Code[5] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	T[6]
	B1	EP 0739115 A2	10-23-1996	AT&T IPM CORP	
	B2	WO 03/036492 A1	05-01-2001	Idetic, Inc.	
	B3	WO 03/073711 A2	09-04-2003	RPOST International Limited	
	B4	WO 04/012415 A1	02-05-2004	Security and Standards Limited	

OTHER REFERENCES – NON-PATENT LITERATURE DOCUMENTS			
Examiner Initials*	Cite No.[1]	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, page(s), volume-issue number(s), publisher, city and/or country where published	T[6]
	C1	BRUSSEE, ROGIER et al, "Content Distribution Networks", Telematica Institute, June 1, 2001, pp. V-IX and 1-88.	
	C2	DAVISON, BRIAN, "Brian Davison's Web-Caching Bibliography", Department of Computer Science, Rutgers, the State University, New Jersey, June 28, 2000, pp. 1-16.	
	C3	GOLDSZMIDT, GERMAN et al., "Load Distribution for Scalable Web Servers: Summer Olympics 1996 – A Case Study", In Proceedings of the 8th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management, Sydney, Australia, October 1997, pp. 1-10.	
	C4	GWERTZMAN, JAMES et al., "The Case for Geographical Push-Caching", *In* Proceedings of the Fifth Workshop on Hot Topics in Operating Systems, IEEE Computer Society, Orcas Island, WA , May 1995, pp. 51-55.	
	C5	GUPTA, G., et al., "Digital Forensics Analysis of E-Mails: A Trusted E-Mail Protocol," International Journal of Digital Evidence, Spring 2004, Vol. 2, Is. 4, 11 Pages.	
	C6	MAMBO, M., et al., "Proxy Signatures for Delegating Signing Operation," Conference on Computer and Communications Security, Proceedings of the 3rd ACM conference on Computer and Communications Security, 1996, Pages 48-57.	
	C7	STANFORD-CLARK, ANDY, "Atlanta Olympics WOMplex", Presented at Get Connected Technical Interchange '96, IBM, Hursley, U.K., October 1996, pp. IBM 000026-000037.	
	C8	VAN STEEN, MAARTEN et al, "Locating Objects in Wide-Area Systems", IEEE Communications Magazine Vo 36, No. 1, Jan. 1998, pp. 104-109.	
	C9	PCT International Preliminary Report on Patentability and Written Opinion, PCT/GB2005/001845, November 29, 2006, 6 pages.	
	C10	International Search Report and Written Opinion, PCT/GB2005/001845, August 24, 2005, 8 Pages.	
	C11	International Search Report and Written Opinion, PCT/GB2004/004609, February 22, 2005, 9 pages.	

Examiner Signature		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.

[1]Applicant's unique citation designation number (optional). [2]See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3]Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4]For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5]Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST. 16 if possible. [6]Applicant is to place a check mark here if English language Translation is attached.

Substitute for form 1449A/PTO				Complete if Known	
INFORMATION DISCLOSURE STATEMENT BY APPLICANT				**Application No.**	11/623,653
				Filing Date	January 16, 2009
				First Named Inventor	Justin Marston
				Art Unit	2158
				Examiner Name	John P. Hocker
Sheet	4	of	4	Attorney Docket Number	23682-12381

OTHER REFERENCES – NON-PATENT LITERATURE DOCUMENTS			
Examiner Initials*	Cite No.[1]	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, page(s), volume-issue number(s), publisher, city and/or country where published	T[6]
	C12	International Search Report and Written Opinion, PCT/GB2004/000805, June 4, 2004, 9 pages.	
	C13	International Search Report and Written Opinion, PCT/GB2005/001868, August 19, 2005, 8 pages.	
	C14	FADEN, G., "Solaris Trusted Extensions," Architectural Overview, Sun Microsystems, April 2006, [online] [Retrieved on December 19, 2008] Retrieved from the internet <URL:http://www.opensolaris.org/os/community/security/projects/tx/TrustedExtensionsArch.pdf>.	
	C15	FADEN, G., "Multilevel Filesystems in Solaris Trusted Extensions," SACMAT'07, June 20-22, 2007, [online] [Retrieved on December 19, 2008] Retrieved from the internet URL:http://www.opensolaris.org/os/community/security/projects/tx/sacmat04s-faden-1.pdf>.	

Examiner Signature		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.

[1]Applicant's unique citation designation number (optional). [2]See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3]Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4]For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5]Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST. 16 if possible. [6]Applicant is to place a check mark here if English language Translation is attached.

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 0 739 115 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
23.10.1996 Bulletin 1996/43

(51) Int Cl.[6]: **H04L 12/58**

(21) Application number: **96302503.6**

(22) Date of filing: **10.04.1996**

(84) Designated Contracting States:
DE ES FR GB

(30) Priority: **20.04.1995 US 425968**

(71) Applicant: **AT&T IPM Corp.**
Coral Gables, Florida 33134 (US)

(72) Inventors:
- **Baldwin, Michael Scott**
  Plainfield, New Jersey 07060 (US)
- **Lustgarten, Paul C.**
  Westfield, New Jersey 07090 (US)

(74) Representative:
Watts, Christopher Malcolm Kelway, Dr. et al
Lucent Technologies (UK) Ltd,
5 Mornington Road
Woodford Green Essex, IG8 0TU (GB)

(54) Electronic messaging in a wide area network

(57) A method for electronic communication using wide area network (WAN) storage and signalling capabilities. The method and apparatus provides for deposit of a message in message stores in the WAN, storage of the message such that the message is uniformly accessible throughout the WAN and notifying the recipient of the message of the availability and address of the message in the message stores in the WAN so that he can access the message.

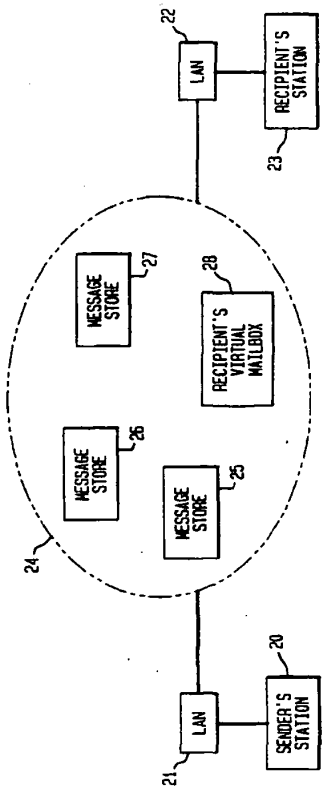


FIG. 2

EP 0 739 115 A2

Printed by Jouve, 75001 PARIS (FR)

Description

TECHNICAL FIELD

This invention relates to a method and apparatus for providing electronic communications.

BACKGROUND OF THE INVENTION

Messaging can be defined as time-delayed communication or communication in which the origination of the communication occurs asynchronously from its reception. For example, letters and public notices are traditional form of messaging. Messaging also includes traditional forms of publishing, whether text/image-based (e.g. memoranda, books and periodicals), audio-based (e.g. tapes and compact discs) or audio/visually-based (e.g. movies). Other forms of messaging include voice mail, including multi-user systems and answering machines, and electronic counterparts of traditional messaging, for example telegrams, cablegrams, faxes, telex, electronic mail, asynchronous computer-based conferencing (i.e. bulletin-boards) and electronic publishing.

Because messaging is communication across time, messaging requires storage of the message content. The way in which the message is stored can be used to distinguish types of messaging.

"Store-and-forward" messaging is one type of electronic messaging system. In store-and-forward messaging, a message is transmitted from its sender to its recipient through at least one "message transfer agent" (MTA). The message can be relayed through as many successive MTAs as are needed for the message to reach its final destination. Each MTA independently chooses the next MTA to which to route the message, based on information carried along with the message, such as its intended recipient, together with information held by the MTA, such as routing tables. If a message is directed to more than one recipient, an MTA can multicast the message by sending copies of the message to separate MTAs for forwarding to each intended recipient. When the message reaches an MTA that has direct access to the recipient's mailbox, the message is deposited in local area network (LAN) storage associated with the recipient's mailbox. The LAN storage associated with recipient's mailbox is comprised of message stores for storing messages deposited therein.

During the messaging transmission process, each MTA temporarily stores the message, until the message is transferred to another MTA or deposited in the LAN storage associated with the recipient's mailbox. Thus, the MTAs constitute a succession of storage locations and the message is forwarded from MTA to MTA until ultimately reaching the recipient's mailbox.

Referring to Fig. 1, the general structure of store and forward messaging as applied to LAN messaging and wide area network (WAN) messaging is shown. With respect to LAN messaging, a message can be sent from a sender at a sender station 1 to a recipient at a recipient station 2 through a LAN 3. The message is transmitted from the sender station 2 through an MTA 6 in LAN 3. MTA 6 has a message store 12 that stores the message temporarily until the MTA, acting on information carried along with the message, routes the message. Because MTA 6 has direct access to the recipient's mailbox 4, MTA 6 deposites the message into LAN 3 storage associated with the recipient's mailbox 4. The recipient can then access the message in his mailbox 4 from the recipient station 2.

With respect to WAN messaging, a series of LAN systems may be connected together (not shown) or LAN systems may be connected through an intermediate WAN (shown in Fig. 1) to permit WAN messaging. As shown in Fig. 1, a message can be sent from the sender at sender station 1 to a recipient at a recipient station 11 through MTA 6 in LAN 3, MTA's 7, 8 and 9 in a WAN 17 and an MTA 10 in a LAN 5. MTA's 6, 7, 8, 9 and 10 have messages stores 12, 13, 14, 15 and 16, respectively, that store the message temporarily until the MTA, acting on information carried along with the message, routes the message. The message is transmitted from sender station 1 to MTA 6 in LAN 3. MTA 6 routes the message to MTA 7 in WAN 17. MTA 7 in turn routes the message to MTA 8 in WAN 17 and MTA 8 in turn routes the message to MTA 9 in WAN 17. MTA 9 routes the message to MTA 10 in LAN 5. MTA 10 has direct access to the recipient's mailbox 18 and deposits the message into the LAN 5 storage associated with the recipient's mailbox 18. The recipient can access the message in his mailbox 18 from the recipient station 11.

Because store-and-forward messaging is an application specific overlay on a network, it has limitations. For example, the costs involved are spread only among the users of the store-and-forward messaging application. Moreover, the scope of messaging is limited to users of that application. Finally, diverse message formats and conversion- from one format to another them are not readily accommodated.

SUMMARY OF THE INVENTION

The above problems are solved in accordance with the principles of the invention by providing a method and apparatus for WAN messaging using WAN storage and signalling capabilities. The invention generally includes three aspects: message origination, message storage and message reception.

Message origination includes the sender's deposit of a message into at least one message store on the WAN, identification of the addresses of the intended recipients and notification of each of the intended recipients of the availability and address of the message on the WAN.

Message storage refers to the characteristics of storage in the WAN. The message stores and the mes-

sages stored have a logical identity in the WAN independent of their physical realization such that uniform accessibility is provided throughout the WAN.

Message reception refers to the recipient's processing of the message. Message reception is provided via "virtual mailboxes" that reside in storage elements on the WAN. The recipient's virtual mailbox does not contain his messages. Instead, it serves as an aggregation and storage point for the signaling information that notifies the recipient of the availability and address of his messages on the WAN. The recipient's station uses the signaling data to alert the recipient of the existence of his messages and to retrieve the messages under the recipient's direction.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram showing the structure of store-and-forward messaging.

Fig. 2 is a block diagram showing the structure of an embodiment of the invention.

Fig. 3 is a flow diagram illustrating the processing of messaging in accordance with an embodiment of the invention.

DETAILED DESCRIPTION

Referring to Fig. 2, the structure of an embodiment of the invention is shown. The sender's station 20 is connected to a LAN 21 which is connected to a WAN 24 and the recipient's station 23 is connected to a LAN 22 which is connected to WAN 24. LAN 21 and LAN 22 do not require message stores because the WAN 24 includes message stores, such as message stores 25, 26 and 27, for storing messages. WAN 24 also includes virtual mailboxes, such as the recipient's virtual mailbox 28, for storing information regarding the availability and address of the recipient's messages.

To deposit a message into message stores in the WAN, the sender at a sender station allocates storage from the WAN by selecting at least one message store from those residing in storage elements integrated into the WAN and fills that message store with the contents of the message.

The storage location of a message may be chosen by the sender's station. For example, if all messages for a given recipient are stored in one location, the sender's station directly deposits the message into those message stores and, if there are multiple recipients, the sender's station deposits the message to the message stores associated with each of the recipients. Alternatively, if the recipients do not have locations dedicated to storage of their messages, the message is deposited in message stores in the WAN selected by the sender and accessible to all of the intended recipients through multiplexed access to the message stores, effectively consolidating the storage required. Consolidated storage is especially advantageous because it provides for administrative simplicity, more efficient use of storage resources and, as explained more fully below, open-ended sets of recipient to facilitate group communications and electronic publishing.

To signal the intended recipients, the recipients are notified of the message's availability and its address on the WAN. The signaling process begins with the determination of the addresses of the recipients. The sender may know the addresses of the intended recipients or may need to refer to network-based interactive directories. Once the recipients' addresses are determined, the sender at a sender station instructs the WAN to notify the recipients of the availability and address of the message via the signaling mechanisms of the WAN.

In addition, the sender may store access authorization information for the message with the message in the message store. The access authorization information may be the list of intended recipients, unlimited in the case of publically available information or specified by identifiable characteristics of a group, for example security clearance, organizational affiliation or subscribers to a service. The access authorization information for a message can be different than the persons signaled of the availability of message. For example, a group may be authorized to access a message, such as any employee of a company, but only certain members of that group being signaled of its availability, such as the members of a particular project.

The sender need not signal all of the recipients at the same time. For instance, if the sender were to think of an additional recipient a day after sending the message, resending it would require only signaling the additional recipient and, perhaps, notifying the message store to add the additional recipient to the list of persons authorized to access the message.

In many messaging applications, the recipients are not explicitly enumerated by the sender, but are determined instead through processes such as self-selection, third-party reference or guided searching. In such cases, the signaling process may be absent. For example, a recipient can select a message by browsing through accessible message stores, by receiving a reference to it from a friend, co-worker or from another message or by searching through messages on the basis of keywords or other identifier provided by the sender. In such situations the new recipient must be authorized to access the message. The access authorization information may be broad enough to already include the new recipient or a recipient may be allowed to extend the access authorization to include a new recipient.

The message stores have the following characteristics. The message stores are located in storage elements integrated into the WAN. This functional integration is independent from questions of ownership or physical location.

The message stores and the messages stored therein have a logical identity in the WAN, independent of their physical realization at any given moment. This

enables greater reliability and efficiency by, for example, distributing the storage of a given message across multiple message stores.

The message stores are uniformly accessible from throughout the WAN. This means that they offer the same functionality independent of their physical location or that of the person accessing the message. As different recipients access a given message store from arbitrary locations, the functionality available to each recipient from the message store varies only with the characteristics of the recipient station used by the recipient to access the message store and its connection to the WAN (e.g. bandwidth).

Access to message stores can be multiplexed, allowing concurrent access by multiple senders and recipients. Thus, the senders and recipients are unlikely to be blocked when sending or receiving messages as can be the case with storage elements having more limited connectivity to the WAN.

The message stores control access to the messages they contain in accordance with the access authorization information provided by the sender.

Message reception is provided for via "virtual mailboxes." A virtual mailbox resides in storage elements in the WAN, but does not store the content of the recipient's messages. Instead, a virtual mailbox serves as an aggregation and storage point for signaling information that notifies a recipient of the availability and address of a message on the WAN. The recipient's station makes use of this signaling information to alert the recipient of the existence and availability of the message and to retrieve the message under the recipient's direction.

Because a recipient's messages can reside in multiple message stores spread throughout the network, the essence of the signaling information contained in the virtual mailbox is the reference to messages available to the recipient. Because of uniform accessibility to the message stores, the recipient's station can make use of those references to retrieve the recipient's messages. Through the use of the information stored in the virtual mailbox, the recipient can access his messages residing in multiple message stores in the WAN as if all of his messages were stored in a single mailbox.

The messages referenced in the virtual mailbox may include any type of information, from any medium. For example, a virtual mailbox could hold references to a text message on one LAN server, an audio message on another LAN server, another audio message on a voice-mail adjunct to a PBX or an audio/video message in another LAN server.

Referring to Fig. 3, a flow diagram illustrating the process of messaging in accordance with an embodiment of the invention is shown. To send a message, the sender instructs his sender station 20 to initiate WAN messaging (step 100). LAN 21 connects sender station 20 to WAN 24 (step 101). The sender station 20 allocates storage from WAN 24 by selecting message stores, e.g. message stores 25, 26 and 27, from available message stores in WAN 24 (step 102). The sender station deposits the message into the selected message stores in the WAN and the message is stored in the selected message stores (step 104).

The sender identifies the intended recipients of the message (step 106). If the sender would like to provide access to additional persons of a specified class, the sender may provide message access authorization information such that additional persons falling within the specified class of authorized persons can access to the message (steps 108, 110).

In response to deposit of the message and identification of intended recipients by the sender, sender's station 20 is disconnected from WAN 24 and the WAN signals the intended recipient of the availability and address of the message in the message stores in the WAN by placing a reference to the message in the recipient's virtual mailbox 28 (step 112).

To receive the message, the recipient instructs the recipient station 23 to access his virtual mailbox 28 (step 116). LAN 22 connects recipient station 23 to WAN 24 (step 117). The virtual mailbox 28 contains a reference notifying the recipient of the availability and address of the message (step 118). The recipient instructs the recipient station 23 to access the message (step 120). Using the message address stored in virtual mailbox, the recipient station accesses and provides the recipient with interactive access to the message stored in message stores in the WAN (step 121). After access to the message is complete the recipient station is disconnected from the WAN 24 (step 122).

If the sender has made the message accessible to additional persons of a specified class, the WAN also makes the message available to persons satisfying the access authorization information (steps 123, 124). Once such a person identifies the address of the message through, for example, the means referred above (step 125) using the address information, he can instruct his user station to access the message (step 126).

When such a person attempts to access the message, the message store 26 verifies that the person satisfies the access authorization information (step 128). If the access authorization information is satisfied, interactive access to the message is provided (step 130). If the access authorization information is not satisfied, access to the message is denied (step 132). Thereafter, the recipient's station is disconnected from the WAN 24 (step 133).

After the message has been accessed by the recipient and/or additional persons of the specified class, the message remains stored in the message stores in the WAN for access by additional intended recipients and other persons satisfying the access authorization information (step 134).

It is to be understood that the above description is only of one preferred embodiment of the invention. Numerous other arrangements may be devised by one skilled in the art without departing from the scope of the

invention. The invention is thus limited only as defined in the accompanying claims.

Claims

1. A method for electronic messaging using wide area network (WAN) signalling and storage capabilities, comprising:

   allocating storage in a WAN having an address in the WAN;
   depositing a message into the allocated storage;
   identifying persons authorized to have access to the message; and
   instructing the WAN to signal persons authorized to have access to the message of an available message and the address of the message in the allocated storage in the WAN.

2. The method of claim 1 wherein allocating storage in a WAN comprises selecting a message store from available message stores in the WAN.

3. The method of claim 1 wherein identifying persons authorized to have access to the message comprises identifying at least one intended recipient of the message.

4. The method of claim 1 wherein identifying persons authorized to have access to the message comprises identifying characteristics of a class authorized to have access to the message.

5. An apparatus for electronic messaging using wide area network (WAN) signalling and storage capabilities, comprising:

   a means for allocating storage in a WAN having an address in the WAN;
   a means for depositing a message into the allocated storage;
   a means for identifying persons authorized to have access to the message; and
   a means for instructing the WAN to signal persons authorized to have access to the message of an available message and the address of the message in the allocated storage in the WAN.

6. The apparatus of claim 5 wherein the means for allocating storage in a WAN comprises a means for selecting a message store from available message stores in the WAN.

7. The apparatus of claim 5 wherein the means for identifying persons authorized to have access to the message comprises a means for identifying at least one intended recipient of the message.

8. The apparatus of claim 5 wherein the means for identifying persons authorized to have access to the message comprises a means for identifying characteristics of a class authorized to have access to the message.

9. A method for electronic messaging using wide area network (WAN) signalling and storage capabilities, comprising:

   identifying an available message and an address of the message in allocated storage in a WAN; and
   using the address of the message in allocated storage in the WAN to interactively access the message in allocated storage in the WAN.

10. The method of claim 9 wherein identifying an available message and an address of the message in allocated storage in a WAN comprises accessing a virtual mailbox comprised of storage elements in the WAN for storing a notification of the availability of the message and the address of the message in allocated storage in the WAN.

11. The method of claim 9 wherein the allocated storage on the WAN comprises a message store.

12. An apparatus for electronic messaging using wide area network (WAN) signalling and storage capabilities, comprising:

   a means for identifying an available message and an address of the message in allocated storage a WAN; and
   a means for using the address of the message in allocated storage in the WAN to interactively access the message in allocated storage in the WAN.

13. The apparatus of claim 12 wherein the means for identifying an available message and an address of the message in allocated storage in a WAN comprises a means for accessing a virtual mailbox comprised of storage elements in the WAN for storing a notification of the availability of the message and the address of the message in allocated storage in the WAN.

14. The apparatus of claim 12 wherein the allocated storage on the WAN comprises a message store.

15. A method for electronic messaging using wide area network (WAN) signalling and storage capabilities, comprising:

providing storage in a WAN for receiving a message and for receiving an identification of persons authorized to have access to the message;
signalling the persons authorized to have access to the message of an available message and an address of the message in the storage in the WAN; and
providing interactive access to the message in the storage in the WAN through the use of the address of the message.

16. The method of claim 15 wherein the storage in a WAN network for receiving a message and for receiving an identification of the persons authorized to have access to the message comprises a message store.

17. The method of claim 15 wherein the identification of the persons having access to the message comprises the identity of at least one intended recipient of the message.

18. The method of claim 15 wherein the identification of the persons having access to the message comprises identifying characteristics of a class authorized to have access to the message.

19. The method of claim 15 wherein signalling the intended recipients of an available message and an address of the message comprises storing in a virtual mailbox comprised of storage elements in the WAN a notification of the availability and address of the message.

20. An apparatus for electronic messaging using wide area network (WAN) signalling and storage capabilities, comprising:

a means for providing storage in a WAN for receiving a message and for receiving an identification of persons authorized to have access to the message;
a means for signalling the persons authorized to have access to the message of an available message and an address of the message in the storage in the WAN; and
a means for providing interactive access to the message in the storage in the WAN through the use of the address of the message.

21. The apparatus of claim 20 wherein the means for providing storage in a WAN for receiving a message and for receiving an identification of persons authorized to have access to the message comprises a message store.

22. The apparatus of claim 20 wherein the means for receiving an identification of the persons having access to the message comprises a means for identifying of at least one intended recipient of the message.

23. The apparatus of claim 20 wherein the means for identification of the persons having access to the message comprises a means for identifying characteristics of a class authorized to have access to the message.

24. The apparatus of claim 15 wherein the means for signalling the intended recipients of an available message and an address of the message comprises a means for storing in a virtual mailbox comprised of storage elements in the WAN a notification of the availability and address of the message.

25. A method of electronic messaging using wide area network (WAN) storage and signalling capabilities, comprising instructing a WAN to signal a virtual mailbox comprised of storage elements in the WAN.

26. The method of claim 25 wherein instructing the WAN to signal a virtual mailbox comprised of storage elements in the WAN comprises instructing the WAN to signal the virtual mailbox of an available message and an address of the message.

27. An apparatus or electronic messaging using wide area (WAN) network storage and signalling capabilities, comprising a means for instructing a WAN to signal a virtual mailbox comprised of storage elements in the WAN.

28. The apparatus of claim 27 wherein the means for instructing a WAN to signal a virtual mailbox comprised of storage elements in the WAN comprises a means for instructing the WAN to signal the virtual mailbox of an available message and an address of the message.

29. A method of electronic messaging using wide area network (WAN) storage and signalling capabilities, comprising signalling a virtual mailbox comprised of storage elements in a WAN.

30. The method of claim 29 wherein signalling a virtual mailbox comprised of storage elements in a WAN comprises signalling the virtual mailbox of an available message and an address of the message.

31. An apparatus for electronic messaging using wide area network (WAN) storage and signalling capabilities, comprising a means for signalling a virtual mailbox comprised of storage elements in a WAN.

32. The apparatus of claim 31 wherein the means for

signalling a virtual mailbox comprised of storage elements in a WAN comprises a means for signalling the virtual mailbox of an available message and an address of the message.

33. A method for electronic messaging using wide area (WAN) network storage and signalling capabilities, comprising retrieving information stored in a virtual mailbox comprised of storage elements in a WAN.

34. The method of claim 33 wherein retrieving information stored in a virtual mailbox comprised of storage elements in a WAN comprises retrieving information regarding an available message and an address of the message stored in the WAN.

35. An apparatus for electronic messaging using wide area network (WAN) storage and signalling capabilities, comprising a means for retrieving information stored in a virtual mailbox comprised of storage elements in a WAN.

36. The apparatus of claim 35 wherein the means for retrieving information stored in a virtual mailbox comprised of storage elements in a WAN comprises a means for retrieving information regarding an available message and an address of the message stored in the WAN.

37. A virtual mailbox comprising storage elements in a wide area network (WAN).

38. The virtual mailbox of claim 37 wherein the signalling information stored comprises information regarding an available message and an address of the message stored in the WAN.

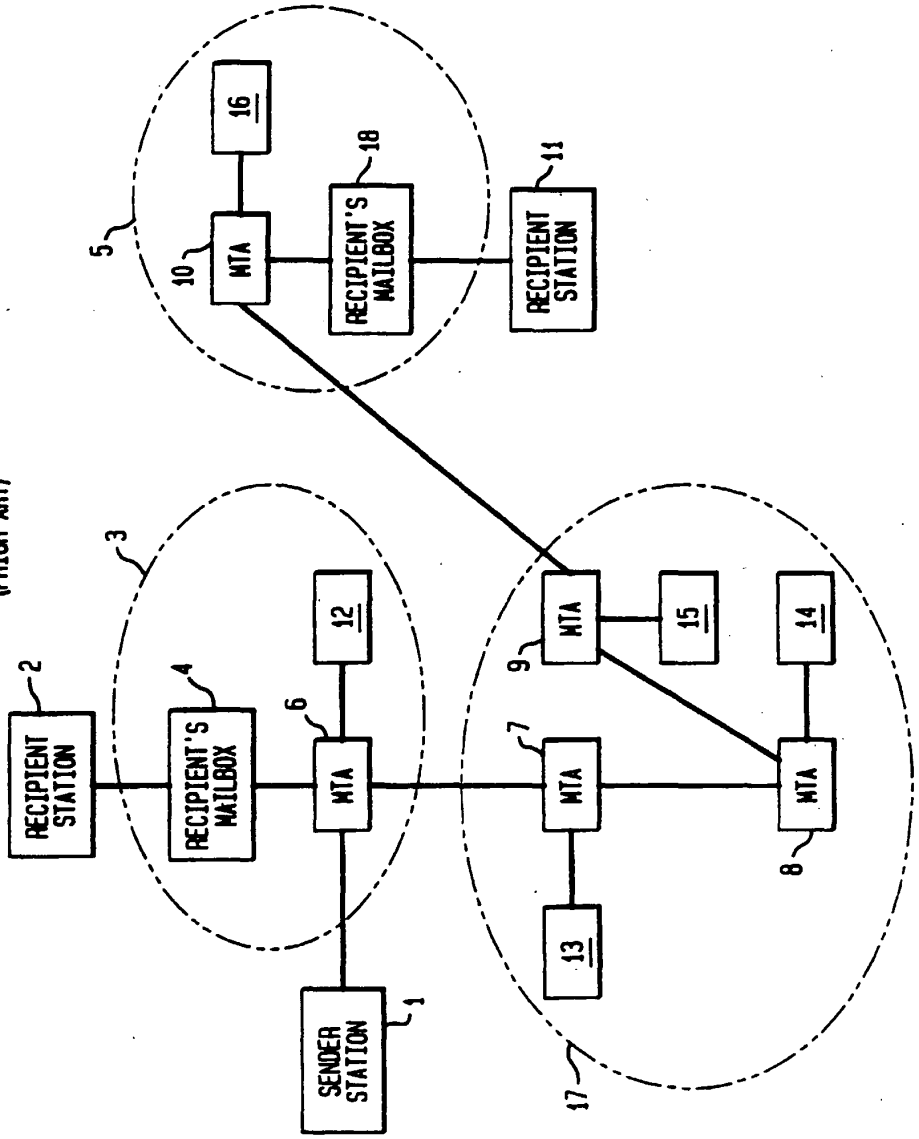


FIG. 1
(PRIOR ART)

FIG. 2

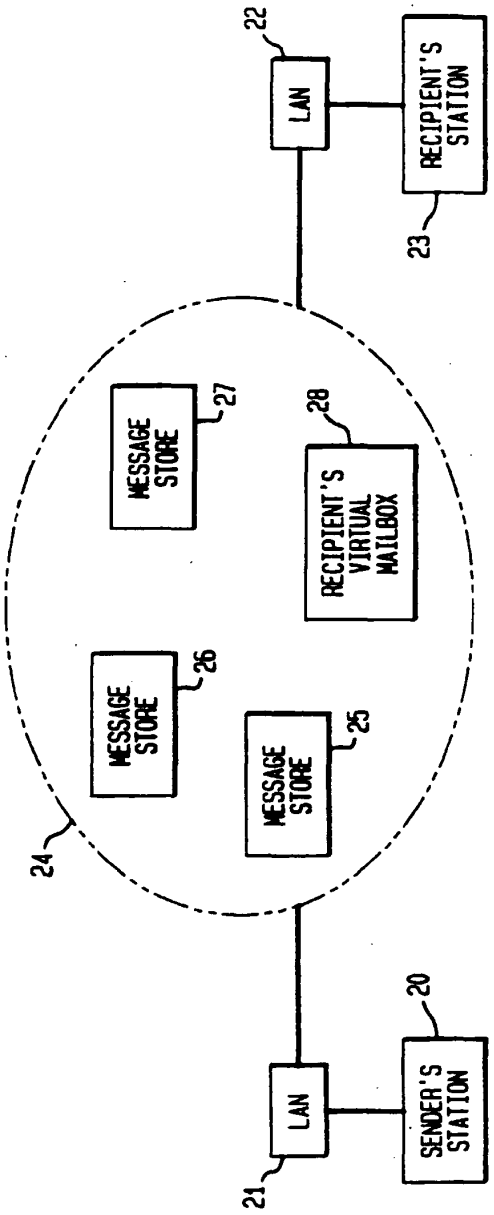

22
LAN
RECIPIENT'S STATION
23
MESSAGE STORE
27
28
RECIPIENT'S VIRTUAL MAILBOX
MESSAGE STORE
26
MESSAGE STORE
25
24
20
LAN
SENDER'S STATION
21

FIG. 3A

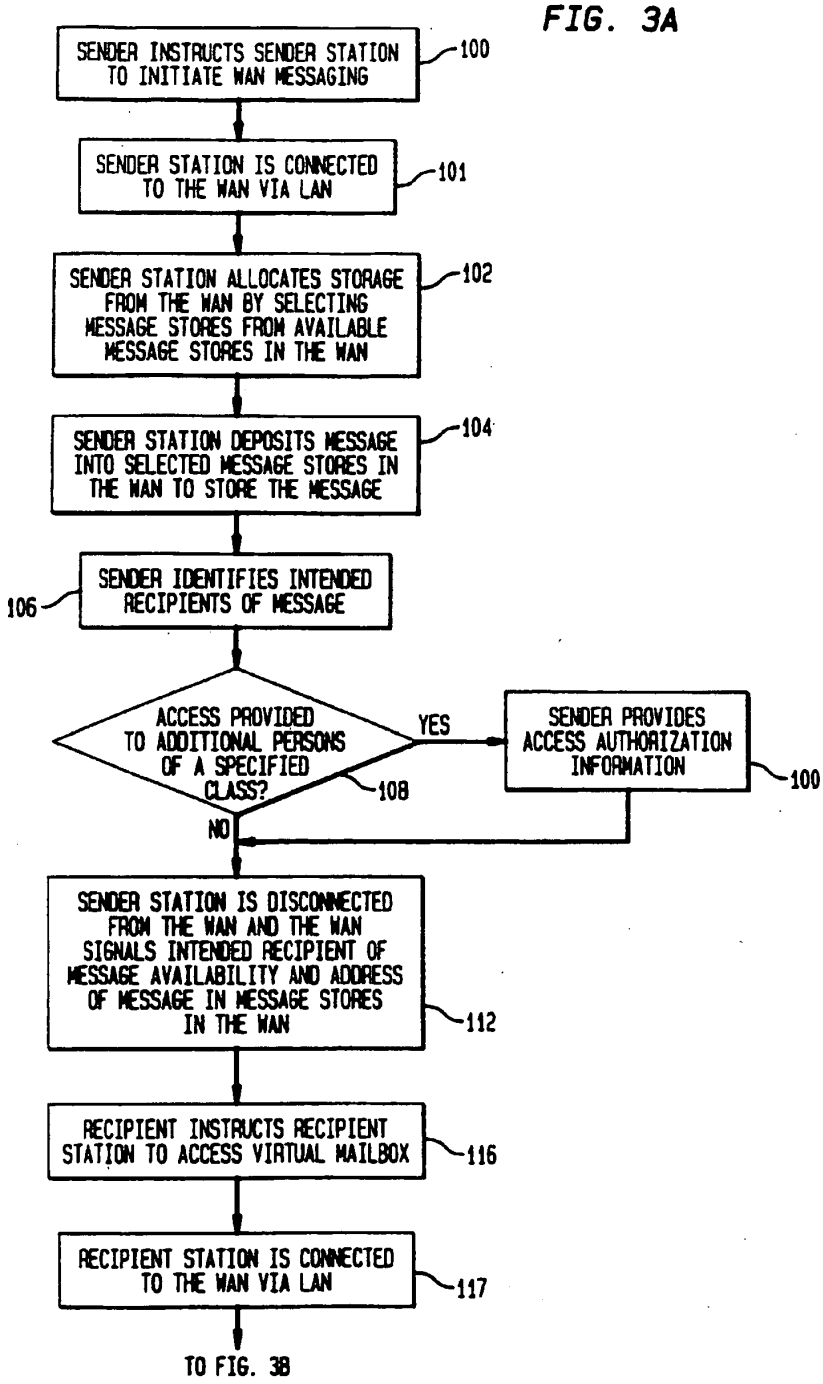
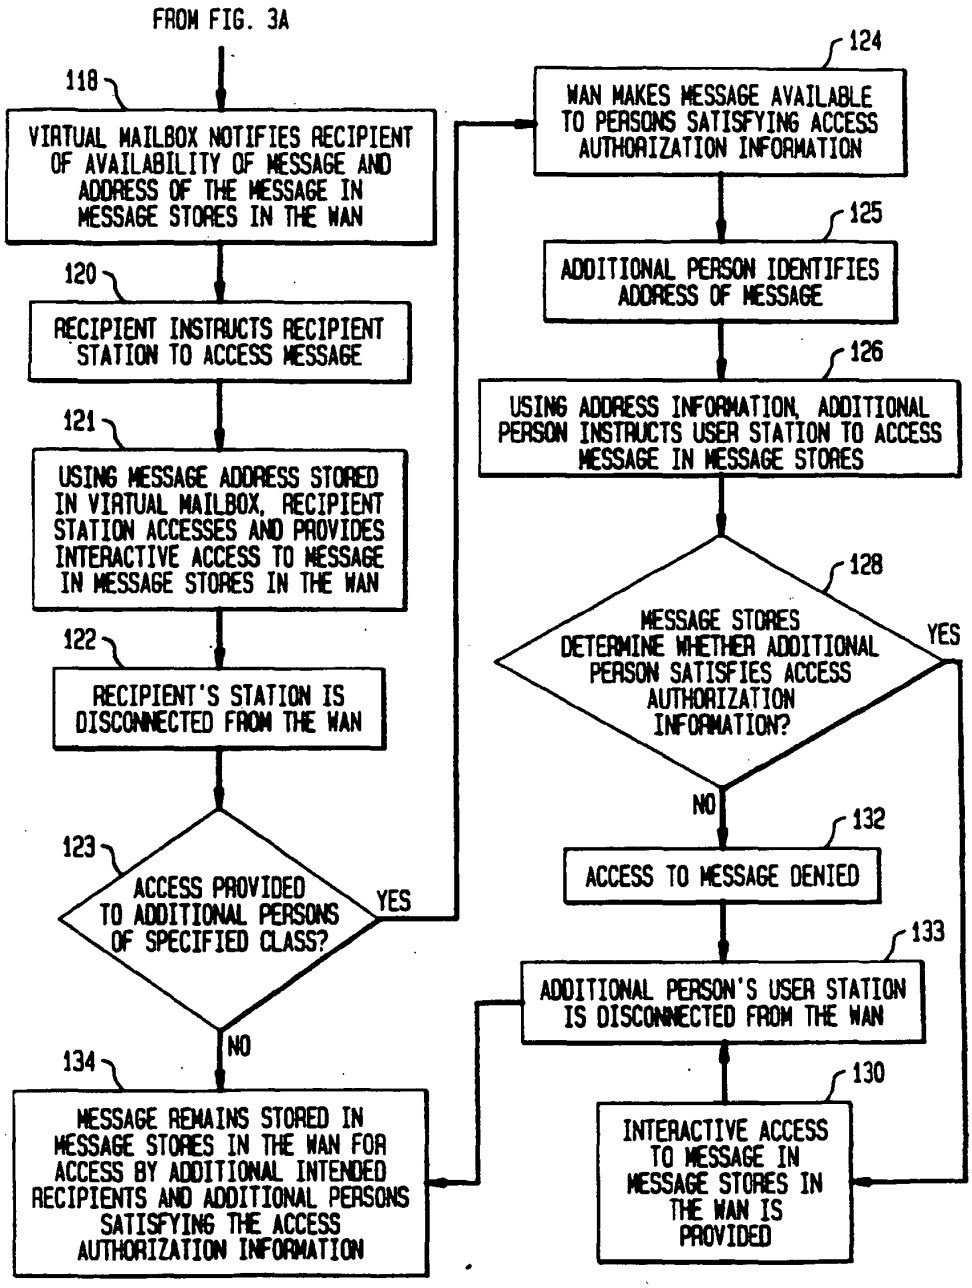

SENDER INSTRUCTS SENDER STATION TO INITIATE WAN MESSAGING
100
SENDER STATION IS CONNECTED TO THE WAN VIA LAN
101
SENDER STATION ALLOCATES STORAGE FROM THE WAN BY SELECTING MESSAGE STORES FROM AVAILABLE MESSAGE STORES IN THE WAN
102
SENDER STATION DEPOSITS MESSAGE INTO SELECTED MESSAGE STORES IN THE WAN TO STORE THE MESSAGE
104
SENDER IDENTIFIES INTENDED RECIPIENTS OF MESSAGE
106
ACCESS PROVIDED TO ADDITIONAL PERSONS OF A SPECIFIED CLASS?
108
YES
SENDER PROVIDES ACCESS AUTHORIZATION INFORMATION
100
NO
SENDER STATION IS DISCONNECTED FROM THE WAN AND THE WAN SIGNALS INTENDED RECIPIENT OF MESSAGE AVAILABILITY AND ADDRESS OF MESSAGE IN MESSAGE STORES IN THE WAN
112
RECIPIENT INSTRUCTS RECIPIENT STATION TO ACCESS VIRTUAL MAILBOX
116
RECIPIENT STATION IS CONNECTED TO THE WAN VIA LAN
117
TO FIG. 3B

FIG. 3B


FROM FIG. 3A
118
VIRTUAL MAILBOX NOTIFIES RECIPIENT OF AVAILABILITY OF MESSAGE AND ADDRESS OF THE MESSAGE IN MESSAGE STORES IN THE WAN
120
RECIPIENT INSTRUCTS RECIPIENT STATION TO ACCESS MESSAGE
121
USING MESSAGE ADDRESS STORED IN VIRTUAL MAILBOX, RECIPIENT STATION ACCESSES AND PROVIDES INTERACTIVE ACCESS TO MESSAGE IN MESSAGE STORES IN THE WAN
122
RECIPIENT'S STATION IS DISCONNECTED FROM THE WAN
123
ACCESS PROVIDED TO ADDITIONAL PERSONS OF SPECIFIED CLASS?
YES
NO
134
MESSAGE REMAINS STORED IN MESSAGE STORES IN THE WAN FOR ACCESS BY ADDITIONAL INTENDED RECIPIENTS AND ADDITIONAL PERSONS SATISFYING THE ACCESS AUTHORIZATION INFORMATION
124
WAN MAKES MESSAGE AVAILABLE TO PERSONS SATISFYING ACCESS AUTHORIZATION INFORMATION
125
ADDITIONAL PERSON IDENTIFIES ADDRESS OF MESSAGE
126
USING ADDRESS INFORMATION, ADDITIONAL PERSON INSTRUCTS USER STATION TO ACCESS MESSAGE IN MESSAGE STORES
128
MESSAGE STORES DETERMINE WHETHER ADDITIONAL PERSON SATISFIES ACCESS AUTHORIZATION INFORMATION?
YES
NO
132
ACCESS TO MESSAGE DENIED
133
ADDITIONAL PERSON'S USER STATION IS DISCONNECTED FROM THE WAN
130
INTERACTIVE ACCESS TO MESSAGE IN MESSAGE STORES IN THE WAN IS PROVIDED

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
1 May 2003 (01.05.2003)

PCT

(10) International Publication Number
WO 03/036492 A1

(51) International Patent Classification[7]: G06F 15/16

(21) International Application Number: PCT/US02/29689

(22) International Filing Date:
18 September 2002 (18.09.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/323,019 18 September 2001 (18.09.2001) US

(71) Applicant *(for all designated States except US)*: IDETIC, INC. [US/US]; 2855 Telegraph Avenue, Suite 510, Berkeley, CA 94705 (US).

(72) Inventor; and

(75) Inventor/Applicant *(for US only)*: STIERS, Todd [US/US]; 3208 King Street, Berkeley, CA 94703 (US).

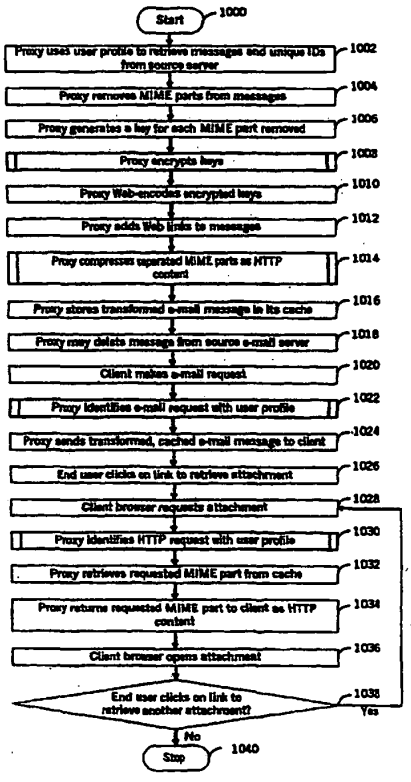
(74) Agent: BERRIER, Mark, L.; Gray Cary Ware & Freidenrich LLP, 1221 S. MoPac Expressway, Suite 400, Austin, TX 78746-6875 (US).

(81) Designated States *(national)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: CLIENTLESS ELECTRONIC MAIL MIME ATTACHMENT RE-DELIVERY SYSTEM VIA THE WEB TO REDUCE NETWORK BANDWIDTH USAGE



(57) Abstract: A method to reduce the network capacity usage of electronic email containing MIME-encoded attachments. A proxy server (100) located between the email client (202) and the source email server separates the MIME parts of the message, removes (1004) one or more MIME attachments, then insert links corresponding to the one or more MIME attachments into the email message. The proxy server (100) transmits the email message to the client using a 7-bit text format. The end user may click on a link corresponding to a MIME attachment. The attachment is transmitted to the client using a non-7-bit format. The proxy server (199) can act as a protocol proxy, retrieving and transmitting messages in real time, or it can retrieve, transform, and cache messages before clients request them.

BEST AVAILABLE COPY

WO 03/036492 A1

BNSDOCID: <WO_____03036492A1_I_>



Declaration under Rule 4.17:

— *as to the identity of the inventor (Rule 4.17(i)) for the following designations AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW, ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)*

Published:

— *with international search report*

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.

DESCRIPTION

CLIENTLESS ELECTRONIC MAIL MIME ATTACHMENT RE-DELIVERY SYSTEM VIA THE WEB TO REDUCE NETWORK BANDWIDTH USAGE

BACKGROUND OF INVENTION

Technical Field

This invention generally relates to processing of electronic mail messages with Multipurpose Internet Mail Extensions (MIME), and more particularly to removing MIME attachments from electronic mail messages and making the attachments accessible through Web links.

Related Art

Internet electronic mail (email) is non-interactive, and most email servers and clients are designed for low-cost, high-bandwidth Internet services. These characteristics mean that email uses resources inefficiently in low-performance, high-cost connection environments, such as wireless networks.

Email gateways require email to be encoded in 7-bit text formats. The first email messages were ASCII text, which required only 7 bits to express the full character set on early mainframe terminals.

With the advent of email attachments and the MIME standard (RFC 1049 and RFC-1341), encoding techniques are required to convert 8-bit multi-media formats into 7-bit formats to guarantee backward compatibility with older email gateways and reader clients.

The base 64 encoding scheme is one popular encoding technique for converting 8-bit content to 7-bit clean MIME attachments. This process adds significant overhead, as every 8-bit byte of content is converted into a 7-bit representation plus remainder. Using only 7 bits out of a potential 8 bits per byte increases the overall size of the message.

Unlike email, the more recent HTTP protocol is capable at handling 8-bit, multi-media content. In addition to supporting content encodings like email's base 64 type, compressed encodings, namely gzip and compress, are available in most modern Web-browsing clients.

Email servers and clients treat MIME-compliant email messages as single, large objects. MIME messages include boundary definitions of the multiple parts of the message, and standard tools (e.g., Perl or Java objects) can parse the messages.

BRIEF SUMMARY OF THE INVENTION

This invention comprises systems and methods for removing electronic mail MIME attachments from an email message and replacing them with Web (HTML) links, thereby reducing the network capacity usage of email containing MIME-encoded attachments. The end user receives the non-attachment portions of an email message with a standard email client and can select the Web links to download the attachments via the Web.

By removing the attachments from email messages, this invention may reduce the bandwidth use between the email gateway and the user in various ways. For example, it may enable faster download of the non-attachment portions of a message. User preferences can determine which portions of a message the proxy server treats as attachments. It may also enable faster download of attached portions of the email message, accomplished both by avoiding base 64 encoding and by applying Web-compliant compression. Further, it may eliminate the need to transmit some data. Because users can control when or if they wish to download attachments, not all attachments may need to be transmitted.

This invention can leverage common clients, standards and protocols for email and the Web, using server-side intelligence to deliver content to the user over poor-performance network environments. It is particularly applicable to the wireless environment where network capacity is expensive and speeds are slow.

One embodiment of the invention comprises a method for increasing the speed and minimizing the bandwidth required for delivering email messages containing MIME attachments to a client. The method of this embodiment comprises receiving an email message having one or more MIME attachments, removing the MIME attachments from the email message, inserting one or more links corresponding to the one or more MIME attachments into the email message, and transmitting the email message containing the one or more links to the client using a 7-bit text format. The client may then select the links to access the attachments, which are delivered to the client using an 8-bit format.

Another embodiment of the invention comprises a system in which a proxy server is located between an email server and a client. The proxy server is configured to receive an email message having one or more MIME attachments from the email server, and to remove the one or more attachments from the email message. The proxy server inserts links corresponding to the one or more attachments into the email message and then transmits the modified email message to the client in a 7-

bit format. In response to the client selecting a link, the corresponding attachment is delivered to the client in an 8-bit format.

Another embodiment of the invention comprises software configured to implement the methods described herein. ("Software" is used herein to refer to software, firmware, and other manners of programming instructions for a computer or other data processor.) The software may be embodied in any medium readable by such a computer or data processor, including floppy disks, magnetic tapes, CD-ROMs, DVD-ROMs, RAM, ROM, and the like. Likewise, a computer or data processor which is configured to execute such software applications, or which is otherwise programmed to perform the methods disclosed herein is intended to be covered by the present application.

Various other embodiments are also possible.

BRIEF DESCRIPTION OF THE DRAWINGS

Other objects and advantages of the invention may become apparent upon reading the following detailed description and upon reference to the accompanying drawings.

FIG. 1 shows the system architecture for an embodiment in which a proxy server collects and caches email messages independently of the client's requests. The proxy server transforms and delivers attachments.

FIG. 2 is a flowchart outlining in detail the steps that the proxy server and client execute in FIG. 1.

FIG. 3 shows the system architecture for an embodiment in which a proxy server scans and reformats email nearly in real time. The proxy server transforms and delivers attachments.

FIG 4. is a flowchart outlining in detail the steps that the proxy server and client execute in FIG. 3.

FIG. 5 shows the system architecture for an embodiment in which an email proxy server scans and reformats email nearly in real time. A separate, stand-alone Web server transforms and delivers attachments.

FIG. 6 is a flowchart outlining in detail the steps that the email proxy server, client, and separate Web server take in FIG. 5.

While the invention is subject to various modifications and alternative forms, specific

embodiments thereof are shown by way of example in the drawings and the accompanying detailed description. It should be understood, however, that the drawings and detailed description are not intended to limit the invention to the particular embodiment which is described. This disclosure is instead intended to cover all modifications, equivalents and alternatives falling within the scope of the present invention as defined by the appended claims.

DETAILED DESCRIPTION OF THE INVENTION

Preferred embodiments of the invention are described below. It should be noted that these and any other embodiments described below are exemplary and are intended to be illustrative of the invention rather than limiting.

This invention comprises systems and methods for removing electronic mail MIME attachments from an email message and replacing the attachments with links to the attachments. The end user receives the non-attachment portions of an email message with a standard email client and can select the links to download the attachments (for example, using an ordinary browser to access the attachments via the Web). The non-attachment portions of the message can be downloaded faster because they do not include the attachments, and the attachments can be downloaded faster because they are converted from the 7-bit MIME format to their original 8-bit format.

This functionality may be implemented in a proxy server positioned between the source email server and the email client. The proxy server is configured to remove the MIME attachments from email messages, decode the attachments, and insert links to the attachments in the email messages.

When the proxy server removes attachments from email messages, the attachments are decoded from the standard base 64, 7-bit clean encoding to the original, 8-bit format. This significantly reduces the attachment size, since bytes encoding only 7-bits each are transcoded into fewer bytes encoding 8 bits each.

The Web links that the proxy server adds to the message make the attachments accessible via a Web server and client, and the Web server's built-in decompression algorithms can compress the attachments to sizes well below the original, non-compressed size. The Web server that transforms and delivers the attachments can be either the email proxy server itself or a separate, stand-alone Web server. The proxy server uses the HTTP protocol to send MIME attachments to client devices, so it can apply all Web-compliant transformations to them. The advantage of using a separate Web server is that an end user can retrieve attachments through a non-proxied channel.

The proxy server combines, and potentially encrypts, several pieces of information to generate a passkey which can be used to uniquely identify and retrieve attachments. The information combined

is: the end user's account name, a unique identifier for the message, and an indicator of which attachment within the message. When the proxy server is configured to allow users to access multiple email servers at the same time, the key also must incorporate the source email server information. To generate these keys, the proxy server must request a unique identifier for each message from the source email server.

In a deployment that uses a separate Web server for attachment traffic, to avoid the need to look up the user's account, mailserver and password when retrieving attachments, the links themselves must include that information. Alternatively, the separate Web server must store that information, potentially using some sort of user profile storage mechanism.

To ensure that the intended recipient is the only one who can retrieve an attachment, the proxy server should encrypt the key that identifies the attachment. To create standards compliant Web links which can be used to retrieve attachments, the proxy server should further Web-encode the keys after they have been encrypted. The proxy server then replaces the attachments within a message with web links, also known as URLs (uniform resource locators), which include the Web-encoded, encrypted keys

In one embodiment of the invention, the end user's email messages and attachments are pre-fetched and cached by the proxy server in a dedicated email spool and Web attachment area, ready for the user to retrieve. The proxy server loads the spools offline periodically at some configurable interval. The advantage of this embodiment is that removing attachments, generating Web links, and compressing attachments can be performed offline or during idle times. This embodiment may, however, use large amounts of disk space on the proxy server and may raise privacy concerns relating to email content re-distribution.

FIG. 1 illustrates a system architecture in accordance with this embodiment, and FIG. 2 is a flowchart showing the steps that the proxy server and client device may execute in this embodiment. Reference numbers for both the architecture diagram of FIG. 1 and the flowchart of FIG. 2 are incorporated into the discussion below. The reference numbers for FIG. 2 are enclosed in parentheses to make them more easily distinguishable from the reference numbers for FIG. 1.

The proxy server 100 collects and stores email messages independently of the client's requests. User profiles 106 stored on the proxy server 100 provide information on end users' source email servers 400, accounts and passwords. The method that the proxy server 100 uses to store and look up user profiles 106 depends on the deployment. The proxy server 100 stores end users' email messages in its cache 104.

The proxy server 100 uses the client's user profile information 106 to issue a retrieve request over the Internet 300 to the end user's source email server 400. When it issues the retrieve request, the proxy server 100 also issues a request for the unique identifier for any message it retrieves. The proxy server 100 retrieves full email messages (and their unique identifiers) from the source email server 400 (1002), leaving the original messages on the source email server 400. The proxy server 100 can retrieve email from multiple accounts for the same user.

The proxy server 100 removes the MIME attachments from the email message (1004). For each attachment removed, the proxy server 100 generates a key (1006). The key combines the following information: the user-identifier, the message-identifier, and the attachment number. When the proxy server 100 is configured to allow users to access multiple email servers at the same time, the key should also include the source email server 400 information.

If the security of the attachments is a concern, the proxy server 100 should encrypt the keys to the attachments (1008). The proxy server 100 can Web-encode the keys (1010), then add the Web-encoded keys to the message as Web link URLs (1012), which connect to the attachments removed from the message. The proxy server 100 may compress separated parts of the message as necessary (1014).

The proxy server 100 stores the transformed email message in its cache 104 (1016), then may then delete the email message from the source email server 400 (1018).

The end user's email application 202 makes an email request (1020). The proxy server's email port 102 receives the request. The proxy server 100 identifies the request with a user profile 106 (1022), retrieves the end user's email from its cache 104, then sends the transformed, cached email message to the client 200 (1024).

When the end user clicks on a link to retrieve an attachment (1026), a browser 204 window opens automatically and issues an HTTP request for the attachment (1028). (The success of this step depends on the configuration of the client device.) The proxy server's HTTP port 108 receives the request. The proxy server 100 identifies the HTTP request with a user profile 106 (1030), retrieves a copy of the requested MIME part from its cache 104 (1032), then returns the requested MIME part to the client 200 as HTTP content (1034). The client device's Web browser 204 opens the attachment (1036). If the browser 204 does not recognize the content type of the attachment, it opens a dialog box allowing the user to save the file to disk. (The success of this step depends on the configuration of the client device.) The end user may retrieve another attachment by clicking on another link (1038).

A second embodiment of this invention is to use a protocol proxy (for example, POP or IMAP)

to make client requests to the source email server and return server responses to the client. The advantage of this system is that the proxy server scans and reformats email messages in near real time, using an off-site email repository (the source email server). The only information stored on (or network accessible to) the proxy is the user profile, which is only required in some embodiments. No part of an email message is stored on the proxy server, so the system does not require the implementation overhead or raise the privacy issues of queuing mail and attachments in third party spools. The main concern that this embodiment raises is email box synchronization, since, if an email message is deleted from the source email server, the Web links to the message's attachments will not work.

FIG. 3 illustrates the system architecture of this second embodiment, and FIG. 4 is a flowchart showing the steps that the proxy server and client device execute. Again, the reference numbers for the flow diagram of the method are enclosed in parentheses, while the reference numbers for the architecture diagram are not.

Before making a request to a source email server 400, the proxy server 100 must receive a request from a client 200. The proxy server 100 may alter this request before sending it over the Internet 300 to the source email server 400.

The end user's email application 202 makes an email request (1202). The proxy server's email port 102 receives the request, and the proxy server 100 identifies the request, potentially using a user profile 106 (1204). The proxy server 100 issues a retrieve request on behalf of the client 200 to the source email server 400. When it issues the retrieve request, the proxy server 100 also issues a request for the unique identifier for any message it retrieves. The proxy server 100 retrieves the end user's email messages (and their unique identifiers) from the source email server 400 (1206), including a full copy of any MIME messages, leaving the original messages on the email server 400.

The proxy server 100 removes attachments from the email message (1208). For each attachment removed, the proxy server 100 generates a key. (1210). The key combines the following information: the user-identifier, the message-identifier, and the attachment number. When the proxy server is configured to allow users to access multiple email servers at the same time, the key also must include the source email server information.

If the security of the attachments is a concern, the proxy server 100 should encrypt the keys to the attachments (1212). The proxy server 100 Web-encodes the encrypted keys (1214), then adds the Web-encoded keys to the message as Web link URLs (1216), which connect to the attachments removed from the message.

The proxy server 100 sends the transformed email message to the email client 202 (1218).

When the end user clicks on a link to retrieve an attachment (1220), a browser 204 window opens automatically and issues an HTTP request for the attachment (1222). (The success of this step depends on the configuration of the client device.) The proxy server's HTTP port 108 receives the request. The proxy server 100 identifies the HTTP request with a user profile 106 (1224), then retrieves a copy of the full email message from the source email server 400 (1226), leaving the original message on the email server 400. The proxy server 100 separates the MIME part requested (1228), compresses it as necessary (1230), then returns the requested MIME part the client 200 as HTTP content (1232). The client device's Web browser 204 opens the attachment (1234). If the browser 204 does not recognize the content type of the attachment, it opens a dialog box allowing the user to save the file to disk. (The success of this step depends on the configuration of the client device.)

The end user may retrieve another attachment by clicking on another link (1236).

The proxy server 100 may delete an email message from a source email server 400 (1238).

A third embodiment of this invention is to use a proxy server for email traffic only, and to use a separate, stand-alone Web server to handle end user requests for attachments and the source server's responses. The email proxy generates attachment links containing enough information to allow the separate Web server to identify and retrieve attachments. (In an alternative embodiment, the separate Web server could store the same user profiles as the email proxy server.) The email proxy server scans and reformats email messages in near real time, using an off-site email repository (the source email server). In addition to offering the advantages of the second embodiment, this third embodiment allows end users to access their attachments through a non-proxied channel. The concerns that this embodiment raises are the same as those of the second embodiment.

FIG. 5 illustrates the system architecture of this third embodiment, and FIG. 6 is a flowchart showing the steps that the email proxy server, client device, and the separate Web server execute.

Before making a request to a source email server 400, the email proxy server 100 must receive a request from a client 200. The email proxy server 100 may alter this request before sending it over the Internet 300 to the source email server 400.

The end user's email application 202 makes an email request (1402). The proxy server's email port 102 receives the request, and the proxy server 100 identifies the request with a user profile 106 (1404). The proxy server 100 issues a retrieve request on behalf of the client 200 to the source email server 400. When it issues the retrieve request, the proxy server 100 also issues a request for the

unique identifier for any message it retrieves. The proxy server 100 retrieves the end user's email messages (and their unique identifiers) from the source email server 400 (1406), including a full copy of any MIME messages, leaving the original messages on the email server 400.

The proxy server 100 removes parts of the email message (1408), For each attachment removed, the proxy server 100 generates a key. (1410). The key combines the following information: the user-identifier, the message-identifier, and the attachment number. When the proxy server 100 is configured to allow users to access multiple email servers at the same time, the key also must include the source email server 400 information. In cases where the attachment is to be retrieved from a stand-alone server 500 that does not have a user profile containing the user's mail server or mail account password, then that information also must be combined into the key. Any reversible method of combination can be used, for example, concatenation.

If the security of the attachments is a concern, the proxy server 100 should encrypt the keys to the attachments (1412), Web-encodes the encrypted keys (1414), then adds the Web-encoded keys to the message as Web link URLs (1416), which connect to the parts removed from the message.

The proxy server 100 sends the transformed email message to the email client 202 (1418).

When the end user clicks on a link to retrieve an attachment (1420), a browser 204 window opens automatically and issues an HTTP request for the attachment (1422). (The success of this step depends on the configuration of the client device.) The separate Web server's HTTP port 502 receives the request. The separate Web server 500 uses the key in the encoded Web link to request the attachment (1424), then retrieves a copy of the full email message from the source email server 400 (1426), leaving the original message on the email server 400. The separate Web server 500 separates the MIME part requested (1428), compresses it as necessary (1430), then returns the requested MIME part the client 200 as HTTP content (1432). The client device's Web browser 204 opens the attachment (1434). If the browser 204 does not recognize the content type of the attachment, it opens a dialog box allowing the user to save the file to disk. (The success of this step depends on the configuration of the client device.)

The end user may retrieve another attachment by clicking on another link (1436).

The proxy server 100 or the separate Web server 500 may delete an email message from a source email server 400 (1438).

The benefits and advantages which may be provided by the present invention have been described above with regard to specific embodiments. These benefits and advantages, and any elements or limitations that may cause them to occur or to become more pronounced are not to be

construed as critical, required, or essential features of any or all of the claims. As used herein, the terms 'comprises,' 'comprising,' or any other variations thereof, are intended to be interpreted as non-exclusively including the elements or limitations which follow those terms. Accordingly, a process, method, article, or apparatus that comprises a list of elements does not include only those elements but may include other elements not expressly listed or inherent to the claimed process, method, article, or apparatus.

While the present invention has been described with reference to particular embodiments, it should be understood that the embodiments are illustrative and that the scope of the invention is not limited to these embodiments. Many variations, modifications, additions and improvements to the embodiments described above are possible. It is contemplated that these variations, modifications, additions and improvements fall within the scope of the invention as detailed within the following claims.

CLAIMS

1. A method comprising:

receiving an email message having one or more MIME attachments;

removing the MIME attachments from the email message;

inserting one or more links corresponding to the one or more MIME attachments into the email message; and

transmitting the email message containing the one or more links to a client.

2. The method of claim 1, further comprising transmitting one of the one or more attachments to the client in response to the client selecting the corresponding one of the one or more links.

3. The method of claim 2, further comprising converting the one of the one or more attachments to a non-7-bit-text format prior to transmitting the one of the one or more attachments to the client.

4. The method of claim 2, further comprising storing the attachments in a cache in a proxy server and retrieving the attachments from the cache prior to transmitting the one or more attachments to the client.

5. The method of claim 2, further comprising, in response to the client selecting one of the one or more links, retrieving the email message from an email source, removing from the email message the one of the one or more attachments corresponding to the selected link, and converting the one of the one or more attachments corresponding to the selected link to a non-7-bit-text format prior to transmitting the one of the one or more attachments to the client.

6. The method of claim 5, wherein the email message containing the one or more links is transmitted to the client by a proxy server and the one of the one or more attachments is transmitted to the client by a Web server which is separate from the proxy server.

7. The method of claim 2, further comprising retrieving a user profile corresponding to the client and retrieving the one of the one or more attachments based upon the user profile.

8. The method of claim 1, The method of claim 1, wherein the email message containing the one or more links is transmitted to the client using a 7-bit text format.

9. The method of claim 1, further comprising generating a key for each of the attachments removed from the email message and incorporating each key in the link to the corresponding attachment.

10. The method of claim 9, wherein the key includes a user identifier, a message identifier, and an attachment identifier.

11. The method of claim 10, wherein the key also includes an email source identifier.

12. The method of claim 9, further comprising encrypting the keys.

13. The method of claim 1, wherein the links comprise Web URLs.

14. A system comprising:

a proxy server,

wherein the proxy server is configured to receive an email message from a mail server, wherein the email message has one or more MIME attachments, and

wherein the proxy server is configured to remove the one or more attachments from the email message and to insert one or more links corresponding to the one or more attachments into the email message, and to transmit the modified email message to a client.

15. The system of claim 14, wherein the proxy server is configured to provide one of the one or more attachments in response to the client selecting a corresponding one of the one or more links.

16. The system of claim 15, wherein, in response to the client selecting the one of the one or more links, the proxy server is configured to retrieve the original email message from a mail server, extract the corresponding one of the one or more attachments from the original email message, and transmit the extracted attachment to the client.

17. The system of claim 16, further comprising a Web server which is separate from the proxy server, wherein the proxy server is configured to transmit the email message containing the one or more links to the client and the Web server is configured to transmit the one of the one or more attachments to the client.

18. The system of claim 15, further comprising a cache, wherein the proxy server is configured to store the attachments in the cache and wherein, in response to the client selecting the one of the one or

more links, the proxy server is configured to retrieve the corresponding one of the one or more attachments from the cache, and transmit the retrieved attachment to the client.

19. The system of claim 15, wherein the proxy server is configured to transmit the extracted attachment is over a non-seven-bit protocol.

20. The system of claim 15, further comprising a memory configured to store one or more user profiles, wherein the proxy server is configured to retrieve a user profile corresponding to the client and to retrieve the one of the one or more attachments based upon the retrieved user profile.

21. The system of claim 14, wherein the proxy server is configured to transmit the email message containing the one or more links to the client in a 7-bit text format.

22. The system of claim 14, wherein the proxy server is configured to generate a key for each of the attachments removed from the email message and to incorporate each key in the link to the corresponding attachment.

23. The system of claim 22, wherein each key includes a user identifier, a message identifier, and an attachment identifier.

24. The system of claim 23, wherein each key further includes an email source identifier.

25. The system of claim 22, wherein the proxy server is configured to encrypt each key.

26. The system of claim 14, wherein the links comprise Web URLs.

FIG. 1

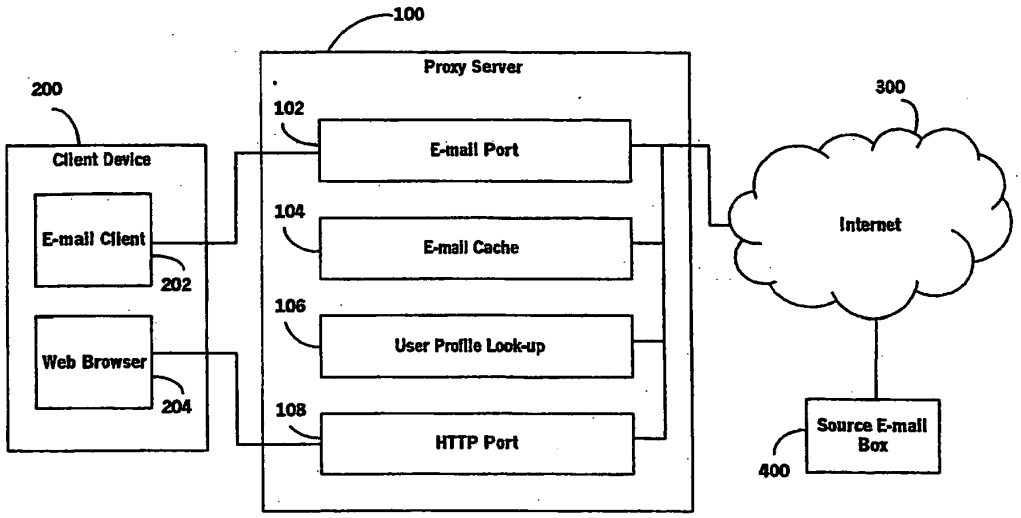
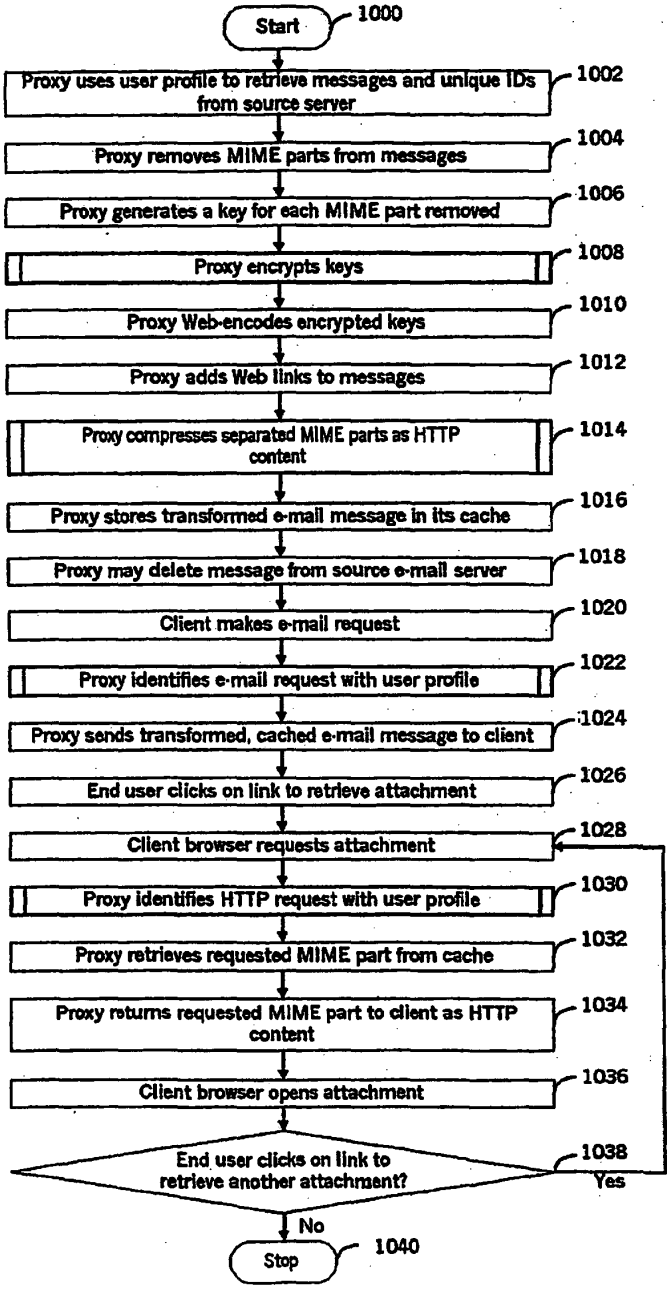

100
Proxy Server
200
Client Device
E-mail Client
202
Web Browser
204
102
E-mail Port
104
E-mail Cache
106
User Profile Look-up
108
HTTP Port
300
Internet
Source E-mail Box
400

FIG. 2



BNSDOCID: <WO_____03036492A1_I_>

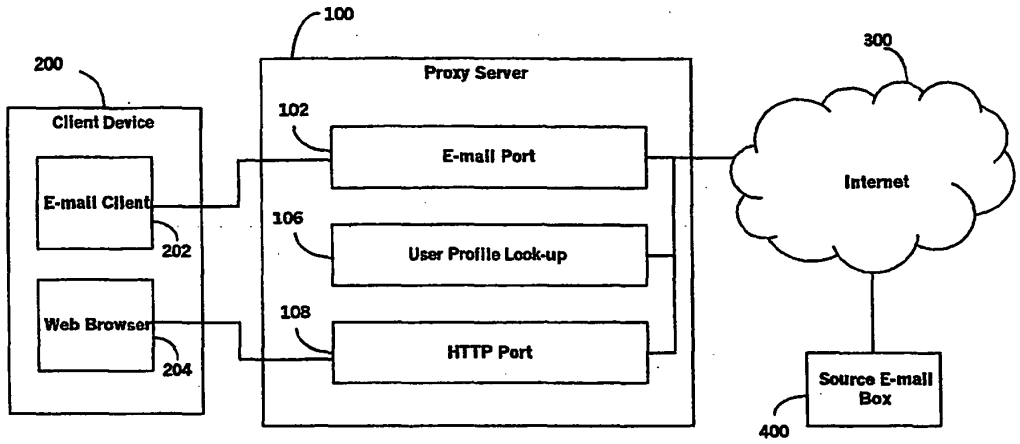
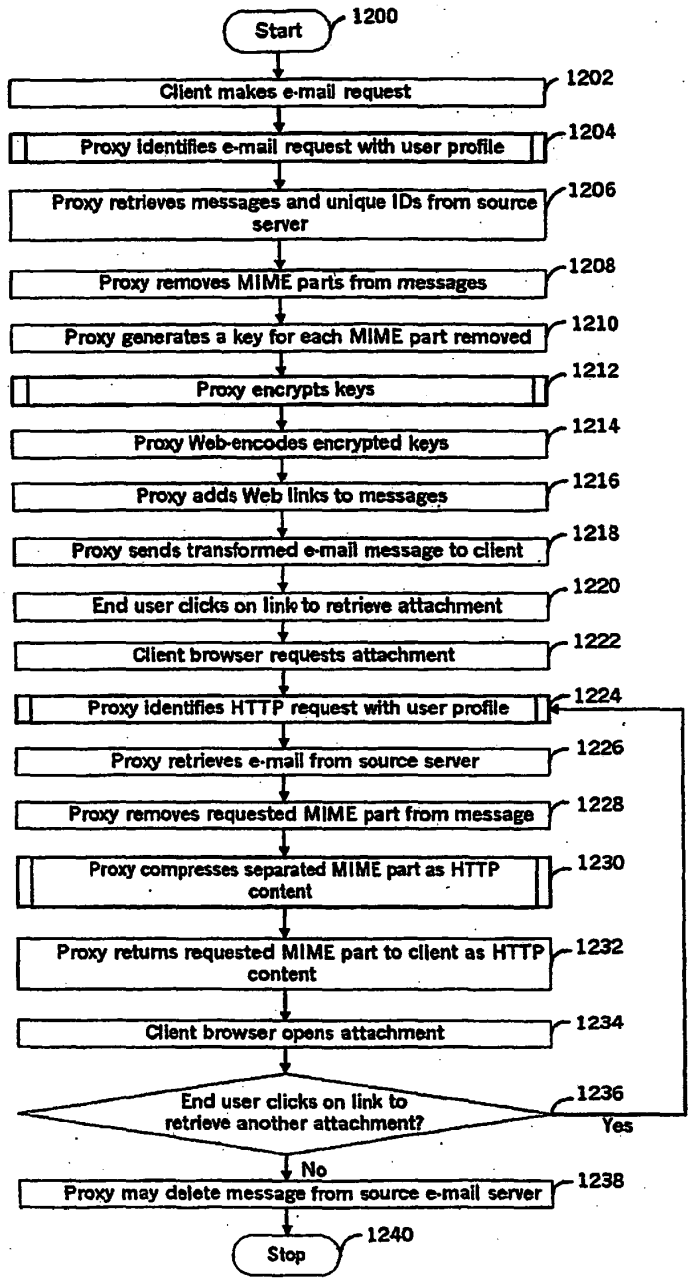

100
200
Client Device
E-mail Client
202
Web Browser
204
Proxy Server
102
E-mail Port
106
User Profile Look-up
108
HTTP Port
300
Internet
Source E-mail Box
400

FIG. 3

FIG. 4



BNSDOCID: <WO_____03036492A1_I_>

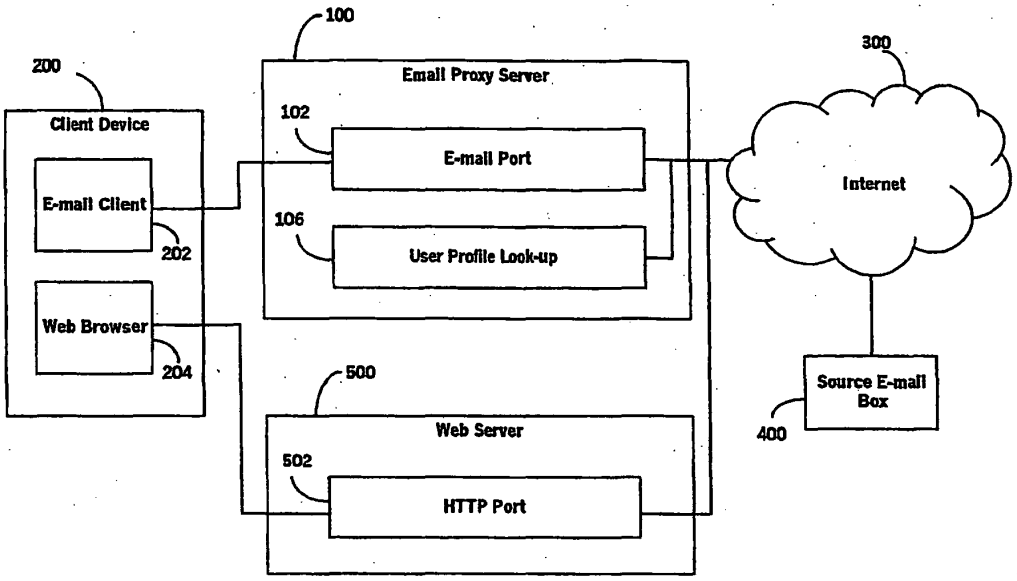
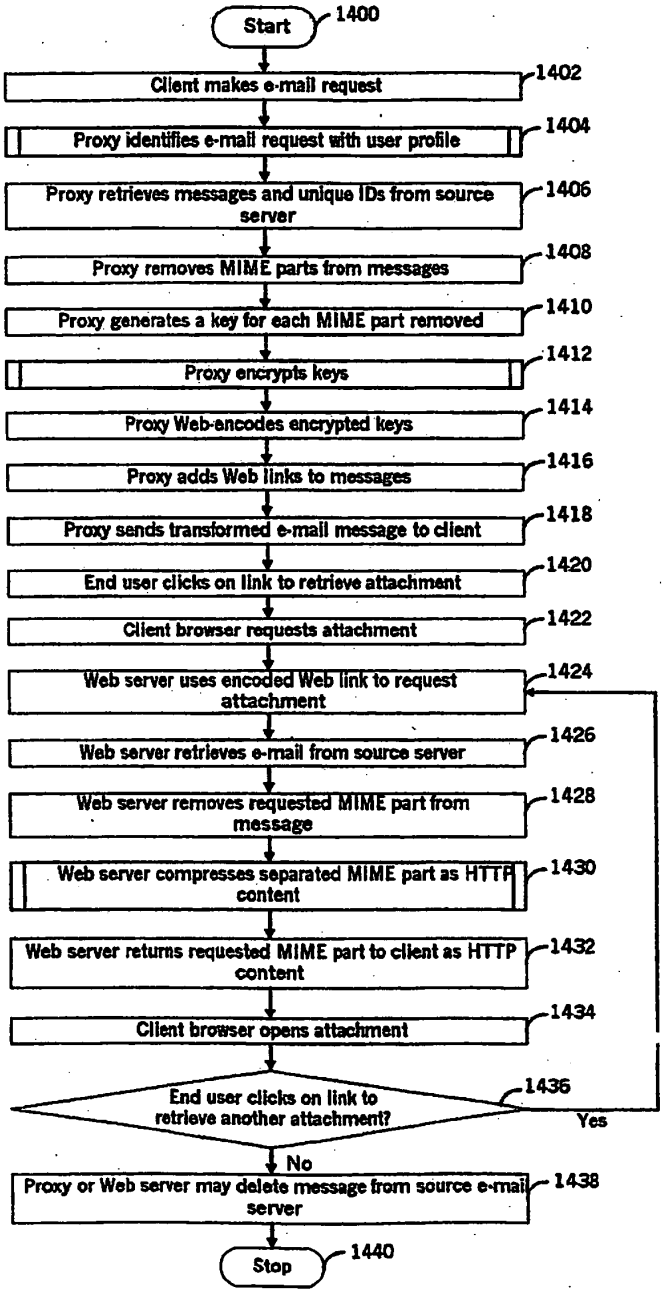

100
Email Proxy Server
200
Client Device
102
E-mail Port
E-mail Client
202
106
User Profile Look-up
Web Browser
204
300
Internet
500
Web Server
502
HTTP Port
Source E-mail Box
400

FIG. 5

FIG. 6



BNSDOCID: <WO_____03036492A1_I_>

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US02/29689

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F 15/16
US CL : 709/206, 207

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
U.S. : 709/206, 207

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 6,275,848 B1 (ARNOLD) 14 August 2001 (14.08.2001), see column 2, line 20 to column 6, line 25.	1-3,5,7-8,13-16,20-21,26
Y	US 6,256,666 B1 (SINGHAL) 03 July 2001 (03.07.2001), see column 3, line 45 to column 7, line 57.	1-3,5,7-8,13-16,20-21,26
A	US 6,275,850 B1 (BEYDA et al.) 14 August 2001 (14.08.2001), see the whole reference.	1-26
A	US 6,249,807 B1 (SHAW et al.) 19 June 2001 (19.06.2001), see the whole reference.	1-26
A	US 5,890,163 A (TODD) 30 March 1999 (30.03.1999), see the whole reference.	1-26

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
14 December 2002 (14.12.2002)	14 JAN 2003
Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703)305-3230	Authorized officer Meng Ai An Telephone No. 703-305-3900

Form PCT/ISA/210 (second sheet) (July 1998)

This Page is Inserted by IFW Indexing and Scanning Operations and is not part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☐ FADED TEXT OR DRAWING
- ☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☒ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.
As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
4 September 2003 (04.09.2003)

PCT

(10) International Publication Number
WO 03/073711 A2

(51) International Patent Classification[7]: H04L 12/58, 29/06

(21) International Application Number: PCT/US03/04964

(22) International Filing Date: 21 February 2003 (21.02.2003)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/991,201 22 February 2002 (22.02.2002) US

(71) Applicant: RPOST INTERNATIONAL LIMITED [US/US]; c/o Deloitte & Touche, Third Floor, Corner House, 20 Parliament Street, Hamilton (BM).

(72) Inventor: TOMKOW, Terrence, A.; 1301 Carmona Avenue, Los Angeles, CA 90019 (US).

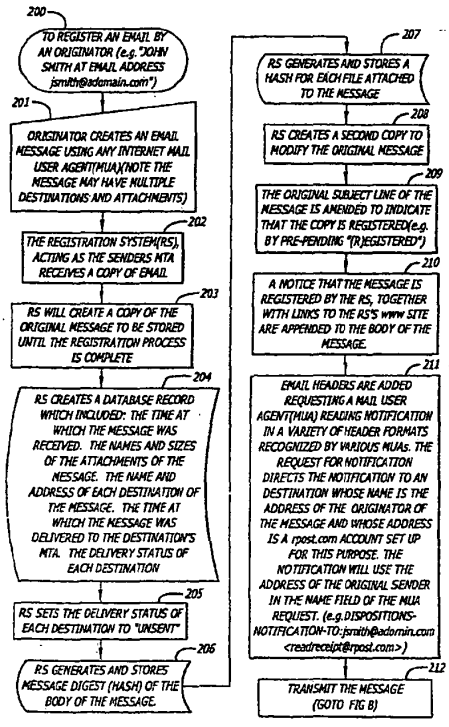
(74) Agents: PARKHURST, David, G. et al.; Fulwider Patton Lee & Utecht, LLP, Howard Hughes Center, 6060 Center Drive, Tenth Floor, Los Angeles, CA 90045 (US).

(81) Designated States *(national)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR VERIFYING DELIVERY AND INTEGRITY OF ELECTRONIC MESSAGES



(57) Abstract: A server receives a message from a sender and transmits the message through the Internet to a recipient. The server normally transmits the message in a first path through the Internet to the recipient. When the sender indicates at a particular position in the message that the message is registered, the server transmits the message in a second path through the Internet to the recipient. The sender can also provide additional indications in the message to have the server handle the message in other special ways not normally provided by the server. After learning from the receipt or the recipient's agent through the Internet that the message was successfully received, the server creates, and forwards to the sender, an electronic receipt. The receipt includes at least one, and preferably all: the message and any attachments, a delivery success/failure table listing the receipts, and the receipt times, of the message by the recipient's specific agents, and the failure of other agents of the recipient to receive the message and a digital signature of the message and attachments subsequently. By verifying that the digital signature on the sender's receipt matches the digital receipt at the server, the server can verify, without retaining the message, that the receipt is genuine and that the message is accurate.

WO 03/073711 A2

Published:
— *without international search report and to be republished upon receipt of that report*

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.

SYSTEM AND METHOD FOR VERIFYING DELIVERY AND INTEGRITY OF ELECTRONIC MESSAGES

BACKGROUND OF PREFERRED EMBODIMENTS OF THE INVENTION

I. FIELD OF THE INVENTION:

This invention relates generally to a system and method for verifying delivery and content of an electronic message and, more particularly, to a system and method of later providing proof regarding the delivery and content of an e-mail message. More specifically, the invention relates to a system and method for sending registered mail through the internet while verifying delivery and content of an electronic message and later providing proof regarding the delivery and content of the registered e-mail message.

II. DESCRIPTION OF RELATED ART:

In recent years e-mail has become an indispensable business tool. E-mail has replaced "snail mail" for many business practices because it is faster, cheaper and generally more reliable. But there remain some mail applications where hard copy is still dominant, such as registered and certified mail. For example, when a letter is sent by certified mail the sender is provided with a receipt to prove that the letter was mailed. A returned registered mail receipt adds the Postal Service's confirmation that the letter was successfully delivered to the addressee or the addressee's authorized agent. Additionally, private couriers such as Federal Express® and United Parcel Service® (UPS) provide some type of delivery confirmation. Since every piece of courier mail is, in effect, registered it is natural for consumers to turn to these services when they want proof of delivery.

Many existing e-mail systems and e-mail programs already provide for some form of proof of delivery. For instance, some e-mail systems today allow a sender to mark a message with "request for notifications" tags. Such tags allow a sender to request notification that the message was delivered and/or when the message was opened. When a sender requests delivery notification, the internet e-mail system may provide the sender with an e-mail receipt that the message was delivered to the mail server or electronic in-box of the recipient. The receipt message may include the title of the message, the destination address, and the time of delivery. It may also include (depending on the types of "flags" that are provided and activated in the mailing software) a list of all the internet "stations" that the message passed through en route to its destination. This form of reporting is built into some of the rules and protocols which implement e-mail. Furthermore, when a message is sent with a "read notification" request, the recipient's e-mail program may send to the sender an e-mail notification that the recipient opened that message for reading. Many electronic mail clients can and do support this kind of reporting; however, Internet protocols do not make this mandatory.

However, this does not mean that an e-mail sent with a notification request is as effective in all respects as registered mail. People certify and register letters because they want proof of delivery, e.g., proof that can be used in a civil or criminal proceeding, or proof that will satisfy a supervisor or a client or a government agency that a message has been sent, a job has been done, an order placed, or a contract requirement satisfied.

A registration receipt from the United States Postal Service (USPS) constitutes proof of delivery because the USPS stands behind it. The receipt represents the Post Office's confirmation that the letter or package in question was actually delivered to the addressee or his authorized representative. On the other hand, with the e-mail receipt various hurdles exist to an e-mail receipt being admitted and relied upon as persuasive evidence in a court of law as a proof that the message was delivered. After all, the receipt may be just another e-mail message that could have been altered or created by anyone, at any time.

There exists a need for an e-mail system and/or method that can provide reliable proof of the content and delivery of an e-mail message in order to take fuller advantage of the convenience and low cost of communicating via e-mail.

To meet this need some systems have been established whereby senders may receive third party proof of delivery by enrolling in services whereby:

a) The sender transmits an electronic message to a third party together with a list of the document's intended recipients.

b) The third party sends a notification to each of the message's intended recipients inviting them to visit the third party's web site where the message can be viewed.

c) If the intended recipient visits the third party's web site to view the message, the third party records this visit so that the sender may know that his message has been read by the recipient.

The drawbacks of such systems are manifold. In the first place, they rely essentially on the co-operation of the recipient of the e-mail to collect his messages from the third party's service. But the circumstances in which a sender may want proof of delivery of a message are often ones in which it cannot be assumed that the intended recipient will co-operate in receiving the message. In such cases, e.g. where acknowledging receipt of the message would place a financial or legal burden on the recipient, the recipient can simply ignore the notification that mail is available for him to receive. Note that there is nothing in such a system to guarantee that the intended recipient has received notification of waiting mail. In the second place, such systems are cumbersome and slow to use as compared to regular e-mail insofar as it can require the sender and/or the recipient to connect to a World Wide Web site to send, collect and verify the delivery of each message. Moreover, transmission of documents by such methods may require both sender and receiver to upload and download files to a web site. Finally, because these methods require the third party to retain a copy of the whole of each message until such time as they are collected or expired, the methods can require its provider to devote substantial computational resources to data storage and data tracking over an

extended period of time. As an alternative method of providing proof of delivery, some systems provide proprietary e-mail clients or web-browser plug-ins that will notify senders when a message has been received provided that a recipient uses the same e-mail client. The obvious disadvantage of such systems is that they require both sender and recipient to use the same e-mail client.

Therefore, there exists a need for an e-mail system/method that can provide reliable proof of the content and delivery of electronic messages which does not require the compliance or co-operation of the recipient, which requires no special e-mail software on the part of sender or recipient, which operates with the same or nearly the same convenience and speed of use as conventional e-mail, and which can be operated economically by a service provider. A general object of the invention disclosed and claimed in co-pending non-provisional application 09/626,577 (attorneys file RPOST-57228), filed by Terrance A. Tomkow on July 27, 2000 and assigned of record to the assignee of record in this application, is to provide a system and method for reliably verifying via secure and tamper-proof documentation the content and delivery of an electronic message such as an e-mail. Ideally, the invention disclosed and claimed in co-pending application 09/626,577 (attorneys file RPOST-57228) will give e-mail and other electronic messages a legal status on a par with, if not superior to, that of registered United States mail. However, it is not necessary to the invention that any particular legal status is accorded to messages sent according to the methods taught herein, as the invention provides useful information and verification regardless.

The invention disclosed and claimed in co-pending non-provisional application 09/626,577 includes an electronic message system that creates and records a digital signature of each electronic message sent through the system. An originator may send a copy of the electronic message to the system or generate the electronic message within the system itself. The system then forwards and delivers the electronic message to all recipients (or to the designated message handlers associated with the recipients), including "to" addressees and "cc" addressees. Thereafter, the system returns a receipt of delivery to the originator of the electronic

message. The receipt includes, among other things: the original message, the digital signature of the message, and a handshaking and delivery history including times of delivery to the recipients. To later verify and authenticate information contained in the receipt, the originator or user sends a copy of the receipt to the system. The system then verifies that the digital signature matches the original message and the rest of the receipt. If the two match, then the system sends a letter or provides other confirmation of authenticity verifying that the electronic message has not been altered.

The system disclosed and claimed in co-pending non-provisional application 09/626,577 may include a form of e-mail server connected to the Internet, which can be utilized in many ways. For instance, individual users can register their electronic messages, such as e-mails, by sending a "carbon copy" (cc:) to the system or composing the message within the system itself. For corporate or e-commerce users, these users can change their server to a server incorporating the present invention and have all of their external electronic messages registered, with the option of having the system retain and archive the receipts. The system can accept and verify encrypted electronic messages and manage the electronic messages within and/or outside a "fire wall." For web-based users, i.e., individuals or corporations using web-based e-mails, such as MSN Hotmail® or Yahoo Mail®, such users could check a box or otherwise set a flag within their e-mail programs to select on a case-by-case basis whether to make the e-mails of record and/or to archive the messages using the system disclosed and claimed in co-pending non-provisional application 09/629,577.

The digital signature can be created using known digital signature techniques, such as by performing a hash function on the message to produce a message digest and then encrypting the message digest. Separate digital signatures can be created for the body of the message, any attachments, and for the overall message including the body, the attachments, and the individual message digests. The encrypted message digest provides one type of message authentication or

validation code, or secure documentation. Other message authentication and/or validation codes may also be generated and used.

In one aspect, the invention disclosed and claimed in co-pending application 09/626,577 is a method of providing proof regarding the delivery and content of an electronic message, comprising: receiving from a sender across a computer network an electronic message, the message having a delivery address associated therewith; computing a message digest according to the message; encrypting the message digest; sending the message electronically to a destination corresponding to the delivery address; recording the Simple Mail Transport Protocol (SMTP) or Extended SMTP (ESMTP) dialog which effects the delivery of the message; receiving Delivery Status Notification information associated with the message and the delivery address; providing to the sender an electronic receipt, the receipt comprising: a copy of the message, the encrypted message digest, the (E)SMTP transcripts, and at least a subset of the Delivery Status notification information, and, at a future date, receiving electronically the electronic receipt from the sender, verifying that the encrypted message digest corresponds to the message, and verifying that the message was received by an electronic message handler associated with the delivery address.

In another aspect, the invention disclosed and claimed in co-pending application 09/626,577 includes a method of verifying delivery of an electronic message, comprising: in a wide area network computer system, receiving an electronic message from a message sender for routing to a destination address; establishing communication with an electronic message server associated with the destination address, the server defining a destination server; querying the destination server to determine whether the destination server supports Delivery Status Notification (DSN) functionality; receiving a response to the query, the query and response together defining an SMTP dialog; requesting Delivery Status notification information from the destination server according to results of the SMTP dialog; transmitting the electronic message to the destination address; receiving DSN information from the destination server with respect to delivery of

the electronic message; and providing to the message sender at least a portion of the SMTP dialog, and at least a portion of the DSN information.

In yet another aspect, the invention disclosed and claimed in co-pending application 09/626,577 includes a method of verifying content of a received electronic message, comprising: receiving the electronic message; generating a digital signature corresponding to the content of the received message; providing the message and the digital signature to a designated addressee; and, at a later time, verifying that the digital signature corresponds to the message.

In accordance with still another aspect of the invention disclosed and claimed in co-pending application 09/626,577, the method includes establishing whether a message was electronically received by a recipient, comprising: providing a message to be dispatched electronically along with a recipient's address from a sender; creating a signature associated with the message; dispatching the message electronically to the recipient's address; tracking the message to determine a final Delivery Status of the message dispatched to the recipient's address; upon receiving final Delivery Status of the message, generating a receipt, the receipt including a copy of the message, the signature, and the final Delivery Status for the message; and providing the receipt to the sender for later establishing that the message was electronically received by the recipient.

In accordance with yet another aspect of the invention disclosed and claimed in co-pending application 09/626,577, a method is provided for proving that an electronic message sent to a recipient was read, comprising: providing an electronic message along with a recipient's address; calculating a digital signature corresponding to the electronic message; dispatching the electronic message electronically to the recipient's address; requesting a Mail User Agent (email client "reading") notification from the recipient; upon receiving the reading notification, generating a reading receipt, the reading receipt including a copy of the message, the digital signature for the corresponding electronic message, and a second digital signature for the reading receipt from the recipient; and providing the reading receipt for later verification that said message was received by the recipient.

In accordance with another aspect of the invention disclosed and claimed in co-pending application 09/626,577, a method is provided for validating the integrity of a purported copy of an electronic message, comprising: receiving the purported electronic message copy, said purported copy including an encrypted message digest associated therewith; decrypting the message digest; generating a second message digest based on content of the purported copy; and validating the purported copy by comparing the decrypted message digest and the second message digest to determine whether the two message digests match.

In accordance with a still further aspect of the invention disclosed and claimed in co-pending application 09/626,577, a method is provided for validating a received registered e-mail, comprising: receiving an electronic receipt, said receipt including a base message and an encrypted message digest; decrypting the encrypted message digest; generating a second message digest from the base message; and validating the e-mail if the decrypted message digest matches the second message digest.

In yet another aspect, the invention disclosed and claimed in co-pending application 09/626,577 includes a website at which users can go to send and receive secure messages, with the website host acting as an independent third party which will send and receive the messages and provide secure documentation regarding the content and delivery of the messages.

The above-described objects of the invention disclosed and claimed in co-pending application 09/626,577 and other features and benefits of the present invention will become clear to those skilled in the art when read in conjunction with the following detailed description of a preferred illustrative embodiment and viewed in conjunction with the attached drawings in which like numbers refer to like parts, and the appended claims.

BRIEF DESCRIPTION OF PREFERRED EMBODIMENTS OF THE INVENTION

A server receives a message from a sender and transmits the message through the Internet to a recipient. The server normally transmits the message in a first path through the Internet to the recipient. When the sender indicates at a particular position in the message that the message is registered, the server transmits the message in a second path through the Internet to the recipient. The sender can also provide additional indications in the message to have the server handle the message in other special ways not normally provided by the server.

After learning from the receipt or the recipient's agent through the Internet that the message was successfully received, the server creates, and forwards to the sender, an electronic receipt. The receipt includes at least one, and preferably all: the message and any attachments, a delivery success/failure table listing the receipts, and the receipt times, of the message by the recipient's specific agents, and the failure of other agents of the recipient to receive the message and a digital signature of the message and attachments subsequently. By verifying that the digital signature on the sender's receipt matches the digital receipt at the server, the server can verify, without retaining the message, that the receipt is genuine and that the message is accurate.

BRIEF DESCRIPTION OF THE DRAWINGS

A detailed description of the preferred embodiment of the invention will be made with reference to the accompanying drawings:

FIG. 1 is a system diagram of a first embodiment of the invention disclosed and claimed in co-pending application 09/626,577, in which embodiment outgoing messages are made of record by being transmitted by a special Mail Transport Agent (MTA).

FIGS. 2-2F constitute a representative flow diagram for making an outgoing e-mail of record according to the embodiment of FIG. 1.

FIG. 3 is a system diagram of a second embodiment of the invention disclosed and claimed in co-pending application 09/626,577, in which embodiment senders may direct a Mail Transport Agent to transmit selected messages through a separate Mail Transport Agent constructed to make the selected messages of record.

FIG. 4 is a system diagram of a third embodiment of the invention disclosed and claimed in co-pending application 09/626,577, in which carbon copies (cc's) of outgoing messages are sent to a special server to be made of record.

FIG. 5 is a system diagram of a fourth embodiment of the invention disclosed and claimed in co-pending application 09/626,577, in which embodiment users compose outgoing messages to be made of record at a designated website.

FIG. 6 is a system diagram of a fifth embodiment of the invention disclosed and claimed in co-pending application 09/626,577 in which users may send e-mails made of record and store receipts from within a Web Based Mail User Agent (MUA).

FIG. 7 is a flow diagram for validating an e-mail receipt made of record.

FIG. 8 is a system diagram of an embodiment of the present invention for making of record incoming messages.

FIG. 9 is a flow diagram for making of record incoming messages.

FIG. 10 is a flow diagram for validating received messages made of record.

FIG. 11 is a system diagram showing an exemplary use of the invention disclosed and claimed in co-pending application 09/626,577 by an e-business to make of record and acknowledge incoming and outgoing communications.

FIG. 12 is a block diagram showing a flow chart for a method of making of record mail in a system such as shown in individual ones of the different embodiments shown in the previous Figures and showing how a message from a sender to a server with an indication representing made of record mail provides for the transmission of the message by the server to a recipient through a special route

different from the route in which the message is normally transmitted by the server to the recipient;

FIG. 13 is a block diagram showing a flow chart similar to that shown in Figure 12 but with additional blocks for providing special functions in addition to the functions provided by the flow chart in Figure 12; and

FIG. 14 is a partial view of a form which is used by the sender to make of record a message to be sent by the server to the recipient.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS OF THE INVENTION

This description is not to be taken in a limiting sense, but is made merely for the purpose of illustrating the general principles of the invention. The section titles and overall organization of the present detailed description are for the purpose of convenience only and are not intended to limit the present invention. Accordingly, the invention will be described with respect to e-mail messaging systems that use the Internet network architecture and infrastructure. It is to be understood that the particular message type and network architecture described herein is for illustration only; the invention also applies to other electronic message protocols and message types using other computer network architectures, including wired and wireless networks. For convenience of discussion, messages that are processed according to the invention disclosed and claimed in co-pending application 09/626,577 may be referred to herein as being "made of record" messages. In the discussion which follows, the term "RPost" will refer in general terms to a third party entity which creates and/or operates software and/or hardware implementing the present invention, and/or acts as a disinterested third party message verifier. Messages that are processed according to the present inventions are referred to as "registered" inventions. The term is used for convenience of exemplary discussion only, and is not to be understood as limiting the invention.

Rpost as outgoing mail server embodiment

FIG. 1 is a system diagram of a first embodiment of the present invention, wherein outgoing e-mails are made of record according to the invention disclosed and claimed in co-pending non-provisional application 09/626,577 . In this embodiment, the RPost server 14 serves as the primary outgoing Mail Transport Agent (MTA) for a message sender's Mail User Agent (MUA) 13. Although message recipient 18 is technically the addressee and is therefore merely the intended recipient or intended destination at this point in time, for simplicity of discussion this entity will be referred to herein as the recipient, addressee, or destination. Note that a single message may have many different destinations and that each of these may be reached through a different MTA. The method of sending messages made of record may divided into three parts:

1) *Preprocessing*: Steps to be taken before a message is transmitted;
2) *Transmission*: The method of delivering messages to addressees;
3) *Post Processing*: Procedures for gathering information about messages after their delivery, the creation of receipts, and the validation of receipts.

1.1 Preprocessing

On receiving a message for transmission, the Rpost server 14 will create records in a database for each message that will be used to store such information as:

a) the time at which the message was received;
b) the names of the attachments of the message; and
c) the number of addresses of the message.

For each destination of the message the database will record:

a) the name of the destination (if available);
b) the internet address of the destination;

c) the time at which the message was delivered to the destination's Mail Server; and

d) the *Delivery Status* of this destination.

Recipient Delivery Statuses used by the system will include:

UNSENT

This status indicates that the message has not been sent.

DELIVERED-AND-WAITING-FOR-DSN

This status indicates that the message has been delivered to an ESMTP compliant MTA that supports Delivery Status Notification (DSN) so that a success/failure notification can be expected.

DELIVERED

This status signifies that the copy of the message sent to this recipient has been successfully delivered to a server that does not support ESMTP DSN.

DELIVERED-TO- MAILBOX

This status signifies that a DSN message has been received which indicates that the copy of the message sent to this recipient was delivered to the mailbox of the recipient.

RELAYED

This status signifies that an MTA DSN has been received which indicates that the copy of the message sent to this recipient has been relayed onward to another server.

UNDELIVERABLE

This status indicates that after repeated attempts RPost has been unable to connect to an MTA to deliver the messages to this recipient.

FAILED

This status signifies that an MTA DSN has been received that indicates a failure to deliver a copy of the message to this recipient.

At this time the system will also perform *hashing* functions on the message's contents.

RPost server 14 employs a hash function and an encryption algorithm. The hash function may be one of any well-known hash functions, including MD2, MD5, the Secure Hashing Algorithm (SHA), or other hash functions which may be developed in the future. Hash algorithms and methods are described in Bruce Schneider, *Applied Cryptography: Protocols, Algorithms, and Source Code in C*, John Wiley & Sons, Inc. (New York) 1993; *Federal Information Processing Standard Publication 180-1 (FIPS PUB 180-1) Secure Hash Standard*, National Institute of Standards and Technology; and U.S. Pat. No. 5,530,757 issued to Krawczyk, entitled "Distributed Fingerprints for Information Integrity Verification," which are hereby incorporated by reference for their teachings of hash functions, encryption, and methods and systems for implementing those functions. Other known or new methods of detecting whether the contents of the message have been altered may be used.

A good hash function H is one-way; that is, it is hard to invert where "hard to invert" means that given a hash value h, it is computationally infeasible to find some input x such that $H(x) = h$. Furthermore, the hash function should be at least weakly collision-free, which means that, given a message x, it is computationally infeasible to find some input y such that $H(x) = H(y)$. The consequence of this is that a would-be forger who knows the algorithm used and the resulting hash value or message digest will nevertheless not be able to create a counterfeit message that will hash to the same number. The hash value h returned by a hash function is generally referred to as a message digest. The message digest is sometimes referred to as a "digital fingerprint" of the message x. Currently, it is recommended that one-way hash functions produce outputs that are at least 128 bits long in order to ensure that the results are secure and not forgeable. As the current state of the art advances, the recommended length for secure hash functions may increase.

RPost server 14 computes a message digest for the message body, and a separate message digest for each of the attachments of the message and stores these in a manner in which they may be later included in a receipt for the message.

Before the message is altered in the ways that registration will require, a copy of the original message and its attachments are stored in a manner in which they can be later retrieved by the system.

The RPost server 14 may alter a message in several ways before transmission to the recipient's MTA.

Although such is not necessary to the practice of the invention, the message may be tagged to denote the fact that the message has been made of record, such as by inserting the words "Made of Record" or at the beginning of the "subject" line of the message, by appending a tag such as,

"This message has been made of record with RPost. Visit our web site at www.RPost.com for additional information."

at the end of the original message or other tagging. Additionally, the tag may contain instructions, World Wide Web addresses, or links that invite and allow the recipient to send a reply made of record to the message by linking to a Web Page from which messages made of record may be composed and sent.

Although tagging is optional, the delivered message will generally be referred to herein as the tagged message.

Internet protocols provide two forms of receipt for e-mail messages:

MTA NOTIFICATIONS

These are e-mails that are sent by a recipient's MTA notifying the nominal sender of the message that various events have occurred. MTAs that conform to the SMTP protocol will typically only send a notification in the event that the mailer cannot deliver a message to the mailbox of the addressee (as might happen if the address is not valid or if the addressee's mailbox has exceeded its allotted storage quota).

With the introduction of the Extended SMTP standard it became possible for sending MTAs to request notices of success and failure in the delivery of messages.

These Delivery Status Notifications (DSNs) are e-mails which are sent by a receiving MTA to the nominal sender of the message when certain events occur: e.g. the message has been successfully deposited into the mailbox of the recipient; the message cannot be delivered to the recipient's mailbox for some reason; the recipient's message has been relayed on to another server which does not give DSN receipts.

Note that only e-mail servers that support the Extended SMTP (ESMTP) protocol support this form of DSN and that support for this function is optional for ESMTP servers and depends on the configuration selected by the server's administrators.

Although DSN is a term that only came into use with the advent of ESMTP, we will, in what follows, use 'DSN' to refer to any MTA generated message relating to the status of a received message whether or not it is in conformity to the ESMTP protocol.

MUA NOTICES (READING NOTIFICATIONS)

These are emails that are sent to the (nominal) author of a message by the recipient's Mail User Agent (MUA) (e-mail program) when certain events occur: e.g. the message is opened for reading, or deleted from the system without being read. By Internet convention (RFC 1891), no MUA program can be forced to generate such notifications. Whether an MUA will generate these receipts will depend upon the configuration chosen by its user.

The RPost server 14 will configure and transmit messages in a way that attempts to elicit both MTA DSNs and MUA notices from compliant MTAs and MUAs. In order to elicit a Reading Receipt from compliant MUAs, certain headers should be included in the header section of an e-mail message. Different MUAs respond to different headers; hence Server 14 will add several different headers to each message requesting a read notification in a form recognized by various MUAs. These headers all take the form:

```
Header label: user name <user address>
```

For example:

```
Disposition-notification-to: john smith
```

<jsmith@adomain.com>
read-notification-to: john smith
<jsmith@adomain.com>

where 'john smith' is the name of the user to whom an MUA notification is to be sent and '<jsmith@adomain.com>' is that user's internet address. Normally such headers would refer to the author of the message but in the case of the present method it is required that the notification be returned to RPost so that the notification can be processed by RPost. To assure that this is so Server 14 will insert headers that request that MUA receipts be sent to an address where they can be processed by the RPost server, for example: "readreceipt@RPost.com". This will direct any compliant recipient MUAs to send their notifications to an RPost address for processing.

The task of processing returned MUA notifications raises another problem that must be dealt with at this stage. There are no standards governing the format or content of MUA notifications. Often they will quote the subject of the original message and the time of the event (e.g. "opened for reading") that they are reporting. But even if this information is included in the notification it is rarely sufficient to uniquely identify the message that prompts it or to identify the author of that message. When the system receives a MUA notification it should be able to identify the message that prompts it, so as to include the notification information in the receipt that RPost will generate for the sender. Alternatively, the system should at least be able to reliably identify the sender of the message to which the MUA notification refers so that the notification information can be passed on to the sender in the form of an RPost Reading receipt (see below).

To accomplish the latter goal, the system can take advantage of the fact that internet addresses have two components: a name field and an address field, where the address field is set off by corner quotes "<>". Most MUAs will include both fields in the destination address of their MUA notifications. In composing its requests for MUA receipts, the RPost system will include the server 14 read receipt-handling address as the address for the notification but will use the *address* of the original sender in the name field of the header. For example, where the original

sender of the message is user John Smith with internet address jsmith@adomain.com, the RPost server 14 will include headers of the form:

```
Disposition-notification-to: jsmith@adomain.com
<readreceipts@RPost.com>
```

This will typically result in the compliant MUA sending a notification to readreceipts@RPost.com addressed as:

```
jsmith@adomain.com <readreceipts@RPost.com>
```

On receipt of such a notification at the address *"readreceipts@RPost.com"*, the server can, by parsing the addressee's field, determine that the notification concerns a message originally sent by *jsmith@adomain.com*, even if this could not be determined by any examination of the contents of the notification. With this information in hand, the server can then package the contents of the notification in a digitally signed RPost Reading receipt and send the receipt to the address jsmith@adomain.com

The RPost system will also endeavor to elicit and collect MTA DSN notices generated by recipient MTAs. Since such notifications are always sent to the address listed in the "FROM:" field of the message header, the server 14 will alter each message header so that the message is received as "FROM:" an RPost address at which DSNs may be processed.

However the problem of processing DSNs raises another issue, which will be dealt with at this stage. DSNs do not have any standard content or format; often it is impossible to determine, merely by examining the contents of these e-mails, what message their contents are giving notification of. This problem was supposed to have been addressed for DSNs generated in compliance with the ESMTP protocol by the use of DSN envelope ID numbers (see RFC 1869). According to the protocol, a transmitting MTA can include a reference number along with its request for a DSN. This number would be quoted in any returning DSN, allowing the sender to identify the subject message of the DSN. However, as a matter of fact, many MTAs that report themselves as supporting ESMTP DSN do not return a DSN envelope ID or any other information sufficient to reliably identify the subject message. Finally, even where a DSN does return information sufficient to identify

the message it is giving notice of, it often will not contain sufficient information to identify the specific addressee of the message that has prompted the notification. Thus, a single message might be sent to two addressees at a domain; one might be successfully delivered to the addressee's mailbox; the other, not. The MTA for the domain may report these events in a DSN in ways that provide no way for the recipient of the DSN to determine which addressee was successfully delivered and which was not (as, for example, may happen if the DSN reports the recipient's addresses as their local alias names rather than by the addresses contained in the original message).

The present invention solves this problem in four steps:

1) A unique identification number is generated for each outgoing message (e.g. based upon a time stamp). This number is stored in a database.

2) The recipients of each message are enumerated and the identifying numbers are stored in a database.

3) The message is sent separately to each intended recipient's MTA. (Even when two recipients have a common domain name and MTA, the server 14 will transmit the message to that MTA in two separate SMTP telnet sessions.)

4) When the server 14 transmits the message to a recipient's MTA it augments the message's "FROM" field to show the message as having been sent from an address which incorporates the message's unique ID and the identifying number of the sender. The address also contains a substring (e.g. "rcpt") that enables the server to identify return messages as DSNs.

Thus, a single message denominated "mmyyddss" by the server 14, from the sender named John Smith, might be sent to its first intended recipient (denominated "a" by the system) with a header reading:

```
From: John Smith <rcptmmddyyssa@RPost.com>
```

The same message would be sent to the second recipient with a header reading:

From: John Smith <rcptmmddyyssb@RPost.com>

Many e-mail MUAs will only display the name of the sender of a message and thus the special address will be unseen by most recipients.

The upshot of this form of addressing is that when the recipient MTAs issue DSNs (whether ESMTP compliant or not) they will address those DSNs to different RPost addressees. On receiving these DSNs, the server 14 can identify them as DSN messages by their "RCPT" prefix and, by parsing the addressees, can determine which message and which recipient is the subject of the DSN.

The server 14 will alter the 'FROM' field of each message to refer to a recipient of the message each time it attempts to transmit the message to that recipient's MTA.

To insure that recipient replies to transmitted messages are directed properly the server 14 will add an explicit "reply-to:" message header into the message listing the original sender's name and internet address. In the case of the present example this would be:

Reply-to: john smith <jsmith@adomain.com>

This will lead recipient MUAs to address replies to a received message to the actual sender's address, rather than the constructed RPost address.

I.2 Transmission

As noted above, it is part of the method that the RPost server 14 transmits a separate copy of an outgoing message to each addressee of that message. Moreover RPost will attempt to make each such delivery through a direct SMTP connection with a mail eXchanger (MX) of record for each destination.

Note: Each valid internet e-mail address includes an internet domain name or IP address. Each domain name/address has associated with it an e-mail server(s) authorized to receive mail for addresses in that domain. It will be noted that some domains have more than one server. The Domain Name Server responsible for each domain broadcasts the identity of its mail servers across the Internet. This information is publicly available and is managed and transmitted in ways that conform to rules and conventions which govern internet e-mail and Domain Name service.

Before transmitting a copy of a message to any destination the RPost server 14 will perform an Internet Name Server Lookup to identify an MTA associated with the destination's domain. Having identified the MTA responsible for receiving mail on behalf of a destination address, the system will attempt to open a telnet connection with the destination's local MTA.

It is common practice for internet e-mails to be relayed from MTA to MTA until they reach their final destination. The primary purpose for providing a direct connection between the RPost server 14 and the destination's MTA is so that the RPost server can record delivery of the message, (this record taking the form of an SMTP dialogue) with the e-mail server which has proprietary responsibility for receiving e-mail for the recipient domain name.

The existence of this record provides helpful evidence that the message was delivered, in much the same way that a registered mail receipt provides evidence of delivery. USPS Registered mail is treated as verifiably delivered if it can be proved to have been delivered to the addressee's authorized agent (e.g. a secretary, or mail room clerk). In the event of any legal challenge to the evidentiary merit of an RPost delivery receipt, it will be recognized that in selecting an internet e-mail service provider, the recipient has authorized that provider to collect electronic messages on his behalf. In turn, that service provider has acknowledged its status as the authorized agent for e-mail recipients of that domain name by broadcasting the address of its MTAs as the receptive e-mail servers for this domain.

Accordingly, having delivered messages directly to the mail server responsible for receiving the recipient's e-mail, RPost will have delivered the message to an agent the recipient has legally authorized to receive his mail. By recording the delivery transaction (that transaction taking the form of an SMTP dialogue) RPost can claim to have proof of delivery to the recipient's authorized agent.

Note that while the method herein described attempts to collect other forms of proof of delivery to each destination, whether or not these attempts succeed will depend upon factors that will not be in the control of RPost, (e.g. the form of SMTP

support deployed on the recipient's mail server). On the other hand, every successful delivery direct to a recipient's mail server will always generate an SMTP record. Recording this record allows RPost to provide proof of delivery to any valid internet destination that complies with the minimum protocols (SMTP) for internet mail. This represents an important advantage of the current method over other methods that might attempt to prove delivery by reliance on ESMTP DSN.

Having identified the MTA for a destination of a message, the RPost server 14 will attempt to open an ESMTP connection with the destination MTA by issuing an "EHLO" handshake in compliance with RFC 1869. If SERVER 16 supports ESMTP, it will respond by listing which ESMTP services it supports.

If SERVER 16 supports ESMTP, the RPost server 14 will first determine if SERVER 16 supports the ESMTP Service "VERIFY". The Verify service allows a calling SMTP server to determine, among other things, if an address in an MTA's domain is genuine. If the RPost server 14 determines by these means that the address it is attempting to deliver its message to is not valid, it will terminate the connection, cease attempting to deliver a message to this addressee, and record, in its database, the status of this message destination as UNDELIVERABLE.

Whatever its result, the RPost server 14 will record the ESMTP VERIFY dialogue in a file and store it so that it may be later attached to or included in the Delivery Receipt for this message. It should be noted that, out of concern for security, few ESMTP servers support the VERIFY function.

If System 16 does not support the VERIFY method, then the RPost server 14 will nevertheless attempt to deliver the message to System 16. Typically an MTA will accept messages for any address nominally in its domain and will later send a DSN if the address is invalid.

The RPost server 14 will then attempt to determine if the destination server supports the ESMTP service DSN. If it does, RPost will transmit the message with a request that SERVER 16 notify the sender of the message with an ESMTP DSN if the delivery to the addressee succeeds or fails. After the successful transmission of

the message to this destination the system will record the Delivery Status of this destination as DELIVERED-WAITING-FOR-DSN.

If Server 16 replies to the "EHLO" handshake in a way that indicates that it does not support ESMTP, the RPost server 14 will issue a "HELO" message to initiate an SMTP connection. If this connection is achieved, the RPost server 14 will transmit the message in compliance with the SMTP protocol and will record the Delivery Status of the destination as DELIVERED.

Whether the connection is SMTP or ESMTP, the RPost server 14 will record the entire protocol dialogue between the two servers. Typically this dialogue will include protocol messages in which, among other things, the destination server identifies itself, grants permission to upload a message for a named recipient, and acknowledges that the message was received. RPost will save the record of this transaction in such way that it may be later retrieved and included in or attached to the RPost Delivery Receipt for this message.

For various reasons RPost may not be able to achieve an SMTP connection with an MTA of a recipient or it may achieve such a connection but be denied permission to transmit the message by the recipient. In that case, if the internet DNS lookup reveals that the destination address is served by multiple MTAs, the RPost server 14 will attempt to deliver its message to each of these in turn. RPost will continue to attempt to deliver to an appropriate MTA as often as system resources permit. If, after a length of time, RPost cannot deliver the message to an address, it will mark the status of this recipient of this message as "UNDELIVERABLE" and stop attempting to send this message to this destination address.

When the RPost server 14 succeeds in transmitting a message to a destination Server that explicitly supports ESMTP DSN, RPost will record the status of this recipient for this message as "DELIVERED-AND-WAITING-FOR-DSN".

When the RPost server 14 succeeds in transmitting a message to the destination Server via a connection that does not explicitly support ESMTP DSN, RPost will record the status of this recipient for this message as "DELIVERED."

I.3. Postprocessing

DSN Processing

MTA DSNs will be returned to the RPost server 14 addressed to fictitious addresses in its proprietary domain (e.g. "RPost.com"), these addresses having been constructed as described above. The RPost server 14 will scan all inbound mail addressed to the domain and detect DSN messages by their identifying substring (e.g. "rcpt"). By parsing these addresses in the manner described above, the system can identify the message and the recipient that has prompted the DSN notification.

There is no standard format for DSN messages; neither is there any standard lexicon in which they report their results. To evaluate a received DSN the system should look in the subject line and the body of DSN messages for words and phrases that express the DSN's meaning. For example, such phrases as "successful delivery" or "delivered to mailbox" or "was delivered" normally signal that the message the DSN concerns was deposited to the mailbox of the destination. When it detects such phrases the System will change the Delivery Status of this destination of the message to "DELIVERED TO MAILBOX".

Phrases such as "could not be delivered", "fatal error", "failure" and "unsuccessful" typically signal a DSN that reports a failure by the MTA to deliver the message to the destination. When it detects phrases such as these in the DSN, the system will change the record of the recipient's Delivery Status to "FAILED".

Though the system always delivers mail to a proprietary MTA for the destination's domain, these MTAs will sometimes relay the message to a different server (as may be the case, for example, if the receiving MTA sends mail behind a firewall). In this case the DSN will contain such phrases as "relayed" or "relayed onward". In such cases the system will change the recipient's Delivery Status to "RELAYED".

Having evaluated the DSN and updated the recipient's Delivery Status accordingly, the system will save the DSN and any attachments it may contain in such a way that this message(s) may be included in and/or attached to an RPost Delivery Receipt.

Message Management

From time to time the system will scan each sent message and examine the status of each destination of that message in order to determine if the system has completed processing of that destination's delivery. The criteria for completion depend upon the destination's Delivery Status:

DELIVERED: This status indicates that a copy of the message for this recipient has been delivered to an MTA that does not support ESMTP DSN. Such an MTA may nevertheless send a form of Delivery Status Notification in the event that the message could not be delivered to the Mailbox of the addressee (as might happen, for example, if the destination address does not correspond to a valid account within the domain). Accordingly, the system will not treat the delivery for such a recipient as completed until a period of time has elapsed since the delivery to the recipient MTA. This time period--typically two to twenty-four hours--represents an estimate of the maximum time required for a majority of servers to return a notification of a failure to deliver and it may be adjusted if the specific destination domain is remote or known to be prompt or tardy in producing such notifications.

RELAYED: This status signifies that a DSN has been received that indicates that the recipient MTA has forwarded the message to another MTA that does not support ESMTP DSN. In this case it is nevertheless possible that the MTA to which the message has been delivered will send a notification of failure to deliver in due course. Accordingly recipients with this status are treated as complete under the same conditions as recipients with the status DELIVERED.

DELIVERED-AND-WAITING-FOR-DSN: This status indicates that the recipient's MTA supports ESMTP DSN and that a DSN has been solicited but not yet received. It may sometimes happen that although an MTA identifies itself as supporting this service it will nevertheless not provide DSNs even in the event of successful delivery. Accordingly, the system will regard deliveries to a destination with this status as completed even if no DSN is received after an interval of time. This interval--typically six to twenty-four hours--represents an estimate of the maximum time typically required for a compliant MTA to return a DSN.

DELIVERED-TO-MAILBOX: This status indicates that a DSN indicating successful delivery has been received for this recipient and hence the delivery of the message to this destination is completed.

FAILED, UNDELIVERABLE: Deliveries to recipients with this status are always treated as complete.

When the system finds that delivery to all recipients of a message has been completed the system will construct a Delivery Receipt for the message.

Creation of Delivery Receipts

Delivery receipts are e-mails sent to the original sender of the made-of-record message. The receipt 20 may contain:

1. an identifier for administrative purposes. This identifier may be or may include reference to the sender's ID and/or the value of the Internet Message-ID of the sender's message as received by the system;
2. the date and time at which the receipt was generated;
3. the quoted body of the original message together with the e-mail addresses of its intended recipients;
4. the date and time at which the RPost server received the message;
5. a table for each destination listing:

(i) the time at which the recipient's MTA received the message and/or the time at which the system received a DSN report from the recipient's MTA;

(ii) a Delivery Status of the message for that destination. The Delivery Status quoted in a Delivery Receipt is based upon the system's internal record of the destination's Delivery Status. They may be transcribed as follows:

- Deliveries to destinations whose status is FAILED or UNDELIVER ABLE will be recorded in the receipt as "failed".
- Deliveries to destinations whose status is DELIVERED or DELIVERED-AND-WAITING-FOR DSN will be recorded in the receipt as "delivered to mail server".
- Deliveries to recipients whose status is DELIVERED-TO-MAILBOX will be recorded in the receipt as "delivered to mail box".

The purpose of these reports is to accurately apprise the reader of the form of verification of delivery the system has been able to achieve.

6. a list of the original attachments of the e-mail together with the separate message digests of those attachments;
7. copies of the attachments to the original message, each original attachment being attached as an attachment to the receipt;
8. transcripts, summaries, or abstractions of the transcripts of all of the SMTP dialogs involved in the delivery of the message to each destination;
9. quotations from the bodies and the attachments of all received DSN reports including whatever details of

delivery or disposition of the message that they might reveal; and

10. any files that were returned to the system as attachments to DSN reports.

All of these separate elements of the receipt may have their own message digests or digital signatures included within the receipt. Additionally, the receipt may include a single overall encrypted message digest or digital signature computed and appended as part of the receipt, thus providing a single message authentication code which could be used to authenticate all of the information contained within the receipt. Since the receipt itself and SMTP dialogs and DSN reports within the receipt contain time-stamps, the receipt includes a non-forgeable record of the message recipient(s), the message content, and the time(s) and route(s) of delivery.

MUA Notification Processing

MUA Notifications could be collected and incorporated within RPost Delivery receipts in the same manner as MTA DSNs. However, MTA notifications are typically issued by receiving MTAs within a few hours of delivery whereas MUA Notifications will not be generated, if ever, until the recipient opens his MUA e-mail client and takes some action with respect to the received mail. For this reason, in this embodiment of the invention MUA notifications are collected separately from MTA notifications and reported in "RPost Reading Receipts" separate from RPost Delivery Receipts.

MUA notifications elicited by message headers constructed in the manner described above will be returned to a common RPost address (e.g. "readreceipts@RPost.com") and each notification will contain-- in the name field of its address-- the address of the original sender of this message. Because this is the only information required to generate an RPost reading receipt in the manner described below, the system can deal with MUA notices whenever these notices

may arrive and without any need to have stored any information about the original message in its databanks.

MUA notices may report, among other things, that a message has been read by a recipient, that a message has been displayed on the recipient's terminal (whether or not read), that a message has been deleted without having been opened. There is no protocol-governed standard for the content or format of MUA messages. The system could be configured so as to examine the text of MUAs to interpret their reports in the same fashion as the system uses for MTA DSNs. However, in the current embodiment of the invention, MUAs are not evaluated or interpreted by the RPost server 14 but are, instead, passed on to the sender for his own evaluation in a form that can be authenticated by RPost. To accomplish this the system will create an e-mail message styled as an "RPost Reading Notice" which may include, among other items:

1. subject line of the received MUA notice;
2. the body of the received MUA notice quoted as the body of the Reading Notice;
3. the received MUA notice included as an attachment;
4. any attachment(s) to the received MUA notice included as an attachment(s).
5. message digests of the received MUA notice and of any attachment(s) to that notice;
6. a date and time stamp;
7. an encrypted hash of at least items 5 and 6 providing an authenticatible date stamped digital signature for the document and all of its contents.

Receipt Disposition

In the case of the current embodiment of the invention, both RPost delivery receipts and Reading Notices are sent to the original sender of the made-of-record message. Since these receipts are digitally signed with an encrypted hash, RPost

can authenticate the information contained in these messages any time they are presented to RPost for this purpose, in the manner described below. This means that once it has transmitted a copy of the receipt to its sender (with instructions to the sender to retain the receipt for his records), RPost has no further need to retain any data concerning the message or its delivery and may expunge all such records from its system. Thus, RPost need not keep any copy of the original message or the receipt. This economy of archival memory gives the present invention an advantage over various prior art message authentication systems that required large amounts of data storage at the service provider side.

In this case the burden of retaining receipt data falls on the original sender of the message. Alternatively or additionally, third party verifier RPost may, perhaps for an additional fee, store a permanent copy of the receipt or of some or all receipt data. The receipt or part(s) thereof may be kept on any suitable archival storage devices including magnetic tape, CD ROM, or other storage device types. Additionally or alternatively, RPost may return receipts or parts thereof to a storage system devoted to this purpose within the control of the sender or the sender's organization.

As described above, RPost receipt information includes all of the data from the original sender's message and its attachments. There are circumstances in which users of the system might not wish to undertake the burden of retaining receipts in their records (e.g., out of fear of accidental data loss) but might also not wish to have the contents of their message in the hands of the RPost third party. Accordingly RPost might discard the contents of messages but store in its database only such information (e.g. sender, date of composition, message digests, destinations and Delivery Statuses) as might be provided by RPost to authenticate and verify the delivery of a message when presented with a copy of the message retained by the sender.

Verification

In the event that the originator of a message requires evidence at a later date that an e-mail was sent, delivered, and/or read, the originator presents the receipt(s) for the message to the operators of the system.

For example, in order to prove that a particular message was sent from sender 10 to recipient 18, sender 10 sends to RPost a copy of receipt 20 with a request to verify the information contained within the receipt. This could be done by sending the receipt to a predefined mailbox at RPost, e.g., verify@RPost.com. RPost then determines whether or not the receipt is a valid receipt. A receipt is a valid receipt if the digital signature matches the remainder of the receipt, and the message digests match the corresponding respective portions of the original message. Specifically, RPost performs the hash function on the various portions of the message including the message body, the attachments, and the overall message including the SMTP dialog and DSN reports, to produce one or more message digests corresponding to the purported message copy. RPost compares the message digests in the purported copy, including the overall message digest, with the message digests which RPost has computed from the purported message copy. The overall message digest can be compared by either decrypting the overall message digest received as the digital signature in the purported receipt, or by encrypting the overall message digest which was calculated from the purported message copy. If the message digests including the digital signature match, then the receipt is an authentic RPost-generated receipt. Assuming that a good hash function was used and that the keys used in the cryptographic hash function and the digital signature encryption algorithm have not been divulged to others, it is virtually impossible that the receipt has been "forged" by the person presenting the receipt. That is, the receipt must have been a receipt that was generated by RPost, and therefore the message contained in the receipt, the to/from information, the date and time of delivery, the fact of successful delivery, the route by which the message traveled, and any DSN information contained within the receipt, must be a true copy of that

information and is accurate. RPost can then provide authentication, verification, and confirmation of the information contained within the receipt. This confirmation can take the form of an e-mail confirmation, affidavit testimony from RPost employees familiar with the methods used by RPost, live sworn testimony in depositions and in court, and other forms of testimony. RPost can charge sender 10, recipient 18, or any other entity, fees for the various respective confirmation services. RPost can also provide testimony or other confirmation with regard to the non-authenticity of a purported receipt. Testimony may be provided in accordance with Federal Rules of Evidence 901(9), 901(10), 803(6), 803(7), 1001-1004, 1006, 702-706, corresponding state rules of evidence, and other applicable rules.

In sum, the system provides reliable evidence based on the testimony of a disinterested third party that a particular message having a particular content was sent, when it was sent, who sent it, who received it, when it was opened for reading, and when it was deleted. This evidence can be presented any time a dispute arises regarding the content and delivery of messages, as for example in contract formation, the timing of stock buy or sell orders, and many other applications. The operators of the system can attest to the accuracy of the information contained in the receipt itself without the need for the operators to preserve any record or copy of the information contained in the receipt.

A significant advantage of the system is that it can be used by existing MUAs without any change thereto. Because all the computation, encryption, ESMTP interrogation and dialog, DSN report collection, and receipt compilation, are performed by the third party RPost server 14, none of these functions need to be implemented within any of the user's equipment. Thus, users can take advantage of the system quickly and easily.

In the embodiment of the invention described above, the RPost server 14 makes of record the delivery of all messages passing through it. Alternatively, an RPost server 14 might make of record only those messages having certain destinations (e.g. external to an organization) or from certain senders (e.g. a customer relations group). Alternatively or additionally, the RPost server 14 might

make of record only those messages that had distinguishing characters or strings in the subject or body of the message. For example, the server might make of record only messages that the sender had included "(Make of Record)" or "(MR)" in the subject of the message. All other messages might be delivered by the RPost server 14 or some other server function as an ordinary internet MTA.

In this embodiment, RPost can raise revenue in a variety of ways. For instance: RPost can charge message sender 10 or her organization a fee on a per-message basis, on a per-kilobyte basis, on a flat fee periodic basis such as monthly, or on a combination of the above. RPost can also charge fees for authenticating and verifying a receipt, with a schedule of charges depending on whether the verification sought is a simple return e-mail, a written affidavit or declaration, sworn fact testimony in deposition or in court, or sworn expert testimony in deposition or in court. If the users opt to have RPost retain copies of the receipts, RPost can charge per item and /or per-kilobyte per month storage fees.

FLOW DIAGRAM FOR MAKING OF RECORD AN OUTGOING MESSAGE

FIGS. 2A-2G constitute a flow chart showing an exemplary operation of the first embodiment of the system. Modifying this flow chart to apply to other embodiments is within the skill of one familiar with software and e-mail protocols.

FIG 3A, Pre-processing, illustrates the steps taken with a message before it will be transmitted by the Making of Record Server (the System).

To make of record an e-mail message, in step 201 an originator/sender/user creates an e-mail message using any internet Mail User Agent (MUA). Possible MUAs include: (1) client side e-mail programs; (2) server based e-mail programs; (3) web based e-mail services; and (4) HTML forms submitted through web pages. The message may contain attached files as described in the Requests for Comments (RFCs) 822, 2046, and 2047, which are hereby incorporated by reference. RFCs are a series of notes regarding the Internet that discuss many aspects of computer

communication, focusing on networking protocols, procedures, programs, and concepts.

In this embodiment, the system functions as the sender's outgoing mail server and hence the sender's message will be directly transferred to the RPost server by the sender's MUA (step 202).

In step 203, the system creates a copy of the original message to be stored for later processing.

In step 204, the system creates a record in a database which may include such information as: the time at which the message was received by the server, the names and size(s) of the file attachment(s) of the message, the name (if known) of each destination of the message; the internet address of each destination; the time at which the message was delivered to the destination's MTA (initially this value is null) and a unit which records the Delivery Status of each destination.

In step 205, the Delivery Status of each destination is set to "UNSENT".

In step 206, the system generates and stores a message digest or digital fingerprint generated from the message body.

In step 207, the system generates and stores a hash or message digest for each attachment included in the message.

In step 208, the system may create a modified copy of the original message. In this second copy (step 209), the original subject line of the message may be amended to indicate that this copy is made of record (e.g. by pre-pending "Made of Record").

In step 210, a notice that the message is made of record by the system together with links to the system's Word Wide Web site may be appended to the body of the message.

In step 211, the e-mail headers may be added requesting reading notification in a variety of header formats recognized by various MUAs. The requests for notification direct the return notification to an address associated with the system: for example, "readreceipt@RPost.com". These headers will also include the

address of the original sender of the message in the name field of the address to which the MUA notification should be sent.

Preprocessing having been completed, the system will now transmit a copy of the message to each of its destinations as illustrated in FIG 2B.

Fig 2B illustrates the steps provided to transmit a message made of record. As step 220 indicates, the process provides a separate transmission for each recipient of the message.

In step 221, the system changes the header field of its working copy of the message to show the message as being "FROM:" a sender whose name is the original sender of the message but whose address is an "RPost.com" address constructed from:

a) a string used to identify returning MTA notifications (e.g. "RCPT");

b) a string which uniquely identifies the message being sent;

c) a tag which uniquely identifies the destination this copy of the message is being sent to.

In step 222, using the domain name of the destination currently being sent to, the system does a Domain Name Server Mail exchange lookup to find the address of the MTA(s) responsible for collecting mail for addresses in this domain.

In step 223, the system attempts to make a direct telnet connection to the MTA of the destination. If the connection fails, the system will try to make the connection again. Provided that the system has not exceed a maximum number of retries (227) for this destination, the system will try to remake the connection perhaps using another MX server for the destination's domain (228).

If, after a maximum number of retries, the system cannot connect to an MTA for this destination, the system will, as in step 226, record this destination's Delivery Status as "UNDELIVERABLE" and cease attempting to deliver this message to this destination.

On connecting to the destination's MTA, the system will begin making a record of its (E)SMTP dialog with the MTA (225).

In step 229, the system attempts to initiate an Extended SMTP (ESMTP) exchange with the destination MTA by issuing an "EHLO" greeting.

If the destination's MTS supports ESMTP, the system will then (230) determine if the destination MTA supports the SMTP function VERIFY. If the MTA supports VERIFY, the system will attempt to determine if the destination address is a valid address within the domain (231).

If the address is not valid, then, as in step 232, the system will record the Delivery Status of this destination as "FAILURE" and will cease attempting to deliver this message to this destination.

If the address is valid or if the ESMTP server does not support VERIFY, the system will then (233) determine if the receiving MTA supports the ESMTP service DSN (Delivery Status Notification).

If the MTA does support ESMTP DSN, the system will transmit the message with ESMTP requests to notify the nominal sender of the message of delivery success or failure (234). Having transmitted the message, the system will record the Delivery Status of this destination as "DELIVERED-AND-WAITING-FOR-DSN" (235).

If the receiving MTA does not support Extended SMTP, the system will transmit the message using SMTP (236) and record the destination's status as "DELIVERED" (237).

Having delivered the message, the system will then store the (E)SMTP dialog, recording the delivery in a manner in which it can later be recovered (238) and attempt to send the message to another destination.

Having transmitted a message to its destination(s), the system must perform several functions in order to gather information about the message's disposition. Fig. 2C illustrates the process by which the system processes MTA Notifications returned by recipient MTAs.

Because of the format used in the headers of sent messages illustrated in Fig 2B step 221, MTA message notifications will be delivered to a fictional local address at the server. The system will be able to detect these notifications by a

string (e.g. "rcpt") embedded in their addresses (241). By parsing the address, as illustrated in 242, the system can determine which message to which destination prompted the received notification.

In step 243, the system scans the subject line and the body of received MTAs for phrases that indicate whether the MTA is reporting a successful delivery, a failed delivery, or that the message has been relayed to another server.

In the event that the process at step 243 reveals that the notification is reporting a successful delivery, the system will, as illustrated in step 245, change the Delivery Status of the relevant destination of the relevant message to "DELIVERED-TO-MAILBOX".

If the system determines that the MTA notice is reporting a delivery failure, the system will (247) change the Delivery Status of the relevant destination of the relevant message to "FAILURE".

In the event that the system determines that the MTA notification indicates that the message was relayed to another server, the system will, as illustrated in step 249, change the Delivery Status of the relevant destination of the relevant message to "RELAYED".

Having processed the MTA Notification, the system will save this message and all of its attachments in such manner that they may be later recalled and used in construction of a receipt for this destination (250).

From time to time, as illustrated in Fig. 2D, the system will examine the status of each message to determine if the system has recovered all of the MTA notifications it is likely to receive for each destination of message and may hence proceed to construct a receipt for the message.

The system will examine the Delivery Status of each destination of the message.

If any destination has the Delivery Status "UNSENT", then the processing of the message is not complete. (252).

If the Delivery Status of a destination is "DELIVERED-AND-WAITING-FOR-DSN", then the system will not regard the processing for this destination as

complete unless, as is illustrated in step 254, the time since delivery of the message has exceeded the system's waiting period (e.g. 24 hrs.).

If the Delivery Status of a destination is "DELIVERED", (257) then the system will regard the processing of this destination as complete provided (258) that a period of time has elapsed which the operators of the system treat as sufficient to have received notice of delivery failure from the destination's MTA. (e.g. 2 hours).

Any other destination Delivery Status (e.g. "FAILED", "UNDELIVERABLE", "DELIVERED TO MAILBOX") is treated as having completed processing.

If processing of any of a message's destinations is not complete the system takes no action but moves to consider other messages in the system (step 255).

However, as illustrated in step 259, if processing of every destination of the message is complete, the system will generate a Delivery Receipt for the message.

As illustrated by way of example in FIG. 2E, the receipt may include:

An identifier for administrative purposes as in block 271. This identifier may be, or may include, reference to the sender's ID and/or the value of the Internet Message-ID of the sender's message as received by the system.

As in block 272, the quoted body of the original message 12 together with the e-mail addresses of its intended recipients may also be included.

As in block 273, a table for each recipient listing may include:

a) the time at which the recipient's MTA received the message and/or the time at which the system received DSN from the recipient's MTA;

b) the Delivery Status report of the message for that destination, i.e., "Delivered to Mail Server", "Delivered to Mail Box", Relayed", "Delivery Failure", or "Undeliverable".

As in block 274, a list of the original attachments of the e-mail together with their separate hash values or message digests.

As in block 275, transcripts or abstractions of the transcripts of all of the SMTP dialogs involved in the delivery of the message to each destination.

As in block 276, quotations from the bodies and the attachments of all received DSNs including whatever details of delivery or disposition of the message that they might reveal.

As in block 277, the system may attach to the receipt copies of all of the attachments of the original message, and, as in block 278, the system may additionally attach files returned to the system as attachments to DSNs.

In step 279, having generated the text of the receipt so far, the system then generates a first hash for the e-mail message and a second hash(es) for any attachments to the body of the receipt and calculates a digital signature for each of the hash(es) using an encryption key known only to the operators of the system. Encryption can employ, for example, the Data Encryption Standard described in Federal Information Processing Standard Publication 4-2 (FIPS PUB 46-2), the Data Encryption Standard, National Institute of Standards and Technology, which is hereby incorporated by reference. Alternatively, other known or new methods of encrypting the hash value may be used.

In step 280, the encrypted hash is then appended to the end of the message as the "document digital signature".

In step 281, the receipt 20, now being complete, may be sent by e-mail to the sender with the advice that it be kept for the sender's records. In step 282, the system may now delete all copies of the original message, attachments, and DSNs. Alternatively, rather than sending the receipt to the sender, the system may store the receipt, or both the sender and system can store the receipt.

Because MUA notifications are returned only at the option of the recipient and only when the recipient takes some action with respect to the received message, embodiments of the system may choose to treat these return messages differently than MTA notifications.

FIG. 2F illustrates how these MUA notifications may be treated by the system. MUA notifications are solicited by the system by including various headers in outgoing messages in the manner of Fig 2A, step 211. These headers direct compliant MUAs to send notifications to a system address (e.g.

"readreciept@RPost.com") set aside for this purpose. The headers also use, in the "name" field of this return address, the e-mail address of the original sender of the message. Accordingly, in step 286, when MUA notifications are returned to readreceipt@RPost.com the system can, by examining the address of the notification, determine the address to which a reading notification should be sent.

Upon the arrival of a read receipt from a destination's MUA, the system, in step 287, generates a reading receipt that contains the subject of the received MUA notification as its subject and incorporates, in its message body, the body of the received MUA Notification.

In step 288, the system attaches to the receipt any files that may accompany the MUA's receipt (typically these may include details of delivery or disposition and identifying references to the original e-mail.)

In step 289, the system generates a hash for any files attached to the receipt and records this hash in the body of the receipt.

In step 290, the system generates a hash for the body of the receipt and its attachments, encrypts this hash, and appends the result to the message as a "document digital signature".

In step 291, the system sends the resulting receipt to the sender of the message. In step 292, having sent this receipt, the system may delete all internal records of the transaction.

III. RPOST AS SECONDARY MAIL SERVER EMBODIMENT

FIG. 3 is a system diagram of a second embodiment of the present invention wherein the RPost server 14 does not serve as the user's primary MTA but rather works in collaboration with another MTA. In this embodiment, the sender may elect to make of record a particular outgoing message by including some form of flag in an outgoing message, message subject, or message addresses. For example, if and only if a sender includes the symbol "(Made of Record)" or '(MR)" in the subject of the message the sender's MTA will direct the message to be transmitted through the RPost server 14 to generate a receipt.

In this embodiment the operators of RPost receive revenues from the operator of the sender's MTA per message and/or per kilobyte transmitted.

IV. CC TO RPOST EMBODIMENT

FIG. 4 is a system diagram of a third embodiment in which a carbon copy ("cc") is sent to the RPost server 14. In this embodiment, the user or message sender 10 can use a standard MUA and standard MTA without modification. Message sender 10 composes the e-mail having a message body and any number of attachments, and addresses it to message recipient 18, along with any carbon copies (cc's) and blind carbon copies (bcc's) as desired. Additionally, message sender 10 addresses a cc to RPost. RPost server 14 tags the message as before, and sends the tagged message including attachments to the recipient's MTA 16 and any designated cc's. On receipt of such a copy RPost server 14 may send an e-mail acknowledging receipt of the copy.

Recipient 18 and other destinations of the message will now receive two versions of the same message: a first version of the message received directly from sender 10, and a second and tagged version which was forwarded from RPost. Once RPost receives confirmation from recipient MTA 16 that the tagged version of the message was successfully received by recipient MTA 16, RPost server 14 composes message receipt 20 as before and sends the receipt to sender 10 for his records.

Revenue can be generated by establishing accounts for message originating domains or individual message senders, and charging the users' accounts per message, per kilobyte, per month, or a combination of these. Revenue can also be generated for the placement of advertisements on receipts and from authentication and verification services as previously described.

V. WEBSITE EMBODIMENT

FIG. 5 is a system diagram of a fourth embodiment. In this embodiment, RPost server 14 is associated with a website at which a user composes messages. Message sender 10 visits the RPost Website and composes his message at the

website by entering the desired "to", "cc", "bcc", "Subject", and message text information. Attachments can be added by using features available on standard browsers and web servers. In this embodiment, the sender additionally provides an address to which the made-of-record receipt may be sent. RPost server 14 sends the receipt to sender 10 through sender's MTA.

Revenue can be generated by establishing accounts for message originating domains or individual message senders, and charging the users' accounts per message, per kilobyte, per month, or a combination of these. Revenue can also be generated for the placement of advertisements on receipts and from authentication and verification services as previously described.

VI. WEB BASED MUA EMBODIMENT

FIG. 6 is a system diagram of a fifth embodiment. In this embodiment, the RPost server 14 is associated with a web based Mail User Agent. In addition to allowing users to compose mail through a web browser, such a MUA provides subscribers with browser viewable mailboxes that display messages stored on the Web server site. Subscribers to such a service gain access to mail accounts with user names and passwords. In this embodiment, message sender 10 visits the RPost Website, accesses a Web Based e-mail account by entering a user name and password, and composes his message which is transported for delivery to RPost server 14. Receipts generated by the RPost server 14 are returned to a web based mailbox associated with the subscriber's account.

In addition to the revenue sources available in other embodiments, in this embodiment the operators can charge storage fees for receipts held in the web based mailbox.

In all of these embodiments, the receipt may serve as evidence that:

(1) the originator sent an e-mail message;

(2) the message was sent at a certain time;

(3) the e-mail was addressed to certain recipient(s);

(4) the e-mail was delivered to the e-mail mailbox of each of its intended recipient(s);

(5) the e-mail was delivered at a certain time;

(6) the e-mail was delivered by a certain network route; and

(7) the e-mail message and its attachments had the specific content recorded in the receipt.

Furthermore, the system under certain circumstances generates a separate receipt, which may be used as evidence that:

(1) the e-mail was inspected through the recipient's Mail User Agent (MUA); and

(2) the recipient took certain actions in response to the message, e.g., read or deleted the e-mail, at a particular time.

As with the other embodiments, this embodiment produces documented evidence which may be attested to and verified by the disinterested third party operators of the system concerning the delivery and integrity of an electronic message. In other words, the system can be thought of as transforming the e-mail to a made-of-record e-mail that can later be used to prove that a particular e-mail message was sent, that it was successfully delivered, and when and how.

Should a dispute ever arise, the dispute can be resolved through the receipt generated by the system because the receipt is so encoded that the operators of the system can determine the authenticity of the receipt as the product of the system. Thereafter, operators of the system can attest to the accuracy of the information contained in an authentic receipt, relying only on information contained in the receipt itself and without the need for the operators to preserve any record or copy of the information contained in the receipt.

In addition to these benefits, the receipts generated by the system may also be useful as evidence of the existence and authorship of such materials as might be transmitted through the system. Moreover, the system is easy to use, as the system can be used from any internet e-mail client program/MUA, so that there is no additional software required.

FLOW DIAGRAM FOR VALIDATING A RECEIPT

FIG. 7 is a flow diagram illustrating an exemplary method for validating a receipt. In the event that the sender of a message should require evidence that an e-mail was sent and delivered (and/or read) the sender presents the receipt(s) corresponding to the message to the operators of the system in step 700. The operators of the system then, in step 702, detach and decrypt the document digital signature appended to the receipt. In step 703, the operators generate a hash of the balance of the document, including attachments.

In step 704, if the current hash value does not match the decrypted hash value, then the system generates a report stating that RPost cannot authenticate the receipt as an accurate record of the delivery or the contents of the message described in the receipt.

If the decrypted hash is equivalent to the current hash of the message, the system can, as in step 706, warrant that the information contained in the body of the message is unchanged since the receipt passed through the system. If the original message contained no attachments, the system may now generate a report that warrants that the receipt is an accurate record of the message's contents and its delivery by the RPost server.

If the receipt reports that the original message contained attachments, then the receipt will also record the name and hash value of each attachment. In generating the receipt all attachments of the original message are attached unchanged to the receipt. Accordingly, the system will, for each such attached file, generate a hash of the attached file (708) and compare it to the hash value recorded in the body of the receipt (709).

If the calculated hash value of a file matches the value included in the receipt, the system can warrant that the file attached to the receipt is identical to that attached to the message as originally delivered. If the hashes do not match, then the system will report that it cannot warrant that the file attached to the receipt is identical to the file attached to the original message.

Having performed this calculation for each file attached to the original message, the system prepares a report which reports on the authenticity of the receipt and each of its attached files (710) or which reports the failure of validation (712).

Having completed its evaluation, the system will then append a copy of the receipt and all of its attachments to the report it has generated and send it via e-mail to the return address of the user who submitted the report for validation.

Registering Inbound E-mails

FIG. 8 is a system diagram illustrating another embodiment of the invention in which incoming e-mails are made of record. In this embodiment, a message sender 60 sends an e-mail message 70. Sender's MTA 62 sends message 70 onto the internet as usual. However, in this embodiment RPost contracts with service subscriber/recipient 68 to make incoming e-mails of record. According to the agreement, RPost is designated with Network Solutions, Inc. (NSI) or other domain name authority as the mail recipient (MX server) for recipient 68. This causes the Domain Name Service (DNS) request performed by the sender's MTA 62 to return the IP address of RPost as the IP address for the recipient, which in turn causes sender's MTA 62 to send the e-mail message to RPost server 64. RPost server 64 acts as an SMTP, POP, POP3 or IMAP MTA (collectively, "POP mail server") for recipient 68. SMTP, POP and IMAP MTAs are governed by RFC 821, the SMTP protocol, RFC 1939 Post Office Protocol - Version 3 (which obsoleted RFC1725), and RFC 2060 IMAP (Internet Message Access Protocol) Version 4 rev 1 (which obsoleted RFC1730), which are hereby incorporated by reference.

RPost server 64 prepares a made-of-record version 74 of the original message 70, and places this made-of-record version 74 into recipient 68's in-box instead of, or in addition to, the original message 70. The made-of-record version may have all of the verification and informational features and options discussed earlier in connection with e-mail receipts. This information can include, but is not limited to: individual message digests for each of the message body and text, the

to/from information, other header information, each attachment, an overall message digest and digital signature and message routing information and tags. Made-of-record version 74 of message 70 as shown in FIG. 6 includes the message body including the header information, an attachment, separate message digests for each, and a digital signature or encrypted message digest. The hash functions and encryption are performed using private phrases or private keys known only to the operators of the system. The made-of-record version 74 is made available to recipient 68 for inspection or downloading through the recipient's MUA.

RPost server 64 can optionally send a confirming e-mail 72 to message sender 60. Confirmation message 72 can be a simple text message indicating that a message was received and made of record. Confirmation message 72 could also include a message such as, "Your e-mail message was received on March 24, 2000 at 2:05 p.m. The digital signature of the message was [128-bit digital signature]. For more information, visit our website at www.RPost.com." Alternatively, or additionally, confirmation message 72 could include all of the information contained in the made-of-record version 74.

Thus, the system may provide to message recipient 68 a receipt 74 or other verifiable confirmation that:

(1) the recipient received an e-mail message;

(2) the message was received at a certain time;

(3) the e-mail was addressed from a certain sender;

(4) the message purports to be delivered via a certain network route; and

(5) the e-mail message and its attachments had a specific content.

Accordingly, the system provides evidence, which may be attested to by the operators of the system, that particular electronic messages and documents were delivered to recipients having certain content and representing themselves as having come from certain senders.

FIG. 9 is a flow chart illustrating one example of making of record in-bound mail. In step 901, RPost server 64 receives a new e-mail message. In step 902, the system generates a hash/digital signature of the message's contents including the message's headers and attachments. Additionally, the system may generate a separate hash for each message attachment. In step 903, the system encrypts the hash(es) using an encryption key known only to the operators of the system. In step 904, the resulting encrypted hash(es) is then appended to the body of the message. Then, in step 905, the modified message may be made available for inspection or downloading through the recipient's MUA.

FIG. 10 is a flow chart of one example of validating a received made-of-record e-mail message. In step 1000, in the event that the recipient of a message should require evidence that an e-mail with a specific content was received at a particular time, the recipient can present a copy of the made-of-record version 74 (FIG. 8) of e-mail message 70 to the operators of the system for verification. To verify the message, in step 1001 the system detaches and decrypts the document digital signature appended to the message. In step 1002, the system generates a hash of the balance of the document, and one for each file attached to the message. In steps 1003 and 1004, the hashes are compared. If the document hash(es) matches the decrypted hash(es), then the message and its attachments have passed through the system and have not been altered since their delivery to the recipient.

Having determined that the e-mail is unaltered, the operators of the system can warrant that:

(1) the e-mail was received by the system at a certain time;

(2) the e-mail purported to arrive at the system via a certain internet route;

(3) the e-mail purported to be from a certain sender; and

(4) the e-mail and its attachments were delivered with the specific content they currently contain.

On the other hand, in step 1006, if the hash values do not match, then the operator cannot warrant that the e-mail is authentic, i.e., that the e-mail is an accurate version of an e-mail that was received by the system.

FIG. 11 illustrates how the invention may be used by a business which utilizes electronic tools (an "e-business"). E-business 30 can utilize the system to make of record all incoming and outgoing e-mail messages from its customers 34. In this case, the system includes Post Office Protocol (POP) server 36 and Simple Mail Transfer Protocol (SMTP) server 38. For example, the e-business 30 can set up its website to e-mail forms to customers, and to forward queries and complaints 40 from customers 34. The made-of-record queries, complaints, orders, offers to purchase, and other information 46 are sent to the e-business 30 by the system. Receipts are then provided to the customers 34 via SMPT server 38. This way there is no question regarding whether or not the customer sent the communication and what it contained. Moreover, the e-business can set up a web site 32 through the RPost server so that every communication with the customers can be made of record. In other words, through the web site form data orders 42 and automated responses 44 can be made of record through the system server; furthermore, any confirmation, collections notices, customer support, and special offers 48 sent by the e-business to customers 34 can be made of record and the confirmation sent to the customer to eliminate arguments about what was ordered, when, or by whom. If desired, identical receipts can be provided to both the customers 34 and to e-business 30. Alternatively, the functions of POP server 36 and SMTP server 38 may be combined in a single system server.

POP is a protocol used to retrieve e-mail from an e-mail server. Many e-mail applications (sometimes called e-mail clients) use the POP protocol, although some can use the newer Internet Message Access Protocol (IMAP). One version of POP, called POP2, requires SMTP to send messages. A newer version, POP3, can be used with or without SMTP. SMTP is a protocol for sending e-mail messages between servers. Many e-mail systems that send e-mail over the Internet use SMTP

to send messages from one server to another; the messages can then be retrieved with an e-mail client using either POP or IMAP. In addition, SMTP is generally used to send messages from a mail client to a mail server. E-mail servers may use a variety of protocols to communicate with the Internet. Commonly used protocols include SMTP, POP3 and IMAP4. Mail readers are at the opposite end of the server. Since mail servers receive messages via SMTP, e-mail readers send e-mail to a mail server using SMPT. Likewise, since mail servers send messages using POP3 and optionally IMAP4, mail readers receive messages from mail servers by using the POP3 or IMAP4 protocol.

Although the above generally describes a system and method of verifying that an e-mail was sent and/or received, the invention disclosed and claimed in application 09/626,577 may apply to any electronic message that can be transmitted through an electronic message network or through any electronic gate. Electronic messages may include text, audio, video, graphics, data, and attachments of various file types. The methods and techniques taught herein can be programmed into servers and other computers, and computer programs implementing the invention can be written onto computer readable media including but not limited to CD ROMs, RAM, hard drives, and magnetic tape. E-mail made-of-record services according to the present invention can be bundled with internet service provider (ISP) services to provide a single provider ISP solution to corporate and other institutional clients. Implementing the above-described invention is within the skill of the ordinary practitioner of the software arts.

As previously indicated, Figures 1-11 show, and the specification describes, systems in which the server receives a message from a sender and transmits this message in a first path to a recipient or to a Mail Transport Agent (MTA) of the recipient. There are times when the sender may wish to have the server send the message to the recipient or to a Mail Transport Agent of the recipient or to the Mail Transport Agent for the recipient by a more traveled path or a less traveled route, or at least a different path, than the first path. To accomplish this, the sender marks a message form 1200 (Fig. 14) with a particular indication such as "(R)" in a

particular position such as the "subject" line of the message. This particular position is indicated at 1202 in the message form 1200 in Figure 14. The step of marking "(R)" in the "subject" line 1202 of the message is shown at 1206 in Figure 12.

The message with the "(R)" in the "subject" line is transmitted by the sender to the server 14 which constitutes the sender's Mail Transport Agent. This is indicated at 1208 in Figure 12. As indicated at 1210, the server scans the "subject" line to determine if there is an "(R)" in the line. If the answer is "No" (see 1211), the server transmits the message to the recipient or the recipient's Mail Transport Agent through the route shown in Figures 1-11 and indicated in Figure 12 as "the ordinary route" and discussed above in the specification. This is indicated at 1212 in Figure 12. If the answer is "Yes" (see 1213), the message is transmitted through a special network route as indicated at 1214 in Figure 14.

Figure 13 is identical in a number of respects to Figure 12. However, Figure 13 includes additional blocks to perform additional functions other than those shown in Figure 12. These include, but are not limited to, the following:

(1) The sender may wish that a copy of the message should be archived. This may be accomplished by adding a coding such as the number "1" after "(R)" in the "subject" line so that the coding is "R1".

(2) The sender may desire that a record of the transmission should be recorded by the server 14 constituting the sender's mail transport agent. This may be accomplished by providing a coding such as "(R2)" in the "subject" line of the message.

(3) The sender may desire that a record of the message transmission should be recorded in a data base. This may be accomplished by providing a coding such as "(R3)" in the "subject" line of the message.

(4) The sender may wish that a record of the message transmission should be recorded in a database with a special annotation or

additional reference. This may be accomplished by providing a coding such as "(R4)" in the "subject" line of the message.

Figure 13 provides a method where the server constituting the sender's Mail Transport Agent processes selected e-mail messages such as those specified in this paragraph.

Figure 13 is particularly limited to a coding "(xyz)" in the "subject" line of the message. In Figure 13, the sender is shown at 1300 as composing an electronic message including "(xyz)" in the "subject" line of the message. As indicated at 1210 in Figures 12 and 13, the server 14 constituting the mail transport agent scans the "subject" line in the outbound message. If the "subject" line in the message does not contain the code "(R)", the server transmits the message though the route shown in Figure 1-11 and discussed above (see 1212 in Figures 12 and 13). If the code "(R)" is detected by the server in the "subject" line of the message, the server transmits the message through the special network route as indicated at 1214 in Figures 12 and 13.

Figure 13 indicates at 1304 that the code "(xyz)" is stripped by the server from the "subject" line of the message. If the delimiter "xyz" is detected, a copy of the message is saved. This is indicated at 1308 in Figure 13. If the delimiter "xyz" is not identified, a copy of the message is not saved.

Although the present invention has thus been described in detail with regard to the preferred embodiments and drawings thereof, it should be apparent to those skilled in the art that various adaptations and modifications of the present invention may be accomplished without departing from the spirit and the scope of the invention. Accordingly, it is to be understood that the detailed description and the accompanying drawings as set forth hereinabove are not intended to limit the breadth of the present invention, which should be inferred only from the following claims and their appropriately construed legal equivalents. In the following claims, those claims which contain the words "means for" are intended to be interpreted in accordance with 35 U.S.C. §112, paragraph 6; those claims which do not include

the words "means for" are intended to not be interpreted in accordance with 35 U.S.C. §112, paragraph 6.

WHAT IS CLAIMED IS:

1. A method of transmitting a message through the internet from a sender to a recipient through a server displaced from the recipient, including the steps of:

receiving the message at the server from the sender and receiving an indication that the sender wishes to send the message in a manner special to the sender and not normally provided by the server,

transmitting, through the internet from the server to an agent of the recipient, the message in the special manner, in accordance with the indication from the sender, an identification and an internet address of the server and the identity of the sender,

receiving from the agent at the server through the internet the identity of the agent and an indication of the receipt of the message by the agent and the identification and internet address of the server and the identity of the sender, and

sending to the sender from the server through the internet a copy of the message and the information received by the server from the agent.

2. A method as set forth in claim 1 wherein

the transmission through the internet from the server to the agent of the recipient is in the normal manner when the sender does not provide an indication through the internet that the sender wishes to transmit the message through the internet in the special manner to the agent of the recipient

3. A method as set forth in claim 1 wherein

the indication received by the server through the internet from the agent of the recipient includes an identification of the agent and any transferring agents through whom the message has passed between the server and the agent of the recipient.

4. A method as set forth in claim 1, wherein

a digital fingerprint of the message is provided by a plurality of digits in a unique sequence and is sent by the server to the sender.

5. A method as set forth in claim 3 wherein

the transmission through the internet from the server to the agent of the recipient is in the normal manner when the sender does not provide an indication through the internet that the sender wishes to transmit the message through the internet in the special manner to the agent of the recipient and wherein

a digital fingerprint of the message is provided by a plurality of digits in a unique sequence and is sent by the server to the sender.

6. A method as set forth in claim 1 wherein

an additional indication is provided to the server with the message from the sender that a high priority should be provided by the server to the sending of the message by the server to the agent of the recipient and wherein

the server provides the high priority in sending the message to the agent of the recipient in accordance with the additional indication.

7. A method as set forth in claim 5 wherein

an additional indication is provided to the server with the message from the sender that a high priority should be provided by the server to the sending of the message by the server to the agent of the recipient and wherein

the server provides the high priority in sending the message to the agent of the recipient in accordance with the additional indication.

8. A method as set forth in claim 1 wherein

an additional indication is provided to the server with the message from the sender that the sending of the message by the server to the agent of the recipient should be recorded by the server and wherein

the server records the sending of the message by the server to the agent of the recipient in accordance with the additional indication.

9. A method as set forth in claim 5 wherein

an additional indication is provided to the server with the message from the sender that the sending of the message by the server to the agent of the recipient should be recorded by the server and wherein

the server records the sending of the message to the agent of the recipient in accordance with the additional indication.

10. A method of transmitting a message through the internet from a sender to a recipient through a server displaced from the recipient, including the steps at the server of:

receiving the message at the server from the sender,

receiving at the server, with the message from the sender, an indication that the message is to be transmitted by the server in a special manner different from the manner normally provided by the server in transmitting messages,

transmitting from the server through the internet to an agent of the recipient, in the special manner indicated by the sender to the server, the message and an identification and an internet address of the server and an indication representing the identity of the sender,

receiving at the server from the agent a handshaking and delivery history of the transmission of the message from the server to the agent of the recipient, and

transmitting from the server to the sender through the Internet the message, a digital signature, including a digital fingerprint, of the message and the handshaking and delivery history of the message received by the server from the agent of the recipient.

11. A method as set forth in claim 10, including the steps of:

receiving at the server, with the message from the sender, an additional indication from the server of an additional function to be performed in the transmission of the message from the server to the agent of the recipient,

providing the additional function in the transmission of the message from the server to the agent of the recipient in accordance with the additional indication provided by the sender to the server.

12. A method as set forth in claim 11 wherein

the message is sent by the server to the sender after the server receives from the agent of the recipient the handshaking and the delivery history of the transmission of the message from the server to the agent of the recipient and wherein the server does not retain the message after it sends the message to the sender.

13. A method as set forth in claim 11 wherein

the additional indication from the sender to the server provides for a recording of the transmission of the message and wherein

the transmission of the message is recorded in accordance with the additional indication from the sender.

14. A method as set forth in claim 11 wherein

the additional indication from the sender provides for an archiving of the message and

wherein the message is archived in accordance with the additional indication from the sender.

15. A method as set forth in claim 11 wherein

the additional indication from the sender provides for the message to be sent to the agent for the recipient by a special route and wherein

the message is sent by the special route from the server to the agent of the recipient in accordance with the additional indication from the sender.

16. A method as set forth in claim 11 wherein

the additional indication from the sender provides for the message to be specially handled in the transmission of the message from the server to the agent of the recipient and wherein

the message is specially handled in the transmission of the message from the server to the agent of the recipient.

17. A method as set forth in claim 11 wherein

the additional indication provides for a transmission of the message with a high priority from the server to the agent of the recipient and wherein

the message is transmitted from the server to the agent of the recipient with a high priority in accordance with the additional indication from the sender.

18. A method as set forth in claim 11 wherein

the server retains a copy of the digital signature of the message and the handshaking and delivery history of the message, but not the message unless requested to do so by the sender, after the server transmits to the sender through the internet the message, the digital signature of the message and the handshaking and delivery history of the message.

19. A method as set forth in claim 10 wherein

the server retains a copy, except for the message, of the information received by the server from the agent of the recipient and sent to the sender and wherein

when the sender wishes to authenticate that the message was sent by the server to the agent of the recipient, the server matches the information, except for the message, sent by the server to the sender relating to the message with the information retained by the server relating to the message.

20. A method as set forth in claim 10 wherein

the server requests a delivery status notification from the agent of the recipient relating to the message when it transmits the message to the agent and wherein

the server receives the delivery status notification from the agent of the recipient when it receives the digital signature of the message from the agent.

21. In a method of transmitting a message through the internet to a recipient through a

server displaced from the recipient, the steps at the server of:

receiving the message at the server from the sender,

generating a hash constituting a synopsis of the message in coded form,

encrypting the hash with a particular encryption code to generate a digital fingerprint of the message,

receiving from the sender an indication with the message from the sender that the message is to be handled by the server in a particular manner different from a normal handling of the message by the server, and

handling the message at the server in the particular manner, in accordance with the indication from the sender, to transmit the message and the digital fingerprint to the recipient.

22. In a method as set forth in claim 21, the steps of:

generating, for any attachment to the message, a hash constituting a synopsis of the attachment in coded form,

encrypting the hash from the attachment with a particular encryption code to generate a digital fingerprint of the attachment, and

transmitting the attachment and the digital fingerprint of the attachment to the sender through the internet at the same time, and in the same manner, that the message and the digital fingerprint of the message are transmitted from the server to the sender through the internet.

23. In a method as set forth in claim 21 wherein

the message is handled by the server in the normal manner when the indication is not provided by the sender to the server with the message and wherein

the message is handled by the server in the particular manner when the indication is provided by the sender to the server with the message.

24. In a method asset forth in claim 23 wherein

the message is processed by the server in a first path when the indication is not provided by the sender to the server with the message and wherein

the message is processed by the server in a second path different from the first path when the indication is provided by the sender to the server with the message.

25. In a method as set forth in claim 23, the steps of:

storing at the server the digital fingerprint of the message, the name of the sender, the identity and internet address of the server and the identity and internet address of the recipient, and

transmitting to the sender for storage by the sender the message, the digital fingerprint of the message, the name of the sender, the identity and internet address of the server and the identity and internet address of the recipient.

26. In a message as set forth in claim 22,

the message is handled by the server in the normal manner when the indication is not provided by the sender to the server with the message and wherein

the message is handled by the server in the particular manner when the indication is provided by the sender to the server with the message and wherein

the message is processed by the server in a first path when the indication is not provided by the sender to the server with the message and wherein

the message is processed by the server in a second path different from the first path when the indication is provided by the sender to the server with the message and wherein

the digital fingerprint of the message, the name of the sender, the identity and internet address of the server and the identity and internet address of the recipient are stored at the server, and

the message, the digital fingerprint of the message, the name of the sender, the identity and internet address of the server and the identity and internet address of the recipient are transmitted to the sender for storage by the sender.

27. A method of transmitting a message through the internet from a sender to an agent

for a recipient through a server displaced from the agent including the steps of:

providing the message from the sender at the server,

providing at the server a digital fingerprint of the message and the identity of the sender and the identity and internet address of the server,

transmitting to the agent in a first route the message and the identity of the sender and the identity and internet address of the server,

providing an indication at the server from the sender that the message from the sender should be transmitted by the server to the sender through a second route different from the first route,

transmitting the message from the server to the agent of the recipient through the second route in accordance with the indication from the sender,

providing at the agent of the recipient an indication of the status of the reception at the agent of the transmittal from the server to the agent of the message and the identity of the sender and the identity and the internet address of the server, and

transmitting to the server from the agent of the recipient the identity and internet address of the gent and the status of the reception at the agent of the message and the identity of the sender and the identity and internet address of the server.

28. A method as set forth in claim 27 wherein

the digital fingerprint of the message includes a digital digest of the message and an encryption of the digital digest and wherein

the message and the digital fingerprint of the message and the identity of the sender and the identity and internet address of the server and the identity and

the internet address of the agent of the recipient and the status at the agent of the reception at the agent of the message are transmitted by the server to the sender.

29. A method as set forth in claim 27 wherein

the sender provides at the server for an additional function to be preformed at the server and wherein

the server performs the additional function in accordance with the indication from the sender.

30. A method as set forth in claim 29 wherein

the additional indication provides for the message to be specially handled in the transmission of the message from the server to the agent of the recipient and wherein

the message is specially handled in the transmission of the message from the server to the agent of the recipient.

31. A method as set forth in claim 28 wherein

the sender provides at the server for an additional function to be performed at the server and wherein

the additional function represented by the additional indication provides for the message to be specially handled in the transmission of the message from the server through the second route to the agent of the recipient and wherein

the message is specially handed in the transmission of the message from the server through the second route to the agent of the recipient.

32. A method of transmitting a message through the internet from a sender to an agent

for a recipient through a server displaced from the agent, including the steps at the server of:

providing at the server a digital fingerprint of the message and the identity of the sender and the identity and the internet address of the server,

transmitting to the agent of the recipient through a first route the message and the identity of the sender and the identity and internet address of the server,

receiving from the sender an indication that the message should be sent by the server to the agent of the recipient through a second path different from the first path,

transmitting to the agent of the recipient the message and the identity of the sender and the identity and internet address of the server through the second path different from the first path in accordance with the indication from the sender,

receiving from the agent of the recipient the identity of the sender and the identity and internet address of the server and the identity and internet address of the agent and an indication of the status of the reception of the message at the agent, and

transmitting to the sender the message and the information received by the server from the agent of the recipient relating to the message.

33. A method as set forth in claim 32 wherein

the server stores the information relating to the message, but not the message, transmitted by the server to the sender and wherein

the server authenticates the message by comparing the information stored by the server relating to the message with the information transmitted from the server to the sender relating to the message.

34. A method as set forth in claim 32 wherein

the server receives additional information from the sender relating to additional functions to be performed by the server on the message in the transmission of the message from the server to the agent of the recipient and wherein

the server performs the additional functions on the message, in accordance with the additional information received by the server from the sender, in the transmission of the message from the server to the agent of the recipient.

35. A method as set forth in claim 34 wherein

the indication received by the server from the sender constitutes a first coding of the message from the sender and wherein

the additional information received by the server from the sender of the additional function to be performed by the server constitutes a second coding, added to the first coding, of the message from the sender.

36. A method set forth in claim 33 wherein

the server receives additional information from the sender relating to additional functions to be performed by the server on the message in the transmission of the message from the server to the agent of the recipient and wherein

the server performs the additional functions on the message, in accordance with the additional information received by the server from the sender, in the transmission of the message from the server to the agent of the recipient and wherein

the indication received by the server from the sender constitutes a first coding of the message from the sender and wherein

the additional information received by the server from the sender of the additional function to be performed by the server constitutes a second coding, added to the first coding, of the message from the sender.

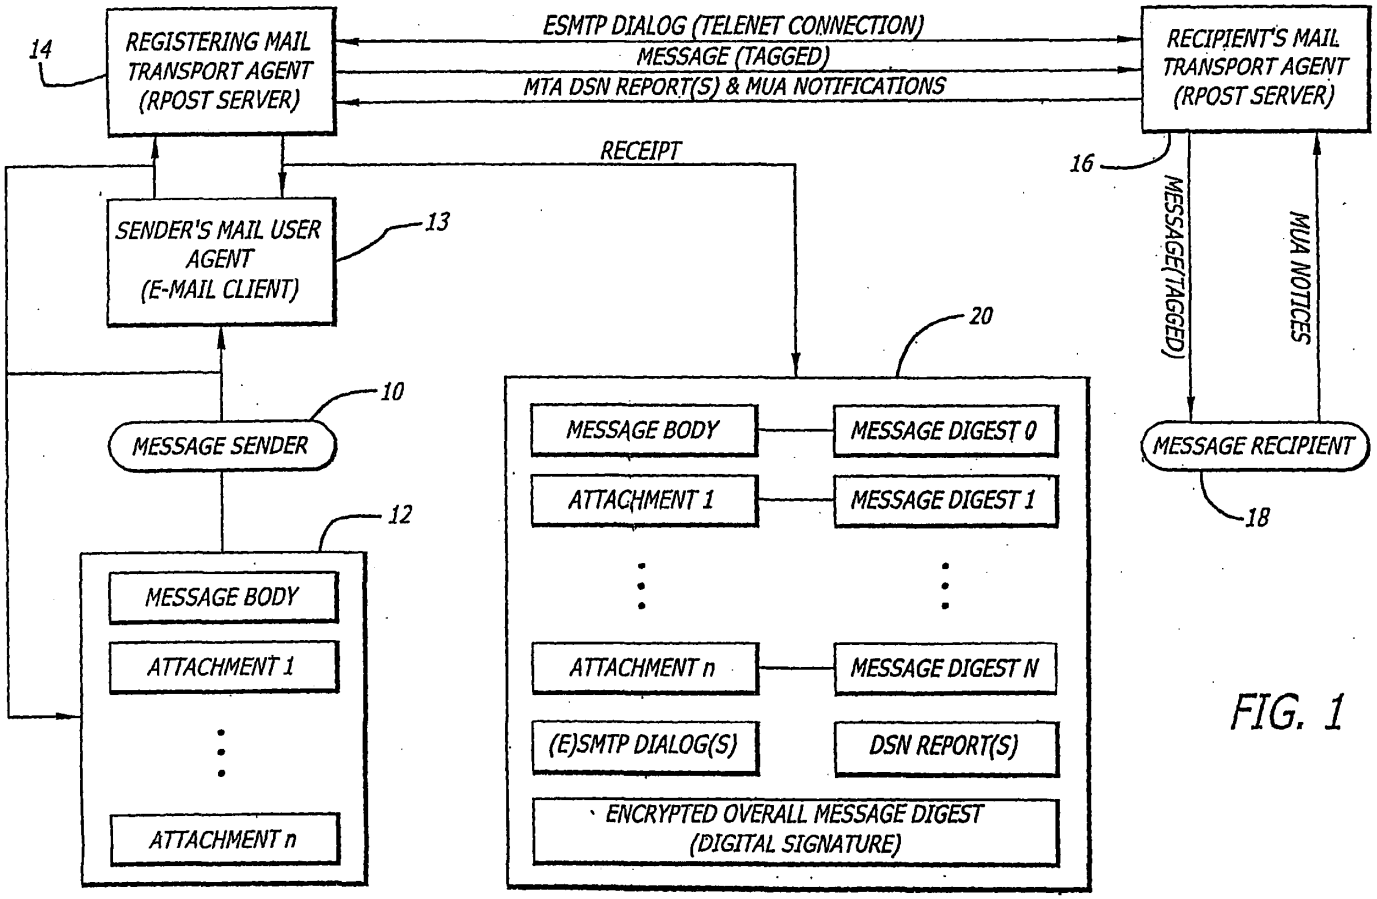
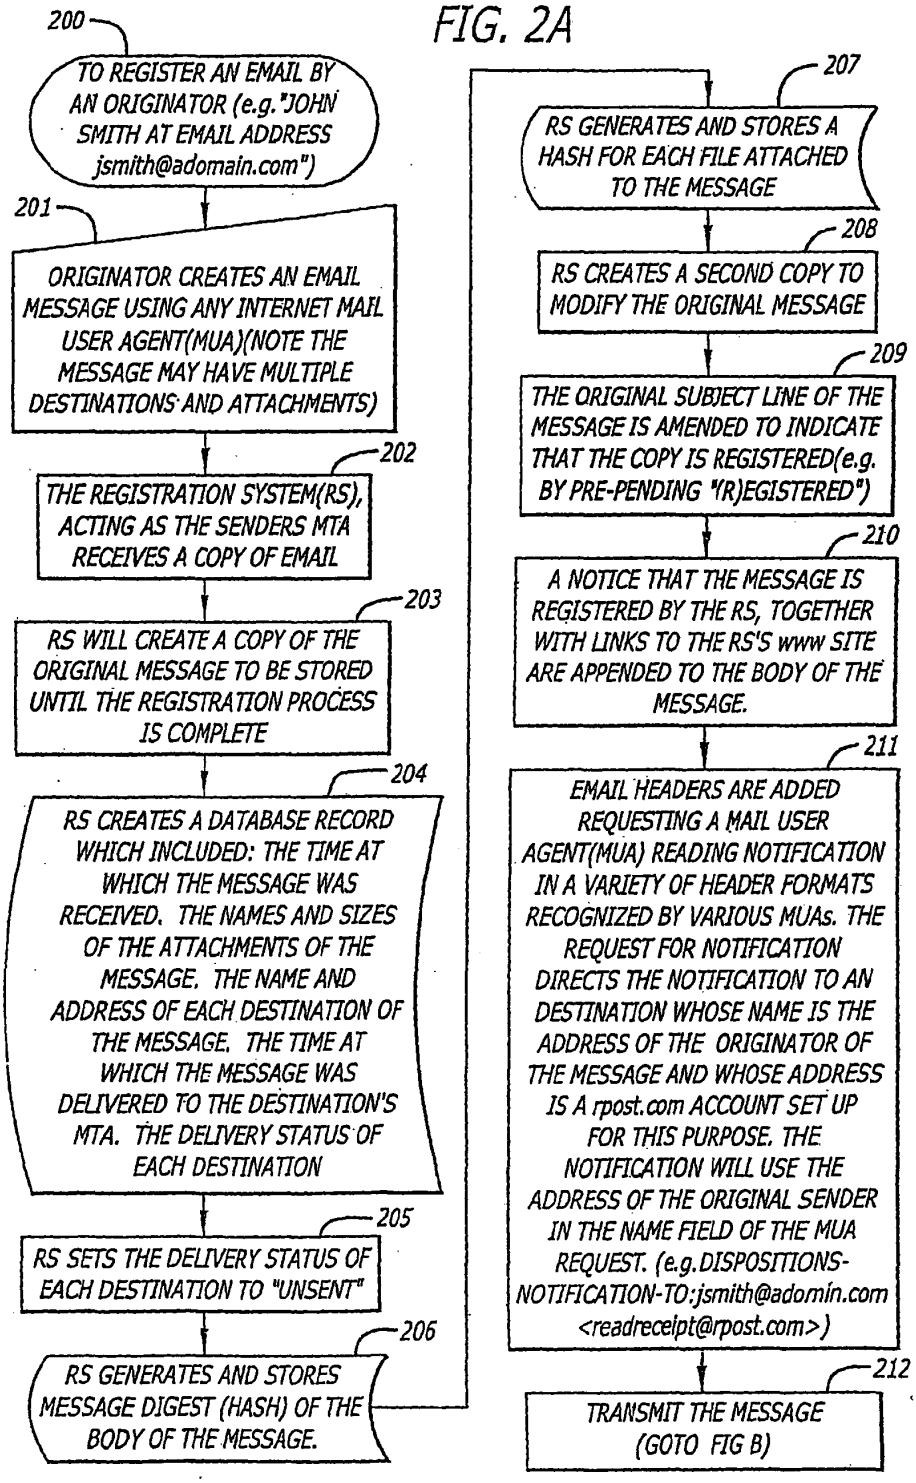

14
REGISTERING MAIL TRANSPORT AGENT (RPOST SERVER)
ESMTP DIALOG (TELENET CONNECTION)
MESSAGE (TAGGED)
MTA DSN REPORT(S) & MUA NOTIFICATIONS
RECIPIENT'S MAIL TRANSPORT AGENT (RPOST SERVER)
16
RECEIPT
MESSAGE(TAGGED)
MUA NOTICES
SENDER'S MAIL USER AGENT (E-MAIL CLIENT)
13
20
10
MESSAGE SENDER
MESSAGE BODY
MESSAGE DIGEST 0
ATTACHMENT 1
MESSAGE DIGEST 1
MESSAGE RECIPIENT
18
12
MESSAGE BODY
ATTACHMENT 1
ATTACHMENT n
MESSAGE DIGEST N
(E)SMTP DIALOG(S)
DSN REPORT(S)
ENCRYPTED OVERALL MESSAGE DIGEST (DIGITAL SIGNATURE)

FIG. 1

FIG. 2A


200
TO REGISTER AN EMAIL BY AN ORIGINATOR (e.g."JOHN SMITH AT EMAIL ADDRESS jsmith@adomain.com")
201
ORIGINATOR CREATES AN EMAIL MESSAGE USING ANY INTERNET MAIL USER AGENT(MUA)(NOTE THE MESSAGE MAY HAVE MULTIPLE DESTINATIONS AND ATTACHMENTS)
202
THE REGISTRATION SYSTEM(RS), ACTING AS THE SENDERS MTA RECEIVES A COPY OF EMAIL
203
RS WILL CREATE A COPY OF THE ORIGINAL MESSAGE TO BE STORED UNTIL THE REGISTRATION PROCESS IS COMPLETE
204
RS CREATES A DATABASE RECORD WHICH INCLUDED: THE TIME AT WHICH THE MESSAGE WAS RECEIVED. THE NAMES AND SIZES OF THE ATTACHMENTS OF THE MESSAGE. THE NAME AND ADDRESS OF EACH DESTINATION OF THE MESSAGE. THE TIME AT WHICH THE MESSAGE WAS DELIVERED TO THE DESTINATION'S MTA. THE DELIVERY STATUS OF EACH DESTINATION
205
RS SETS THE DELIVERY STATUS OF EACH DESTINATION TO "UNSENT"
206
RS GENERATES AND STORES MESSAGE DIGEST (HASH) OF THE BODY OF THE MESSAGE.
207
RS GENERATES AND STORES A HASH FOR EACH FILE ATTACHED TO THE MESSAGE
208
RS CREATES A SECOND COPY TO MODIFY THE ORIGINAL MESSAGE
209
THE ORIGINAL SUBJECT LINE OF THE MESSAGE IS AMENDED TO INDICATE THAT THE COPY IS REGISTERED(e.g. BY PRE-PENDING "(R)EGISTERED")
210
A NOTICE THAT THE MESSAGE IS REGISTERED BY THE RS, TOGETHER WITH LINKS TO THE RS'S www SITE ARE APPENDED TO THE BODY OF THE MESSAGE.
211
EMAIL HEADERS ARE ADDED REQUESTING A MAIL USER AGENT(MUA) READING NOTIFICATION IN A VARIETY OF HEADER FORMATS RECOGNIZED BY VARIOUS MUAs. THE REQUEST FOR NOTIFICATION DIRECTS THE NOTIFICATION TO AN DESTINATION WHOSE NAME IS THE ADDRESS OF THE ORIGINATOR OF THE MESSAGE AND WHOSE ADDRESS IS A rpost.com ACCOUNT SET UP FOR THIS PURPOSE. THE NOTIFICATION WILL USE THE ADDRESS OF THE ORIGINAL SENDER IN THE NAME FIELD OF THE MUA REQUEST. (e.g.DISPOSITIONS-NOTIFICATION-TO:jsmith@adomin.com <readreceipt@rpost.com>)
212
TRANSMIT THE MESSAGE (GOTO FIG B)

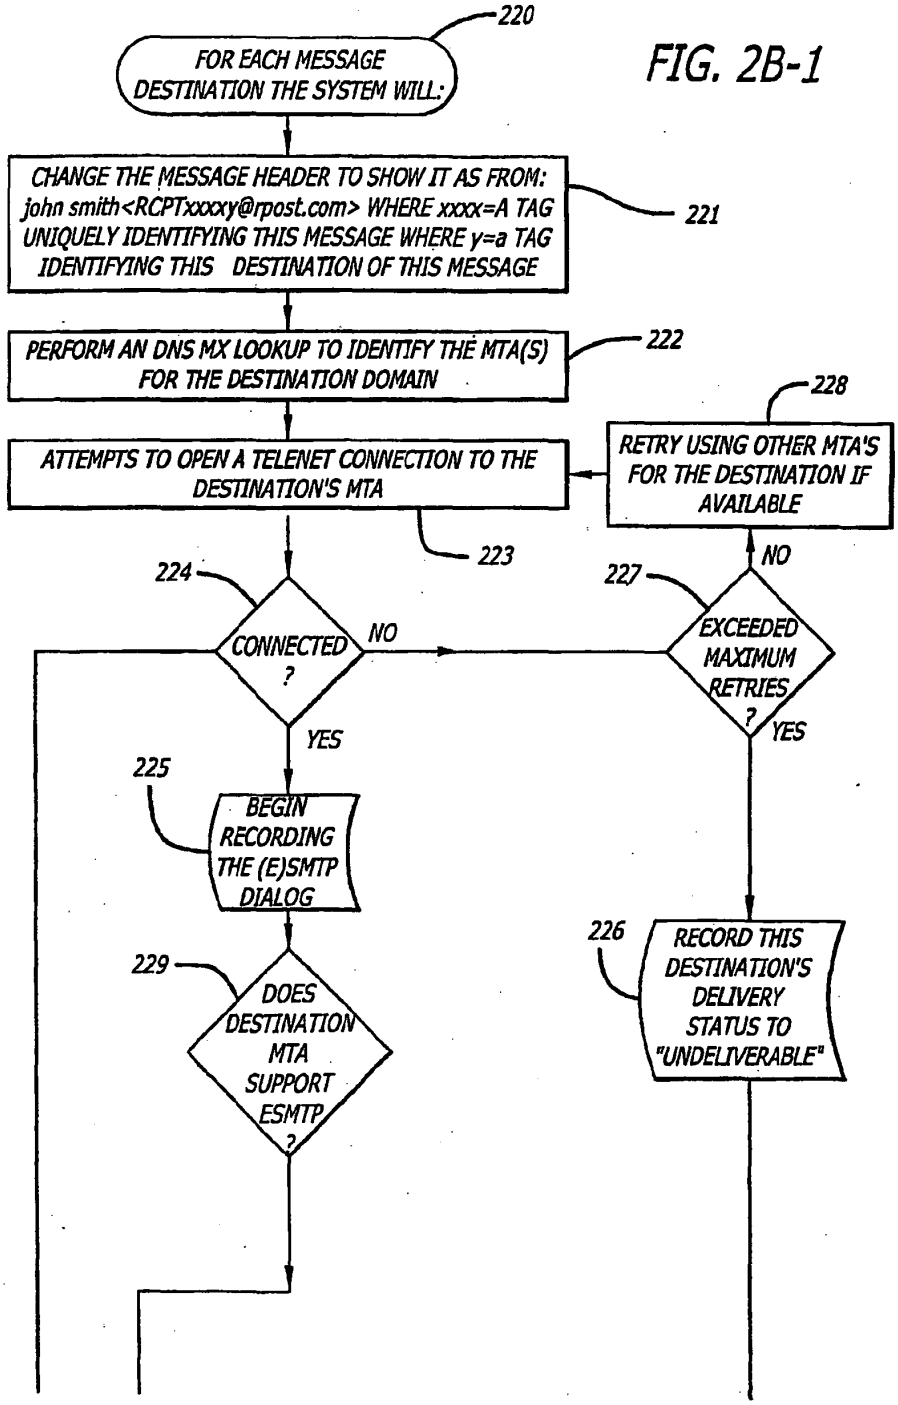

220
FOR EACH MESSAGE DESTINATION THE SYSTEM WILL:
CHANGE THE MESSAGE HEADER TO SHOW IT AS FROM: john smith<RCPTxxxxy@rpost.com> WHERE xxxx=A TAG UNIQUELY IDENTIFYING THIS MESSAGE WHERE y=a TAG IDENTIFYING THIS DESTINATION OF THIS MESSAGE
221
PERFORM AN DNS MX LOOKUP TO IDENTIFY THE MTA(S) FOR THE DESTINATION DOMAIN
222
ATTEMPTS TO OPEN A TELENET CONNECTION TO THE DESTINATION'S MTA
223
224
CONNECTED ?
NO
YES
225
BEGIN RECORDING THE (E)SMTP DIALOG
229
DOES DESTINATION MTA SUPPORT ESMTP ?
227
EXCEEDED MAXIMUM RETRIES ?
NO
YES
228
RETRY USING OTHER MTA'S FOR THE DESTINATION IF AVAILABLE
226
RECORD THIS DESTINATION'S DELIVERY STATUS TO "UNDELIVERABLE"

FIG. 2B-1

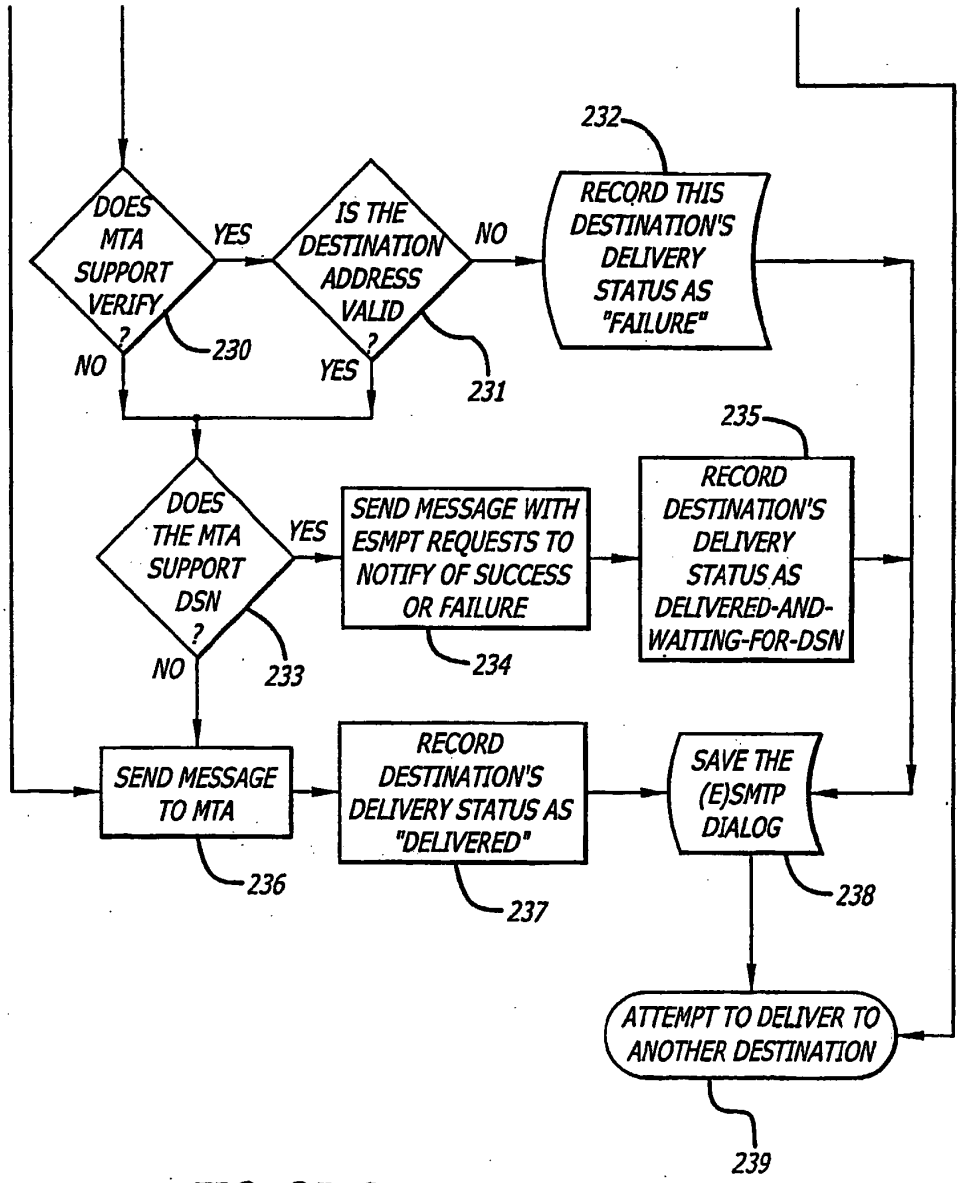

DOES MTA SUPPORT VERIFY ?
YES
NO
230
IS THE DESTINATION ADDRESS VALID ?
YES
NO
231
232
RECORD THIS DESTINATION'S DELIVERY STATUS AS "FAILURE"
DOES THE MTA SUPPORT DSN ?
YES
NO
233
SEND MESSAGE WITH ESMPT REQUESTS TO NOTIFY OF SUCCESS OR FAILURE
234
235
RECORD DESTINATION'S DELIVERY STATUS AS DELIVERED-AND-WAITING-FOR-DSN
SEND MESSAGE TO MTA
236
RECORD DESTINATION'S DELIVERY STATUS AS "DELIVERED"
237
SAVE THE (E)SMTP DIALOG
238
ATTEMPT TO DELIVER TO ANOTHER DESTINATION
239

FIG. 2B-2

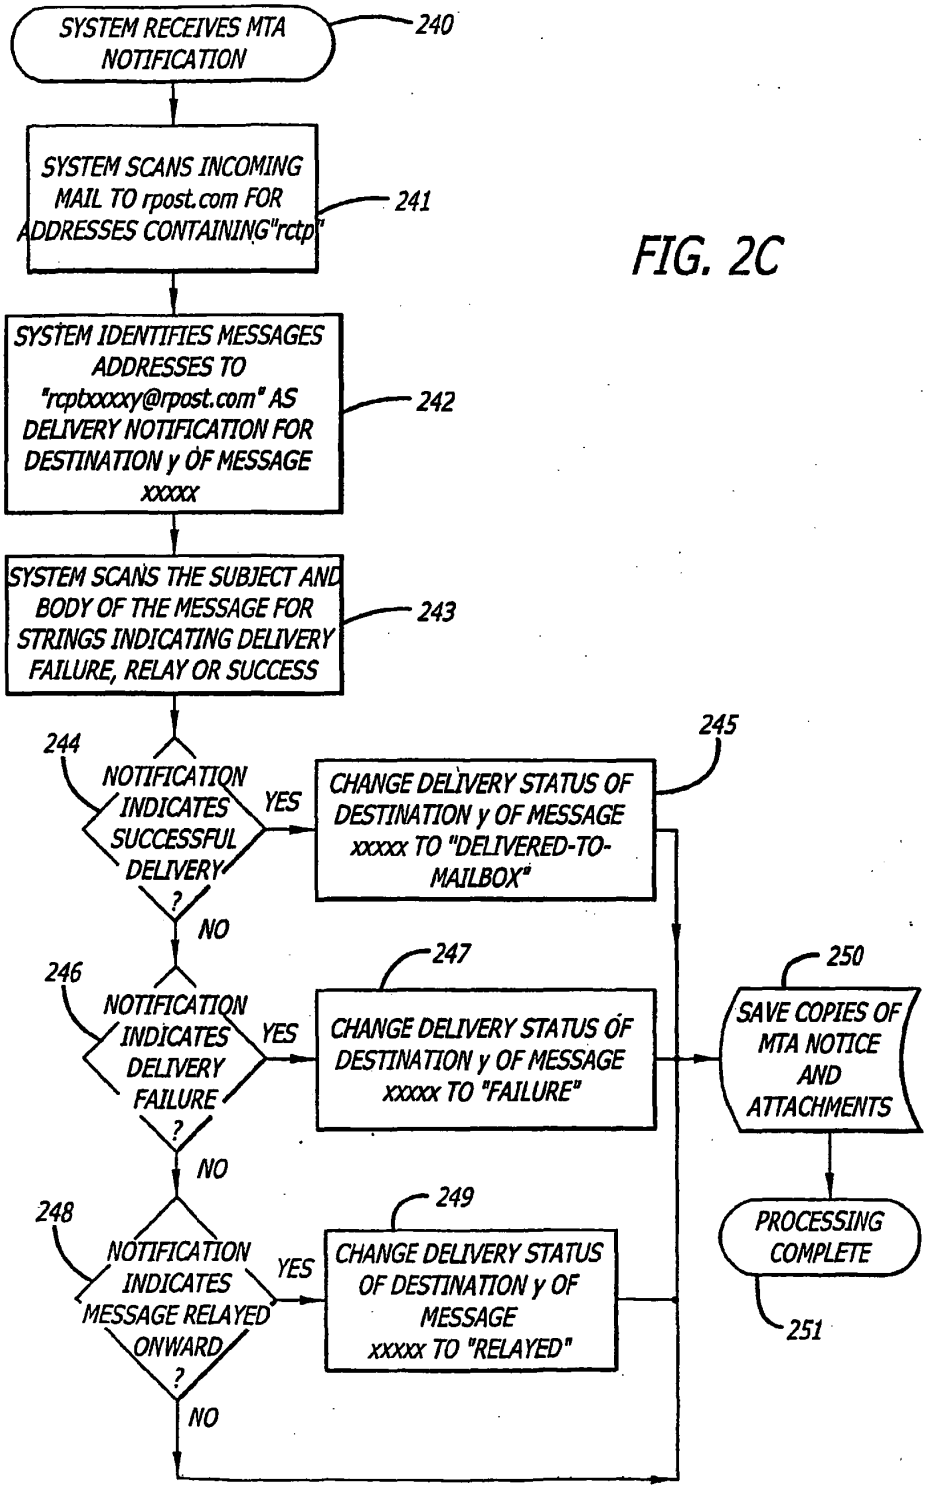

SYSTEM RECEIVES MTA NOTIFICATION
240
SYSTEM SCANS INCOMING MAIL TO rpost.com FOR ADDRESSES CONTAINING"rctp"
241
SYSTEM IDENTIFIES MESSAGES ADDRESSES TO "rcptxxxxy@rpost.com" AS DELIVERY NOTIFICATION FOR DESTINATION y OF MESSAGE xxxxx
242
SYSTEM SCANS THE SUBJECT AND BODY OF THE MESSAGE FOR STRINGS INDICATING DELIVERY FAILURE, RELAY OR SUCCESS
243
244
NOTIFICATION INDICATES SUCCESSFUL DELIVERY ?
YES
CHANGE DELIVERY STATUS OF DESTINATION y OF MESSAGE xxxxx TO "DELIVERED-TO-MAILBOX"
245
NO
246
NOTIFICATION INDICATES DELIVERY FAILURE ?
YES
247
CHANGE DELIVERY STATUS OF DESTINATION y OF MESSAGE xxxxx TO "FAILURE"
250
SAVE COPIES OF MTA NOTICE AND ATTACHMENTS
NO
248
NOTIFICATION INDICATES MESSAGE RELAYED ONWARD ?
YES
249
CHANGE DELIVERY STATUS OF DESTINATION y OF MESSAGE xxxxx TO "RELAYED"
PROCESSING COMPLETE
251
NO

FIG. 2C

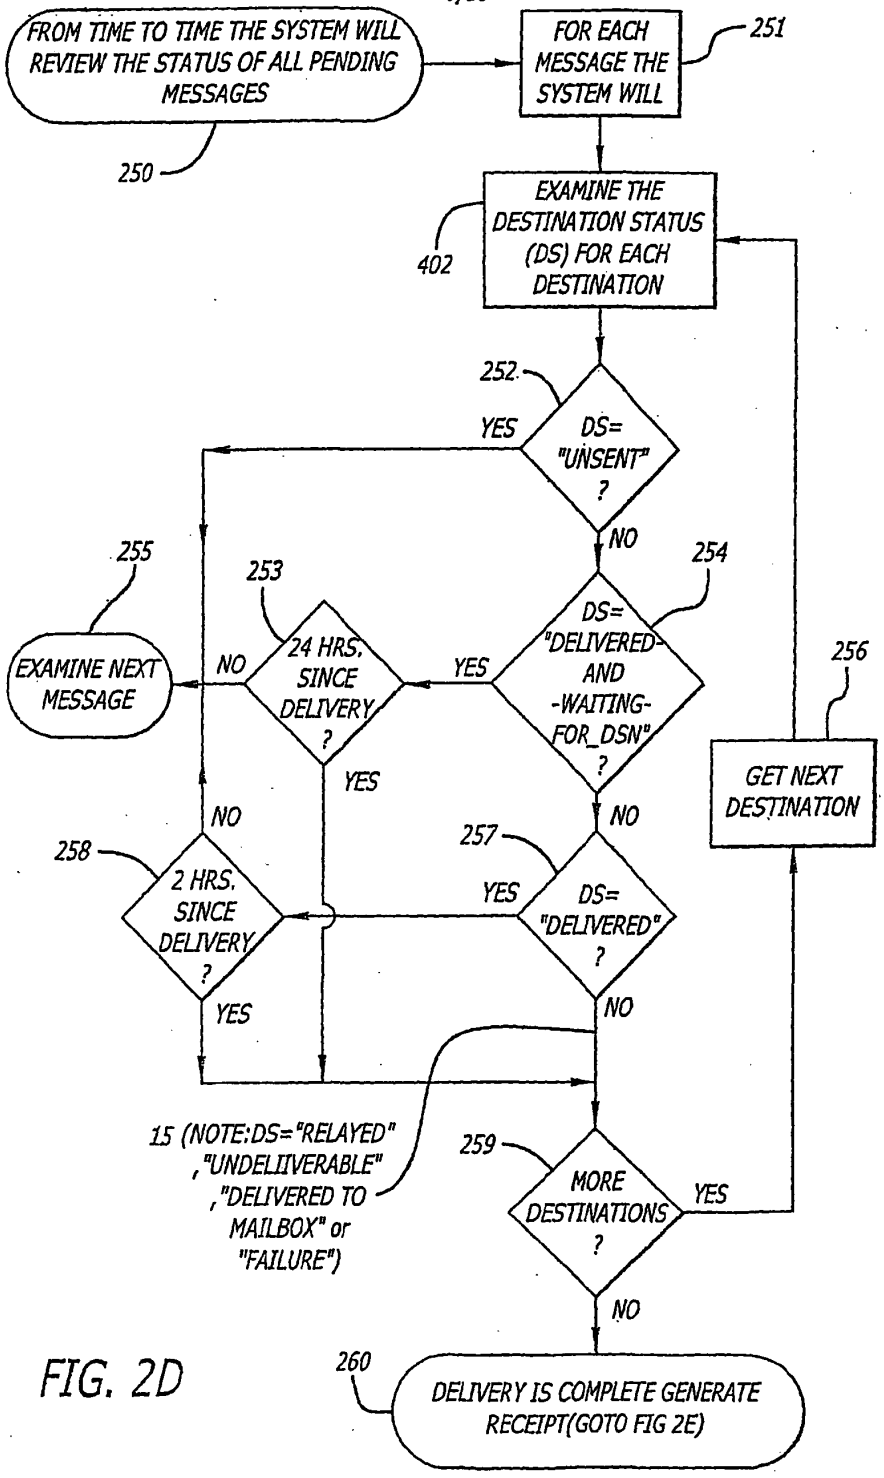
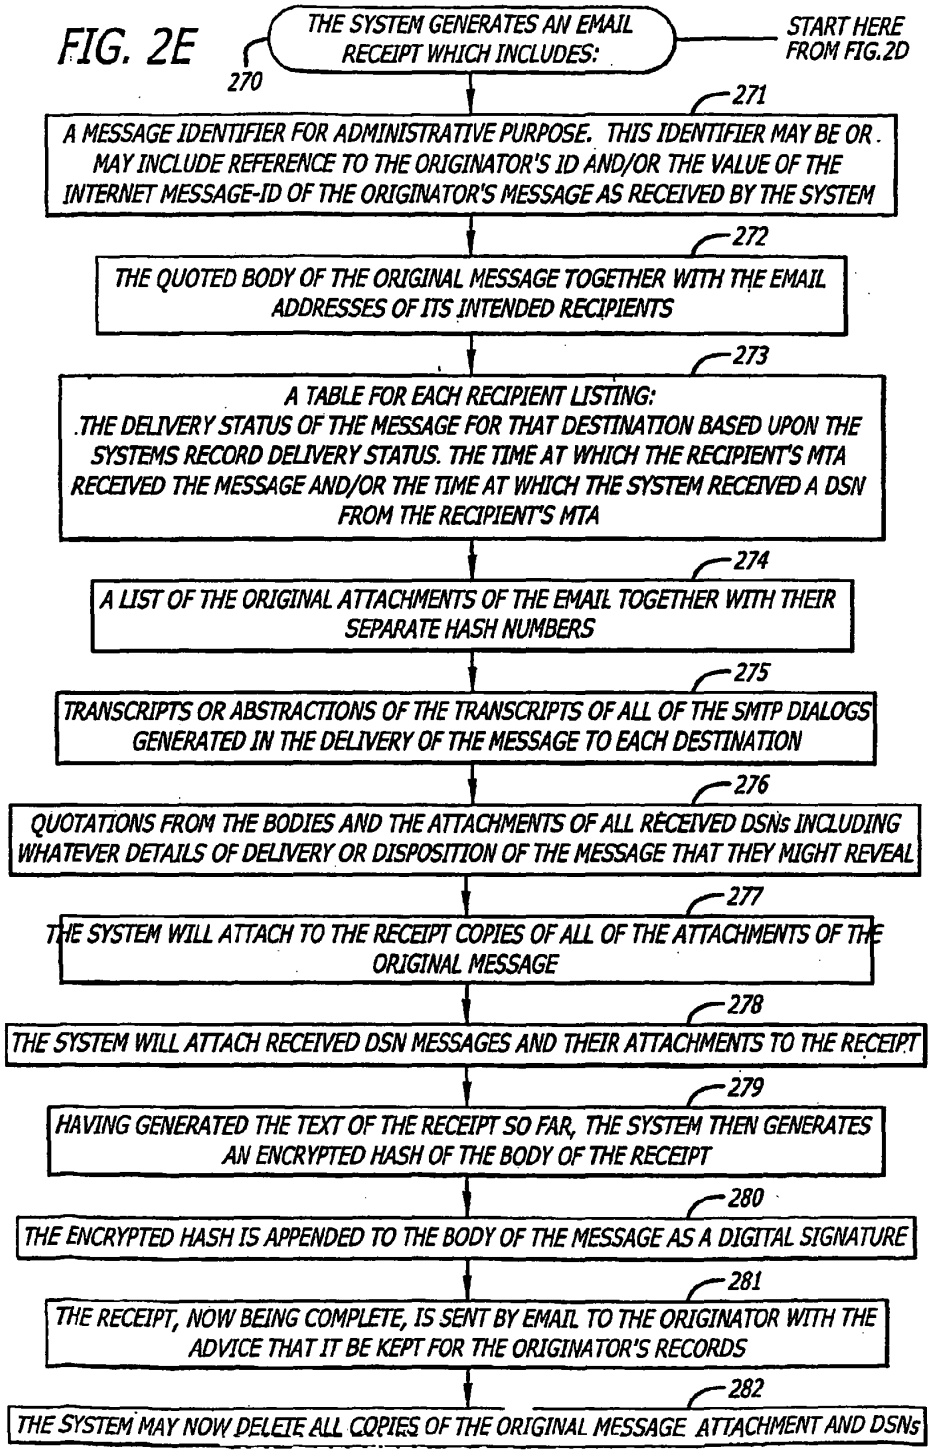

FROM TIME TO TIME THE SYSTEM WILL REVIEW THE STATUS OF ALL PENDING MESSAGES
250
FOR EACH MESSAGE THE SYSTEM WILL
251
EXAMINE THE DESTINATION STATUS (DS) FOR EACH DESTINATION
402
252
DS= "UNSENT" ?
YES
NO
254
DS= "DELIVERED-AND -WAITING-FOR_DSN" ?
253
24 HRS. SINCE DELIVERY ?
255
EXAMINE NEXT MESSAGE
256
GET NEXT DESTINATION
257
DS= "DELIVERED" ?
258
2 HRS. SINCE DELIVERY ?
15 (NOTE:DS="RELAYED" ,"UNDELIIVERABLE" ,"DELIVERED TO MAILBOX" or "FAILURE")
259
MORE DESTINATIONS ?
260
DELIVERY IS COMPLETE GENERATE RECEIPT(GOTO FIG 2E)

FIG. 2D

FIG. 2E


THE SYSTEM GENERATES AN EMAIL RECEIPT WHICH INCLUDES:
START HERE FROM FIG.2D
270
271
A MESSAGE IDENTIFIER FOR ADMINISTRATIVE PURPOSE. THIS IDENTIFIER MAY BE OR MAY INCLUDE REFERENCE TO THE ORIGINATOR'S ID AND/OR THE VALUE OF THE INTERNET MESSAGE-ID OF THE ORIGINATOR'S MESSAGE AS RECEIVED BY THE SYSTEM
272
THE QUOTED BODY OF THE ORIGINAL MESSAGE TOGETHER WITH THE EMAIL ADDRESSES OF ITS INTENDED RECIPIENTS
273
A TABLE FOR EACH RECIPIENT LISTING:
.THE DELIVERY STATUS OF THE MESSAGE FOR THAT DESTINATION BASED UPON THE SYSTEMS RECORD DELIVERY STATUS. THE TIME AT WHICH THE RECIPIENT'S MTA RECEIVED THE MESSAGE AND/OR THE TIME AT WHICH THE SYSTEM RECEIVED A DSN FROM THE RECIPIENT'S MTA
274
A LIST OF THE ORIGINAL ATTACHMENTS OF THE EMAIL TOGETHER WITH THEIR SEPARATE HASH NUMBERS
275
TRANSCRIPTS OR ABSTRACTIONS OF THE TRANSCRIPTS OF ALL OF THE SMTP DIALOGS GENERATED IN THE DELIVERY OF THE MESSAGE TO EACH DESTINATION
276
QUOTATIONS FROM THE BODIES AND THE ATTACHMENTS OF ALL RECEIVED DSNs INCLUDING WHATEVER DETAILS OF DELIVERY OR DISPOSITION OF THE MESSAGE THAT THEY MIGHT REVEAL
277
THE SYSTEM WILL ATTACH TO THE RECEIPT COPIES OF ALL OF THE ATTACHMENTS OF THE ORIGINAL MESSAGE
278
THE SYSTEM WILL ATTACH RECEIVED DSN MESSAGES AND THEIR ATTACHMENTS TO THE RECEIPT
279
HAVING GENERATED THE TEXT OF THE RECEIPT SO FAR, THE SYSTEM THEN GENERATES AN ENCRYPTED HASH OF THE BODY OF THE RECEIPT
280
THE ENCRYPTED HASH IS APPENDED TO THE BODY OF THE MESSAGE AS A DIGITAL SIGNATURE
281
THE RECEIPT, NOW BEING COMPLETE, IS SENT BY EMAIL TO THE ORIGINATOR WITH THE ADVICE THAT IT BE KEPT FOR THE ORIGINATOR'S RECORDS
282
THE SYSTEM MAY NOW DELETE ALL COPIES OF THE ORIGINAL MESSAGE ATTACHMENT AND DSNs

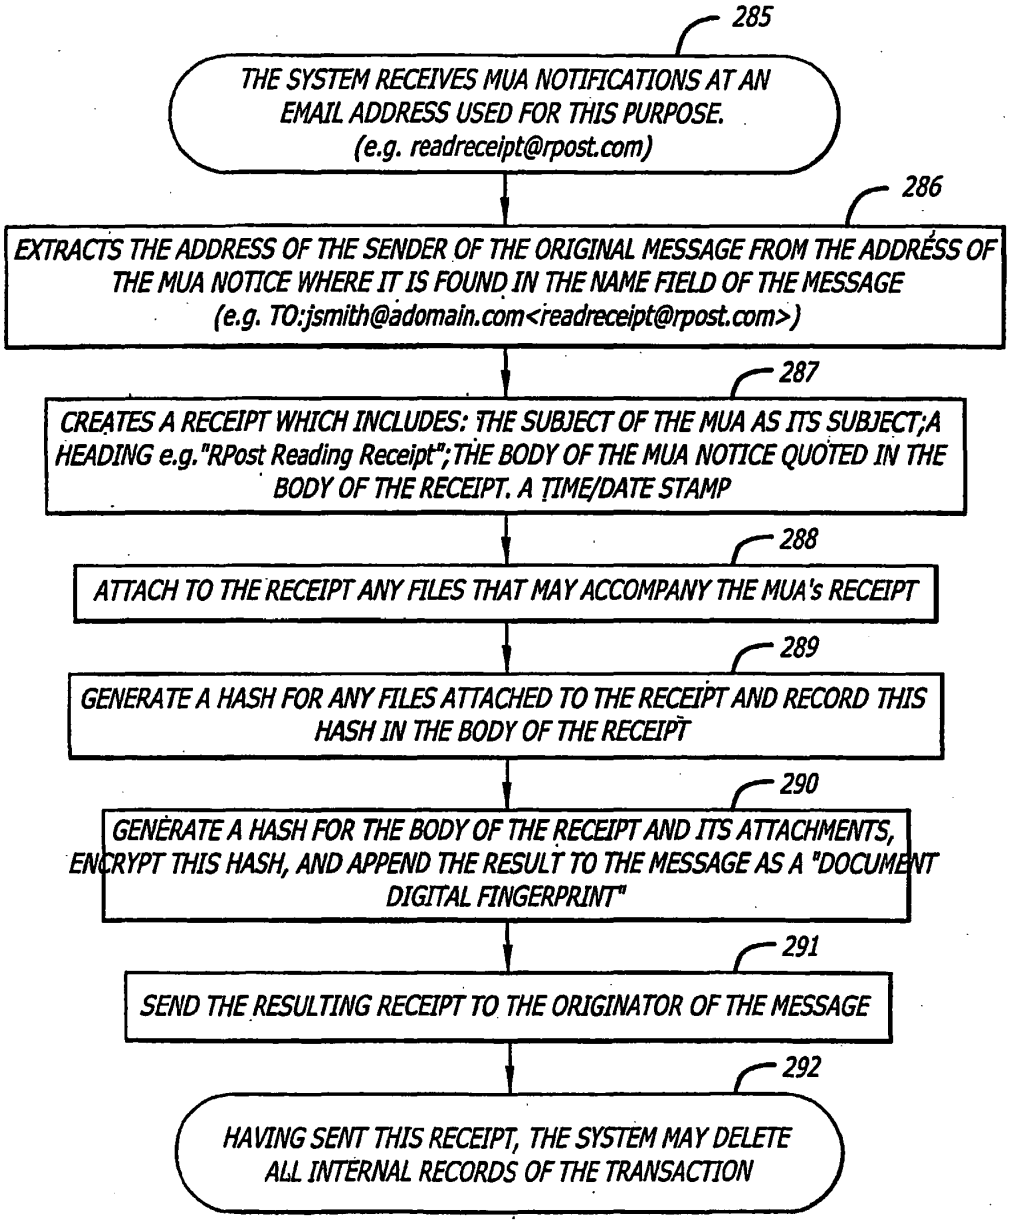

285
THE SYSTEM RECEIVES MUA NOTIFICATIONS AT AN EMAIL ADDRESS USED FOR THIS PURPOSE. (e.g. readreceipt@rpost.com)
286
EXTRACTS THE ADDRESS OF THE SENDER OF THE ORIGINAL MESSAGE FROM THE ADDRESS OF THE MUA NOTICE WHERE IT IS FOUND IN THE NAME FIELD OF THE MESSAGE (e.g. TO:jsmith@adomain.com<readreceipt@rpost.com>)
287
CREATES A RECEIPT WHICH INCLUDES: THE SUBJECT OF THE MUA AS ITS SUBJECT;A HEADING e.g."RPost Reading Receipt";THE BODY OF THE MUA NOTICE QUOTED IN THE BODY OF THE RECEIPT. A TIME/DATE STAMP
288
ATTACH TO THE RECEIPT ANY FILES THAT MAY ACCOMPANY THE MUA's RECEIPT
289
GENERATE A HASH FOR ANY FILES ATTACHED TO THE RECEIPT AND RECORD THIS HASH IN THE BODY OF THE RECEIPT
290
GENERATE A HASH FOR THE BODY OF THE RECEIPT AND ITS ATTACHMENTS, ENCRYPT THIS HASH, AND APPEND THE RESULT TO THE MESSAGE AS A "DOCUMENT DIGITAL FINGERPRINT"
291
SEND THE RESULTING RECEIPT TO THE ORIGINATOR OF THE MESSAGE
292
HAVING SENT THIS RECEIPT, THE SYSTEM MAY DELETE ALL INTERNAL RECORDS OF THE TRANSACTION

FIG. 2F

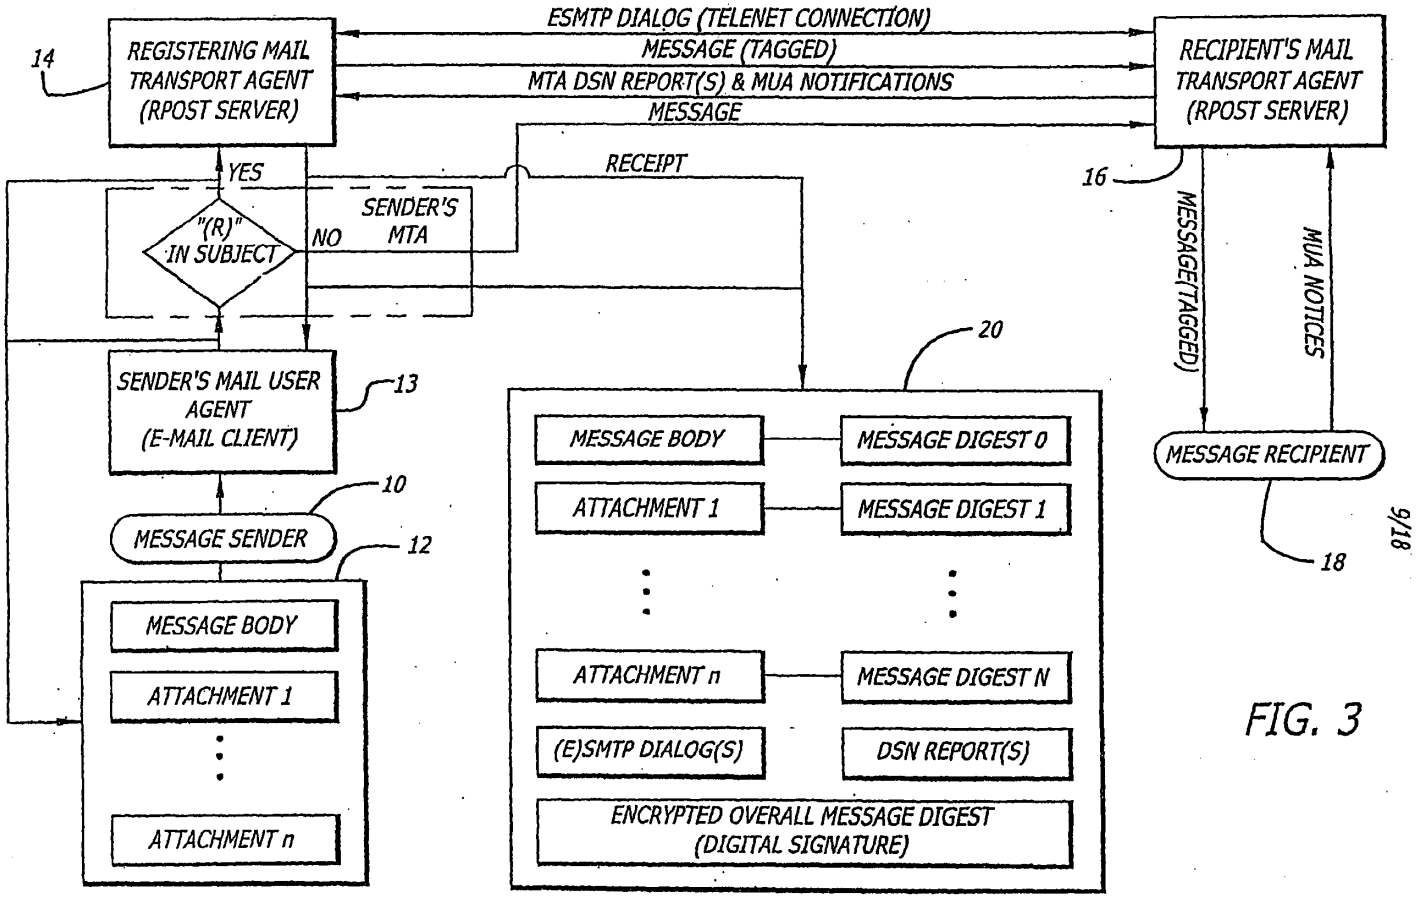

ESMTP DIALOG (TELENET CONNECTION)
MESSAGE (TAGGED)
MTA DSN REPORT(S) & MUA NOTIFICATIONS
MESSAGE
RECEIPT
14
REGISTERING MAIL TRANSPORT AGENT (RPOST SERVER)
RECIPIENT'S MAIL TRANSPORT AGENT (RPOST SERVER)
16
YES
NO
"(R)" IN SUBJECT
SENDER'S MTA
MESSAGE(TAGGED)
MUA NOTICES
20
SENDER'S MAIL USER AGENT (E-MAIL CLIENT)
13
MESSAGE BODY
MESSAGE DIGEST 0
ATTACHMENT 1
MESSAGE DIGEST 1
MESSAGE RECIPIENT
10
MESSAGE SENDER
12
18
MESSAGE BODY
ATTACHMENT 1
ATTACHMENT n
MESSAGE DIGEST N
(E)SMTP DIALOG(S)
DSN REPORT(S)
ENCRYPTED OVERALL MESSAGE DIGEST (DIGITAL SIGNATURE)

FIG. 3

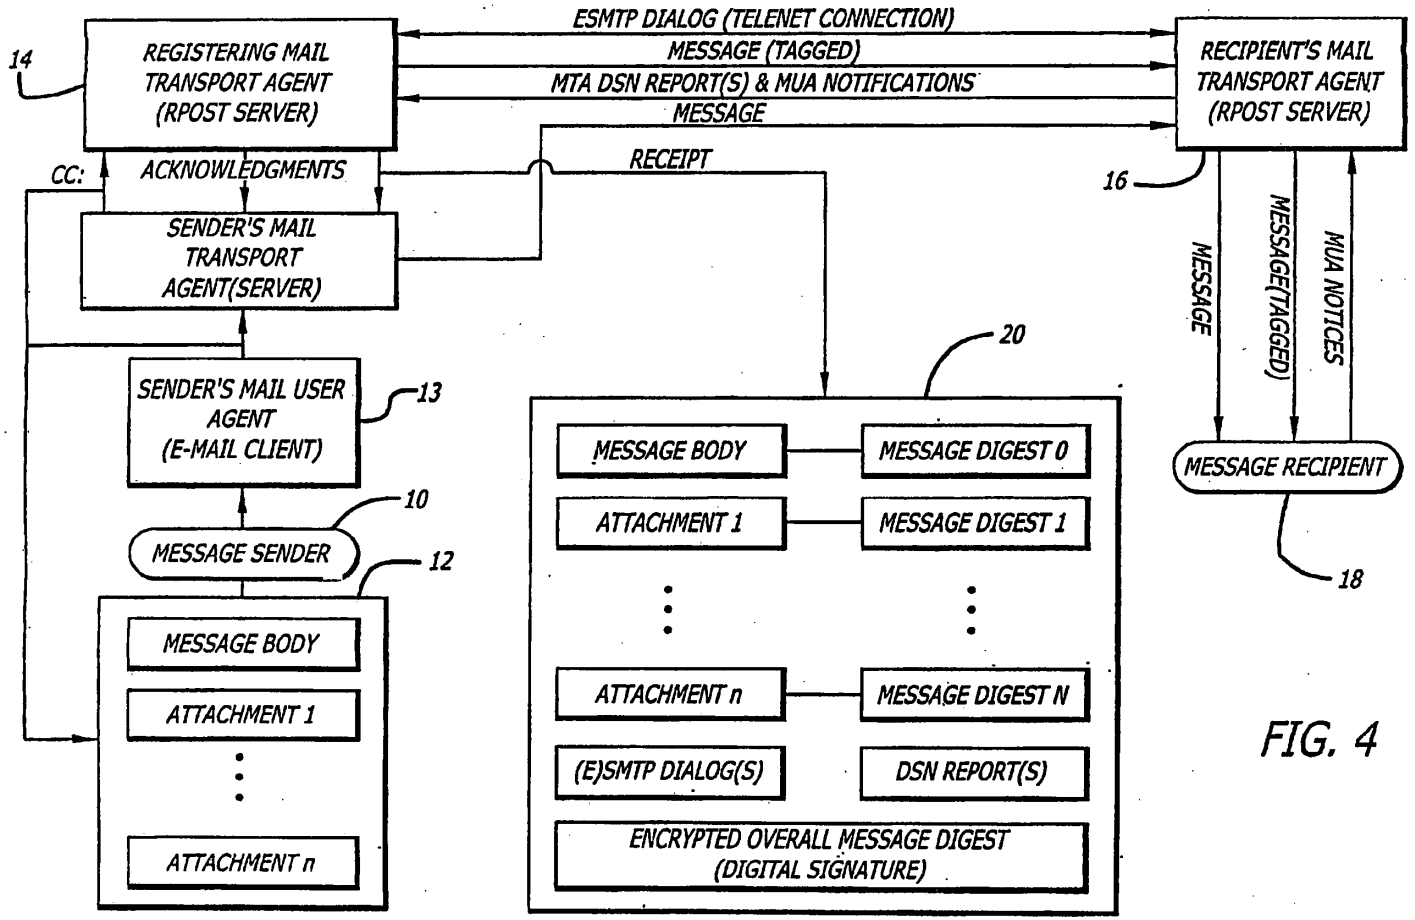

ESMTP DIALOG (TELENET CONNECTION)
MESSAGE (TAGGED)
MTA DSN REPORT(S) & MUA NOTIFICATIONS
MESSAGE
RECEIPT
14
REGISTERING MAIL TRANSPORT AGENT (RPOST SERVER)
RECIPIENT'S MAIL TRANSPORT AGENT (RPOST SERVER)
16
CC:
ACKNOWLEDGMENTS
SENDER'S MAIL TRANSPORT AGENT(SERVER)
MESSAGE
MESSAGE(TAGGED)
MUA NOTICES
20
SENDER'S MAIL USER AGENT (E-MAIL CLIENT)
13
MESSAGE BODY
MESSAGE DIGEST 0
ATTACHMENT 1
MESSAGE DIGEST 1
10
MESSAGE SENDER
MESSAGE RECIPIENT
12
18
MESSAGE BODY
ATTACHMENT 1
ATTACHMENT n
ATTACHMENT n
MESSAGE DIGEST N
(E)SMTP DIALOG(S)
DSN REPORT(S)
ENCRYPTED OVERALL MESSAGE DIGEST (DIGITAL SIGNATURE)

FIG. 4

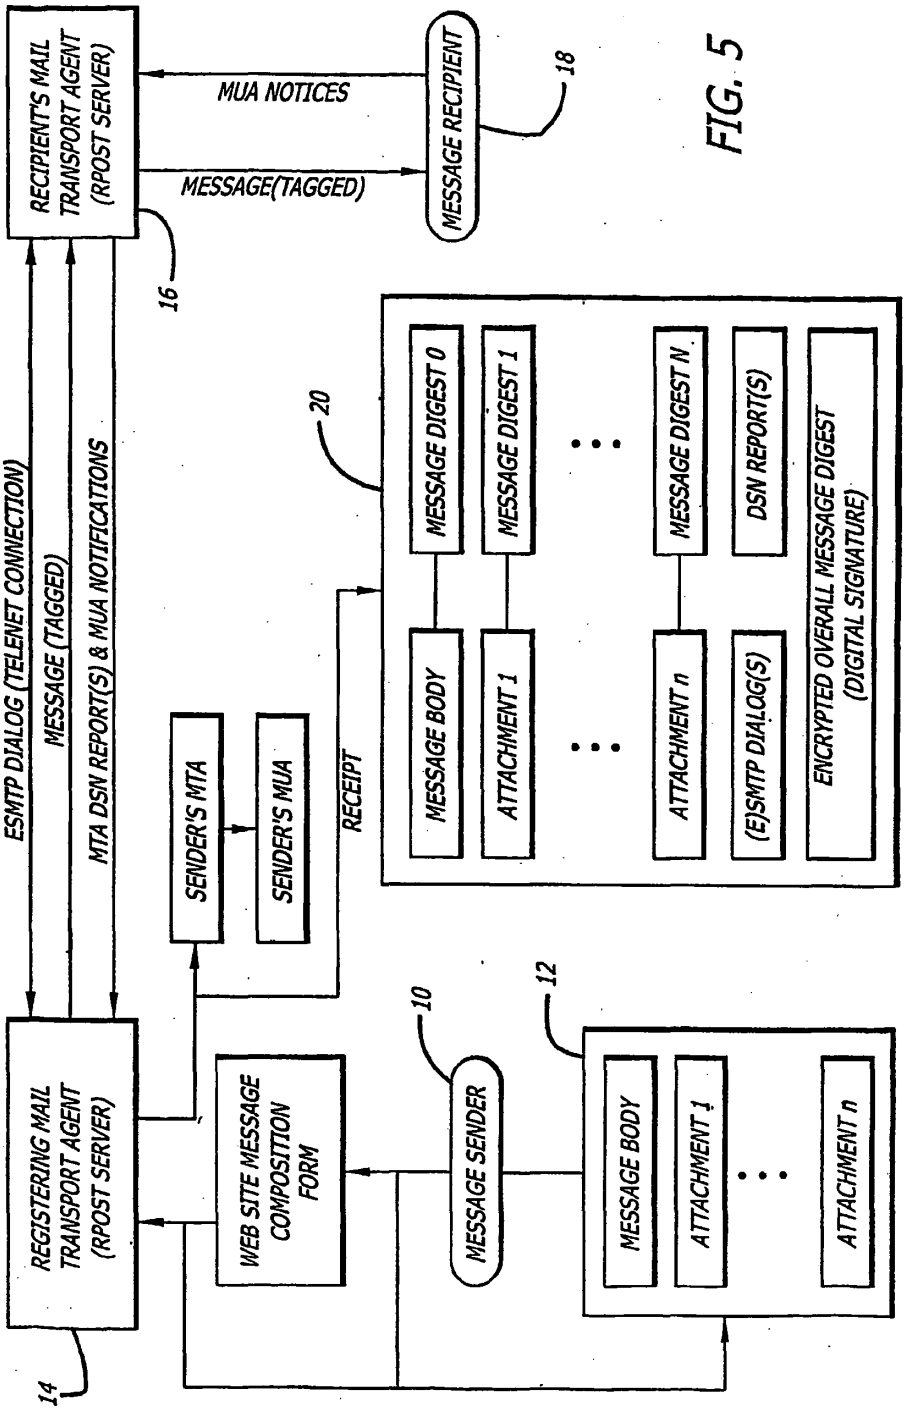

RECIPIENT'S MAIL TRANSPORT AGENT (RPOST SERVER)
MUA NOTICES
MESSAGE(TAGGED)
MESSAGE RECIPIENT
18
16
ESMTP DIALOG (TELENET CONNECTION)
MESSAGE (TAGGED)
MTA DSN REPORT(S) & MUA NOTIFICATIONS
20
MESSAGE DIGEST 0
MESSAGE DIGEST 1
MESSAGE DIGEST N
DSN REPORT(S)
MESSAGE BODY
ATTACHMENT 1
ATTACHMENT n
(E)SMTP DIALOG(S)
ENCRYPTED OVERALL MESSAGE DIGEST (DIGITAL SIGNATURE)
SENDER'S MTA
SENDER'S MUA
RECEIPT
REGISTERING MAIL TRANSPORT AGENT (RPOST SERVER)
WEB SITE MESSAGE COMPOSITION FORM
MESSAGE SENDER
10
12
MESSAGE BODY
ATTACHMENT 1
ATTACHMENT n
14

FIG. 5

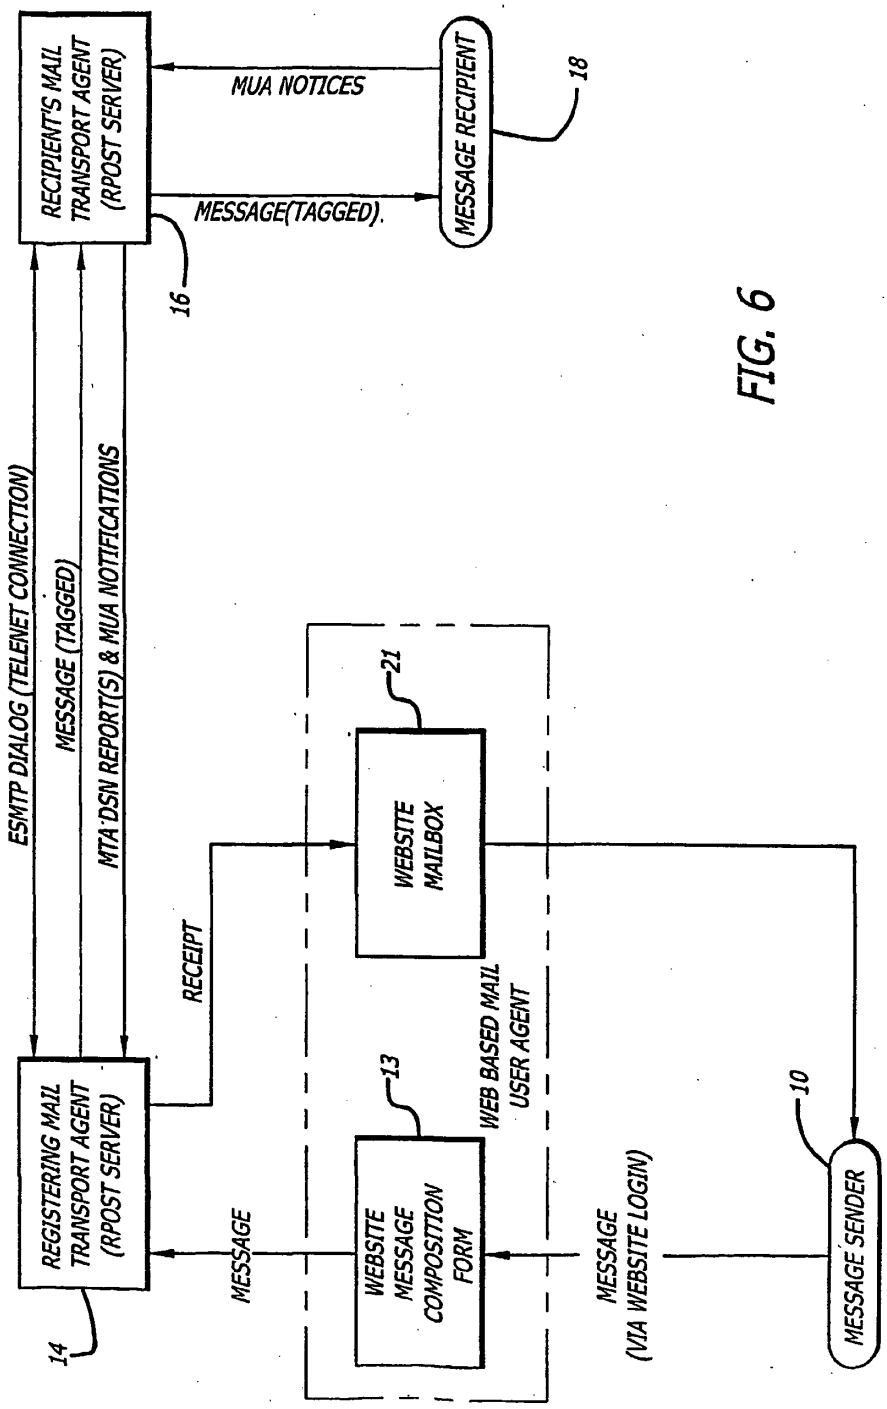

RECIPIENT'S MAIL TRANSPORT AGENT (RPOST SERVER)
MUA NOTICES
MESSAGE RECIPIENT
18
MESSAGE(TAGGED).
16
ESMTP DIALOG (TELENET CONNECTION)
MESSAGE (TAGGED)
MTA DSN REPORT(S) & MUA NOTIFICATIONS
21
WEBSITE MAILBOX
RECEIPT
WEB BASED MAIL USER AGENT
13
REGISTERING MAIL TRANSPORT AGENT (RPOST SERVER)
WEBSITE MESSAGE COMPOSITION FORM
10
MESSAGE
MESSAGE (VIA WEBSITE LOGIN)
MESSAGE SENDER
14

FIG. 6

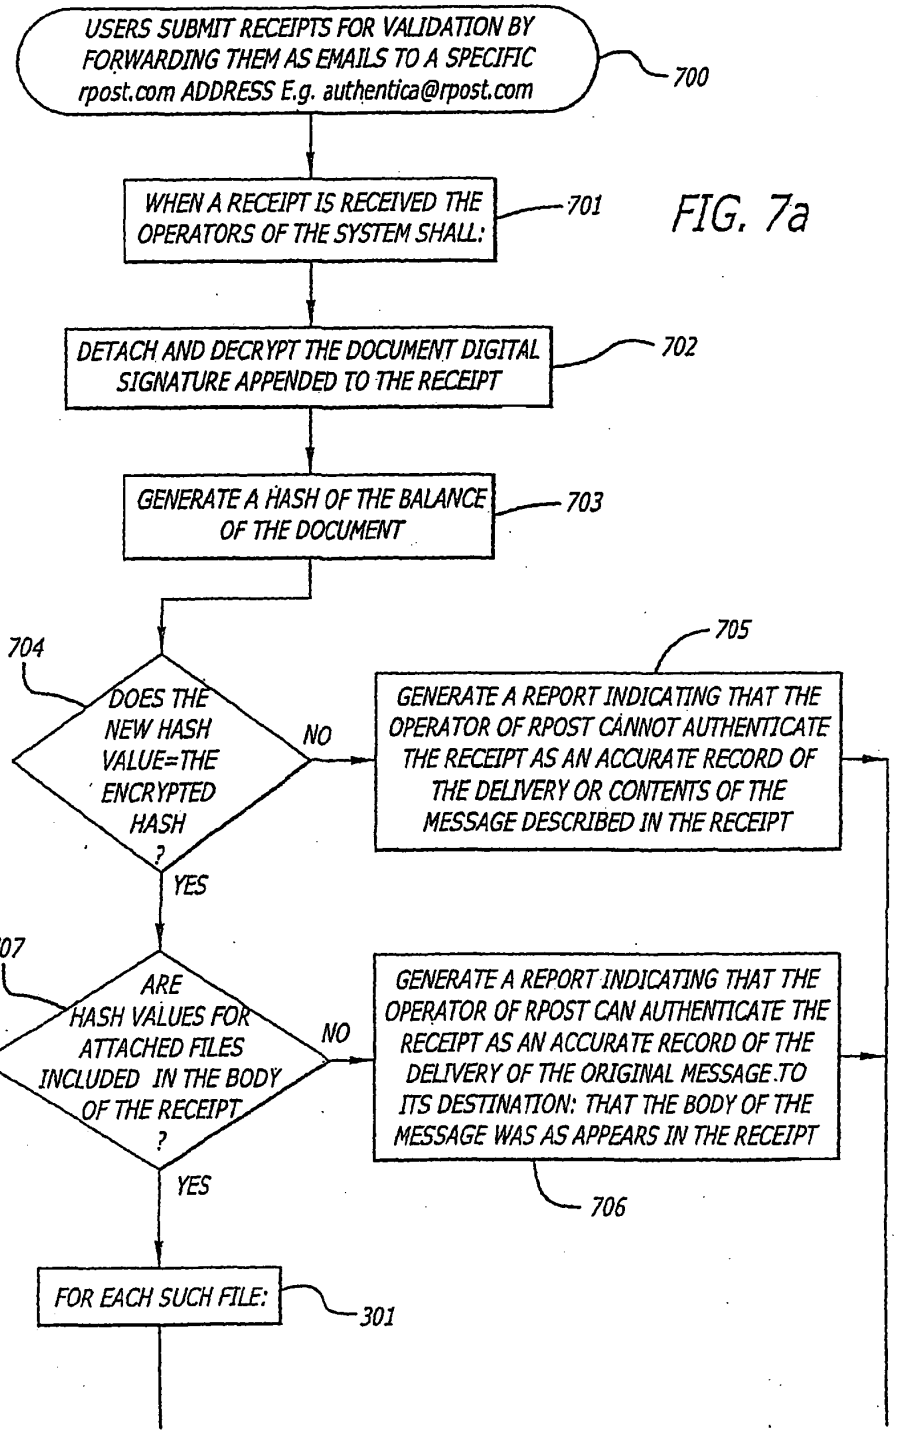

USERS SUBMIT RECEIPTS FOR VALIDATION BY FORWARDING THEM AS EMAILS TO A SPECIFIC rpost.com ADDRESS E.g. authentica@rpost.com
700
WHEN A RECEIPT IS RECEIVED THE OPERATORS OF THE SYSTEM SHALL:
701
DETACH AND DECRYPT THE DOCUMENT DIGITAL SIGNATURE APPENDED TO THE RECEIPT
702
GENERATE A HASH OF THE BALANCE OF THE DOCUMENT
703
704
DOES THE NEW HASH VALUE=THE ENCRYPTED HASH ?
NO
705
GENERATE A REPORT INDICATING THAT THE OPERATOR OF RPOST CANNOT AUTHENTICATE THE RECEIPT AS AN ACCURATE RECORD OF THE DELIVERY OR CONTENTS OF THE MESSAGE DESCRIBED IN THE RECEIPT
YES
707
ARE HASH VALUES FOR ATTACHED FILES INCLUDED IN THE BODY OF THE RECEIPT ?
NO
GENERATE A REPORT INDICATING THAT THE OPERATOR OF RPOST CAN AUTHENTICATE THE RECEIPT AS AN ACCURATE RECORD OF THE DELIVERY OF THE ORIGINAL MESSAGE.TO ITS DESTINATION: THAT THE BODY OF THE MESSAGE WAS AS APPEARS IN THE RECEIPT
706
YES
FOR EACH SUCH FILE:
301

FIG. 7a

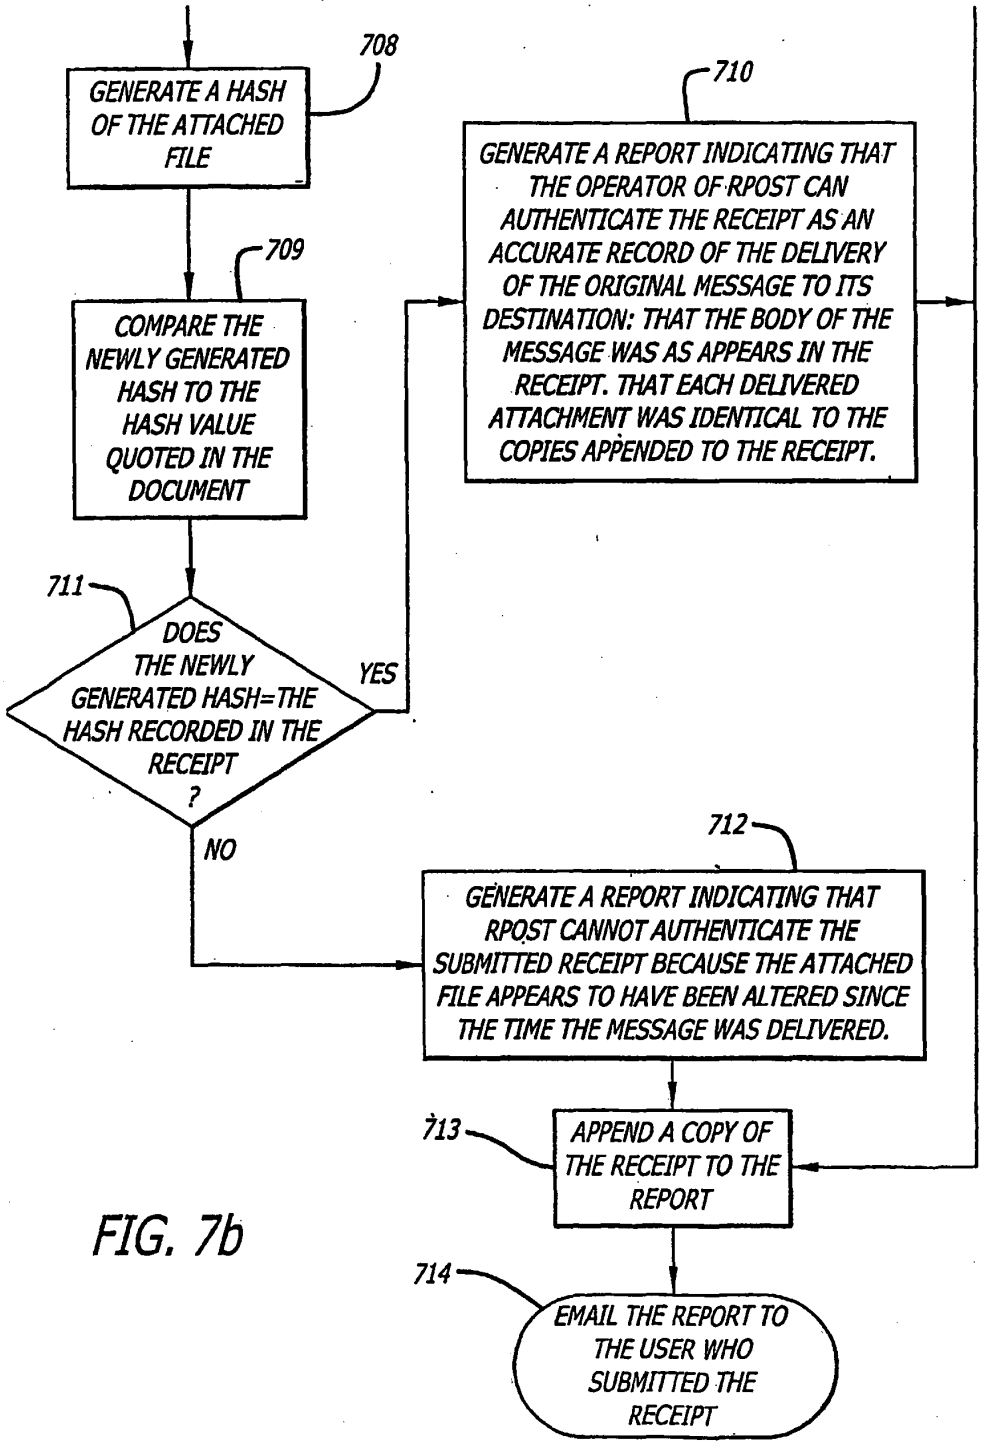

708
GENERATE A HASH OF THE ATTACHED FILE
709
COMPARE THE NEWLY GENERATED HASH TO THE HASH VALUE QUOTED IN THE DOCUMENT
711
DOES THE NEWLY GENERATED HASH=THE HASH RECORDED IN THE RECEIPT ?
YES
NO
710
GENERATE A REPORT INDICATING THAT THE OPERATOR OF RPOST CAN AUTHENTICATE THE RECEIPT AS AN ACCURATE RECORD OF THE DELIVERY OF THE ORIGINAL MESSAGE TO ITS DESTINATION: THAT THE BODY OF THE MESSAGE WAS AS APPEARS IN THE RECEIPT. THAT EACH DELIVERED ATTACHMENT WAS IDENTICAL TO THE COPIES APPENDED TO THE RECEIPT.
712
GENERATE A REPORT INDICATING THAT RPOST CANNOT AUTHENTICATE THE SUBMITTED RECEIPT BECAUSE THE ATTACHED FILE APPEARS TO HAVE BEEN ALTERED SINCE THE TIME THE MESSAGE WAS DELIVERED.
713
APPEND A COPY OF THE RECEIPT TO THE REPORT
714
EMAIL THE REPORT TO THE USER WHO SUBMITTED THE RECEIPT

FIG. 7b

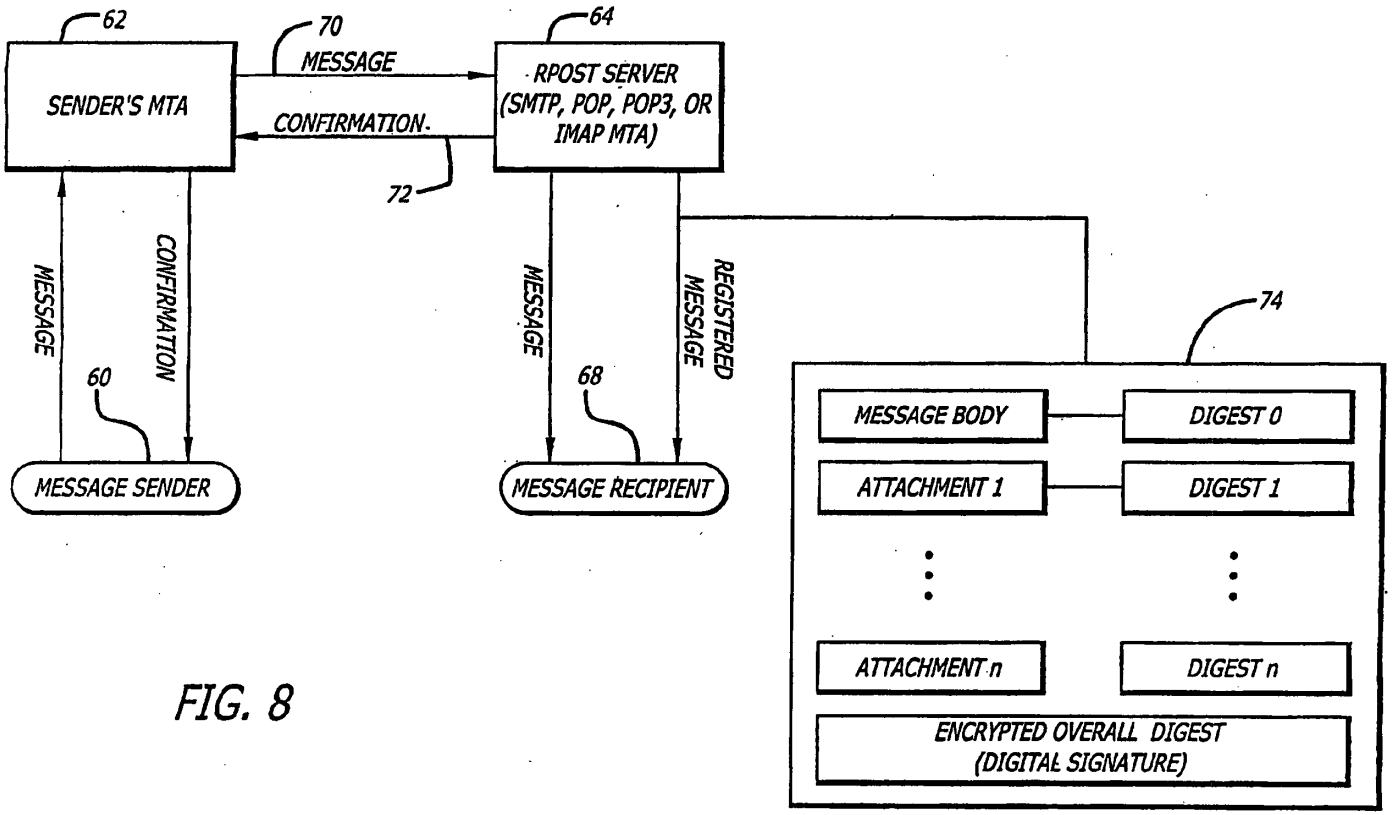
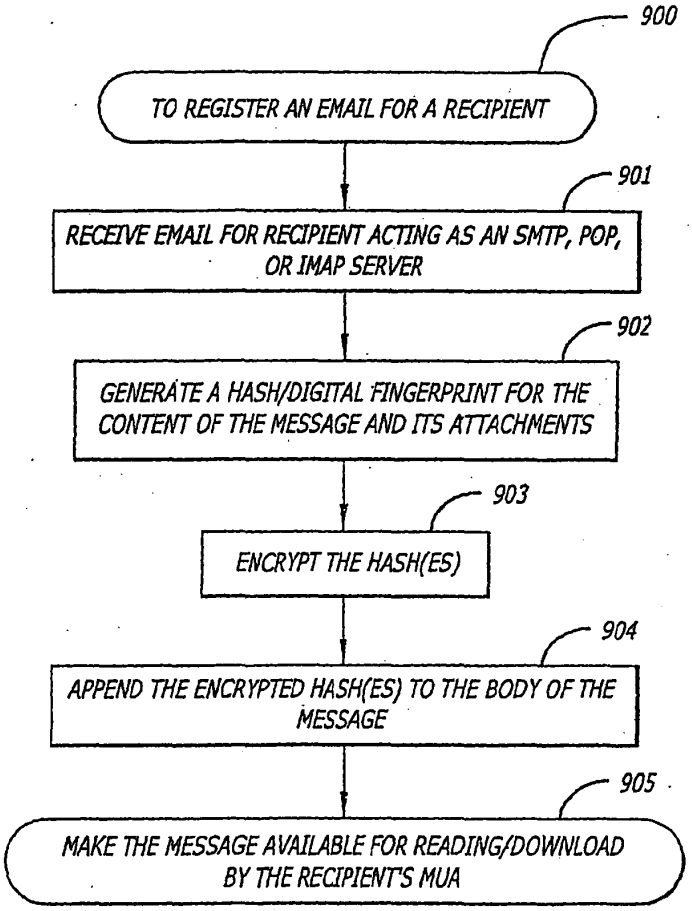


FIG. 8

SUBSTITUTE SHEET (RULE 26)

FIG. 9


900
TO REGISTER AN EMAIL FOR A RECIPIENT
901
RECEIVE EMAIL FOR RECIPIENT ACTING AS AN SMTP, POP, OR IMAP SERVER
902
GENERATE A HASH/DIGITAL FINGERPRINT FOR THE CONTENT OF THE MESSAGE AND ITS ATTACHMENTS
903
ENCRYPT THE HASH(ES)
904
APPEND THE ENCRYPTED HASH(ES) TO THE BODY OF THE MESSAGE
905
MAKE THE MESSAGE AVAILABLE FOR READING/DOWNLOAD BY THE RECIPIENT'S MUA

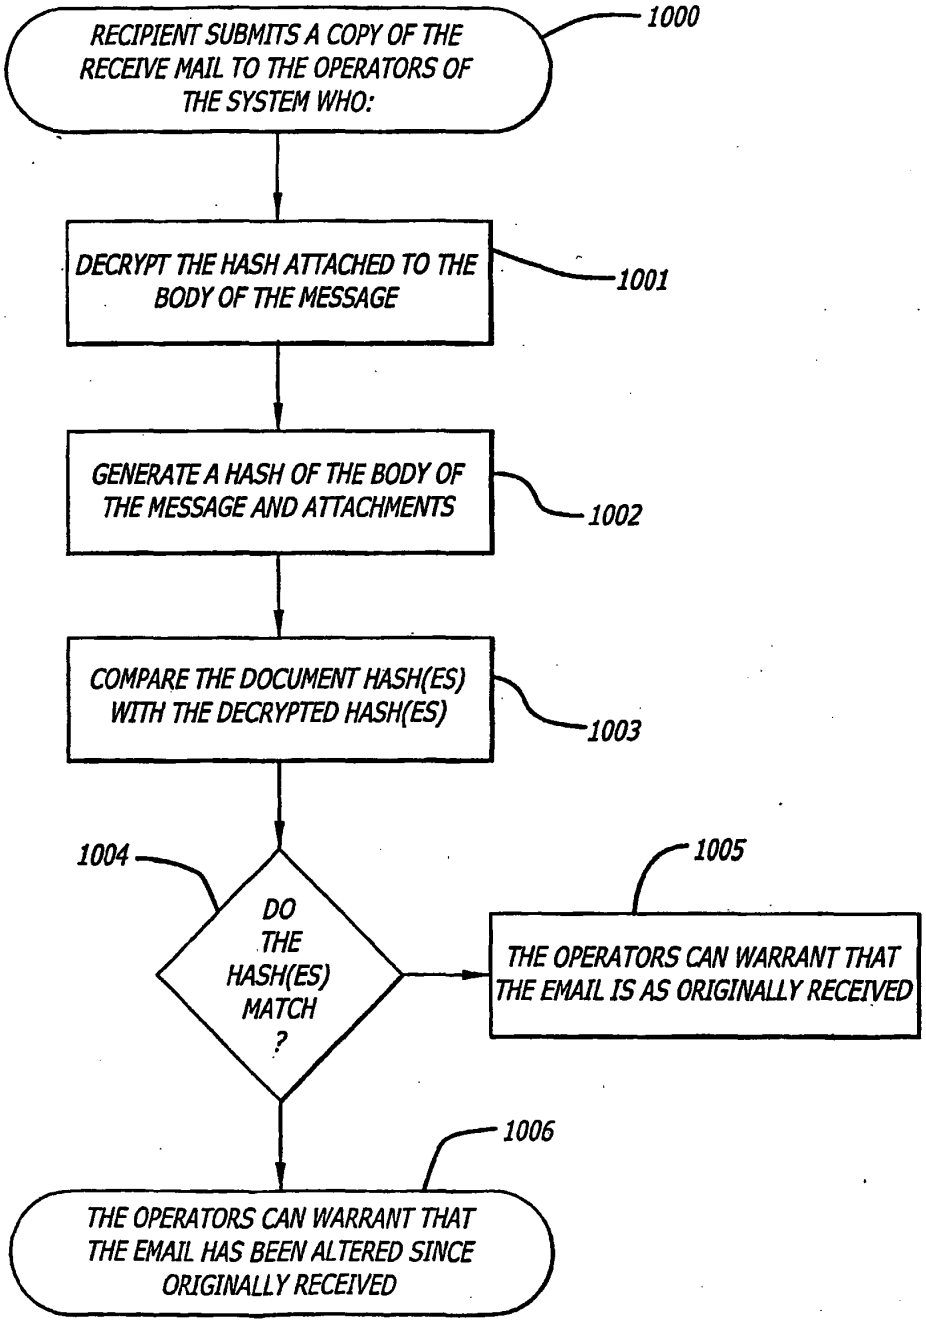

RECIPIENT SUBMITS A COPY OF THE RECEIVE MAIL TO THE OPERATORS OF THE SYSTEM WHO:
1000
DECRYPT THE HASH ATTACHED TO THE BODY OF THE MESSAGE
1001
GENERATE A HASH OF THE BODY OF THE MESSAGE AND ATTACHMENTS
1002
COMPARE THE DOCUMENT HASH(ES) WITH THE DECRYPTED HASH(ES)
1003
1004
DO THE HASH(ES) MATCH ?
1005
THE OPERATORS CAN WARRANT THAT THE EMAIL IS AS ORIGINALLY RECEIVED
1006
THE OPERATORS CAN WARRANT THAT THE EMAIL HAS BEEN ALTERED SINCE ORIGINALLY RECEIVED

FIG. 10

SUBSTITUTE SHEET (RULE 26)

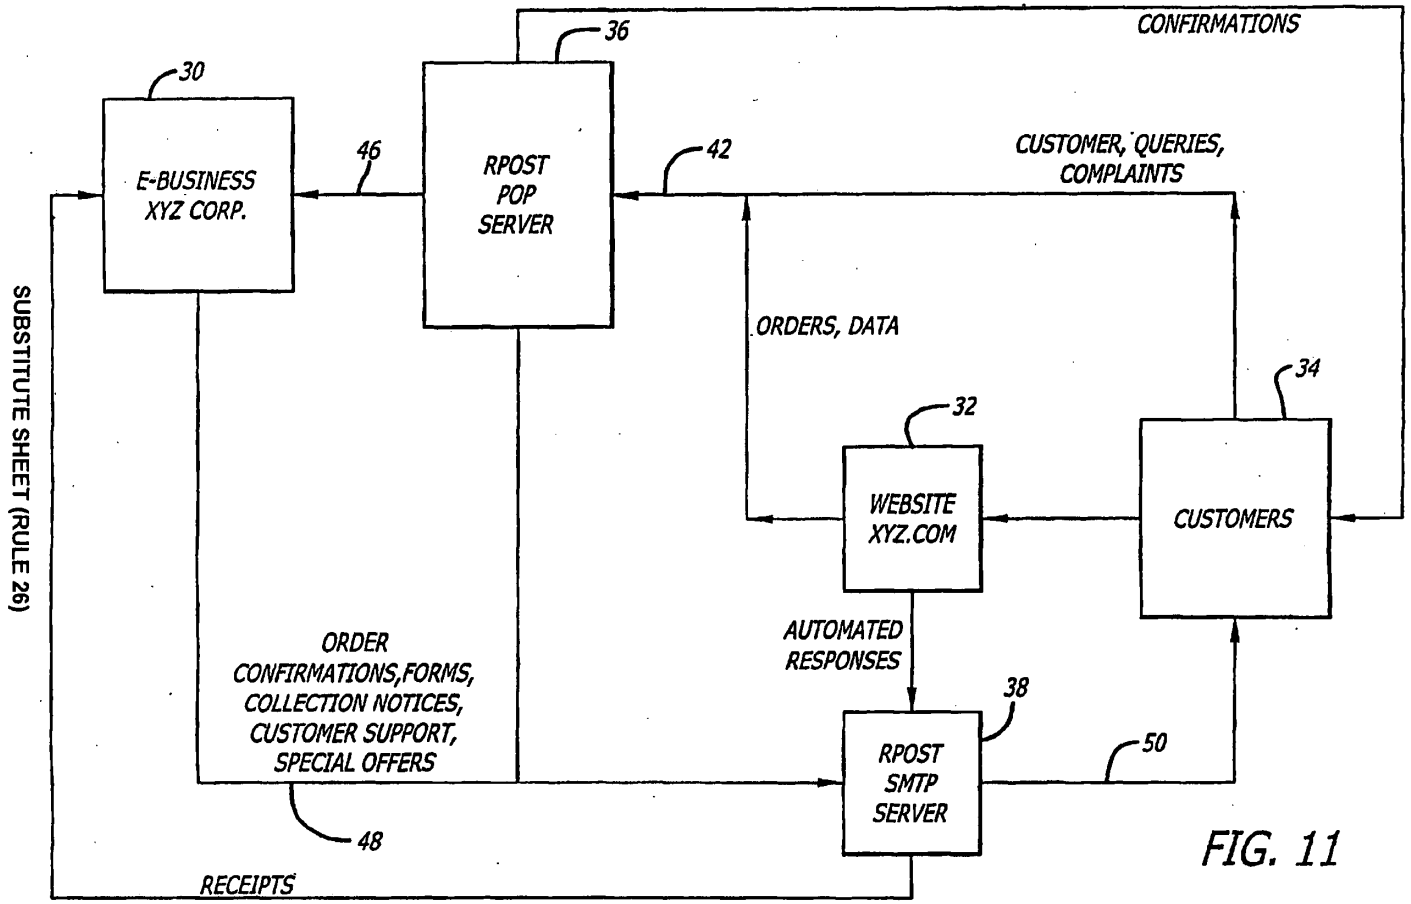


FIG. 11

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
5 February 2004 (05.02.2004)

PCT

(10) International Publication Number
WO 2004/012415 A1

(51) International Patent Classification[7]: H04L 29/06, 12/58, G06F 1/00

(21) International Application Number: PCT/GB2003/003243

(22) International Filing Date: 21 July 2003 (21.07.2003)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
0217610.5 30 July 2002 (30.07.2002) GB
03250672.7 3 February 2003 (03.02.2003) EP

(71) Applicant *(for all designated States except US)*: SECURITY AND STANDARDS LIMITED [GB/GB]; Suite A, 192 Moulsham Street, Chelmsford, Essex CM2 0LG (GB).

(72) Inventors; and
(75) Inventors/Applicants *(for US only)*: POPE, Nicholas, Henry [GB/GB]; 167 Mountnessing Road, Billericay CM12 0EE (GB). ROSS, John, Gordon [GB/GB]; 38 Quilp Drive, Chelmsford CM1 4YA (GB).

(74) Agents: EXELL, Jonathan et al.; Morley House, 26-30 Holborn Viaduct, London EC1A 2BP (GB).

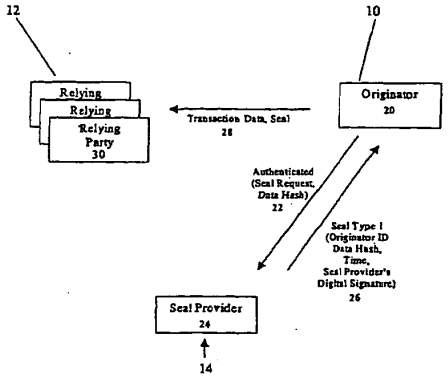
(81) Designated States *(national)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— *with international search report*
— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

[Continued on next page]

(54) Title: ELECTRONIC SEALING FOR ELECTRONIC TRANSACTIONS



(57) Abstract: A sealing method and system based on any data originator authentication mechanism involves an originator (10) which represents the entity requiring the data to be sealed and to have itself identified as the originator of the sealed data. The relying party (12) is an entity requiring to use the sealed data as proof relating to a transaction. The seal provider (14) is an entity trusted to provide seals by the originator and relying parties. The procedure for creating a seal is as follows: the originator (10) creates a hash value or other one-way representation of the data to be sealed. The originator (10) then sends the hash value with a seal request through a secure channel which authenticates the originator (10) to the seal provider (14) and ensures the integrity of the request. The seal provider (14) then determines if the authentication of the request is correct and, if so, creates a seal which contains an identifier for the originator (10), the time of the request, the hash value and a digital signature or other similar mechanism which authenticates the data unit as coming from the seal provider (14). The preferred form of digital signature is one using public key cryptography such as specified in ITU-T X.509 or Internet RFC 2560.

WO 2004/012415 A1

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.

ELECTRONIC SEALING FOR ELECTRONIC TRANSACTIONS

The present invention relates to an electronic sealing method and to an electronic registration method for use in electronic transactions.

Existing electronic signature methods are commonly based on use of public/private key (asymmetric) cryptography to create a digital signature supported by electronic certificates (e.g. as defined in ITU-T X.509). This requires use of certificate status methods such as Certificate Revocation Lists (as defined in ITU-T X.509) or on-line certificate status checking (as defined in Internet RFC 2560) to ensure the validity of certified key (revoked or not). As the validity of the certified key used to create a digital signature may change over time, it is often considered necessary to know the time at which the signature was created. This can be achieved by applying a secure time-stamp method (Internet RFC 3161, US patent RE 34,954) over the digital signature (ETSI TS 101 733).

Secure time-stamping methods (Internet RFC 3161; US patent RE 34,954) can also be used on their own (without the digital signature of the originator) to provide integrity of transaction data and proof of transaction time. This does not, however, provide any proof of the originator.

These electronic signature methods have two major disadvantages. First, it can be difficult to gain assurance of the validity of the public key used to verify the digital signature (that is the signature verification key). When verifying a signature the relying party needs to obtain information on the status of certificate used to validate the signing key. This can involve a significant overhead in obtaining the appropriate (revocation) status information. It can be particularly difficult, if not impossible, to get this status information as it applied when the signature was created when there is a significant period between the creation of the signature and its verification (as in the case of a subsequent dispute over a transaction).

Secondly, the certificate, which relates a name to a signing key, is often created by a party which is not directly involved with the business using that certificate. As the subject of the certificate is not already "known" to the authority issuing the certificate (the

certification authority), there can be difficulty in checking the subject's name. This can result in significant additional costs and also increases the probability of errors.

The present invention seeks to provide an electronic sealing and an electronic registration method for use in electronic transactions.

According to an aspect of the present invention, there is provided a sealing method for sealing the details of an electronic transaction including the step of obtaining authentication information from an originator, checking the validity of that originator's authentication at the time of the transaction and creating a sealed record related to the transaction and validity of the originator's authentication information.

Preferably, the sealed record includes details of the originator of transaction data and the time of the transaction and thereby proves the integrity of the transaction data. The seal is preferably digitally signed by a trusted seal provider. In the preferred embodiment, the transaction data itself is not revealed to the seal provider enhancing its confidentiality. The seal may include other information relating to a transaction originator such as a copy of originator's public key certified by the seal provider.

The method advantageously forms an electronic signature having non-repudiation properties. The method may use any data originator authentication mechanism or more specific form of originator authentication using digital signatures based on public key cryptography.

The preferred sealing method overcomes the first disadvantage mentioned above by checking the validity of originator's authentication information (for example signing key) at the time of the transaction. The registration method overcomes the second disadvantage mentioned above by involving a party involved in the business functions in the registration process.

According to a preferred embodiment of the present invention. The method includes the steps of:

the originator digitally signing a transaction data or other data,

transmitting the transaction data or other data or a hash value or another representation of the transaction, the digital signature and a seal request to the or a seal provider;

providing for the seal provider to determine whether the authentication of the request is correct and whether the digital signature is currently valid and if so creating a seal.

The seal may contain an identifier for the originator, the time of the request, the transaction data or other data, hash value or other one-way representation, by the originator, and at least one digital signature or other authentication of the seal as coming from the seal provider. The seal may also contain a digital signature provided by the originator which has been checked by the seal provider.

The originator may be a transaction peer. The sealed record may include a text header and binary data encoded as a text string.

The method may preferably further comprise the step of encoding the transaction data or other data and the sealed record, the encoding being arranged such that upon accessing the encoded data, the sealed record is displayed in a first area of a computer interface and the transaction data or other data is displayed in one or more other areas.

The encoding may comprise MIME and HTML encoding and the areas comprise frames.

According to another aspect of the present invention, there is provided a computer implemented transaction method comprising the steps of accepting one or more user inputs specifying a transaction, submitting authentication data and data representative of transaction data to a system operating the method of any of the preceding claims for creating a sealed record, receiving said sealed record and storing said sealed record and at least selected parts of said data representative of transaction data or the transaction data.

According to another aspect of the present invention, there is provided a system for providing a seal for sealing the details of an electronic transaction or data including means for obtaining authentication information from an originator, means for checking the validity of that originator's authentication at the time of sealing and means for creating a sealed record related to the transaction or data and validity of the originator's authentication information.

According to another aspect of the present invention, there is provided a transaction system arranged to accept one or more user inputs specifying a transaction, submit authentication data and data representative of transaction data to a system for creating a sealed record as defined above, receive said sealed record and store said sealed record and at least selected parts of said data representative of transaction data or the transaction data.

Preferably the system further comprises a client interface, the transaction system being arranged to transmit said sealed record and said at least selected parts of said data representative of transaction data or the transaction data to the client interface.

Preferably, the system further comprises a validation interface, the validation interface being arranged to accept a user input to validate an existing sealed record, the validation interface being operative to transmit said sealed record and said data representative of transaction data to a validation system for validating said sealed record, said validation system being arranged to generate at least part of a sealed record, compare it to said sealed record and output a result of said comparison to said validation interface.

The system may comprise one or more web pages.

In one embodiment, an email system may incorporate or interface with a transaction system as defined above. The email system may be arranged to receive an email message for an intended addressee, submit data from said email or data representative of data from said email to the system to obtain a sealed record, append the sealed record to the email message and transmit the email message and appended sealed record to the intended addressee.

The invention also seeks to provide a sealing method for sealing the details of an electronic transaction or data including the step of obtaining authentication information from an originator, obtaining an identification of the originator at the time of sealing and creating a sealed record related to the transaction or data and validity of the originator's identification.

In the preferred embodiment, a one-way mathematical function of the transaction data, and not the data itself, is revealed to the seal provider. This may be used to ensure that the seal applies to the data in its original form.

Embodiments of the present invention are described below, by way of example only, with reference to the accompanying drawings, in which:

Figure 1 shows an embodiment of sealing method using any originator authentication mechanism;

Figure 2a shows an embodiment of sealing method based on originator digital signatures;

Figure 2b shows an embodiment of sealing method incorporating a certified copy of the originator's public key;

Figure 3 shows an embodiment of sealing method for email transactions;

Figure 4 shows an embodiment of sealing method providing a transaction peer;

Figure 5 shows a modification to the embodiment of sealing method of Figure 4;

Figure 6 shows an embodiment of sealed registration method and in particular an originator set-up stage;

Figure 7 shows the steps of a routine for creating a sealed registration document;

Figure 8 shows a example of how a sealed document may be encoded as a text string;

Figure 9 shows a example of how a sealed document may be displayed; and

Figure 10 shows the steps of how a trusted agent may be used to validate a seal.

Referring to Figure 1, there is shown an embodiment of sealing method which is based on any data originator authentication mechanism. The scenario depicted in Figure 1 involves an originator 10 which represents the entity requiring the data to be sealed and to have itself identified as the originator of the sealed data. The relying party 12 (of which there may be more than one) is an entity requiring to use the sealed data as proof relating to a transaction. This may, for example, include the recipient of the transaction data, an arbitrator or judge in a dispute or even the originator.

The seal provider 14 is an entity trusted to provide seals by the originator and relying parties. The seal provider 14 can be a system operated by any suitable organisation, such a business specialising in the provision of security services, trade organisation or any other suitable organisation.

In this embodiment and scenario depicted in Figure 1, the procedure for creating a seal is as follows. In the first instance, at step 20, the originator 10 creates a hash value (for example, the secure hash algorithm SHA1 defined in FIPS 180-1) or other one-way representation of the data to be sealed (typically the transaction data). At step 22 the originator 10 sends the hash value with a seal request through a secure channel which authenticates the originator 10 to the seal provider 14 and ensures the integrity of the request.

At step 24, the seal provider determines if the authentication and integrity of the request are correct. If so, the seal provider 14 creates a data unit, conveniently referred to

in this description as a seal, which, in this embodiment, contains an identifier for the originator 10, the time of the request, the hash value and a digital signature or other similar mechanism which authenticates the data unit as coming from the seal provider 14. The preferred form of digital signature is one using public key cryptography such as specified in ITU-T X.509 or Internet RFC 2560. The seal provider 14 may also copy the data forming the seal into an audit log file to provide an additional mechanism for assuring the validity of the transaction.

At step 26, the seal is returned to the originator 10. At step 28, the transaction data with its seal is passed to each relying party 12, either directly or indirectly.

At step 30, any relying party 12 which trusts the seal provider 14 can verify the seal and the integrity of the transaction data to which the seal is applied by verifying the digital signature in the seal. Thus, the relying party 12 is provided with proof of the originator of the transaction data, the time of transaction and the integrity of the transaction data.

Figure 2 shows an embodiment of sealing method which is based on a digital signature from the originator and still provides the scenario of one originator 10, one or more relying parties 12 and a seal provider 14.

In this scenario, the preferred procedure for creating a seal involves the following steps. At step 40 the originator 10 digitally signs a transaction data and creates a hash value or other one-way representation of that transaction data. At step 42 the originator 10 sends the hash value and the digital signature(s), with a seal request, through a secure channel which authenticates the originator 10 to the seal provider 14 and ensures the integrity of the request. In this embodiment, the request need not necessarily be authenticated. If the request is not authenticated, the digital signature is used to authenticate the originator 10. Other parties may later submit the signed data for sealing but the seal with the earliest time can be taken as the one representing the time of the original transaction.

At step 44 the seal provider 14 determines whether the authentication and integrity of the request is correct and also whether the digital signature is currently valid, using the signature verification key of the originator 10. If this is the case, the seal provider 14 creates a data unit called a seal which, in this example, contains: an identifier for the originator 10, the time of the request, the hash value, the digital signature provided by the originator 10, and at least one digital signature or other similar mechanism which authenticates the data unit as coming from the seal provider 14. The seal provider 14 may also copy the data forming the seal into an audit log file to provide an additional mechanism of assuring the validity of the transaction.

At step 46 the seal is returned to the originator 10. At step 48 the transaction data with its seal is passed to each relying party 12.

At step 50, any relying party 12 which trusts the seal provider 14 can verify the seal and the integrity of the transaction data to which the seal is applied by verifying the digital signature(s) in the seal. Thus, each relying party 12 is provided with proof of the originator 10 of the transaction data, the time of transaction, the integrity of the transaction data and the validity of the digital signature which the originator 10 applied to the transaction data at that given time. In this embodiment, as is depicted in Figure 2a, two digital signatures may be provided, the originator's digital signature and the seal provider's digital signature.

A variation for the sealing method based on digital signatures, as shown in Figure 2b, is one providing a public key that is certified by the sealing authority as being valid at the time of the transaction. In this variation, the originator 10 includes a request for its public key be certified in the seal request 42b to the seal provider 14. The seal provider checks that the originator's public key is valid 44b, for example by checking the validity period in the originator's current X.509 identity certificate and the revocation status of that certificate. If the originator's public key is valid then this is included in the seal certifying the validity of the key at the time of the transaction 46b. The originator sends this seal along with the digital signature to the relying party 48b. The relying party verifies the

digital signature using the public key in the seal 50b that the sealing authority certifies as being valid at the time of the transaction.

This form of seal (i.e. one including the public key) may be used as an alternative a public key certificate (such as defined in X.509) with the validity only for a specific transaction (i.e. a transaction certificate). Furthermore, a seal may be used as a "qualified certificate" meeting the requirements identified in the European Directive 1999/93/EC on a "Community framework for electronic signatures" as described in the following table.

EU Electronic Signature Directive 1999/93/EC – Annex I Requirement on Qualified Certificate Content	Qualified Seal
(a) an indication that the certificate is issued as a qualified certificate;	Additional transaction attribute indicating that the seal is a form of qualified certificate;
(b) the identification of the certification-service-provider and the State in which it is established;	Already identified through the Sealing Authority certificate. This could be repeated as in an additional transaction attribute.
(c) the name of the signatory or a pseudonym, which shall be identified as such;	Already included in user name.
(d) provision for a specific attribute of the signatory to be included if relevant, depending on the purpose for which the certificate is intended;	Additional transaction attribute
(e) signature-verification data which correspond to signature-creation data under the control of the signatory;	Additional transaction attribute containing user public key

(f) an indication of the beginning and end of the period of validity of the certificate;	Validity period is for a single transaction at the time already given in the time-stamp.
(g) the identity code of the certificate;	Already in time-stamp serial number
(h) the advanced electronic signature of the certification-service-provider issuing it;	Already in the time-stamp signature
(i) limitations on the scope of use of the certificate, if applicable; and	Additional transaction attribute
(j) limits on the value of transactions for which the certificate can be used, if applicable.	Additional transaction attribute

This further enhanced seal, which is called herein a "qualified seal", has additional advantages over an X.509 based certificate in that:

1) This qualified seal may be used to provide an electronic signature that under the European Directive 1999/93/EC article 5.1 "satisfy the legal requirements of a signature in relation to data in electronic form in the same manner as a hand-written signature satisfies those requirements in relation to paper-based data".

2) A user who already has conventional X.509 certificate, which does not meet the above legal requirements, can use that certificate to obtain a "qualified seal", which does meet these requirements.

3) As the "qualified seal" is valid only for a particular transaction uncertainty over the validity of the certificate as the time of the transaction is minimised.

4) As the "qualified seal" applies to a single transaction the issuing authority has much greater control over its liability.

A variation to the sealing methods described above with respect to Figures 1, 2a and 2b is shown in Figure 3 and, in this example, relates to email transaction environments.

As illustrated in Figure 3, a seal request is passed to the seal provider 14 at step 60, which is then to be subsequently forwarded to the recipient(s) 12 identified in the recipient list. The transaction data is passed to the seal provider 14 together with the seal request and then is forwarded, at step 62, with the resulting seal to the identified recipient(s) 12.

Figure 4 shows another scenario of sealing method for use with transaction peers. As illustrated in Figure 4, at step 70 a seal request from the originator 10 is passed through the peer 15 with whom the originator 10 is carrying out the transaction. The seal request is passed, at step 72, to the seal provider 14 only if the transaction peer 15 agrees with the transaction. At step 74 the seal is returned by the seal provider 14 to the transaction peer 15 and subsequently, at step 76, passed to the originator 10.

Figure 5 shows another scenario of sealing method for with transaction passing through the seal provider. As illustrated in Figure 5, at step 80 references to the originator and transaction data (for example a universal Resource Identifier, see Internet RFC 2396) is passed from the originator 10 to the transaction peer 15. The preferred form of the reference to the transaction data is that the reference value is unpredictable, to avoid unauthorised access to the transaction data. The transaction peer 15 then sends these references to the seal provider 14 at step 82. The seal provider 14 then requests the transaction data from the originator 10 at step 84, which returns with the transaction data through a secure channel which authenticates the originator 10 at step 86. The seal provider 14 then creates the seal which it passes, along with the transaction data, to the transaction peer 15 at step 88. The seal may also be sent back to the originator 10, with the reference to the transaction data, at step 90.

The preferred embodiments of seal registration methods provide two forms of identification, the first being a security system identifier and the second being a business identifier.

The preferred seal registration method is a technique for securely binding a security system identifier to a business identifier.

The security system identifier is allocated and authenticated by a security system, such as one of the implementations of sealing method described above. When allocated, this identifier is unrelated to any "real world" identification but can be used with a security system for identification purposes. A security system identifier may, for example, be a number pair allocated by a seal provider 14 to an originator of sealed data. One number of the number pair identifies a community of users, the other number identifies the user within the community. The number pair is unique within the set of identifiers allocated by the seal provider.

The business identifier provides a set of attributes used by the business to identify an entity (such as a customer) for the purposes of transacting business therewith. This may include, for example, references used within the business (such as a customer reference number) and/or information used to identify the entity more generally (such as the customer's name and home address). The business identifier is agreed between the subject and a registration authority. Some attributes may be allocated by the registration authority (for example, customer number), others may be provided by the subject (for example, first name, surname and address). The registration would either be a system operated by or on behalf of the business with which the subject is trading, or the registration authority would be operated by a third party business whose function is to check information concerning the user.

When being set up to use one of the sealing methods described above, a sealed data originator 21 is allocated a security system identifier by a seal provider 14, as shown in Figure 6. Before allocating a security system to the originator identifier, the seal provider proves ownership of some means for authentication. In the embodiment using public key techniques this involves proof of possession of a private signing key.

Having set up two seal originators 21 which are to be parties in a business transaction, such as a business and its customer, a sealed registration document is created. One party (for example the business) acts as a registration authority for the other party (for example, the customer). The customer acts as a subject of the registration. Between them,

these two parties create a sealed registration document which binds the subject identifier to the business identifier.

In the preferred embodiment, the sealed registration document is allocated a document reference (for example a Universal Resource Identifier, see Internet RFC 2396) by the registration authority. This registration document reference can later be used in subsequent sealed documents produced by the subject, thereby providing reference to its authenticated business identifier.

In the preferred embodiment, the sealed registration document contains:

(a) the business identifier,

(b) the registration document reference or some other information that may be used to identify the registration for a particular business purpose,

(c) a seal from the subject which includes its security system identifier and which indicates the subject's agreement that the business identifier is correct for itself,

(d) a seal from the registration authority which indicates the registration authority's agreement that the business identifier is correct for the subject which created the seal in step (c) above.

An embodiment of routine for creating such a sealed registration document between a business acting as the registration authority and a customer acting as a subject of the registration is shown in Figure 7.

The procedure of the example of Figure 7 is as follows. At step 100, the subject 17 and the registration authority 19 establish the attributes required for the business identifier of the subject 17, such as name, address, customer identifier and so on.

At step 102, the registration authority 19 allocates the registration a reference (for example a Universal Resource Identifier) which is sent to the subject 17. At step 104, the subject 17 sends to the registration authority 19 a sealed document containing the business identifier and the registration document reference. The seal indicates the entity identified

with the security system identifier in the seal and agrees that the business identifier is applicable to itself.

At step 106, if the registration authority 19 has also agreed that the business identifier is applicable to the subject 17 which sent the sealed document, it will add its own seal. This sealed registration document is stored by the registration authority 19 and, if required, also passed to the subject 17 for later use in verifying a claim of identity.

Any sealed document can then include the appropriate registration reference to authenticate the identity of the originator.

The business operating as a registration authority 19 may similarly obtain a sealed registration with the seal provider 14 which, in turn, acts as another registration authority for the business registration authority.

It will be appreciated that the various embodiments described above can be combined with one another, as required in any particular implementation of the systems described herein.

It will also be appreciated that the embodiments described can provide a method of sealing data by a trusted third party to provide independent proof of: the originator, of the transaction data, the time of transaction, and the transaction data. The sealing method produces a form of electronic signature, which has non-repudiation properties and avoids the need for certificate revocation or other authentication status checking. The preferred sealing method builds on existing secure time-stamping methods.

A method of encoding a seal provides a text string containing a descriptive header and binary protection data encoded in Base64 as defined in Internet RFC 2045. The header includes information on: the identity of the sealing authority, the originator's identity (system security identity and business identity) and the time of the transaction. The binary protection data includes binary encoding of the sealing authorities digital signature against a hash of the transaction data with the originator's identity and the time of the transaction.

An example of a text string encoding is given in Figure 8. An alternative encoding is to encode the data as described above in XML as defined in World Wide Web Consortium (W3C) specification Extensible Markup Language (XML) (reference http://www.w3.org/TR/REC-xml).

A method of encoding the transaction data with the seal uses the encoding defined in Internet RFC 2045 to RFC 2049 (commonly called MIME) combined with code for viewing hypertext in multiple frames such as defined in the World Wide Web Consortium (W3C) HTML 4.01 Specification (reference: http://www.w3.org/TR/html4/). The seal itself, encoded as described above, is shown in one frame as a text string, only the header need be immediately visible. The sealed document is in the other frame or frames. This enables sealed documents, including documents incorporating text, images, voice or video, to be viewed through any suitably MIME enabled Web Browser. An example of a sealed document is given in Figure 9.

Referring to Figure 10, an embodiment of routine illustrating how a trusted agent may be used to validate a seal is shown. At step 120 the seal provider 14 issues a seal to the originator 10, which in turn, at step 122, provides the data and seal to the relying party 12. At step 124 the relying party 12 forwards the hash of the data together with the seal to a seal validator 23 which, at step 125, determines whether the seal is valid. At step 126 the seal validator 23 returns to the relying party 12 a valid/invalid reply.

The embodiments described above provide a number of features for the creation of a secure seal for subsequent use. First, the seal protects the authenticity of a data such as a document or transaction. This protection applies to the data both when communicated between systems and in storage. The seal can be applied to a hash of data along with the time and authenticated identity of the data originator. The authenticity is certified using, for example, the digital signature of the seal provider.

The seal provides long term authenticity of the data and in a form which cannot be repudiated, for example by claiming compromise of authentication keys subsequent to the seal being applied. Thus, a seal is provided which establishes the authenticity of the

parties and seals the transaction data in a manner which maintains its integrity over time. A copy of the seal can be kept by the seal provider in an audit log to provide further proof of authenticity. As the seal only contains a hash of the data being sealed, from which the original data cannot be re-created, the privacy of the data is maintained.

The seal can be used by the originator and any other party relying on the authenticity of data to validate its authenticity. Any party that trusts the seal provider and has a means of authenticating the seal (e.g. using the seal providers public key) can verify the validity of the seal. The hash value within the seal can be checked against a recalculated hash value to check that the seal applies to the given data and that the data has been unchanged. Further assurance of the validity of a seal can be achieved by requesting a copy of the seal from the seal provider's audit log.

In this embodiment, once a transaction record and receipt (seal) has been generated, as well as the seal to the appropriate party, a copy is saved at the seal provider or at a secure location controlled by the seal provider.

In some embodiments, such as those involving a transaction sequence a seal can be applied to each stage of the transaction proving the timing sequence of the transactions. For example, in a transaction involving an offer and an acceptance, data relating to the offer and also to the acceptance can be each be sealed proving the time relationship between the offer and acceptance. This can, for example, have particular application in trading in shares.

It is also envisaged that in some embodiments the generation of the seal also stores data relating to a single event. For example, the event might be the creation of documents or data by an identified originator at a given time. The seal, therefore, would provide evidence of that event.

Thus, in the preferred embodiments, the seal is issued by a trusted authority (the seal provider) which authenticates the document originator, provides proof of time of creation of the document linked to a hash code representing the document and is protected

by a cryptographic code which proves the authenticity of the seal and which can detect any changes in the sealed information. The seal, along with the data it authenticates, can be sent to any party relying on the authenticity of that data.

In this regard, a seal can be checked for validity by the relying party or an agent trusted by the relying party to validate seals (seal validator). As shown in figure 10, the trusted agent can be used to validate a seal simply by forwarding the seal and a recalculated hash of the data to that party (124). This can be initiated, for example, by the relying party clicking on a button or icon displayed as part of the seal. If the seal validator knows and trusts the sealing authority and has the key necessary to verify the cryptographic code (e.g. the sealing authority's public key) it can check the validity of the seal and confirm that the hash in the seal matches that recalculated by the relying party. If the seal is valid and the data hash is the same, the validator would respond by indicating that the sealed document provided by the relying party is valid (126). On the other hand, if there has been any change in the seal, the cryptographic code was not created a trusted sealing authority or the hash in the seal does not match the recalculated hash, the seal provider would indicate that that seal has become invalid.

Thus, the preferred embodiment can provide a seal which is a form of certificate, specific to each transaction which incorporates attributes. This seal certifies the user's signing key, as well as providing a time stamp for relating to the transaction or event. This seal provides a form of certificate which is bound to the transaction, i.e. is transaction specific.

The seals can also be used for voice and video data and indeed can be applied to any digitised media such as web pages, video files, audio files, digital photographs and so on. In the preferred embodiment, the seals are simple to view and verify using, for example, an Internet browsing software. Furthermore, the seals can enable disputes to be resolved quickly, saving time and money. It can also protect all parties against denial of the existence or content of a document or event and can provide evidence of a transaction or data when the information is contested, even years after the event.

Three variations of the sealing method have been described, one using originator authentication, the second adding originator digital signatures, the third adding public keys certified to be valid at the time of the transaction. Further variations of the sealing methods are envisaged for alternative transaction environments including but not limited to e-mail, requests passed via transaction peers and transaction data passing through the seal provider.

There is also described a method for registration that applies the sealing method to the registration of a business identity between business partners. The registration method can be used as an alternative to existing electronic certification methods. The registration method enables alternative business related registrations to be applied to the same entity for different business applications.

CLAIMS

1. A sealing method for sealing the details of an electronic transaction or data including the step of obtaining authentication information from an originator , checking the validity of that originator's authentication at the time of sealing and creating a sealed record related to the transaction or data and validity of the originator's authentication information.

2. A method according to claim 1, wherein the sealed record includes details of the originator of the transaction or data and the time of sealing.

3. A method according to claim 1 or 2, including the step of digitally signing the seal.

4. A method according to claim 3, wherein the step of digitally signing the seal is carried out by a trusted seal provider.

5. A method according to claim 4, wherein the transaction data or other data itself is not revealed to the seal provider.

6. A method according to any preceding claim, wherein a result of a one-way mathematical function applied to the data, is revealed to the seal provider.

7. A method according to any preceding claim, including the step of storing a copy of the data forming the seal into an audit log file held remotely from parties to the transaction or involved in the data.

8. A method according to any preceding claim, wherein the seal is passed to a relying party directly or indirectly.

9. A method according to any preceding claim, including the step of providing for checking of a seal, and that it belongs to data which is identical to that to which the seal was originally applied, by a relying party or an agent trusted by the relying party.

10. A method according to any preceding claim, including the steps of:

the originator digitally signing a transaction data or other data,

transmitting the transaction data or other data or a hash value or another representation of the transaction, the digital signature and a seal request to the or a seal provider;

providing for the seal provider to determine whether the authentication of the request is correct and whether the digital signature is currently valid and if so creating a seal.

11. A method according to claim 10, wherein the seal contains: an identifier for the originator, the time of the request, the transaction data or other data, hash value or other one-way representation, by the originator, and at least one digital signature or other authentication of the seal as coming from the seal provider.

12. A method according to claim 11, wherein the seal also contains a digital signature provided by the originator which has been checked by the seal provider.

13. A method according to any preceding claim, wherein the originator is a transaction peer.

14. A method according to claim 13, wherein the sealed record includes a text header and binary data encoded as a text string.

15. A method according to any preceding claim, further comprising the step of encoding the transaction data or other data and the sealed record, the encoding being arranged such that upon accessing the encoded data, the sealed record is displayed in a first area of a computer interface and the transaction data or other data is displayed in one or more other areas.

16. A method as claimed in claim 15, wherein the encoding comprises MIME and HTML encoding and the areas comprise frames.

17. A computer implemented transaction method comprising the steps of accepting one or more user inputs specifying a transaction, submitting authentication data and data representative of transaction data to a system operating the method of any of the preceding claims for creating a sealed record, receiving said sealed record and storing said sealed record and at least selected parts of said data representative of transaction data or the transaction data.

18. A computer program comprising computer program code means for performing all of the steps of any of claims 1 to 17 when said program is run on a computer.

19. A computer program as claimed in claim 18 embodied on a computer readable medium.

20. A system for providing a seal for sealing the details of an electronic transaction or data including means for obtaining authentication information from an originator (10), means for checking the validity of that originator's authentication at the time of sealing and means for creating a sealed record related to the transaction or data and validity of the originator's authentication information.

21. A transaction system arranged to accept one or more user inputs specifying a transaction, submit authentication data and data representative of transaction data to a system for creating a sealed record according to claim 20, receive said sealed record and store said sealed record and at least selected parts of said data representative of transaction data or the transaction data.

22. A transaction system according to claim 21, further comprising a client interface, the transaction system being arranged to transmit said sealed record and said at

least selected parts of said data representative of transaction data or the transaction data to the client interface.

23. A transaction system according to claim 21 or 22, further comprising a validation interface, the validation interface being arranged to accept a user input to validate an existing sealed record, the validation interface being operative to transmit said sealed record and said data representative of transaction data to a validation system for validating said sealed record, said validation system being arranged to generate at least part of a sealed record, compare it to said sealed record and output a result of said comparison to said validation interface.

24. A transaction system according to any of claims 21 to 23 comprising one or more web pages.

25. An email system including a system according to any of claims 20 to 24, the email system being arranged to receive an email message for an intended addressee, submit data from said email or data representative of data from said email to the system to obtain a sealed record, append the sealed record to the email message and transmit the email message and appended sealed record to the intended addressee.

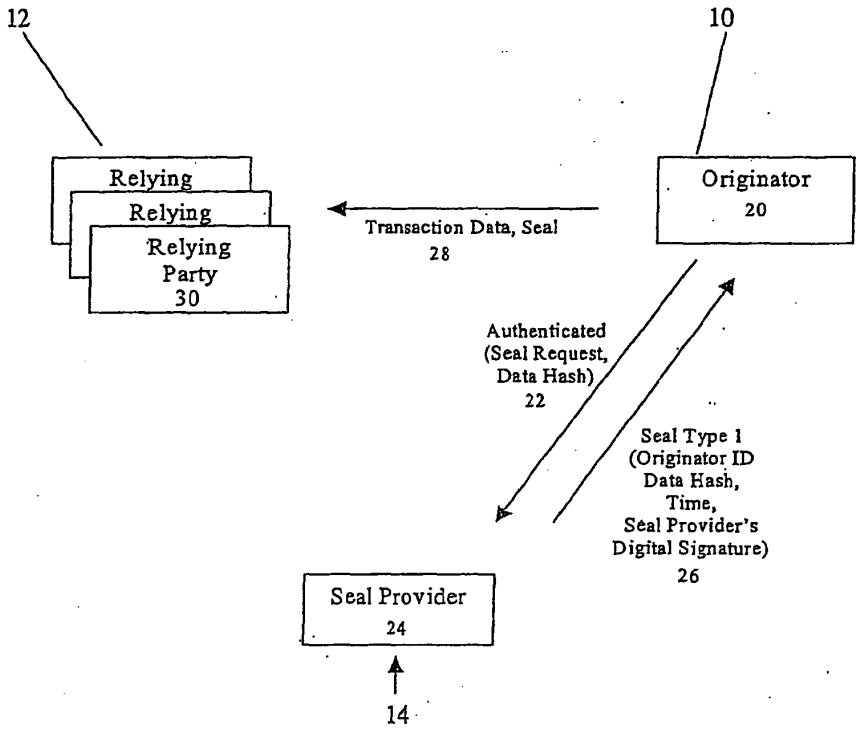

12
10
Relying
Relying
Relying
Party
30
Transaction Data, Seal
28
Originator
20
Authenticated
(Seal Request,
Data Hash)
22
Seal Type 1
(Originator ID
Data Hash,
Time,
Seal Provider's
Digital Signature)
26
Seal Provider
24
14

Figure 1

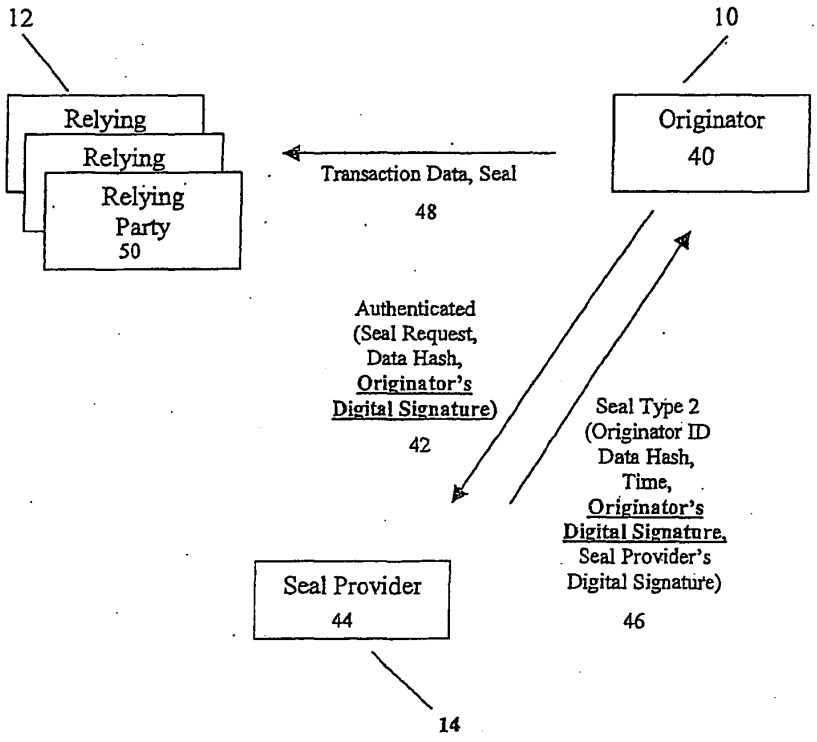

12
10
Relying
Relying
Relying
Party
50
Originator
40
Transaction Data, Seal
48
Authenticated
(Seal Request,
Data Hash,
Originator's
Digital Signature)
42
Seal Type 2
(Originator ID
Data Hash,
Time,
Originator's
Digital Signature,
Seal Provider's
Digital Signature)
46
Seal Provider
44
14

Figure 2 a

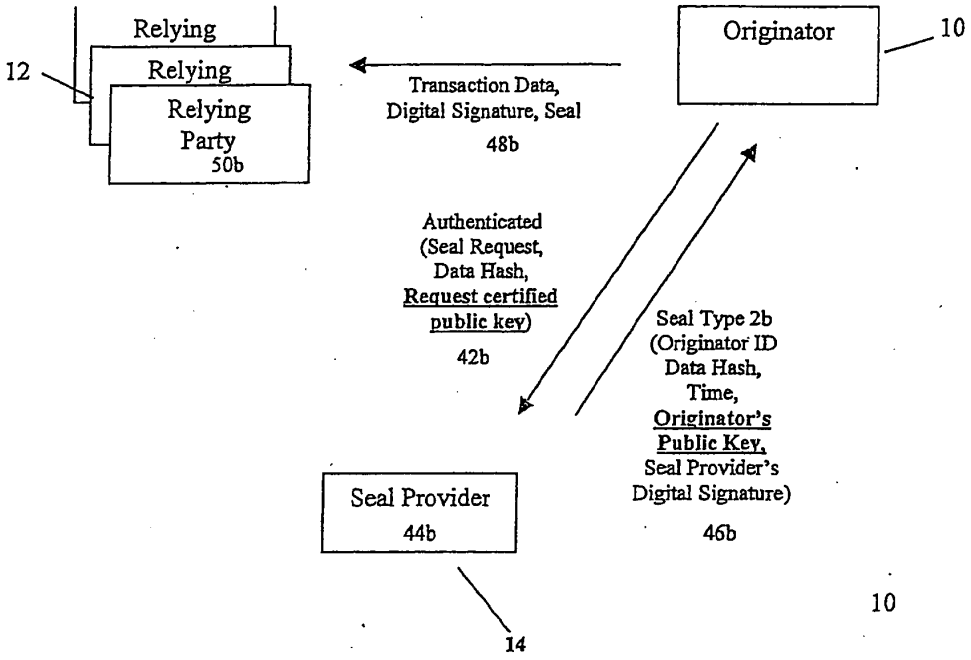

Relying
Relying
Relying
Party
50b
12
Transaction Data,
Digital Signature, Seal
48b
Originator
10
Authenticated
(Seal Request,
Data Hash,
Request certified
public key)
42b
Seal Type 2b
(Originator ID
Data Hash,
Time,
Originator's
Public Key,
Seal Provider's
Digital Signature)
46b
Seal Provider
44b
14
10

Figure 2b

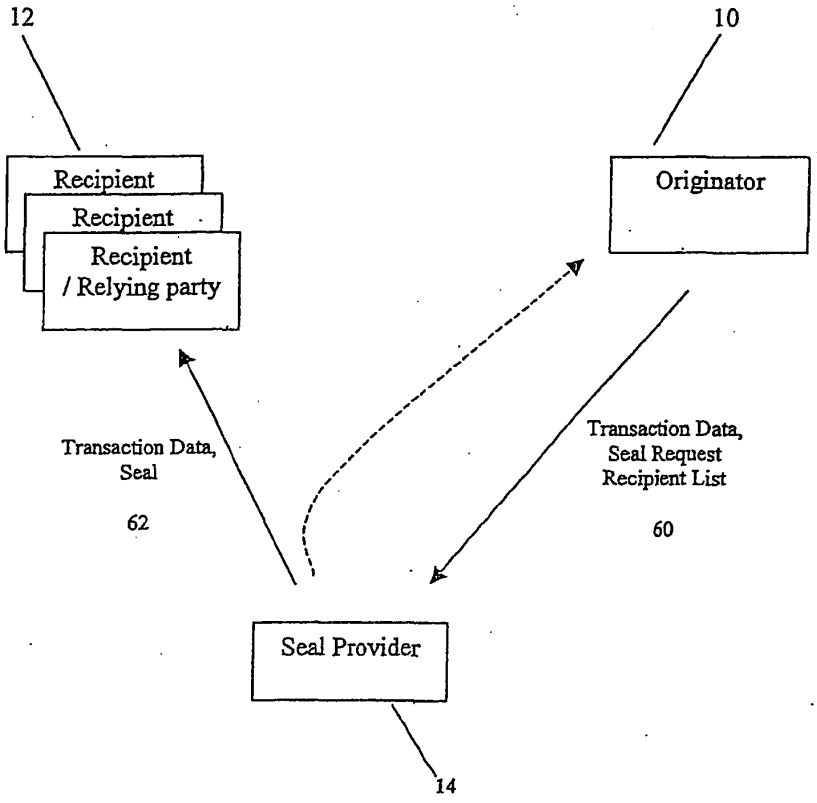

12
10
Recipient
Recipient
Recipient
/ Relying party
Originator
Transaction Data,
Seal
62
Transaction Data,
Seal Request
Recipient List
60
Seal Provider
14

Figure 3

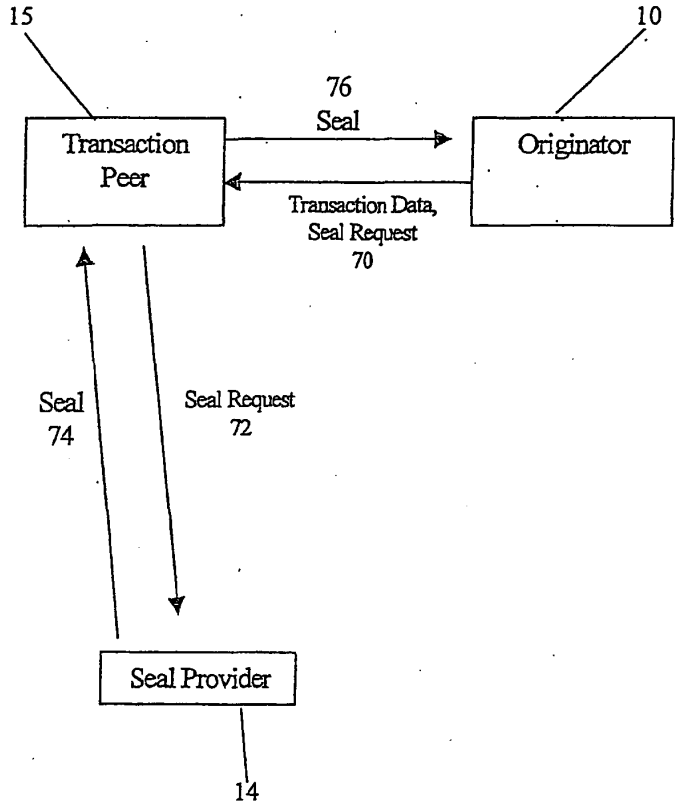

15
10
76
Seal
Transaction Peer
Originator
Transaction Data, Seal Request 70
Seal 74
Seal Request 72
Seal Provider
14

Figure 4

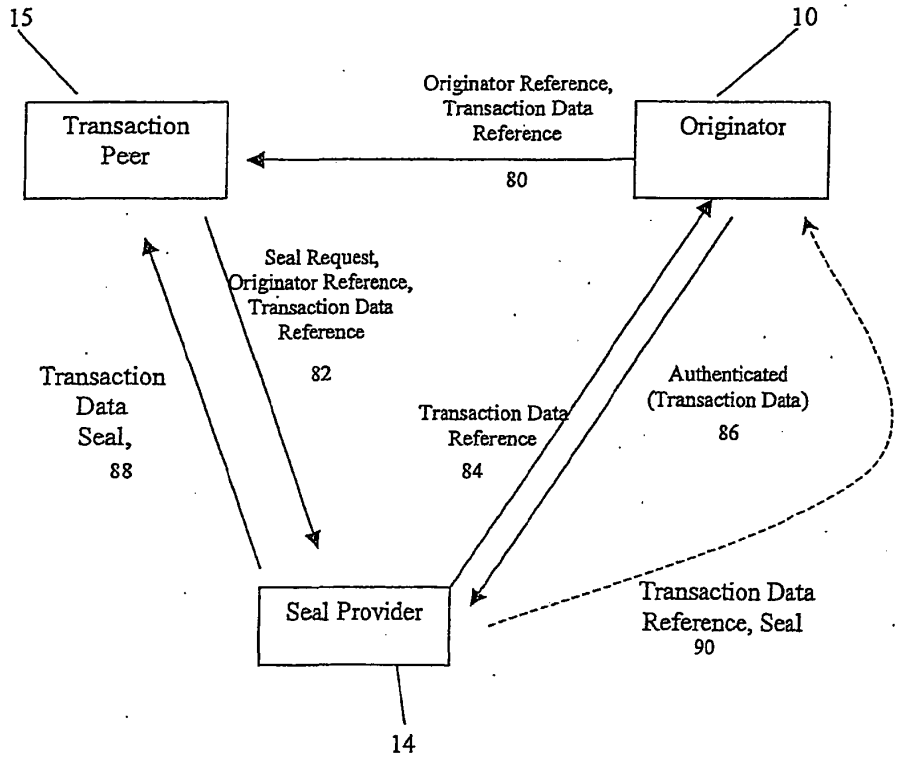

15
10
Transaction Peer
Originator Reference, Transaction Data Reference
80
Originator
Seal Request, Originator Reference, Transaction Data Reference
82
Transaction Data Seal,
88
Transaction Data Reference
84
Authenticated (Transaction Data)
86
Seal Provider
Transaction Data Reference, Seal
90
14

Figure 5

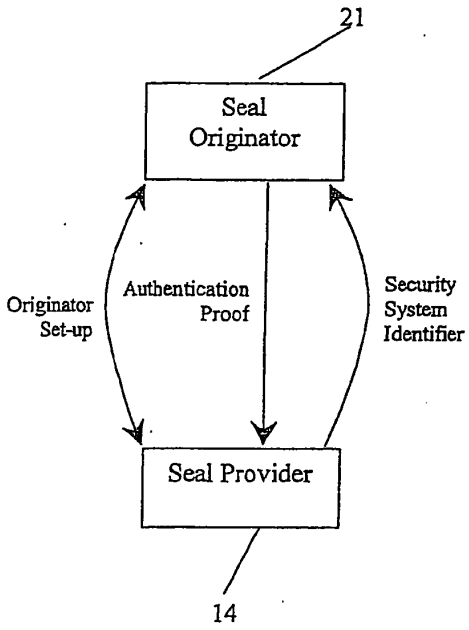

21
Seal Originator
Originator Set-up
Authentication Proof
Security System Identifier
Seal Provider
14

Figure 6

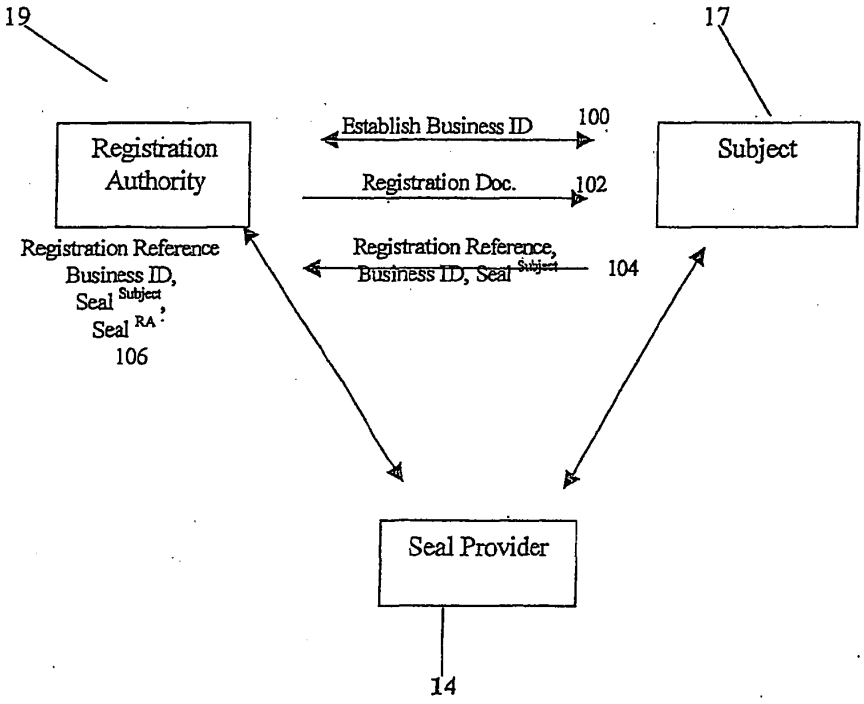

19
17
Registration Authority
Establish Business ID
100
Subject
Registration Doc.
102
Registration Reference Business ID, Seal Subject, Seal RA
106
Registration Reference, Business ID, Seal Subject
104
Seal Provider
14

Figure 7

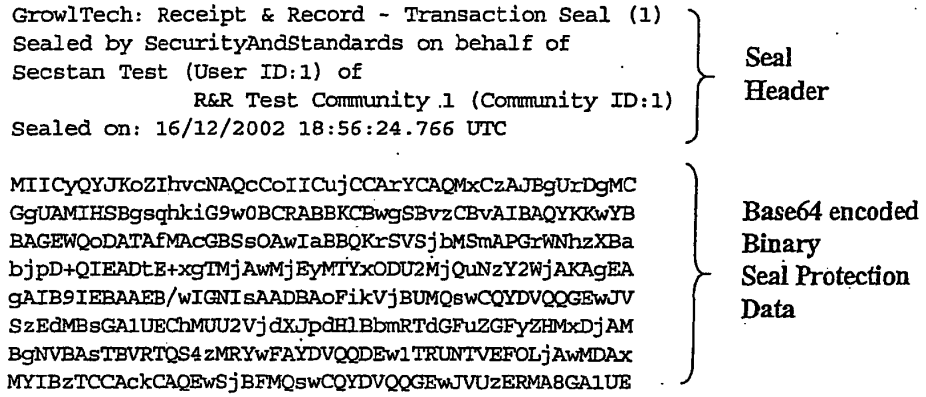

GrowlTech: Receipt & Record - Transaction Seal (1)
Sealed by SecurityAndStandards on behalf of
Secstan Test (User ID:1) of
R&R Test Community 1 (Community ID:1)
Sealed on: 16/12/2002 18:56:24.766 UTC
Seal Header
MIICyQYJKoZIhvcNAQcCoIICujCCArYCAQMxCzAJBgUrDgMC
GgUAMIHSBgsqhkiG9w0BCRABBKCBwgSBvzCBvAIBAQYKKwYB
BAGEWQoDATAfMAcGBSsOAwIaBBQKrSVSjbMSmAPGrWNhzXBa
bjpD+QIEADtE+xgTMjAwMjEyMTYxODU2MjQuNzY2WjAKAgEA
gAIB9IEBAAEB/wIGNIsAADBAoFikVjBUMQswCQYDVQQGEwJV
SzEdMBsGA1UEChMUU2VjdXJpdHlBbmRTdGFuZGFyZHMxDjAM
BgNVBAsTBVRTQS4zMRYwFAYDVQQDEw1TRUNTVEFOLjAwMDAx
MYIBzTCCAckCAQEwSjBFMQswCQYDVQQGEwJVUzERMA8GA1UE
Base64 encoded Binary Seal Protection Data

Figure 8

BEST AVAILABLE COPY

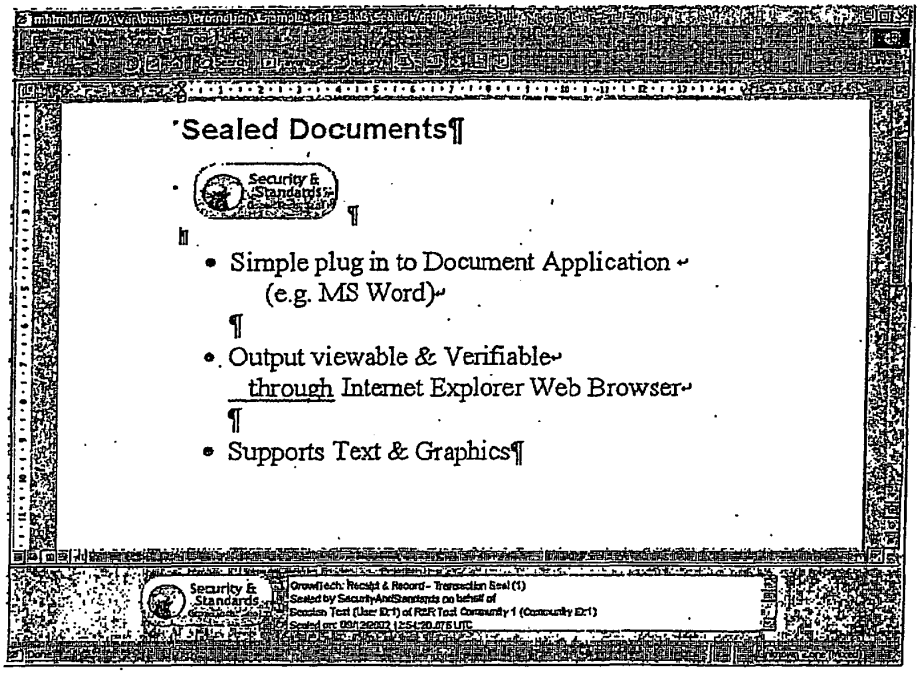


Figure 9

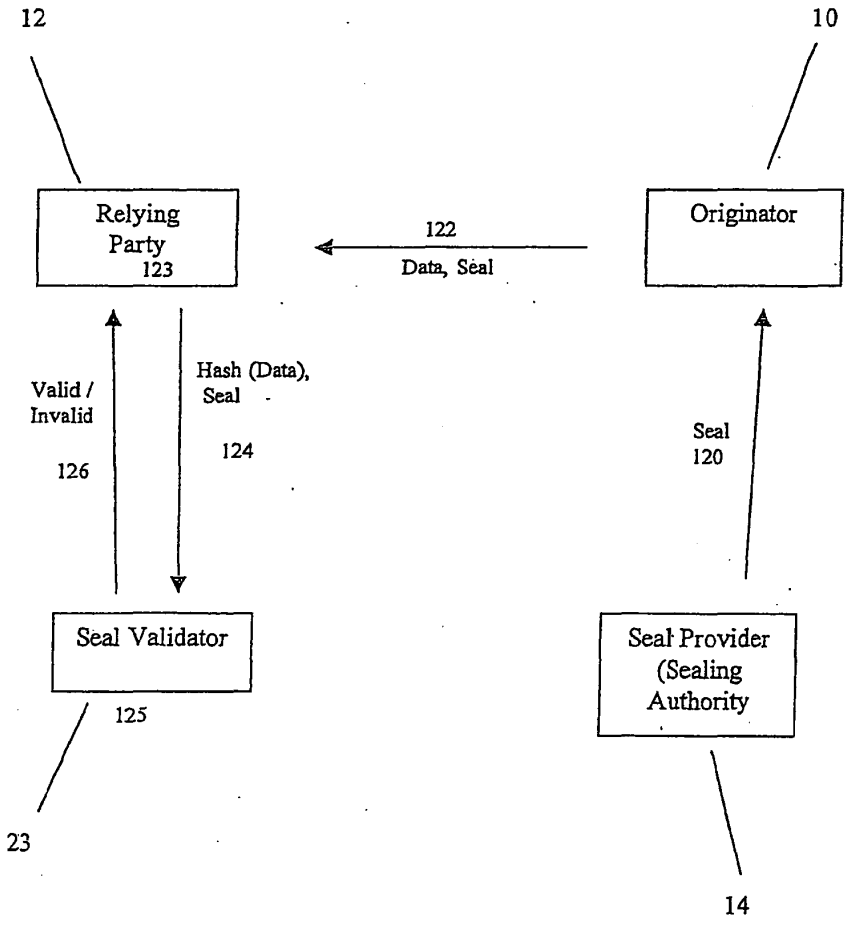

12
10
Relying Party
123
122
Data, Seal
Originator
Valid / Invalid
126
Hash (Data), Seal
124
Seal
120
Seal Validator
125
Seal Provider (Sealing Authority
23
14

Figure 10

INTERNATIONAL SEARCH REPORT

Internatio pplication No
PCT/GB 03/03243

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L29/06 H04L12/58 G06F1/00

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)
EPO-Internal, WPI Data, PAJ

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	ZIXMAIL: "Zixmail technology" INTERNET CITATION, 'Online! 6 September 2000 (2000-09-06), XP002260773 Retrieved from the Internet: <URL:http://www.zixmail.com/zixmail.pdf> 'retrieved on 2003-11-07! the whole document	1-25
X	US 2002/032865 A1 (CHOU CHENG-FU ET AL) 14 March 2002 (2002-03-14) abstract paragraph '0061! - paragraph '0065!; figures 4,5	1-25

-/--

[X] Further documents are listed in the continuation of box C.

[X] Patent family members are listed in annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
27 November 2003	12/12/2003

Name and mailing address of the ISA	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Bertolissi, E

Form PCT/ISA/210 (second sheet) (July 1992)

INTERNATIONAL SEARCH REPORT

Internatic plication No
PCT/GB 03/03243

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2002/078355 A1 (SAMAR VIPIN) 20 June 2002 (2002-06-20) abstract paragraph '0044! - paragraph '0050! figures 1,2	1-25
P,X	US 2002/143710 A1 (LIU GARY) 3 October 2002 (2002-10-03) paragraph '0069! - paragraph '0090!; figures 3,4A,4B	1-25
A	US 6 393 566 B1 (LEVINE JUDAH) 21 May 2002 (2002-05-21) abstract column 2, line 8 -column 3, line 3; figures 1-5	1-25
A	WO 99 35785 A (KONINKL PHILIPS ELECTRONICS NV ;PHILIPS SVENSKA AB (SE)) 15 July 1999 (1999-07-15) abstract	1-25
A	PATENT ABSTRACTS OF JAPAN vol. 013, no. 477 (M-885), 27 October 1989 (1989-10-27) -& JP 01 188322 A (TOYOBO CO LTD;OTHERS: 01), 27 July 1989 (1989-07-27) abstract	1-25

Form PCT/ISA/210 (continuation of second sheet) (July 1992)

INTERNATIONAL SEARCH REPORT

Information on patent family members

Internati Application No
PCT/GB 03/03243

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2002032865 A1	14-03-2002	NONE	
US 2002078355 A1	20-06-2002	AU 3662202 A	24-06-2002
		CA 2426828 A1	20-06-2002
		WO 0248848 A2	20-06-2002
US 2002143710 A1	03-10-2002	NONE	
US 6393566 B1	21-05-2002	NONE	
WO 9935785 A	15-07-1999	US 6601172 B1	29-07-2003
		CN 1254464 T	24-05-2000
		EP 0963637 A2	15-12-1999
		WO 9935785 A2	15-07-1999
		JP 2001515612 T	18-09-2001
JP 01188322 A	27-07-1989	JP 2108598 C	21-11-1996
		JP 7106599 B	15-11-1995

Form PCT/ISA/210 (patent family annex) (July 1992)

C1

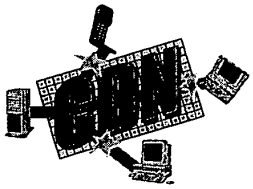
Telematica
Instituut

Content Distribution Networks

State of the art









Telematica
Instituut

Colophon

Date :	June 1, 2001
Version :	1.0
Change :	-
Project reference:	CDN2
TI reference :	TI/RS/2001xx
Company reference :	-
URL :	-
Access permissions :	Anyone
Status :	Final
Editor :	Bob Hulsebosch
Company :	Telematica Instituut
Author(s) :	Rogier Brussee, Henk Eertink, Wolf Huijsen, Bob Hulsebosch, Michiel Rougoor, Wouter Teeuw, Martin Wibbels, Hans Zandbelt.

Synopsis:

This document presents an overview of the current state-of-the-art of Content Distribution Networks.

Copyright © 2001 Telematica Instituut

Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from or via Telematica Instituut (http://www.telin.nl).

Management Summary

World-wide the Internet population is currently estimated at 170 million users. As Internet use continues to grow, greater access speeds create a need for more sophisticated bandwidth to support high-impact sites. E-businesses are experiencing increased pressure to provide fast, reliable content delivery to Web users as site performance can strongly impact a content provider's bottom line in that consumers are more likely to visit and/or purchase from sites that load quickly, reliably and consistently. Moreover, the increased use of rich-media content consisting of audio, video, and images puts a huge load on the storage and network infrastructure. Helping drive the use of such rich-media is a rapid adoption of broadband access technologies that enable applications such as movies-on-demand and videoconference calling. Content Delivery Networks (CDNs) optimise content delivery by putting the content closer to the consumer and shorting the delivery path via global networks of strategically placed servers. CDNs also manage and maintain the network elements that deliver Web content, such as text, images, and streaming audio and video, to the end user, streamlining the entire process. Moreover, a CDN offers unique possibility to provide for value added services like customisation and adaptation of content, virus scanning and ad insertion.

This state of the art investigation presents a state-of-the-art survey of Content Distribution Networks. It gives insight in:

- The current content distribution landscape. It is concluded that there are many emerging competitors in the content-distribution space. To survive, CDN services must expand beyond cache-based delivery to offer application logic and point of interaction customisation. By delivering value-added applications at the edge of the network, content providers are able to develop a more profitable, personalised, and persistent relationship with end-user subscribers.

- The business models for content distribution networks. Based on the value chain of content delivery, we distinguish the following roles (business functions):
  - Content provider (CP, originator, content creator, publisher)
  - Syndicator (content assembler)
  - Distributor (content distribution service provider, CDSP, content networker)
  - Content consumer (consumer, customer, end-user)

  CDN peering allows multiple CDN resources to be combined so as to provide larger scale and / or reach to participants than any single CDN could achieve by itself. Future CDN service scenarios are virus scanning, insertion of ad banner, insertion of regional data, and adaptation of streaming media.

- CDN components, architectures and protocols. I.e., the components that constitute a CDN, the technicalities of finding the most appropriate surrogate server, replication techniques for content caching and distribution, proxy technologies and architectures for streaming and other media, and the protocols that are used within a CDN.

- Content negotiation in a CDN. Content negotiation provides a tool where the client can indicate his preferences and capabilities. It allows CDN providers to offer value-added services based on these negotiation elements. Several protocols for content negotiation are MIME-type based, HTTP, CC/PP, and SDP. The IETF ConNeg

working group has proposed and described a protocol-independent content negotiation framework.

- Content adaptation in a CDN. Besides delivering content, CDNs may also adapt content. For instance by transcoding multimedia streams or by translating from a particular language into another. There is currently new standardisation work being set-up that defines standard mechanism to extend HTTP-intermediates with application-specific value added services (such as virus checking or transcoding). The iCAP protocol for instance facilitates such content adaptation functionality. Middle boxes and media gateways are intermediary devices that may offer additional intelligence for content adaptation or transcoding.

- Authorisation, authentication and accounting. The AAA requirements for a CDN service environment are driven by the need to ensure authorisation of the client, publishing server or administrative server attempting to inject proxylet functionality, to authenticate injected proxylets, and to perform accounting on proxylet functions so the client or publishing server can be billed for the services. In addition, AAA is also required for a host willing to act as a remote callout server. Digital Rights Management (DRM), i.e. the process of protecting and managing the rights of all participants engaged in the electronic commerce and digital distribution of content, will become an important issue in a CDN since original content will be adapted and distributed over the network.

- Related platforms and architectures. In a way, CDN providers offer a (middleware) platform for a wide range of interactive functions, from searching to user profiling to order processing. The Globe middleware platform helps design wide area distributed applications and is in many aspects similar to a CDN platform. Globus, a Grid middleware layer is another example. The areas of distributed operating systems and parallel computing on the one hand (from which Grid comes) and middleware platforms on the other hand (from which CDN comes) seem come closer and might even benefit from each other. Parlay, OSA, and JAIN define standard application programming interfaces that may facilitate rapid deployment of new CDN services.

Based on an analysis of the strengths, weaknesses, opportunities and future threats for a CDN we have observed the following research opportunities:

- ASP and CDN synergy,
- Grid and CDN synergy,
- Broadcasting of streaming media in a CDN,
- Personalisation and localisation (mobility),
- Globalisation.

Table of Contents

1 Introduction **1**
1.1 How do CDNs work? 1
1.2 Reading guide 2
2 Current content delivery landscape **4**
2.1 CDN service providers 4
2.2 CDN market forecasts 6
2.3 Standardisation activities 7
2.3.1 Within the IETF 7
2.3.2 Outside the IETF 7
2.4 Streaming content delivery 8
2.5 Telematica Instituut point of view 9
2.5.1 Bridging distance 10
2.5.2 Bridging time 10
2.5.3 Bridging heterogeneity 10
3 Content distribution services: business models **12**
3.1 Internet developments 12
3.1.1 Internet trends 12
3.1.2 Internet business models 13
3.1.3 Content Distribution Networks 14
3.1.3.1 Functionality 14
3.1.3.2 CDN Business models 14
3.2 Business roles 15
3.2.1 Content provider 16
3.2.2 Syndicator 16
3.2.3 Content distribution service provider 16
3.2.4 Content consumer 17
3.2.5 ISP or local access provider 18
3.2.6 Server capacity provider 18
3.2.7 CDN product manufacturers 18
3.3 Peering CDNs 19
3.4 Future scenarios 21
3.4.1 Virus scanning 21
3.4.2 Insertion of ad banners 21
3.4.3 Insertion of regional data 21
3.4.4 Content adaptation for alternate Web access devices 21
3.4.5 Adaptation of streaming media 22
3.5 Mapping value-added services on business roles 22
4 CDN components, architectures and protocols **23**
4.1 Introduction 23
4.2 Replication 23
4.2.1 Client-Replica protocols 23
4.2.2 Inter-Replica protocols 23
4.3 Caching 24
4.3.1 Proxies 24
4.3.1.1 Filtering Requests 24
4.3.1.2 Sharing Connections 24
4.3.1.3 Improving Performance 25
4.3.2 Caching proxies 25
4.3.3 Web Cache Architectures 26

4.3.4 Caching Protocols 27
4.3.4.1 ICP 27
4.3.4.2 Cache Digests 27
4.3.4.3 HTCP 29
4.3.4.4 CARP 29
4.4 OPES 30
4.5 Streaming Proxies 31
4.5.1 Cached Delivery 32
4.5.2 Replication 32
4.5.3 Unicast Split 33
4.5.4 Multicast Split 33
4.5.5 Pass-Through Delivery 33
4.6 Products 34
5 Content negotiation **35**
5.1 MIME-type based content negotiation 35
5.2 Content negotiation in HTTP 35
5.3 IETF Content Negotiation working group 38
5.4 Transparent Content Negotiation 39
5.5 User (agent) profiles 39
5.5.1 W3C CC/PP (Composite Capability / Preference Profiles) 40
5.6 SDP version 2 41
6 Content adaptation **43**
6.1 ICAP – Internet Content Adaptation Protocol. 45
6.1.1 Benefits of iCAP 45
6.1.2 ICAP architecture 45
6.1.3 Trends and iCAP opportunities 48
6.1.4 ICAP limitations 49
6.2 Middle boxes 49
6.3 Transcoding and media gateways 49
6.4 Transcoding and XML/HTML 50
7 Authorisation, authentication, and accounting **51**
7.1 What is AAA? 51
7.2 AAA definitions 51
7.3 AAA standardisation 52
7.4 AAA in a CDN 53
7.4.1 AAA in the existing Web system model 54
7.4.2 AAA in the service environment caching proxy model 55
7.4.3 AAA in the Remote Callout Server model 55
7.4.4 AAA in the Administrative Server model 57
7.5 Accounting in peered CDNs 57
7.6 DRM 58
7.7 Lack of AAA in current CDNs 61
7.8 Accounting revenue sources 62
8 Other platforms and system architectures **63**
8.1 The Globe middleware, GlobeDoc and the GDN. 63
8.1.1 The Globe system 63
8.1.2 The GlobeDoc System 64
8.1.3 The Globe Distribution Network (GDN). 65
8.1.4 Status 66
8.2 Globus, a Grid middleware layer 66
8.2.1 Grid Security Infrastructure. 67
8.2.2 Globus Resource Management 67

8.2.2.1 QoS Management 67
8.2.2.2 The GRAM resource manager 68
8.2.3 Globus Data Management 68
8.2.3.1 Replica Management 68
8.3 Parlay 69
8.4 3GPP-OSA 71
8.5 JAIN 71
9 Conclusions **73**
9.1 Strength of current CDN approaches 73
9.2 Weakness of current CDN approaches 73
9.3 Opportunities for future CDNs 73
9.4 Threats for future CDNs 74
9.5 CDN research opportunities 74

1 Introduction

The Internet has matured to the point where providing mere connectivity to support Web-browsing and e-mail is no longer the main value. E-business companies, publishers, and content providers view the Web as a vehicle to bring rich content to their customers — wherever they are, whenever they want. Example applications are news and entertainment services, e-commerce transactions, and live sporting events.

The infrastructures supporting this kind of applications is referred to as a CDN: a Content Distribution (or: Delivery) Network. CDNs add management and quality of service to the Internet, e.g., a better performance through caching or replicating content. They offer new possibilities for value-added services, such as localised or personalised content, fast and secure access to content, automatic adaptation of content to increase the 'value of experience', et cetera. It is clear that here lies a potential benefit for content owners, end-users and service providers alike.

1.1 How do CDNs work?

A CDN service is typically accessed through application-specific proxies. Examples are HTTP proxies (for regular Web traffic) and RTSP proxies (for multimedia streaming). These (caching) proxies are located at the edge of the network to which end-users are connected, as depicted in Figure 1.

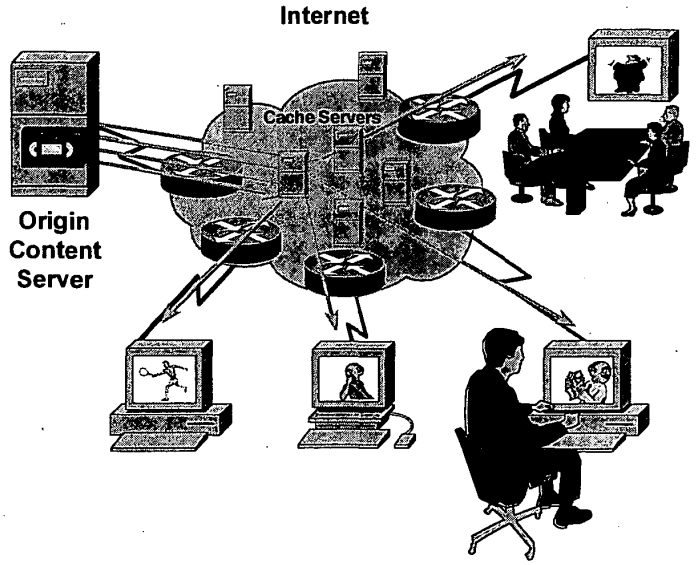


Figure 1: Model of a CDN network.

Each of the nodes in the CDN is located close to the user (access network), which makes it much easier to adapt to varying qualities of end-user equipment or their preferences. Typically, the CDN provides the following functions:

- *Redirection services* to direct a request to the cache server that is the closest and most available.
- *Distribution services*, e.g., a distributed set of surrogate servers that cache content on behalf of the origin server, mechanisms to bypass congested areas of the Internet or technologies like IP-multicast, and replication services.
- *Accounting* services to handle, measure, and log the usage of content.

The CDN infrastructure, consisting of the entirety of services and equipment, provides overlays for the existing Internet, and is necessary to ensure scalable quality management for multimedia content.

Why do CDNs work? A number of advantages are relevant:

1. Less time on the network
    - Reduces network load
    - Less geography to traverse
    - Fewer peers to negotiate with
2. Closer to the clients
    - Fewer hops, traffic stay local
    - Reduces "turns" delay
3. Reduces load on the origin server
    - Reduces processing time
    - Improve consistency and availability of content

In a way, a CDN operates as an intermediary between end-user and content-owner. Content delivery services are typically provided by ISPs, content-owners, or, for internal use, by large companies. They form an interesting vehicle for research, because its operation is invisible for the end-user. That makes it possible to innovate without any need for the user to install new applications, or whatever.

1.2 Reading guide

This CDN state of the art survey is written for anyone involved in content distribution at a (technical) management level. We recommend for this state of the art deliverable to ask your boss for a day off, then go and sit on a beach and start reading. It is the cheapest way to revitalise yourself technically and find out all there is to know about content distribution networks. So once you sink into that comfortable chair overlooking the ocean, we think you will find this deliverable a lot of fun - maybe even relaxing.

This state of the art deliverable focuses on several aspects of content delivery networks.

Chapter 2 and 3 will provide you more general information about CDN networks. Chapter 2 gives you an overview of the current CDN landscape. It presents facts, the actors involved, trends, standardisation activities, and the Telematica Instituut point of view. Chapter 3 tells you something about the existing business models for exploitation of CDN networks en services. Chapter 4 explains the general aspects of content delivery and distribution via the Internet network. An overview of several proxy types, technologies, architectures, and products is given in this chapter. If you do somehow become bored reading Chapter 4, which is sometimes a bit technical, simply jump to the next chapter. In

this chapter, Chapter 5, new and exciting aspects of content negotiation are discussed. Negotiation about e.g. the language of a document or media type format requires knowledge of the user's preferences and system capabilities. Protocols that address the elements for negotiation are described in this chapter.

Another aspect typical for a CDN network is the issue of content adaptation. To provide content-adaptation, the CDN must know what the exact resource-availability of a client (in terms of codecs, drivers, multimedia capabilities, ...) and the intermediate networks (in terms of delay, available bandwidth, ...) are. Examples of content adaptation such as transcoding of multimedia streams, translation of HTML content into WML (for access to mobile WAP devices), insertion advertisements, or translation from a particular language into another are discussed in Chapter 6.

An increasing part of the content of the Internet is access-controlled. Therefore, proper authentication, accounting, and access control is necessary, certainly for third-party content-delivery service providers. There is currently not a single, or standardised, way of doing this. Chapter 7 will address the aspects of Authorisation, Authentication and Accounting for a CDN network.

Chapter 8 addresses platforms and architectures, other then CDNs, for content distribution.

In the concluding Chapter 9, the strengths, weaknesses, opportunities and threats of CDNs are analysed.

2 Current content delivery landscape

This section describes the current CDN landscape. Section 2.1 gives an overview of the traditional CDNs, and shortly describes their business (which will be further elaborated in Chapter 3). Section 2.2 gives an overview of the market forecast for CDN services and products. Section 2.3 describes the current standardisation efforts. Section 2.4 focuses on streaming content, the most important delivery subject of current, performance-based CDNs. Section 2.5, finally, describes the current trends of CDNs to move from caching and replication only towards providing more and more value-added services.

2.1 CDN service providers

Lately, many parties emerge which call themselves 'Content Delivery Networkers' or 'CDNs' or 'CDN service providers'. They vary greatly in infrastructure and business models (see also Chapter 3), but the services they provide are remarkably similar. In general they create a (virtual) world-wide network that speeds ups Web site requests by caching and replicating data. Clearly, the current '1st generation' CDN service providers focus on performance. Today's Internet users are connecting to the Internet at kilobit speeds, and ISPs are having difficulties keeping pace. CDN services provide increased performance to users browsing an Internet content site "*by making Web sites run up to 10 times faster, no matter how many people visit the site*" (Digital Island about its Footprint performance, see www.digitalisland.com).

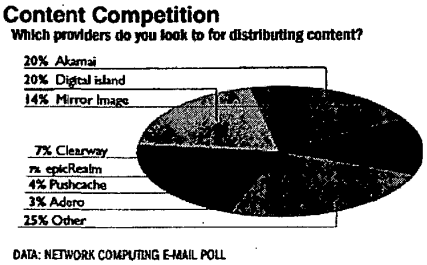


Figure 2: CDN vendors pie (source: Informationweek.com, December 4, 2000 http://www.informationweek.com/815/cdnvendors.htm). Note that the CDN market is rapidly changing; new CDN providers emerge while others disappear or are acquired by other companies.

Figure 2 gives an indication of which CDNs currently share the market for distributing content on behalf of content providers. The annual costs for content providers to subscribe to a CDN service provider are estimated to be 30% cheaper with respect to an in-house hosting model. This includes bandwidth, equipment and labour costs (source: http://www.htrcgroup.com, "The Content Delivery Network Market"). Large CDN service providers are Akamai, Digital Island (including the formal Sandpiper), and Mirror Image. Table 1 gives a characterisation of some example companies providing CDN services (an overview of traditional CDNs, companies delivering CDN software, and satellite-based distribution services can be found on http://www.web-caching.com/cdns.html[1]). From the table it is clear that the different CDNs vary in the way they realise their business.

[1] See also http://www.webreference.com/internet/software/site_management/cdns.html.

Table 1: Examples of Content Distribution Service Providers

COMPANY	DESCRIPTION
Adero (www.adero.com)	Adero provides high performance, quality enhanced content delivery solutions to carriers and hosting providers through the established Adero™ GlobalWiseSM Network and content delivery services. The GlobalWise Network is comprised of strategically placed servers around the world, which redirect content closer to the audience for on-net enhanced services.
Akamai (www.akamai.com)	Akamai sells technology designed to speed up delivery of content over the Internet. Akamai operates a global Internet content-delivery network designed to alleviate Web-server overload and shorten the distances that content travels. Servers are distributed in a wide geographical area, putting Akamai's caching service closer to users at the edge of the network. This decreases network congestion and decreases the response time of its customers' Web sites. The system is distributed with no central control, which allows for self-correction if a part of the system fails, from servers to entire backbones. Akamai's widely distributed caching network has made it a market leader.
CacheWare (www.cacheware.com)	CacheWare specialises in content distribution and caching from an origin server to edge servers. Its CacheWare Content Manager takes the load off an origin server by acting as the intermediary between origin and edge servers.
Cidera (www.cidera.com)	Cidera's network is satellite-based and specialises in transporting data streams. It has more than 300 points of presence in North America and presence in Europe, with expansion into Latin America and Asia later this year.
Clearway (www.clearway.com)	Clearway is a provider of server-based content delivery solutions that provides Web performance services to e-businesses of all sizes. Using Clearway's services, the customers can bypass technical performance barriers. Clearway has been acquired by Xcelera/Mirror Image in January 2001.
Digital Island (www.digitalisland.com), includes the former Sandpiper Networks	Digital Island provides global application hosting and content distribution over a private network that bypasses oversubscribed public networks. Digital Island has Web-hosting facilities in New York, Santa Clara, Honolulu, Hong Kong, Tokyo, and London that provide network access to 27 countries. Streaming content delivery is also provided.
EpicRealm (www.epicrealm.com)	EpicRealm operates a global network to provide prioritised traffic flow control and constant connection with the user, in addition to fast, reliable content delivery. The company launched its global network in April 2000.

COMPANY	DESCRIPTION
IBeam (www.ibeam.com)	Ibeam is a broadcasting company whose technology and infrastructure streams high-quality content to audiences.
Mirror Image (www.mirror-image.com)	Mirror Image exploits a global "Content Access Point" (CAP) infrastructure on top the Internet to provide content providers, service providers and enterprises with a platform that delivers Web content to end users. Besides content distribution, streaming and content access services are provided. As a secure and managed high-speed layer on top of the Internet, each CAP offloads origin servers and networks by intelligently placing content at locations closer to users world-wide.
Pushcache (www.pushcache.com)	Pushcache.com is an Internet software company focused on the development and sale of products based on pushcache, a communication and middleware architecture, based on Web caching, capable of scalable and flexible data delivery.

A detailed list of CDN organisations and the products they offer can be found in Appendix B - Overview of CDN organisations.

2.2 CDN market forecasts



The Internet continues to grow as both the amount of content and the number of online users increases. The growth rate of content on the Internet is significantly increasing as organisations around the world deploy new content types, such as streaming media and dynamic content, to further Web site differentiation. *"The growth of the Internet continues to blaze forward at an incredible rate, fuelling the content delivery network (CDN) market"*, said Greg Howard, principal analyst and founder of the HTRC Group, LLC. *"The world-wide CDN product market will grow from $122M in 2000 to an estimated $1.4B by 2004"* (see Figure 3).

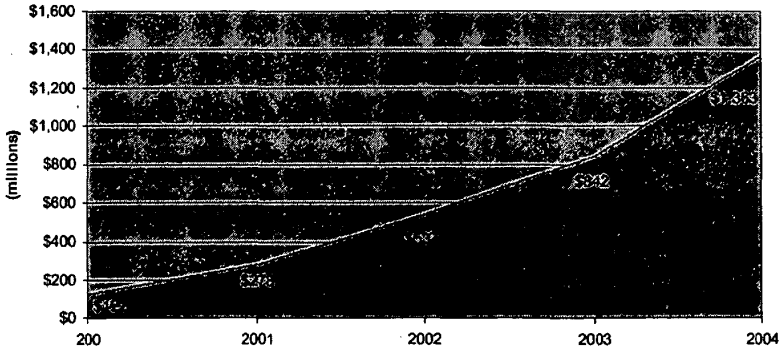


Figure 3: World-wide CDN products forecast (source: http://www.htrcgroup.com).

The new CDN *service* market will grow significantly to an estimated $2.2 billion by 2003 (source: http://www.htrcgroup.com/). The expected growth of the CDN *service* market is shown in Figure 4.

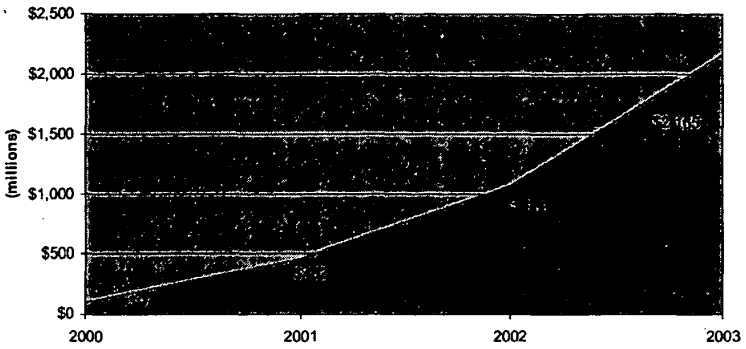


Figure 4: Growth estimation of the CDN service market (source: http://www.htrcgroup.com).

2.3 Standardisation activities

Since CDN is a booming business, vendors are organising themselves and standardisation activities emerge. The Internet Engineering Task Force (IETF) is the official body that defines standard Internet operating protocols. In this section, the most important working groups inside and outside the IETF are listed (a further elaboration on some standardisation issues can be found in Chapter 6).

2.3.1 Within the IETF

The following relevant working groups active within the IETF are:

- Web Replication and Caching (WREC, www.wrec.org, expired): Worked on a taxonomy for Web replication and caching.
- Middlebox Communication (MIDCOM, www.ietf.org/html.charters/midcom-charter.html): Works on protocols allowing applications to communicate their needs to devices in the network (see also section 6.2).
- Reliable Multicast (RMT, www.ietf.org/html.charters/rmt-charter.html): Works on a protocol framework for reliable multicast communication.
- Web Infrastructure (WEBI, www.ietf.org/html.charters/webi-charter.html): Addresses issues specific to intermediaries in the World Wide Web infrastructure.
- Content Distribution Internetworking (CDI, www.content-peering.org/ietf-cdi.html): Addresses issues specific to internetworking of CDNs ("Content Peering").
- Open Pluggable Extension Services (OPES, www.ietf-opes.org): Works on a framework and protocols for extensible content services (see also section 4.4).

2.3.2 Outside the IETF

The following relevant working groups active outside the IETF are:

- World Wide Web Consortium (W3C, www.w3.org): Develops interoperable technologies for the World Wide Web (e.g. HTML, XML, SMIL, etc.).

- Broadband Content Delivery Forum (BCDF, www.bcdforum.org): Consortium of Internet infrastructure, content and service providers to develop standards for delivering broadband content.
- ICAP forum (www.i-cap.org): Consortium of infrastructure and service companies promoting the Internet Content Adaptation Protocol (iCAP); technical work will move into the IETF (OPES group). See also section 6.1 for more information.
- Content Alliance (www.content-peering.org) and Content Bridge (www.content-bridge.com): Consortia of infrastructure and service companies working on issues related to content peering; technical work will move to IETF CDI group.
- Internet Research Task Force - Reliable Multicast (RMRG, www.irtf.org/charters/reliable-multicast.html): Group looking into research issues related to reliable multicast (delivers into IETF RMT Working Group).
- The Wireless Multimedia Forum (WMF, www.wmmforum.com) is an international, multi-vendor forum and gathering point for vendors developing products, services and information focused on the delivery of rich media content to mobile, wireless devices.
- Internet Streaming Media Alliance (ISMA, www.ism-alliance.org or www.isma.tv). The goal of ISMA is to accelerate the adoption of open standards for streaming rich media — video, audio, and associated data — over the Internet.

2.4 Streaming content delivery

For CDNs, the bandwidth consuming streaming applications are important. Often, CDNs only deliver the streaming content whereas 'regular' ISPs deliver the other parts of a Web site.

Delivering quality multimedia over the Internet presents many challenges. Before streaming, audio and video clips had to be downloaded in their entirety before use. Streaming lets a media player start presenting content as it arrives: frame by frame. To speed delivery, media is commonly transported over UDP. But datagrams get lost and arrive out of order. Many players buffer frames to improve quality of the stream presented to the end user. Streaming media therefor form a prime candidate for caching. Early cache extensions for real-time streaming protocols enabled stream splitting. Bandwidth was saved when more than one media player shared the same live broadcast, conveyed just once across the backbone between media server and cache. These caches proxied live streams, but they did not store them.

Streams are, by nature, bandwidth intensive. When many streams compete for resources over a highly variable and lossy medium like the Internet, client-side buffering is not enough. Delaying an occasional HTTP/TCP packet a few extra seconds degrades user experience. Delaying streamed packets, however, can render multimedia unplayable, and dropped UDP packets leave "holes" in the stream.

Furthermore, interaction between client and server is required for "rich" multimedia. End users want VCR-like controls that pause, rewind, forward, and index into streams. Content providers who own media servers want the ability to authenticate and charge users for delivered streams. The Real Time Streaming Protocol (RTSP) enables set-up and control over streams delivered from media server to client. RTSP acts as a network remote control. The content itself is delivered using data protocols like RTP or Real's proprietary RDT. These real time transport protocols allow frames that arrive out of order to be reassembled with the intended timing and sequencing.

These protocols are, of course, used between origin server and media player client. But proxies to relay live broadcasts and on-demand content can also use them. During a live broadcast, an RTSP proxy uses one data session to receive the stream from a media server. It may split the stream to several clients, deliver the stream over IP multicast to many clients, or pass the stream through to a single client. In each case, the proxy accounts for use by establishing an RTSP control session per client. Only authorised clients can receive the stream, and statistics are returned to the media server for each client. Live stream delivery is analogous to pay-per-view TV— consumers join the regularly scheduled program and pay for what they watch.

Another delivery model resembles video-on-demand: Consumers request a movie whenever they want, and have discretionary control over playback (pause, rewind, etc.). On-demand stream delivery from origin server to media client can be impractical, costly, or completely impossible, depending upon network bandwidth, speed, and loss. Delivering on-demand streams across the backbone in volume would quickly gobble up capacity, even if quality of service could be adequately controlled. Often it cannot be.

Clearly, the most economical approach for delivering high-quality on-demand streams is to cache them at the edge of the destination network, for example, at the broadband CLEC (Competitive Local Exchange Carrier) head-end, backbone NAP (Network Access Point), or ISP POP (Point of Presence), with a media cache. Like a Web cache, a media cache records content supplied by origin servers in response to user requests. When the same user, or a different user, requests the same content, it is fetched from the cache instead of the origin server. When delivering cached content, response is faster, quality is higher, and upstream bandwidth consumption and costs are reduced. Like their Web counterpart, media caches must verify content freshness and may operate in transparent or non-transparent mode.

But the content stored by a media cache must be treated as a borrowed asset, made available for resale. The media cache must proxy licensing schemes, authentication, accounting and usage statistics by establishing an RTSP control session to the origin server whenever a client requests previously-cached content. The cache should protect against unauthorised access to stored content. For example, the Real Networks RealProxy 2.0 encrypts content replicated locally, and terminates client streams if the server becomes unreachable during replay.

Media caches can also leverage their role to overcome quality problems that plague live streams (source: ISP-Planet - Equipment - Stream Caching with TeraEDGE - http://www.isp-planet.com/equipment/teraedge1.html).

2.5 Telematica Instituut point of view

Currently, we observe that CDNs start to focus on Content Distribution Networks 'Next Generation', which should focus on integrating value-added services into the CDN, e.g., enabling e-Commerce. In our vision, we are heading towards 'Adaptive Broadband Content Delivery', as shown in Figure 5. Data is provided via the Internet in a multi-channel way (office documents, audio, video, Web pages, etc.), as well as accessed though different access networks (telephony, cable, mobile, etc.). The intermediary Internet bridges distance, time and heterogeneity.

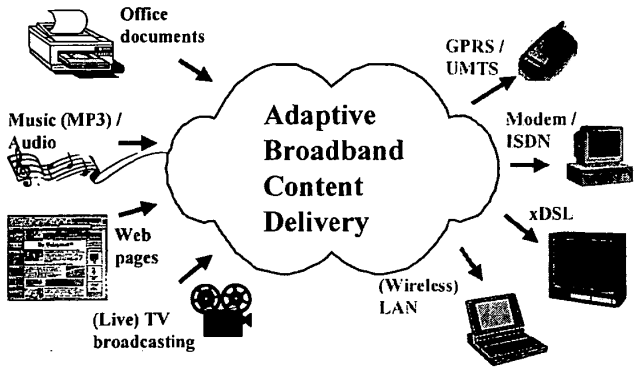


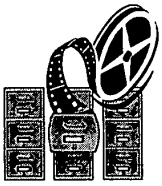
Figure 5: Multi-channel, multi-access content delivery.

2.5.1 Bridging distance

A CDN bridging distance primarily means a CDN delivering content as the traditional CDNs do. For content providers and consumers this means globalisation: content available any time, any place, everywhere. This requires available and scalable solutions for content distribution management and delivery.

Bridging distance also means that is guaranteed what is delivered and how it is delivered. Knowing what is delivered is important for accounting-related issues like monetising content delivery or digital rights management (see section 7.6). Knowing how the content is delivered is a Quality of Service issue. If the content is delivered in a poor quality, the content provider received complaints, even if being 'innocent' with respect to the causes of poor quality. In that case one may cease distributing the data, or take financial measures as agreed in service level agreements.

2.5.2 Bridging time



A CDN bridging time means the Internet becomes a single large archive, a distribution medium for digital data. Data is inserted into the Internet, and may be retrieved years later on. To do so, content management functionality is required. Among them are searching and retrieval issues, e.g., summarising and indexing the data. Indexing the data becomes important. Other issues are storage service provision, data warehousing, and the persistence of digital storage formats in time.

Besides storing data for a long term, data may also be retrieved immediately after being provided by the publisher. Real-time streaming is the example. The CDN should support the entire spectrum from real-time retrieval to long-term archiving.

2.5.3 Bridging heterogeneity

There is a growing diversity and heterogeneity in types and capabilities of client devices as well as the forms of network connections that people use to access the Web. Clients include cell phones and PDAs as well as PCs, TVs (with SetTop boxes), etc. However, these appliances have quite diverse display capabilities, storage, processing power, as

well as slow network access. As a result, Internet access is still constrained on these devices and users are limited to only a small fraction of the total number of Web pages/content available in the Internet today.

Besides heterogeneity in network and terminal characteristics, there is heterogeneity in services/applications, data formats and personal preferences as well. Also, location-based adaptation of content may be desired (local advertisements, language conversions, etc.). Possible adaptations to meet the special requirements of different Web access devices are:

- Conversion of HTML pages to WML (Wireless Markup Language) pages
- Conversion of JPEG images to black and white GIF images
- Conversion of HTML tables to plain text
- Reduction of image quality
- Removal of redundant information
- Stripping of Java applets / JavaScript
- Audio to text conversion
- Video to key frame or video to text conversion
- Content extraction

One has to ensure that the automatic adaptation process will not make changes to a Web page that are unwanted by either the content provider or the recipient. A strategy to achieve this would be to allow the content provider as well as the client to define their preferences as to how they want Web pages to be adapted. The actual adaptation decisions would then be made based on the given preferences and a set of transformation rules. There would have to be a mechanism of resolving potential conflicts between the content provider's and the user's adaptation preferences. If neither the content provider nor the client has expressed his preferences, a default adaptation of the requested Web page may be possible but investigation is needed.

We conclude that there are many emerging competitors in the content-delivery space. To survive, CDN services must expand beyond cache-based delivery to offer application logic and point of interaction customisation. By delivering value-added applications at the edge of the network, content providers are able to develop a more profitable, personalised, and persistent relationship with end-user subscribers.

3 Content distribution services: business models

This section is on business models for content distribution networks. Section 3.1 describes the relevant Internet trends which have led to the emerge of content distribution networks, as well as explain why content distribution networks are a solution. Section 3.2 lists the several types of actors involved in content distribution networks and describe the economic value for them in terms of their benefits in using a CDN. Section 3.3 focuses on CDN peering as a business model. Section 3.4 lists some future scenarios for CDNs that are based on value-added services. Section 3.5 finally links these scenarios to the business roles of section 3.2.

3.1 Internet developments

3.1.1 Internet trends

The Internet has matured to the point where providing mere connectivity to support Web-browsing and e-mail is no longer the main value. E-business companies, publishers, and content providers view the Web as a vehicle to bring rich content to their customers — wherever they are, whenever they want. Example applications are news and entertainment services, e-commerce transactions, and live sporting events. This means a shift in paradigm from connectivity (e-mail, Web browsing) to content (see Figure 6)[2].

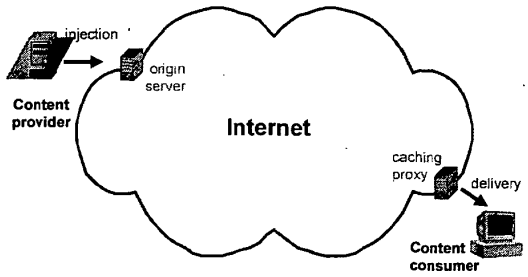


Figure 6: Content delivery value chain

The Internet traffic grows at a considerable rate. Some people claim the Internet traffic is doubling every 3-4 months. Others state a growth rate of 100% per year as more realistic[3]. Even such a moderate rate, however, is still provoking with respect to Internet management and Quality of Service issues. Also, note that *the* Internet does not exist. Internet is a collection of interconnected networks (see Figure 7). The number of networks is increasing as well, and is currently over 7000.

[2] Note that there are also voices that believe the huge sums being invested by carriers in content are misdirected. In particular, see: A. M. Odlyzko, 'Content is not king', *First Monday* 6(2) (February 2001), http://firstmonday.org.

[3] A. M. Odlyzko, 'Internet growth: Myth and reality, use and abuse',. to appear in *J. Computer Resource Management*, April 2001. http://www.research.att.com/~amo

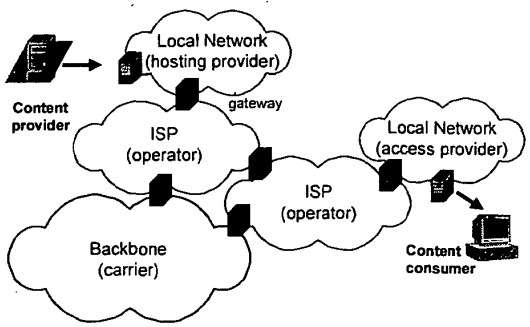


Figure 7: Many networks interconnected.

Finally, end-users want customised services. The trend to personalisation of Web content is noticeable. A good example is the My.Yahoo.com personal Web site. This trend is reflected in the emerging business models for the Internet.

3.1.2 Internet business models

The emerge of the Internet and so-called 'e-commerce' (end-user transactions like the 'Amazons') and 'e-business' (business to business) has led to a shift in paradigm from products to services. Buzzwords are 'apps-on-tap' and 'application hosting'. Portals handle the dynamic brokerage of electronic services. What is needed in this situation is a 'retailer', like defined in the TINA Business Model[4]. A retailer provides end-users with a single point of access, where they can describe to a wide variety of services. Service providers get access to a wide range of end-users. Retailers use broker mechanisms to match user request (required functionality) with (third party) provided services, i.e., to find provided services. Moreover, the retailer may provide 'generic' functionality like access control, authentication, accounting, transaction services, hosting services, etc., which means that the 3rd party service providers can focus on their core business only, the development of application functionality[5].

On the content (rather than service) level, we observe the same kind of development, which is called syndication. Syndication is defined as sale/licensing of the same goods (in particular content, but in principle anything can be syndicated) to multiple customers, who then integrate it with other offerings and redistribute it. Syndication is called the future Internet business model[6]. The most common example of syndication is in newspapers, where such content as wire-service news, comics, columns, horoscopes, and crossword puzzles are usually syndicated content. Newspapers receive the content from the content providers, reformat it as required, integrate it with other copy, print it, and publish it. For many years mainly a feature of print media, today content syndication is the way a great deal of information is disseminated across the Web. Reuters, for example, provides online news content to over 900 Web sites and portals, such as Yahoo and

[4] Mulder, H. (ed.), *TINA Business Model and Reference Points: Version 4.0*. TINA-Consortium, May 1997. http://www.tinac.com/specifications/specifications.htm
[5] See also: W.B. Teeuw *et al.*, 'Samenwerken en zakendoen via het Internet', *Architectuur & Infrastructuur* (tenHagenStam), nr. 1, 2001 (in Dutch).
[6] K. Werbach, 'Syndication: The emerging model for business in the Internet era', *Harvard Business Review*, 78 (3), May-June 2000, 84-93.

America Online. Online content syndication is a growing industry sector, in terms of both content syndication and hardware and software development.

3.1.3 Content Distribution Networks

3.1.3.1 Functionality

In a technical sense, the syndication business model perfectly suites a CDN, which generally distribute modules of content from the content providers to content users, often retailers that assemble the content and deliver it to the end-user. CDNs mainly provide content hosting and distribution capabilities to content providers. They add management and quality of service to the Internet, e.g., a better performance through caching or replicating content. They offer new possibilities for value-added services, such as localised or personalised content, fast and secure access to content, automatic adaptation of content to increase the 'value of experience', et cetera. It is clear that here lies a potential benefit for content owners, end-users and service providers alike.

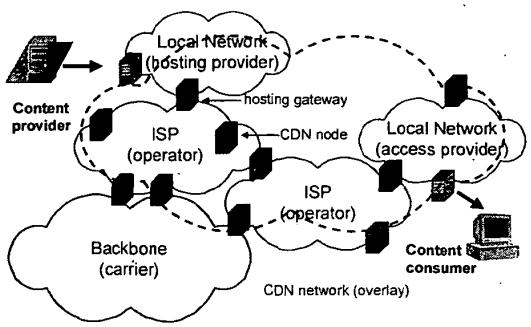


Figure 8: CDN overlay network.

As shown in Figure 8, a content distribution network is an 'overlay' network, consisting of servers being hosted in several networks of possibly different Internet service providers. The typical functionality of a CDN includes:

- *Redirection and delivery services* to direct a request to the cache server that is the closest and most available, e.g., using mechanisms to bypass congested areas of the Internet or technologies like IP-multicast.
- *Distribution services* like a distributed set of surrogate servers that cache content on behalf of the origin server, or alternatively using replication of data.
- *Content negotiation services* that automatically take care of network and terminal capabilities as well as handle user preferences.
- *Content adaptation services* to convert formats or to include (local) advertisements.
- *Management services* to handle, e.g., accounting or digital rights management or services to monitor and report on content usage.

3.1.3.2 CDN Business models

A CDN is an answer to the trends mentioned above because from a business point of view a CDN brings:

1. *Performance* in an era of traffic growth and Internet congestion.
2. *Money* through economies of scale, advertisement and accounting features.

3. *Personalisation* to content customers, including location awareness.
4. *Value-added services* to support, e.g., e-commerce.

We define the role of a Content Distribution Service Provider (CDSP) as the actor who provides the infrastructure for content distribution (or delivery) network services. In section 3.2.3 we show several examples of CDSPs. We observe that there is a large variety of service providers that we may call content networkers. Within the margins of a CDSP, their business may differ in several ways. Most important choices are:

- The CDSP may partly own its private network (like Digital Island), or may place its CDN servers at the edges of as many facilities-based providers as possible, creating an internetworking of CDN servers that cross multiple ISP backbones (like Akamai).
- The CDSP may provide end-to-end content delivery, or alternatively the client (end consumer) will be connected to the CDSP through an ISP or local access provider ('ISP owns the customer'). The latter case is more common and has the advantage that an access provider may use CDNs while still keeping their customer relationships.
- Identically, on the server side the CDSP may 'own the content provider' and have a customer relationship with them, or the CDSP may support the hosting service provider (like Adero does)
- The CDSP may host entire Web pages, or only (streaming) parts of a Web page. Currently the latter model is more common, with CDSPs bringing streaming content to the end-user whereas the service providers deliver the text parts of a Web page. The trend, however, seems that more and more CDSPs host entire pages.
- The CDSP may provide network services across a large geographic area —for example, large ISPs may become CDN service providers— or CDSP may interwork, as shown in Figure 9.

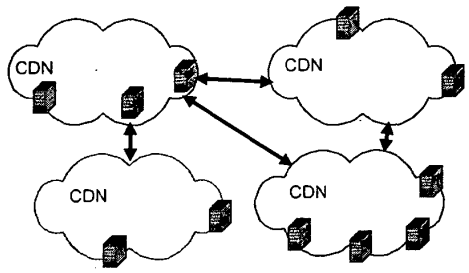


Figure 9: CDN Interworking (CDNI) or peering.

3.2 Business roles

In this section we focus on what CDNs bring to the different stakeholders in terms of economic advantage. Based on the value chain of content delivery, we distinguish the following roles (business functions):

- Content provider (CP, originator, content creator, publisher)
- Syndicator (content assembler)
- Distributor (content distribution service provider, CDSP, content networker)
- Content consumer (consumer, customer, end-user)

Besides we distinguish the following supporting roles, which we grouped into clusters of stakeholders that benefit from CDNs in identical ways:

- Connectivity provider (like a Local access provider or Internet service provider (ISP))

- Server capacity provider (like a Hosting provider or Storage server provider)
- Product manufacturers (both hardware and software)

In the following we will discuss these roles, provide examples and list their benefits if using a CDN. Notice that a single market party may fill in different roles.

3.2.1 Content provider

The core business of content publishers is the creation of content. Examples are Bertelsmann, BBC, and NOS.

The advantages for a content provider of using a CDN are:
- Being relieved of content hosting and distribution issues, among which rights management, accounting and format conversions (= focus on their core business).
- Keep in full control of the content by managing the so-called origin server, which governs content, access rights and policies.
- Get insight into content usage through the usage statistics reported by the CDSP to the content provider.
- Higher uptime: By using a CDN (caching) information may still be available even if the origin server is down.
- Lower costs: Consultancy studies show that CDNs provide the highest performance for the lowest costs compared with in-house sort out or storage service provision[7].

3.2.2 Syndicator

The syndicator assembles and packages content, and manages the relationship between content providers and distributors. This role is 'optional' in the sense that content providers may directly distribute (or let distribute) their own content. Examples of a syndicator are iSyndicate (www.isyndicate.com) and Tanto (www.tanto.de).

The advantages for a syndicator of using a CDN are:
- They need a 'meta service' network like a CDN as distribution channel (seamless distribution of content).
- CDNs allow real-time content assembly (due to their performance).

3.2.3 Content distribution service provider

The CDN service providers are basically in the business of bringing management and quality of service (QoS) to Internet services. Based on service level agreements with their customers (content providers) they distribute and deliver content. Their resources provide caching and replication of data, and refreshing the data if invalidated; they perform content negotiation and adaptation; authorisation, access control and accounting capabilities; and may also use third-party clearing house services. Also they may provide other value-added services, like virus checking and advertisement insertion. Examples of CDSP are shown in Table 1 in section 2.1.

[7] *The Content delivery Network Market*, HTRC Group, San Andreas, CA, http://www.htrcgroup.com

The service level agreement (SLA) between a content distribution service provider and a content provider typically includes:

- Performance and availability over a period, either absolute (network uptime, response times) or relative to base (= non-CDN) pages. Performance should be measured from a statistically valid number of measurement points (end-users).
- Financial/contractual consequences for failure to meet minimum criteria.
- Financial/contractual incentives for exceeding acceptable criteria.
- Response times for reported problems.

Content providers pay content distribution service providers for their delivered services. The general payment model is based on data usage[8]. Every 5-12 minutes the CDSP measures the traffic concerning the data of the content provider (in Mbyte/s). These statistics are gathered during a time-interval (e.g., a month), are topped-off (e.g., the top 5% is left out), and the remaining rate (in Mbyte/s per month) is multiplied with the tariff. Also, payment models exist that charge on a region base, i.e., the number of regions supported by the CDSP (a strategy supported by Adero).

Monetising content distribution services is a hot issue: who pays? Looking to content provision and content distribution on the Internet, there are three possible answers to this question. Money is coming from either sponsoring, or advertisement or the end-user. So far, advertisement has been very important. There is much 'branding' on the Internet. With the emerge of e-commerce and transactions, content networkers expect that more and more the end-user is going to pay, e.g. for value-added services. End-users may pay on a subscription-, usage-, or content-base. Opinions differ about whether users want to pay for content. On the one hand users want quality and will pay for it. On the other hand people do not pay for what you can get for free somewhere else on the Internet. Figure 10 shows the flows of money between several stakeholders.

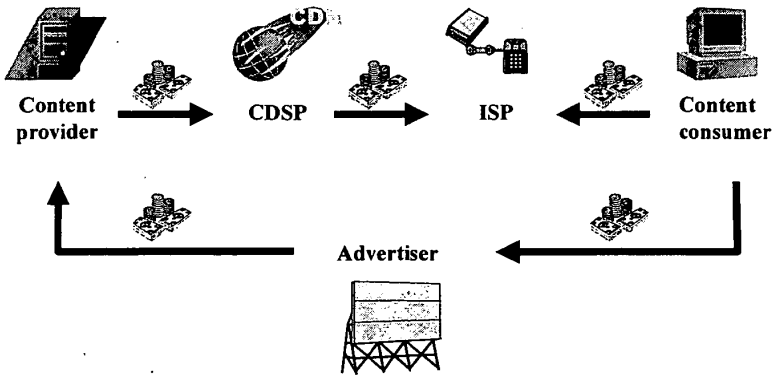


Figure 10: Cash flows in CDN.

3.2.4 Content consumer

Those who use or view the content have the following advantages if using a CDN:

[8] A.M. Pugh, 'Blowing away Web delays', *Data Communications*, October 1999, 30-38, http://www.data.com

- Personalised services, i.e., customisation through personal preferences, tune on terminal equipment and available network bandwidth, et cetera.
- Quality of Service (performance).
- Single point of payment and transparency of costs? This advantage can be achieved if, e.g., a syndicator assembles the requested content or the CDSP acts as a retailer.
- Quality of content?

3.2.5 ISP or local access provider

ISPs and local access providers provide connectivity in either backbone or access networks. Note that the access networks may also be wireless, mobile or cable. Examples of ISPs are KPN, Libertel, Telfort, AT&T, etc. Examples of access providers are World Online, Freeler, Planet Internet, etc.

The advantages for an ISP or access provider of using a CDN are:
- Having the CDNs as bandwidth-consuming customers and sell bandwidth to them.
- Faster response times and higher throughput because bandwidth usage is reduced through CDN caching.
- Being able to provide CDN services to their customers.

3.2.6 Server capacity provider

'Service capacity providers' will provide the storage and server capacity needed for CDNs. Typical examples of server capacity providers or hosts are ISPs who provide this functionality. Other examples are data warehouses or Storage Server Providers (SSP) that lease, e.g., storage area networks (SAN) or network attached storage (NAS) solutions. Because there is a large market for CDNs there is a significant opportunity for server capacity providers as well. An example of an SSP is Managed Storage (www.managedstorage.com).

3.2.7 CDN product manufacturers

CDN product manufactures provide infrastructure, both hardware and software, needed for CDNs. They include accounting and billing product vendors, third-party clearinghouse services, content signalling technologies, caching, load-balancing, redirection appliances and e-commerce products. Relevant examples are:
- Vendors of network infrastructure like Cisco (www.cisco.com), Lucent (www.lucent.com) or Nortel (www.nortel.com).
- Vendors of caching hardware like Cacheflow (www.cacheflow.com), Cisco (www.cisco.com), InfoLibria (www.infolibria.com), or Network Appliance (www.netapp.com).
- Vendors of caching software like Inktomi (www.inktomi.com), Novell (www.novell.com), and NLANR (www.squid-cache.org).

Like for server capacity providers, there is a large market for the makers of CDN network elements as well.

3.3 Peering CDNs

CDN peering allows multiple CDN resources to be combined so as to provide larger scale and/or reach to participants than any single CDN could achieve by itself. At the core of CDN peering are four principal architectural elements that constitute the building blocks of the CDN peering system. These elements are the Request-Routing Peering system, Distribution Peering system, Accounting Peering system, and Surrogates. Collectively, they control selection of the delivery CDN, content distribution between peering CDNs, and usage accounting, including billing settlement among the peering CDNs. In order for CDNs to peer with one another, it is necessary to interconnect several of these core elements of the individual CDNs. The interconnection of CDN core system elements occurs through network elements called CDN Peering Gateways (CPGs). Namely, the core system elements that need to be interconnected are the Request-Routing system, the Distribution system, and the Accounting system. The net result of peering CDNs is that a larger set of surrogates becomes available to the clients. Figure 11 shows a conceptual overview of three CDNs, which have peered to provide greater scale and reach to their existing customers.

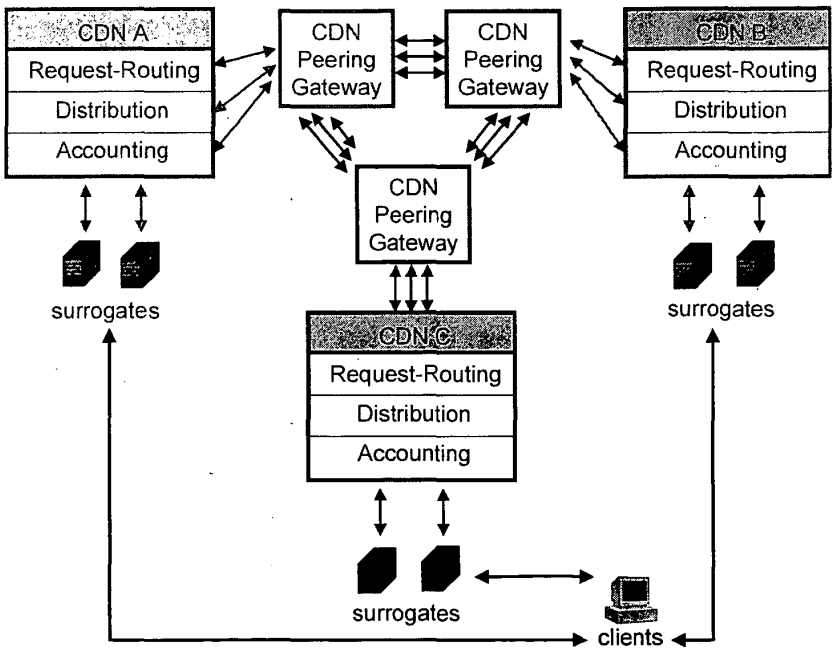


Figure 11: Peering existing CDNs (source: ."CDN Peering Architectural Overview" IETF Internet draft, http://www.ietf.org/internet-drafts/draft-green-cdnp-gen-arch-03.txt). The CDNs are peered through interconnection at CPGs. The result is presented as a virtual CDN to clients for the delivery of content by the aggregated set of surrogates.

The system architecture of a CDN peering system is comprised of seven major elements, three of which constitute the CDN peering system itself. Figure 12 contains a system architecture diagram of the core elements involved in CDN peering. The arrows in the diagram represent the following dynamics:

1. The Origin delegates its URI name space for objects to be distributed and delivered by the peering CDNs to the Request-Routing peering system.
2. The Origin publishes Content that is to be distributed and delivered by the peering CDNs into the Distribution peering system. Note: Content which is to be pre-populated (pushed) within the peering CDNs is pro-actively published, while Content which is to be pulled on demand is published at the time the object is being requested for Delivery.
3. The Distribution peering system moves content between CDN Distribution systems. Additionally this system interacts with the Request-Routing peering system via feedback Advertisements to assist in the peered CDN selection process for Client requests.
4. The Client requests Content from what it perceives to be the Origin, however due to URI name space delegation, the request is actually made to the Request-Routing peering system.
5. The Request-Routing peering system routes the request to a suitable Surrogate in a peering CDN. Request-Routing peering systems interact with one another via feedback Advertisements in order to keep request-routing tables current.
6. The selected Surrogate delivers the requested content to the Client. Additionally, the Surrogate sends accounting information for delivered content to the Accounting peering system.
7. The Accounting peering system aggregates and distils the accounting information into statistics and content detail records for use by the Origin and Billing organisation. Statistics are also used as feedback to the Request-Routing peering system.

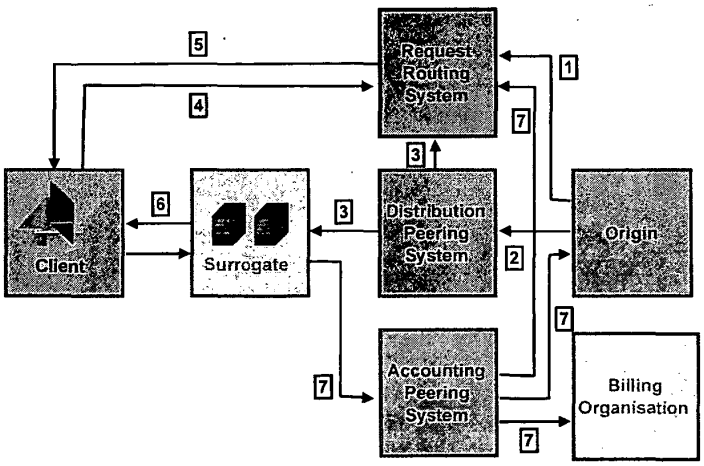


Figure 12: System architecture elements of a CDN peering system (source: ."CDN Peering Architectural Overview" IETF Internet draft, http://www.ietf.org/internet-drafts/draft-green-cdnp-gen-arch-03.txt).

Note that the request-routing peering system is the only mandatory element for CDN peering to function. A distribution peering system is needed when the publisher does not have a negotiated relationship with every peering CDN. Additionally, an accounting peering system is needed when statistical and usage information is needed in order to satisfy publisher and/or billing organisation requirements.

3.4 Future scenarios

This section discusses several service scenarios that could be implemented on top of a CDN platform. For a more detailed description of the scenarios presented below and for other scenarios see the IETF Internet draft: *Example services for network edge proxies* - http://www.ietf.org/internet-drafts/draft-beck-opes-esfnep-01.txt .

3.4.1 Virus scanning

Viruses, Trojan Horses, and worms have always posed a threat to Internet users. Just recently, for instance, a number of e-mail based worms have hit millions of Internet users world-wide within a few hours. With the help of a content scanning and filtering system at the caching proxy level, Web pages and also file transfers could be scanned for malicious content prior to sending them to the user. In Web pages active content like ActiveX, Java and JavaScript could be scanned for harmful code (e.g. code exploiting security holes). File transfers could be scanned for known viruses. If a virus is found, the adaptation server could try to remove it or deny the delivery of the infected content. A general rule could be that the caching proxy might store and/or deliver content only, if the content adaptation server has scanned it and no viruses are found.

3.4.2 Insertion of ad banners

Many Internet companies rely heavily on revenue made by selling advertisement space on their Web pages. Whenever advertisement banners are inserted dynamically depending on who requests the page, they cannot be cached, even when the content of the page itself is static. This behaviour prevents Web pages from being cached, although their static content would allow for it. Therefore it seems reasonable to cache the static part of those Web pages at a caching proxy near the client and to insert ad banners into the cached Web pages before serving them to the client.

3.4.3 Insertion of regional data

If a content provider wants to add user-specific regional information (weather forecasts for certain areas for example) to his Web pages, he has little choice but to have the user select his location from a list of regions. Usually it is not possible for origin servers to reliably detect from where Web users connect to Web sites because user requests can get routed through a number of proxy servers on their way from the client to the origin server. In a network edge caching proxy environment user requests are usually redirected to the nearest proxy that is available to respond to the request. Regional information that is relevant to all users who are likely to connect to a certain proxy could be stored at the corresponding caching proxy. Whenever the proxy receives a user request, a module on the caching proxy could insert the regional information into the requested Web page. If the Web page does not contain any user-specific non-cacheable content other than the inserted regional information, the Web page content can now be cached for future requests.

3.4.4 Content adaptation for alternate Web access devices

Since the number of different access devices is growing constantly content providers cannot be expected to provide different versions of their Web pages for each and every

Web access device that is available in the market. Therefore, if it is possible to transcode the general full-fledged Web pages at some point on their way from the origin server to the user so that they are optimised for (or at least adapted to) the end users' specific requirements, it would provide a valuable service for the end customer, the service provider, and the content provider.

3.4.5 Adaptation of streaming media

In particular, media streams could be adapted to meet the bandwidth of the user's connection. It would also be possible to insert pre- recorded advertisements into audio or video streams. Even content analysis and content filtering could be applied to streaming media.

3.5 Mapping value-added services on business roles

The OPES (see section **Error! Reference source not found.**) group has formulated a taxonomy [http://www.ietf.org/internet-drafts/draft-erickson-opes-taxonomy-00.txt] in which also a number of typical value added services per business role are mentioned. In Table 2, some services are mentioned together with the business roles that are likely to implement these services. Note that in particular a CDN and an Access Network will deploy OPES-like boxes to provide the value-added services; the content provider (or the ISP that hosts the content provider) and the client may use application-specific means to provide these services.

Table 2: Value-added services implemented by several business roles (OPES).

Added value	Content provider	CDN service provider	Access Network	Client
virus scanning			●	●
insertion of ad banners	●	●	●	
insertion of regional data		●	●	
caching of personalised/customised Web pages		●		●
content adaptation for alternate devices		●	●	●
bandwidth adaptation			●	
adaptation of streaming media		●		
request filtering			●	●
request filtering through content analysis				●
creation of user profiles	●	●	●	
search engine index on caches	●	●	●	
language translation	●	●	●	●

4 CDN components, architectures and protocols

This section provides an overview of the components that constitute a CDN network, its architectural properties, and the protocols that are used.

4.1 Introduction

The Internet today has evolved from a client-server model into a complex distributed architecture that has to deal with the scaling problems associated with the exponential growth. Two core infrastructure techniques, described in the following paragraphs, are employed to meet the demands: replication and caching. A more detailed overview of Internet Web Replication and Caching Taxonomies can be found in RFC 3040[9].

4.2 Replication

Replication according to the Free Online Dictionary of Computing means: "Creating and maintaining a duplicate copy of a database or file system on a different computer, typically a server." It typically involves "pushing" content from a master server to (and between) replica servers. Two types of communication in a replication architecture can be distinguished:

1. Client replication communication;
2. Inter replication communication.

4.2.1 Client-Replica protocols

A protocol running between client and replica servers and/or master origin server(s) ensures that the client retrieves data from the server that delivers it in the most efficient way. Examples of such protocols are:

- Navigation Hyperlinks: the content consumer manually selects the link of the replica server that he wants to use.
- Replica HTTP Redirection: clients are redirected to an optimal replica server via the use of the HTTP redirection responses.
- DNS Redirection: a Domain Name Server returns a sorted list of replica servers based on quality of service polices upon a client request for an origin server.

4.2.2 Inter-Replica protocols

A protocol running between replica servers and the master origin server(s) ensures that the replicated data remains valid. Examples are:

- Batch driver replication: the replica server initiates an update session (FTP, RDIST) with the origin server at specified times according to a scheduling policy.

[9] http://www.faqs.org/rfcs/rfc3040.html

- Demand driven replication: the replica server initiates an update session with the origin server whenever a client requests a resource that is not up-to-date. (Notice that the difference between caching here lies in the fact that the inter-replica communication protocol can be different from the client-replica protocol and the master origin server is aware of the replica servers)
- Synchronised replication: replicated origin servers co-operate using synchronised strategies and protocols. Updates occur based upon the synchronisation time constraints and involve deltas only.

4.3 Caching

A caching program controls a local store where it stores, retrieves and deletes response messages based upon client requests. A cache stores cacheable responses in order to reduce the response time and network bandwidth consumption on future, equivalent requests. Web caching is typically done in two places: browsers and proxies.

4.3.1 Proxies

Schematically, a proxy server[10] sits between a client program (typically a Web browser) and some external server (typically another server on the Web). The proxy server can monitor and intercept any and all requests being sent to the external server or that comes in from the Internet connection. This positioning gives the proxy server three key capabilities as described in the following paragraphs.

4.3.1.1 Filtering Requests

Filtering requests is the security function and the original reason for having a proxy server. Proxy servers can inspect all traffic (in and out) over an Internet connection and determine if there is anything that should be denied transmission, reception, or access. Since this filtering cuts both ways, a proxy server can be used to keep users out of particular Web sites (by monitoring for specific URLs) or restrict unauthorised access to the internal network by authenticating users. In this way a proxy can be seen as an application level firewall. Before a connection is made, the server can ask the user to log in. To a Web user this makes every site look like it requires a log in. Because proxy servers are handling all communication, they can log everything the user does. For HTTP (Web) proxies this includes logging every URL. For FTP proxies this includes every downloaded file. A proxy can also examine the content of transmissions for "inappropriate" words or scan for viruses, although this may impose serious overhead on performance (see Value Added Services in section 3.5).

4.3.1.2 Sharing Connections

Some proxy servers, particularly those targeted at small business, provide a means for sharing a single Internet connection among a number of workstations. They do so by performing so-called Network Address Translation between the workstations in the Local Area Network and the Internet. While this has practical limits in performance, it can still

[10] http://serverwatch.internet.com/proxyservers.html - Proxy Server Overview

be a very effective and inexpensive way to provide Internet services, such as e-mail, throughout an office.

4.3.1.3 Improving Performance

The third aspect of proxy servers is improving performance. This capability is usually called proxy server caching. In simplest terms, the proxy server analyses user requests and determines which, if any, should have the content stored temporarily for immediate access. A typical corporate example would be a company's home page located on a remote server. Many employees may visit this page several times a day. Since this page is requested repeatedly, the proxy server would cache it for immediate delivery to the Web browser. Cache management is a big part of many proxy servers, and it is important to consider how easily the cache can be tuned and for whom it provides the most benefit. The following paragraphs describe proxy server caching in detail.

4.3.2 Caching proxies

A proxy cache[11] is an application-layer network service for caching Web objects. Unlike browser caches that cache Web objects locally on a machine on a per-client basis, proxy caches can be simultaneously accessed and shared by many users. Proxy caches often operate on dedicated hardware. These tend to be high-end systems with fast processors, 5--50 GB of disk space, and 64--512 MB of RAM. Proxy caches are usually operated much like other network services (e-mail, Web servers, DNS).

The term proxy refers to an important aspect of their design. The proxy application acts as an intermediary between Web clients and servers. Without a proxy, clients make TCP connections directly to servers. In certain environments, i.e. networks behind firewalls, this is not allowed. To prevent exposure of the internal network, firewalls require all external traffic to pass through gateways. Clients must make their connections to proxy applications (also knows as application-layer gateways) running on the firewall host(s). The proxy then connects to the server and relays data between the client and the server.

Strictly speaking there is a difference between a proxy and a cache. A proxy does not always also cache the replies passing through it. A proxy may be used on a firewall only to allow and monitor internal clients access to external servers. Several commercial firewall proxies exist which only proxy Web requests. A proxy may also be used primarily to check incoming files for viruses without caching them (see Value Added Services in section 3.5).

We use the term proxy cache to mean a Web cache, which is implemented as a HTTP proxy, and just to be clear we are not talking about other types of caches (browser caches, RAM caches). Until recently, all Web caches were implemented as HTTP proxies. Now some new and exciting caching technologies have been developed that enlarge the usability of proxies in the network.

---

[11] http://ircache.nlanr.net/Cache/FAQ/ircache-faq-2.html - What is a proxy cache?

4.3.3 Web Cache Architectures

A single Web cache will reduce the amount of traffic generated by the clients behind it. Similarly, a group of Web caches can benefit by sharing another cache in much the same way. Researchers on the caching protocols envisioned that it would be important to connect Web caches hierarchically. In a cache hierarchy (or mesh) one cache establishes peering relationships with its neighbour caches. There are two types of relationship: parent and sibling. A parent cache is essentially one level up in a cache hierarchy. A sibling cache is on the same level. The terms "neighbour" and "peer" are used to refer to either parents or siblings, which are a single "cache-hop" away. What does it mean to be "on the same level" or "one level up?" The general flow of document requests is up the hierarchy. When a cache does not hold a requested object, it may ask (via ICP) whether any of its neighbour caches has the object. If any of the neighbours does have the requested object (i.e., a "neighbour hit"), then the cache will request it from them. If none of the neighbours has the object (a "neighbour miss"), then the cache must forward the request either to a parent, or directly to the origin server. The essential difference between a parent and sibling is that a "neighbour hit" may be fetched from either one, but a "neighbour miss" may NOT be fetched from a sibling. In other words, in a sibling relationship, a cache can only ask to retrieve objects that the sibling already has cached, whereas the same cache can ask a parent to retrieve any object regardless of whether or not it is cached. A parent cache's role is to provide "transit" for the request if necessary, and accordingly parent caches are ideally located within or on the way to a transit Internet service provider (ISP).

There are several problems associated with a caching hierarchy:

- Every hierarchy level may introduce additional delays,
- Higher level caches may become bottlenecks and have long queuing delays, and
- Multiple document copies are stored at different cache levels.

As an alternative to hierarchical caching a distributed caching scheme has been proposed. In distributed Web caching, no intermediate caches are set up and there only caches at the bottom level of the network which co-operate and serve each other's misses. In order to decide from which cache to retrieve a miss document, the caches keep metadata information about the content of every other co-operating cache. To make distribution of the metadata information more efficient and scalable, a hierarchical distribution can be used. However, the hierarchy is only used to distribute information about the location of the documents and not to store document copies. With distributed caching most of the traffic flows through low network levels, which are less congested and no additional disk space is required at intermediate levels. Large-scale deployment of distributed caching, however, encounters several other problems, such as high connection times, higher bandwidth usage, and administrative issues.

Performance analysis of both caching architectures shows that hierarchical caching has shorter connection times than distributed caching. Placing additional copies at intermediate network levels reduces the retrieval latency for small documents. Moreover, distributed caching provides shorter transmission times and has higher bandwidth usage.

However, the network traffic generated by a distributed scheme is better distributed, using more bandwidth in the lower network levels, which are less congested[12].

4.3.4 Caching Protocols

This section describes some of the most frequently used caching protocols in today's Internet caching architectures. A more complete overview can be found in RFC 3040[13].

4.3.4.1 ICP

The Internet Cache Protocol (ICP, currently version 2) is an informational IETF Request For Comments (RFC2186[14], RFC2187[15]) developed by National Laboratory for Applied Network Research (NLANR) in 1997. ICP is a lightweight message format used for communicating among Web caches. It is used to exchange hints about the existence of URLs in neighbour caches. Caches exchange ICP queries and replies to gather information to use in selecting the most appropriate location from which to retrieve an object.

Although Web caches use HTTP for the transfer of object data, caches benefit from a simpler, lighter communication protocol. ICP is primarily used in a cache mesh to locate specific Web objects in neighbouring caches. One cache sends an ICP query to its neighbours. The neighbours send back ICP replies indicating a "HIT" or a "MISS. In current practice, ICP is implemented on top of UDP, but there is no requirement that it be limited to UDP. It is felt that ICP over UDP offers features important to Web caching applications. An ICP query/reply exchange needs to occur quickly, typically within a second or two. A cache cannot wait longer than that before beginning to retrieve an object. Failure to receive a reply message most likely means the network path is either congested or broken. In either case one would not want to select that neighbour. As an indication of immediate network conditions between neighbour caches, ICP over a lightweight protocol such as UDP is better than one with the overhead of TCP. In addition to its use as an object location protocol, ICP messages can be used for cache selection. Failure to receive a reply from a cache may indicate a network or system failure. The ICP reply may include information that could assist selection of the most appropriate source from which to retrieve an object.

4.3.4.2 Cache Digests

Cache Digests[16] are an exchange protocol and data format developed by the NLANR in 1998. They form a response to the problems of latency and congestion associated with other inter-cache communications mechanisms such as the Internet Cache Protocol (ICP, see section 4.3.4.1) and the HyperText Cache Protocol (HTCP, see section 4.3.4.3). Unlike most of these protocols, Cache Digests support peering between cache servers

---

[12] Pablo Rodriguez, Christian Spanner, Ernst W. Biersack "Web Caching Architectures: Hierarchical and Distributed Caching", 4th International Web Caching Workshop, San Diego, California, 1999.
[13] http://www.faqs.org/rfcs/rfc3040.html
[14] http://www.ircache.net/Cache/ICP/rfc2186.txt
[15] http://www.ircache.net/Cache/ICP/rfc2187.txt
[16] http://www.squid-cache.org/CacheDigest/cache-digest-v5.txt

without a request-response exchange taking place. Instead, other servers who peer with it fetch a summary of the contents of the server (the Digest). Using Cache Digests it is possible to determine with a relatively high degree of accuracy whether a particular server caches a given URL. This is done by feeding the URL and the HTTP method by which it is being requested into a hash function that returns a list of bits to test against in the Cache Digest.

Cache Digests are both a protocol and a data format, in the sense that the construction of the Cache Digest itself is well defined, and there is a well defined protocol for fetching Cache Digests over a network - currently via HTTP. A peer answering a request for its digest will specify an expiry time for that digest by using the HTTP *Expires* header. The requesting cache thus knows when it should request a fresh copy of that peer's digest. Requesting caches use an If-Modified-Since request in case the peer has not rebuilt its digest for some reason since the last time it was fetched.

It's possible that Cache Digests could be exchanged via other mechanisms, in addition to HTTP, e.g. via FTP. The Cache Digest is calculated internally by the cache server and can exist as (for instance) a cached object like any other - subject to object refresh and expiry rules. Although Cache Digests as currently conceived are intended primarily for use in sharing summaries of which URLs are cached by a given server, this capability can be extended to cover other data sources. For example, an FTP mirror server might make a Cache Digest available that indicated matches for all of the URLs by which the resources it mirrored may be accessed. This is potentially a very powerful mechanism for eliminating redundancy and making better use of Internet server and bandwidth resources.

A Cache Digest is a summary of the contents of an Internet Object Caching Server. It contains, in a compact (i.e. compressed) format, an indication of whether or not particular URLs are in the cache. A "lossy" technique is used for compression, which means that very high compression factors can be achieved at the expense of not having 100% correct information.

Cache servers periodically exchange their digests with each other. When a request for an object (URL) is received from a client, a cache can use digests from its peers to find out which of its peers (if any) have that object. The cache can then request the object from the closest peer.

The checks in the digest are very fast and they eliminate the need for per-request queries to peers. Hence:

- Latency is eliminated and client response time should be improved.
- Network utilisation may be improved.

Note that the use of Cache Digests (for querying the cache contents of peers) and the generation of a Cache Digest (for retrieval by peers) are independent. So, it is possible for a cache to make a digest available for peers, and not use the functionality itself and vice versa.

4.3.4.3 HTCP

The Hyper Text Caching Protocol[17] (HTCP) is a protocol for discovering HTTP caches and cached data, managing sets of HTTP caches, and monitoring cache activity. It was developed as an Internet Draft by the ICP working group in 1999.

HTTP 1.1 permits the transfer of Web objects from origin servers, possibly via proxies (which are allowed under some circumstances to cache such objects for subsequent reuse) to clients which consume the object in some way, usually by displaying it as part of a Web page. HTTP 1.0 and later permit headers to be included in a request and/or a response, thus expanding upon the HTTP 0.9 (and earlier) behaviour of specifying only a URI in the request and offering only a body in the response. ICP was designed with HTTP/0.9 in mind, such that only the URI (without any headers) is used when describing cached content and the possibility of multiple compatible bodies for the same URI had not yet been imagined. HTCP permits full request and response headers to be used in cache management. It expands the domain of cache management to include monitoring a remote cache's additions and deletions, requesting immediate deletions, and sending hints about Web objects such as the third party locations of cacheable objects or the measured uncacheability or unavailability of Web objects.

4.3.4.4 CARP

The Cache Array Routing Protocol[18] [19] (CARP) is an IETF draft co-developed (and implemented) by Microsoft. It divides URL-space by hashing among an array of loosely coupled proxy servers. Proxy servers and client browsers can route requests to any member of the Proxy Array. Due to the resulting sorting of requests through these proxies, duplication of cache contents is eliminated and global cache hit rates are improved. According to Microsoft it has the following advantages:

- CARP doesn't conduct queries. Instead it uses hash-based routing to provide a deterministic "request resolution path" through an array of proxies. The result is single-hop resolution. The Web browser or a downstream proxy will know exactly where each URL would be stored across the array of servers.
- CARP has positive scalability. Due to its hash-based routing, and hence, its freedom from peer-to-peer pinging, CARP becomes faster and more efficient as more proxy servers are added.
- CARP protects proxy server arrays from becoming redundant mirrors of content. This vastly improves the efficiency of the proxy array, allowing all servers to act as a single logical cache.
- CARP automatically adjusts to additions or deletions of servers in the array. The hashed-based routing means that when a server is either taken off line or added, only minimal reassignment of URL cache locations is required.
- CARP provides its efficiencies without requiring a new wire protocol. It simply uses the open standard HTTP. One advantage of this is compatibility with existing firewalls and proxy servers.

[17] http://www.ircache.net/Cache/ICP/htcp.txt
[18] http://www.microsoft.com/proxy/guide/CarpWP.asp?A=2&B=3
[19] http://www.ircache.net/Cache/ICP/carp.txt

- CARP can be implemented on clients using the existing, industry-standard client Proxy Auto-Config file (PAC). This extends the systemic benefits of single hop resolution to clients as well as proxies. By contrast, ICP is only implemented on Proxy servers.

4.4 OPES

The Open Pluggable Edge Services (OPES) architecture, under development in the IETF, enables construction of services executed on application data by participating transit intermediaries. Caching is the most basic intermediary service, one that requires a basic understanding of application semantics by the cache server. Because intermediaries divert data temporarily over a pathway different from the transit pathway, one can think of the service path as being orthogonal to the main transit path. The purpose of the IETF OPES working group is to define the protocols and API's for a broad set of services that facilitate efficient delivery of complex content or services related to content. The architecture supports services that are either co-located with the transit intermediary or located on other (auxiliary) servers.

The current System Architecture of an 'OPES box' looks as follows:

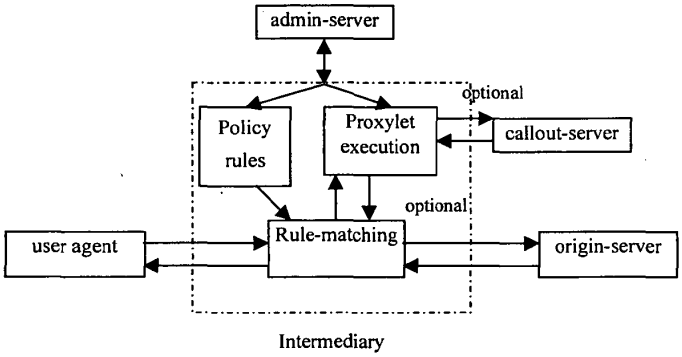


Figure 13: OPES system architecture

This architecture shows that the admin-server is responsible for setting some 'policy rules' that are used to match either requests from the client to the origin server, or responses from the origin server to the client. When the rules match, a specified action takes place. This action is typically the execution of a (co-located) 'proxylet' (application-specific code executed on the intermediate system), that may contain a call to a callout server (a remote entity). The protocol between intermediary and callout server will probably be based on ICAP (see section 6.1). The ICAP protocol is being developed for carrying HTTP headers and data to co-operating servers; other protocols for carrying SMTP or other protocols to co-operating servers will be supported by the framework, as they exist or become available.

The security model for intermediary services involves defining the administrator roles and privileges for the application client, application server, intermediary, and auxiliary server. The working group will use the Policy Configuration Information Model to define the security attributes and the enforceable policy.

4.5 Streaming Proxies

The increased presence of streaming media places new demands on the management of network bandwidth[20]. Streaming media not only consumes more bandwidth than do Web pages; it also requires a continuous uninterrupted flow of data to yield the best possible end-user experience, because the client/server connections are persistent. Unlike HTML, images, or downloadable files, streaming media depends critically on consistent and reliable packet delivery over complicated network paths spanning many segments, routers, and switches. Temporary delays (network congestion) and packet loss are more than inconvenient—they affect the smoothness of playback for the end user, and they may also lead to audio dropouts or poor video quality that compromise the overall user experience.

The use of streaming media content presents several unique challenges:

- *It is bandwidth intensive.* The transmission rate of streamed content can be as low as 28 kilobits per second (Kbps), but is now being encoded as high as 1 megabyte per second (MBps) on the Internet—and even higher in controlled environments. Unmanaged, this vast range of speeds has significant network implications.
- *It requires an uninterrupted flow of data.* For content to be worth watching, the end user must be able to receive a continuous flow of bits. The shorter the distance those bits have to travel, the better the end user's experience will be.
- *It requires different ports than does Web content.* Streaming media travels by means of a streaming protocol. This can be the Real Time Streaming Protocol (RTSP), a standard that has been submitted for acceptance to the Internet Engineering Task Force (IETF). RTSP defines UDP data channels that require additional ports to be opened, which affects issues such as firewalls, authentication, and security. In addition to that Microsoft has defined a proprietary streaming protocol named Microsoft Media Streaming (MMS).
- *It needs to protect broadcasters' rights.* Broadcasters of streaming media require that their content be protected and managed. This prevents streaming media caching from being as loosely managed as Web caching currently is.

In the case of broadband-streamed media on-demand, CDN surrogates need to be as close to the edge of the network as possible and content will mostly likely be pushed to the surrogates in advance.

Rejaie *et al.* describe a proxy caching mechanism for multimedia playback streams in the Internet[21].

The following sections discusses five basic delivery methods that streaming proxies can support for streaming content to connected clients: cached delivery, replication, splitting streams by means of either UDP or TCP unicasting or IP multicasting and pass-through delivery.

---

[20] http://service.real.com/help/library/whitepapers/rproxy/proxy.html

[21] Rejaie, R., Handley, M., Yu, H., and Estrin, D., "Proxy caching mechanism for multimedia playback streams in the Internet", 4th International Web Caching Workshop, San Diego, California, March 31 - April 2, 1999, http://workshop99.ircache.net/Papers/rejaie-html/ .

4.5.1 Cached Delivery

A proxy may be equipped with a streaming media cache This enables on-demand content to be dynamically replicated locally, perhaps in an encrypted format. The proxy may attempt to store all cacheable media files upon first request.

When a proxy receives a client request for on-demand media, it determines whether the content is cacheable. Then it checks to see whether the requested media already resides in its local cache. If the media is not already in the cache, the proxy acquires the media file from the source server and simultaneously delivers it to the requesting client. Subsequent requests for the same media clip can be served without repeatedly pulling the clip across the network from the source Server.

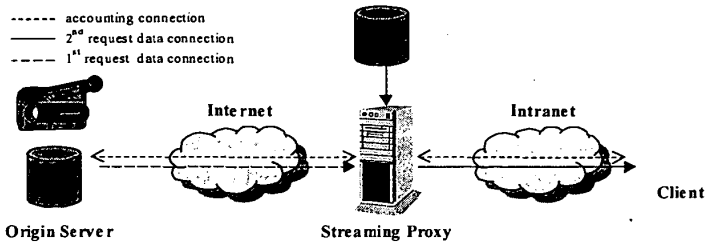


Figure 14: cached stream delivery

4.5.2 Replication

Using replication techniques, one or more copies of a single streaming media asset or even a whole file-system, containing multiple streaming media assets or databases, can be maintained on one or more different servers, called 'replica origin servers'. Clients discover an optimal replica origin server for clients to communicate with. Optimality is a policy based decision, often based upon proximity, but may be based on other criteria such as load.

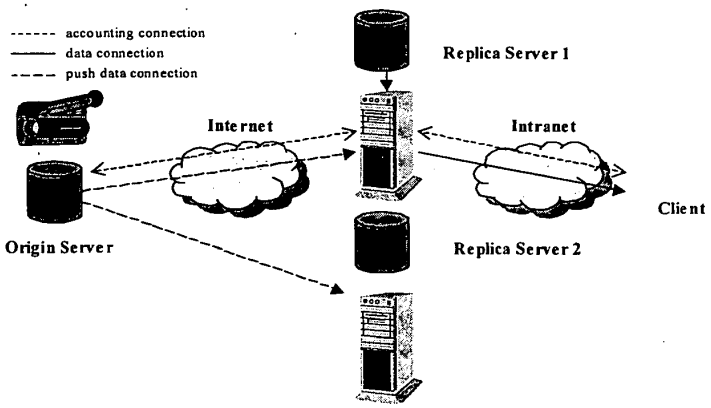


Figure 15: Replication stream delivery

4.5.3 Unicast Split

After initiating a single data-channel connection to a source server, the proxy splits live broadcasts for any clients connected to it. Subsequent requests for the same live stream are then delivered from the proxy, without pulling redundant live data from the source server. For each client requesting the live stream, the proxy establishes an accounting connection back to the source server. This accounting ensures that the client is permitted to access the stream and that it forwards unique session statistics back to the source server. This sort of splitting is also known as "application level multicast".

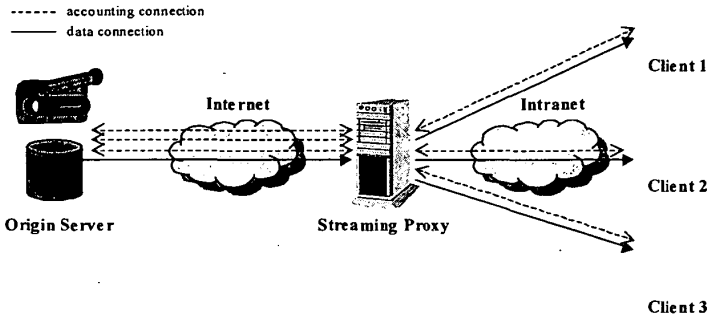


Figure 16: a stream split into multiple unicast streams

4.5.4 Multicast Split

The proxy can rebroadcast a unicast live split stream to its connecting clients by way of IP multicast. IP multicast requires that the network between the client and the proxy is IP multicast-enabled and so are the clients.

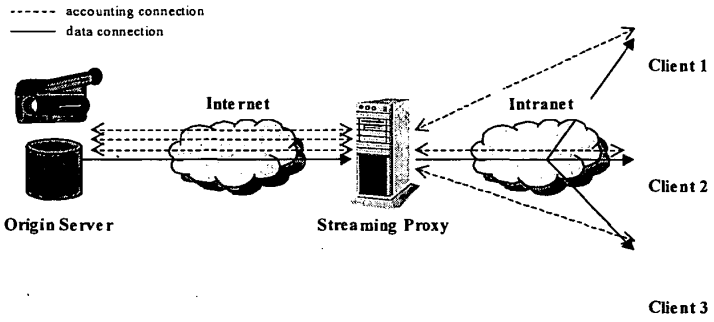


Figure 17: a stream rebroadcast into a multicast stream

4.5.5 Pass-Through Delivery

If live content cannot be split, a proxy can simply pass the stream through to each connecting client, establishing both an accounting connection and a data connection for each client that has requested the live stream from the proxy.

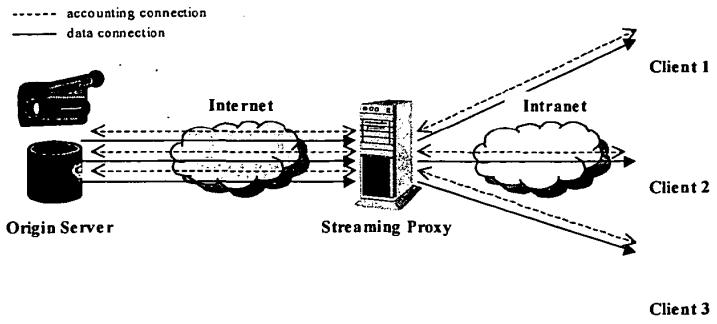


Figure 18: pass-through delivery of a stream

For a detailed description of the state of the art of streaming media caching and replication techniques the reader is referred to the Telematica Instituut project Video over IP (VIP) deliverable 3.1:
http://alpha.sec.nl/vip/Streaming_media_Caching_and_replication_techniques.pdf .

4.6 Products

A list of companies that sell CDN related products can be found in Appendix B - Overview of CDN organisations. Several resources that focus on caching products can be found on the Internet. Some links:

http://directory.google.com/Top/Computers/Software/Internet/Servers/Proxy/Caching/Vendors/.

http://www..web-caching.com/proxy-caches.html

http://www.web-caching.com/proxy-comparison.html

5 Content negotiation

Content negotiation is a very powerful tool where the client can indicate what type of information he can accept, and the server decides what (if any) type of information to return. The term *type* is used very loosely here, because negotiation can apply to several aspects of the information. For example, it can be used to choose the appropriate human language for a document (say, French or German), or to choose the media type that the browser of the client can display (say, GIF or JPEG).

A general framework for content negotiation requires a means for describing the meta-data or attributes and preferences of the user or his/hers/its agents, the attributes of the content and the rules for adapting content to the capabilities and preferences of the user.

Content negotiation covers three elements:
1. expressing the capabilities of the sender and the data resource to be transmitted (as far as a particular message is concerned),
2. expressing the capabilities of a receiver (in advance of the transmission of the message), and
3. a protocol by which capabilities are exchanged.

These negotiation elements are addressed by several protocols.

5.1 MIME-type based content negotiation

MIME means **M**ultipurpose **I**nternet **M**ail **E**xtensions, and refers to an official Internet standard that specifies how messages must be formatted so that they can be exchanged between different e-mail systems. MIME is a very flexible format, permitting one to include virtually any type of file or document in an e-mail message. Specifically, MIME messages can contain text, images, audio, video, or other application-specific data.

The MIME format is also very similar to the format of information that is exchanged between a Web browser and the Web server it connects to. This related format is specified as part of the Hypertext Transfer Protocol (HTTP). With HTTP content negotiation the client and server actually negotiate, via MIME types, on what that particular browser can accept, and what that particular server can give. Once they reach an agreement, the request is granted. Unfortunately, only a few servers support this feature (Apache and W30) and even fewer browsers fully support it.

A number of Internet application protocols have a need to provide content negotiation for the resources with which they interact. MIME media types provide a standard method for expressing several negotiation activities. However, resources vary in ways which, cannot always be expressed using currently available MIME headers.

5.2 Content negotiation in HTTP

Web users speak many languages and use many character sets. Some Web resources are available in several *variants* to satisfy this multiplicity. HTTP/1.0 includes the notion of

content negotiation, a mechanism by which a client can inform the server which language(s) and/or character set(s) are acceptable to the user.

In order for the server to deliver the correct representation of the data, the client must send some information about what he can accept. A browser used on a French-language machine, for instance, should indicate that it can accept data in French (of course, this should also be user-configurable).

To use negotiation, two things are needed. Firstly, a resource is needed that exists in more than one format (for example, a document in French and German, or an image stored as a GIF and a JPEG), and secondly a configurable server is needed that knows that each of these files are actually the same resource. Two methods are available to achieve these things:

- Using a Variants File
- Using file extensions

- *Using a Variants File*

This method involves creating a **variants** file, usually referred to as a **var**-file. This lists each of the files, which contains the same resource, along with details of what representation it is. Any request for this **var**-file causes the server to return the best file, based on the contents of the **var**-file and the information supplied by the browser.

As an example, say there is a file in English and a file in German containing the same information. The files could be called english.html and german.html (they are both HTML files). So create a **var**-file listing each of these files, and specifying which languages they are in. Create a **var**-file called (say) `info.var` containing:

```
URI: english.html
Content-Language: en

URI: german.html
Content-Language: de
```

This file consists of a series of sections, separated by blank lines. Each section contains the name of the file (on the URI: line) and header information used in the negotiation.

Now, when a request for `info.var` is received, the server will read the **var**-file and return the best file, based on which languages the browser has said it can accept. Similarly, the **var**-file could be used to select files based on content type (using Content-Type or content encoding (using Content-Encoding, or any combination.

The Content-Type: line in a variants file can also give any other content type parameters, such as the subjective qualify factor. This will be used in the negotiation when picking the 'best' match. For example, an image available as a JPEG might be regarded as having higher quality then the same image in GIF format. To tell this to the server, the following .var contents could be used:

```
URI: image.jpg
Content-Type: image/jpeg; qs=0.6

URI: image.gif
Content-Type: image/gif; qs=0.4
```

Here the **qs** parameters give the 'source quality' for these two files, in the range 0.000 to 1.000, with the highest value being the most desirable. For instance, a browser than can handle both GIF and JPEG files equally well, can indicate a preference (qs) to see the JPEG version rather than the GIF.

Using variant files gives complete control over the scope of the negotiation, however, it does require the file to be created and maintained for each resource. An alternative interface to the negotiation mechanism is to get the server to identify the negotiation parameters (language, content type, encoding) from the file extensions.

- *Using File Extensions*

Instead of using a **var**-file, file extensions can be used to identify the content of files. For example, the extension `eng` could be used on English files, and `ger` on German files. Then the `AddLanguage` directive can be used to map these extensions onto the standard language tags.

After enabling the `multiview` option, the directives, which map extensions onto representation types can be given. These are `AddLanguage`, `AddEncoding` and `AddType` (content types are also set in the `mime.types` file). For example:

```
AddLanguage en                          .eng
AddLanguage de                          .ger
AddEncoding x-compress                  .Z
AddType         application/pdf         pdf
```

When a request is received, the server looks at all the files in the directory, which **start with** the same filename. So a request for /about/info would cause the server to negotiate between all the files names /about/info.*.

For each matching file, the server checks its extensions and sets the content type, language and encodings appropriately. For example, a file called `info.eng.html` would be associated with the language tag **en** and the content type **text/html**. The source quality, to express the importance or degree of acceptability of various negotiable parameters, is assumed to be 1.000 for all.

The extensions can be listed in any order, and the request itself can include one or more extensions. For example, the files *info.html.eng* and *info.html.ger* could be requested with the URL `info.html`. This provides an easy way to upgrade a site to use negotiation without having to change existing links.

HTTP/1.0 provided a few features to support content negotiation. The HTTP/1.1 specification specifies these features with far greater care, and introduces a number of new concepts. HTTP/1.1 provides two orthogonal forms of content negotiation, differing in where the choice is made:

1. In server-driven negotiation, the more mature form, the client sends hints about the user's preferences to the server, using headers such as Accept-Language, Accept-Charset, etc. The server then chooses the representation that best matches the preferences expressed in these headers.

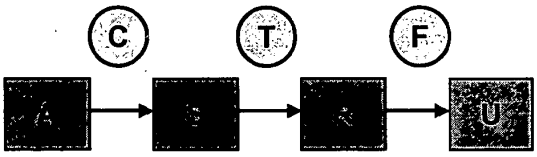
2. In agent-driven negotiation, when the client requests a varying resource, the server replies with a 300 (Multiple Choices) response that contains a list of the available representations and a description of each representation's properties (such as its language and character set). The client (agent) then chooses one representation, either automatically or with user intervention, and resubmits the request, specifying the chosen variant.

Although the HTTP/1.1 specification reserves the `Alternates` header name for use in agent-driven negotiation, the HTTP working group never completed a specification of this header, and server-driven negotiation remains the only usable form.

5.3 IETF Content Negotiation working group

The Content Negotiation working group (ConNeg, http://www.imc.org/ietf-medfree) has proposed and described a protocol-independent content negotiation framework. The negotiation framework provides for an exchange of negotiation meta-data between the sender and receiver of a message, which leads to determination of a data format which the sender can provide and the recipient can process. The subjects of the negotiation process and whose capabilities are described by the negotiation meta-data thus are: the sender, the transmitted data file format, and the receiver.

The life of a data resource may be viewed as:



Where: [A] = author of document, (C) = original document content, [S] = message sending system, (T) = transmitted data file (representation of C), [R] = receiving system, (F) = formatted (rendered) document data (presentation of C), [U] = user or consumer of a document. Source: "Protocol-independent Content Negotiation Framework", Request for Comment 2703, ConNeg working group draft, http://www.imc.org/rfc2703 .

Here, it is [S] and [R] who exchange negotiation meta-data to decide the form of (T). Negotiation meta-data provided by [S] would take account of available document content (C) (e.g. availability of resource variants) as well as its own possible ability to offer that content in a variety of formats. Negotiation meta-data provided by [R] would similarly take account of the needs and preferences of its user [U] as well as its own capabilities to process and render received data.

Negotiation between the sender [S] and the receiver [R] consists of a series of negotiation meta-data exchanges that proceeds until either party determines a specific data file (T) to be transmitted. If the sender makes the final determination, it can send the file directly. Otherwise the receiver must communicate its selection to the sender who sends the indicated file.

5.4 Transparent Content Negotiation

Transparent Content Negotiation (TCN) uses a model in which one of the following happens:

- The recipient requests a resource with no variants, in which case the sender simply sends what is available.
- A variant resource is requested, in which case the server replies with a list of available variants and the client chooses one variant from those offered.
- The recipient requests a variant resource, and also provides negotiation meta-data (in the form 'Accept' headers) which allows the server to make a choice on the client's behalf.

For more information about transparent content negotiation see www.gewis.win.tue.nl/~koen/conneg/rfc2295.html .

Another, simpler example is that of fax negotiation: in this case the intended recipient declares its capabilities, and the sender chooses a message variant to match.

5.5 User (agent) profiles

For service providers, content providers and network operators it would be ideal if they could make use of information about the user and/or its user-agent (terminal/application/whatever). It would make solutions for personalisation or terminal-adaptation quite easy when this information would be available. Of course, privacy-issues are the most prominent aspects that need to be taken into account in this respect. Therefore, the end-user must always be able to control the (access to) the profile-content.

It is important in this discussion to separate 'profiling' or 'usage profiles' from 'user profiles'. The first category of profiles are automatically generated due to interaction of end-users with particular systems or services, and are managed and controlled by the system-operator or service-provider; typically for marketing purposes or personalisation. These are out-of-scope for this discussion. User (agent) profiles, as used in this section, are collections of attributes that are managed by the users (or their user agents) themselves, and selectively provided on behalf of that user to other parties (users, service providers, or whoever). They can be used by CDNs to select automatically the proper content on behalf of that user.

The following aspects are relevant for content negotiation with user (agent) profiles:

- Terminal capabilities and preferences of the active terminal in the current session. This is handled in W3C's CC/PP activity (see section 5.5.1). The CC/PP has the advantage that the information is transferred along with HTTP requests, and, therefore, has end-to-end significance. It can be used by all entities in the CDN chain of responsibility.

- User-specific information. Access network providers may be able to identify end-users, and/or the location of these end-users. For instance, using protocols like RADIUS[22]/DIAMETER[23] during the establishment of an access-session (see also Chapter 7), a unique coupling between IP-address and end-user can be determined. This enables them to derive service-specific usage parameters, but they could also allow end-users to provide more specialised user-profiles (with varying access permissions or degrees of anonymity). Benefits of such an approach are outlined in a recent Internet Draft: http://www.ietf.org/internet-drafts/draft-penno-cdnp-nacct-userid-02.txt . Also, this may be a subject for the definition of an extended Parlay-like API (http://www.parlay.org, see also section 8.3). An organisation like the PAM-forum (http://www.pamforum.org) is currently active in defining presence-like environments. This is closely linked to user profile management. Because it is not likely that an end-user would want to have its identity communicated to all service-providers; it is likely that a new business role will emerge, a 'user profile provider' or 'presence-information provider' that actually manages the on-line state and user preferences of individual users.
- Current resource availability. This is quite difficult to determine. Access providers typically know the maximum capacity of the access network. However, current availability is typically not available. This information can be derived (ideally) using terminal co-operation (e.g. using quality feedback agents on the terminal; this allows for end-to-end resource availability feedback), or otherwise by measuring (again, by the access provider) current capacity on the end-users access network. This information can be made available in, e.g., the 'proxylet' execution environment (or even the rule matching engine) of the OPES architecture (see also section 4.4), which makes it possible to run filters based on resource availability.

5.5.1 W3C CC/PP (Composite Capability / Preference Profiles)

The W3C Composite Capability / Preference Profile (CC/PP; http://www.w3.org/TR/NOTE-CCPP/) specifies client capabilities and user preferences as a collection of URIs and Resource Description Framework (RDF) text, which is sent by the client along with a HTTP request. The URIs point to an RDF document which contains the details of the clients capabilities. RDF provides a way to express meta-data for a Web document. The CC/PP scheme allows proxies and servers to collect information about the client, from the client directly, and to make decisions based on this information for content adaptation and delivery. The CC/PP is the encoding of profile information that needs to be shared between a client and a server, gateway or proxy. CC/PPs are intended to provide information necessary to adapt the content and the content delivery mechanisms to best fit the capabilities and preferences of the user and its agents.

[22] RADIUS (Remote Authentication Dial-In User Service) is a client/server protocol and software that enables remote access servers to communicate with a central server to authenticate dial-in users and authorise their access to the requested system or service. RADIUS is a de facto industry standard and is a proposed IETF standard.

[23] Like RADIUS, Diameter is a "triple-A" protocol - it authenticates and authorises users and performs basic back-end accounting services for bookkeeping purposes.

5.6 SDP version 2

SDP[24] allows specifying multimedia sessions (i.e. conferences) by providing general information about the session as a whole and specifications for all the media streams to be used to exchange information within the multimedia session. Currently, media descriptions in SDP are used for two purposes:

1. to describe session parameters for announcements and invitations; the original purpose of SDP,
2. to describe the capabilities of a system (and possibly to provide a choice between a number of alternatives). Note that SDP was not designed to facilitate this.

A distinction between these two "sets of semantics" is only made implicitly. The IETF Multiparty Multimedia Session Control (MMUSIC) working group (http://www.ietf.org/html.charters/mmusic-charter.html) is currently defining a Next Generation SDP protocol to initiate the development of a session description and capability negotiation framework. In a new IETF draft[25], a language for describing multimedia sessions with respect to configuration parameters and capabilities of end systems is defined to allow for content negotiation. MMUSIC also defines terminology and lists a set of requirements that are relevant for a framework for session description and endpoint capability negotiation in multiparty multimedia conferencing scenarios[26].

The SDP concept of a capability description language addresses various pieces of a full description of system and application capabilities in four separate "sections":

1. Definitions (elementary and compound): specification of a number of basic abstractions that are later referenced to avoid repetitions in more complex specifications and allow for a concise representation. Definition elements are labelled with an identifier by which they may be referenced. They may be elementary or compound (i.e. combinations of elementary entities). Examples of definitions that sections include (but are not limited to) codec definitions (<audio-codec name="audio-basic" encoding="PCMU sampling_rate="8000 channels="1"/>), redundancy schemes, transport mechanisms and payload formats.
2. Potential or Actual Configurations: all the components that constitute the multimedia conference (IP telephone call, multi-player gaming session, etc.). For each of these components, the potential and, later, the actual configurations are given. Potential configurations are used during capability exchange and/or negotiation; actual configurations to configure media streams after negotiation or in session announcements.
3. Constraints: Constraints refer to potential configurations and to entity definitions and express and use simple logic to express mutual exclusion, limit the number of instantiations, and allow only certain combinations.

---

[24] Handley, M. and V. Jacobsen, "SDP: Session Description Protocol", RFC 2327, April 1998, http://www.faqs.org/rfcs/rfc2327.html .

[25] Kutscher, Ott, Bormann, "*Session Description and Capability Negotiation*", IETF draft, http://www.ietf.org/internet-drafts/draft-ietf-mmusic-sdpng-00.txt .

[26] Kutscher, Ott, Bormann, "*Requirements for Session Description and Capability Negotiation*", IETF draft, http://www.ietf.org/internet-drafts/draft-ietf-mmusic-sdpng-req-01.txt .

4. Session attributes: description of general meta-information parameters of the communication relationship to be invoked or modified. It also allows to tie together different media streams or provide a more elaborate description of alternatives (e.g. subtitles or not, which language).

6 Content adaptation

Besides delivering content, CDNs may also *adapt* the content. For instance by transcoding multimedia streams, or by translating HTML content into WML (for access to mobile WAP devices), or by inserting advertisements, or by translating from a particular language into another. Hence, CDN-providers may also provide multi-channelling support to the content-owners. To provide content-adaptation, the CDN must know what the exact resource-availability of a client (in terms of codecs, drivers, multimedia capabilities, ...) and the intermediate networks (in terms of delay, available bandwidth, ...) are. The W3C CC/PP group is currently working on standardising such negotiations between Web-browsers and Web-servers, in close co-operation with the WAP-forum (for mobile terminals). However, these negotiations are only 'static', which means that current resource availability is not taken into account, but only the capabilities of the client. Furthermore, actual network-load is not taken into account either. That is reasonable for their purpose of course, because most Web content has not so many real-time requirements. However, for streaming services and multimedia content more realistic adaptation is needed. The actual capability-information that is required for such adaptation can be stored in either specific servers (say, capability management servers), but it may also provide useful to store this information in presence-servers. That means that the resource availability can be made available on the basis of a personal identification (e.g. using presence URI like presence:H.Eertink@telin.nl). This matches with the standardisation efforts of the *impp* group within the IETF, and relates this work with more synchronous communication means.

There is currently new standardisation work being set-up that defines standard mechanism to extend HTTP-intermediates with application-specific value added services (such as transcoding, or virus checking, or whatever). This is quite similar to the major objective of this project. The standardisation takes place in the *IETF opes-group* (open proxy extended service), and is still in a requirements phase. Proposed results are both an API and protocol descriptions. The protocol will probably be based on iCAP (Internet Content Adaptation Protocol; a non-standard protocol developed by NetAppliances; see http://www.i-cap.org/). An outline of the scope of this group is given in ftp://ftp.isi.edu/internet-drafts/draft-tomlinson-epsfw-00.txt.

Transcoding *inside* the network (as can be done by proxies) is a disputed subject, certainly when it becomes 'standardised practice'. It has all kinds of implications:

- It breaks end-to-end security. For instance, digital signatures are worthless when the content that has been signed is being manipulated by an intermediate system.
- It may result in quality degradation that is not acceptable for the content owner.
- It has all kinds of possibilities for eavesdropping and other privacy-related infringements.

Therefore, there is a common understanding that this should only be done either under control of the end-user (e.g. on the client itself, or within the scope of an SLA with his (access?) provider) or under control of the content-owner (e.g. via negotiated policies between content-owner and proxy-manager).

Adaptive content delivery technologies transform Web content and provide delivery schemes according to viewers' heterogeneous and changing conditions to enable universal access. The goal of adaptive content delivery is to take into account these heterogeneous and changing conditions and provide the best information accessibility and perceived quality of service over the Internet. The improved perceived quality of service by adaptive content delivery means for instance for e-commerce applications that shoppers are more likely to stay and return, thus resulting in a greater profit. Most Web content has been designed with desktop computers in mind, and they often contain rich media. This media-rich content may not be suitable for Internet appliances with relatively limited display capabilities, storage, processing power, as well as slow network access. Several content adaptation techniques are:

- Information abstraction. The goal of information abstraction is to reduce the bandwidth requirement for delivering the content by compressing the data, while preserving the information that has highest value to the user.
- Modality transformation. Modality transform is the process of transforming content from one mode to another so that the content can become useful for a particular client device. For instance the transformation of video into sets of images for handheld computers.
- Data transcoding. Data transcoding is the process of converting data format according to client device capability. For example, GIF images to JPEG images or audio format conversion such as WAV to MP3.
- Data prioritisation. The goal of data prioritisation is to distinguish the more important part of the data from the less important part so that different quality of service levels can be provided when delivering the data through the network. For example, to allow less important data to be dropped under network constraints. Or send the more important data first.
- Purpose classification. By classification of the purpose of each media object (e.g. images of banners, logos, and advertisements) in a Web page, one can improve the efficiency of information delivery by either removing redundant objects or prioritising them according to their importance.
- Proxy-based adaptation. In a proxy-based adaptation, the client connects through a proxy, which then makes the request to the server on behalf of the client. The proxy intercepts the reply from the server, decides on and performs the adaptation, and then sends the transformed content back to the client. A proxy-based architecture makes it easy to place adaptation geographically close to the clients. Adapting the proxy means that there is no need to change existing clients and servers, and it achieves economy of scale more than a server-based adaptation architecture since each proxy can transform content for many servers. The proxy can transform existing Web content so that existing content does not have to be re-authored. The issue of copyright infringement becomes significant in a proxy-based system, since an author has little control for performing adaptation. Transcoding proxies are used as intermediaries between generic WWW servers and a variety of client devices in order to adapt to the greatly varying bandwidths of different client communication links and to handle the heterogeneity of possibly small-screened client devices.

Some of these techniques are explained in the remainder of this section.

6.1 ICAP – Internet Content Adaptation Protocol.

Content delivery caching systems have dramatically improved the speed and reliability of the Internet, benefiting Web sites and users alike. But while Internet-based applications and services continue to flourish, no one has defined a way for network-based applications to communicate with the latest content delivery systems. The iCAP Forum, a consortium of Internet businesses covering a wide array of services (www.i-cap.org), introduced in 1999 the iCAP protocol to enable such communication. This protocol is currently being drafted via the IETF (http://www.i-cap.org/icap/media/draft-opes-icap-00.txt). ICAP is an open protocol designed to facilitate better distribution and caching for the Web. It distributes Internet-based content from the origin servers, via proxy caches (iCAP clients), to dedicated iCAP servers. These iCAP servers are focussed on specific value-added services such as access control, authentication, language translation, content filtering, virus scanning, and ad insertions. Moreover, iCAP enables adaptation of content in such a way that it becomes suitable for other less powerful devices such as PDAs and mobile phones.

Since iCAP allows the removal of the value-added services from the critical path, i.e. from the origin server to the iCAP server, it reduces the load on the origin server and the network. Additionally, iCAP-enabled devices are able to store modified data, which eliminates repeated adaptation.

6.1.1 Benefits of iCAP

ICAP is useful in a number of ways. For example it might be used to scale Internet services such as:

- Simple transformations of content can be performed near the edge of the network instead of requiring an updated copy of an object from an origin server.
- Avoiding proxy caches or origin servers to perform expensive operations by shipping the work off to other (iCAP) servers. This helps distribute load across multiple machines.
- Allowing firewalls or proxy caches to act as iCAP clients that send outgoing requests to a service that checks to make sure the URI in the request is allowed.

Another advantage of iCAP is that it creates a standard interface for adaptation of HTTP messages, allowing interoperability.

6.1.2 ICAP architecture

ICAP is in essence an HTTP-based remote procedure call protocol that empowers an edge device, like a cache, to forward HTTP messages to an application server, without overloading the cache and slowing response times. In other words, iCAP allows its clients to pass HTTP based (HTML) messages (content) to iCAP servers for adaptation. Adaptation refers to performing particular value-added services (content manipulation) to the associated client request/response.

There are three ways in which iCAP can work:

- Request modification method. In this mode, a client sends a request to an origin server [1]. This request is redirected to an iCAP server by the intervening proxy server (cache) [2]. The iCAP server modifies the message and sends it back to the proxy server [3]. The proxy server parses the modified message and forwards it to the origin server to fulfil the client's request [4]. The origin server then executes the request and the response is, via the proxy server [5], delivered to the client [6] (see Figure 19).

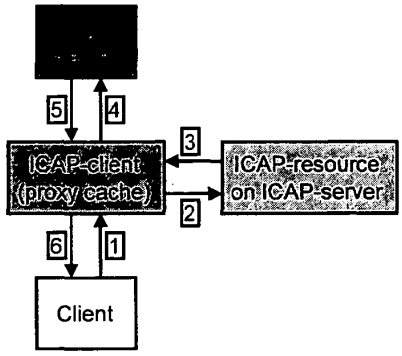


Figure 19: ICAP request modification method.

- Request satisfaction method. Here, the client's request to an origin server is redirected to an iCAP server by the intervening proxy cache [1]. The iCAP server modifies the message [2], and sends it straight to the origin server for fulfilment [3]. After processing the request, the origin server sends it back to the client via the iCAP server and proxy server [4], [5], [6] (see Figure 20).

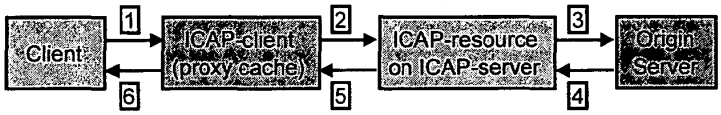


Figure 20: ICAP request satisfaction method.

- Response modification method. In this mode, a client makes an HTTP request to an iCAP capable proxy intermediary [1]. The intermediary forwards the HTTP request to the origin server [2]. The response [3], however, is redirected by the proxy server to the iCAP server [4]. The iCAP server executes the requested iCAP service and sends the possibly modified response back [5]. The proxy server sends the reply (possibly modified from the origin server's response) to the client [6] (see Figure 21).

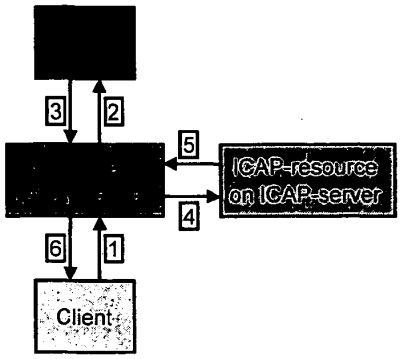


Figure 21: ICAP response modification method.

In each case, the iCAP client and server exchange standard HTTP GET and POST requests and responses. The iCAP client for instance uses an HTTP POST in which "client request" and "propose origin server response" are encapsulated within the first part of the HTML body.

ICAP is a request/response protocol similar in semantics and usage to HTTP/1.1. Despite the similarity, iCAP is not HTTP, nor is it an application protocol that runs over HTTP. ICAP communication usually takes place over TCP/IP connections. Two examples are shown in the table below.

Example	ICAP Request	ICAP Response
A proxy cache receives a GET request from a client. The proxy cache, acting as an iCAP client, then forwards this request to an iCAP server for modification. The iCAP server modifies the request headers and sends them back to the iCAP client. The iCAP server in this example modifies several headers and strips the cookie from the original request.	REQMOD icap://icap-server.net/server ICAP/1.0 Host: icap-server.net Encapsulated: req-hdr=0 GET / HTTP/1.1 Host: www.origin-server.com Accept: text/html, text/plain Accept-Encoding: compress Cookie: ff39fk3jur@4ii0e02i If-None-Match: "xyzzy", "r2d2xxxx"	ICAP/1.0 200 OK Date: Mon, 22 Jan 2001 09:55:21 GMT Server: ICAP-Server-Software/1.0 Connection: close Encapsulated: req-hdr=0 GET /modified-path HTTP/1.1 Host: www.origin-server.com Accept: text/html, text/plain, image/gif Accept-Encoding: gzip, compress" If-None-Match: "xyzzy", "r2d2xxxx"
An iCAP server returning an error response when it receives a Request Modification request.	REQMOD icap://icap-server.net/content-filter ICAP/1.0 Host: icap-server.net Encapsulated: req-hdr=0 GET /naughty-content HTTP/1.1 Host: www.naughty-site.com Accept: text/html, text/plain Accept-Encoding: compress	ICAP/1.0 200 OK Date: Mon, 22 Jan 2001 09:55:21 GMT Server: ICAP-Server-Software/1.0 Connection: close Encapsulated: res-hdr=0, res-body=198 HTTP/1.1 403 Forbidden Date: Thu, 25 Nov 2001 16:02:10 GMT Server: Apache/1.3.12 (Unix) Last-Modified: Fri, 25 Jan 2001 13:51:37 GMT Etag: "63600-1989-3a017169" Content-Length: 62 Content-Type: text/html Sorry, you are not allowed to access that naughty content.

Generally, iCAP does not specify the *when*, *who*, or *why* for content manipulation, but only *how* to make content available to an application server that will perform the adaptation. For example, if iCAP is the tool to allow content translation/adaptation, one will still need an adaptation engine (iCAP server) to decide when, who, or why.

6.1.3 Trends and iCAP opportunities

Currently there is considerable demand for value-added services in three major areas: ad insertion, virus detection, and wireless. ICAP has the potential to give birth to such value-added services by enabling Web sites to offer Web applications closer to the user. Usually, ad insertion is based on either the origin Web site of the ISP, the hosting provider, or the site itself signing up for direct advertising. With iCAP, ad insertion will become more focussed and targeted to individuals based on the originating IP address (location), click behaviour of the customer (keyword adaptation), or user profiles. Target advertising is much more valuable when localised. An implementation of targeted ad insertion is shown in Figure 22.

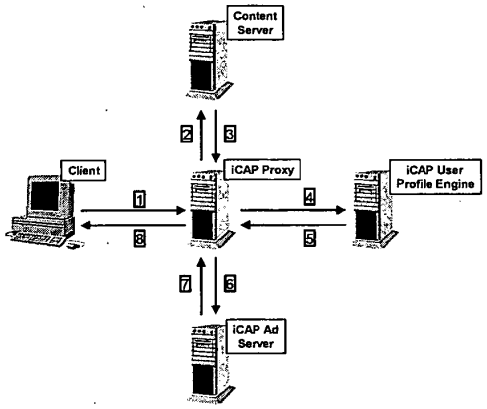


Figure 22: Possible implementation of targeted ad insertion. A client request for content [1]. The iCAP Proxy forwards the request to the Content Server [2]. The Content Server returns the requested content [3]. Based on e.g. the IP address of the Client, the iCAP Proxy request the Client's profile at the iCAP Profile Engine [4]. The Profile Engine returns the profile [5]. Both the Content and the Client's profile are forwarded to the iCAP Ad Server [6]. The Ad Server returns the Content with the proper ads inserted [7]. The iCAP Proxy returns the adapted content to the Client [8].

With concern growing over Internet security – or lack thereof – efficient, effective virus protection is another hot item. Virus scanning is usually left to the receiving network (or PC) to accomplish, and every object has the potential to be scanned many times, causing waste of resources. There is no historical "on-the-fly" method for virus scanning prior to delivery. Under iCAP, virus scanning allows previously scanned (and unaffected) objects to be cached and provided virus free. These objects never have to be scanned again.

For wireless devices such as PDAs and cell phones iCAP provides a means to adapt content for such heterogeneous devices. A cache can handle all client requests through redirects to such translation iCAP servers and maintain cached copies of multiple formatted objects for faster response to the client.

Tweaking Web content to speak the user's natural language might be another value-added service that becomes feasible under iCAP.

6.1.4 ICAP limitations

ICAP is not the perfect solution for content adaptation. It has several limitations:

- ICAP defines a method for forwarding HTTP messages only; it has no support for other protocols and for streaming media (e.g. audio / video),
- ICAP only covers the *transaction semantics* ("How do I ask for adaptation?") and not the *control policy* ("When am I supposed to ask for which adaptation from where?"),
- The current iCAP version relies on some form of encryption on the link or network layer for security,
- There are many different "flavours" of iCAP implementations (e.g. version 0.9, version 0.95, version 1.0 with modifications, etc.).

6.2 Middle boxes

There are a variety of intermediate devices in the Internet today that require application intelligence for their operation. Many of these devices enforce application specific policy based functions such as packet filtering, differentiated Quality of Service, tunnelling, intrusion detection, security and so forth. Network Address Translators, on the other hand, provide routing transparency across address. A Firewall is a policy-based packet filtering Middlebox, typically used for restricting access to/from specific devices and applications. There may be other types of devices requiring application intelligence for their operation. A middlebox is an intermediate device requiring application intelligence to implement one or more of the functions described. The discussion scope of this document is, however, limited to middleboxes implementing Firewall and NAT functions only. MIDCOM (see section 2.3.1) and MIDTAX are two IETF working groups that are currently solving several middlebox issues.

6.3 Transcoding and media gateways

Transcoding typically refers to the adaptation of streaming content. It is not very common to do that, because of the large performance penalties. Typical scenarios exploit content-negotiation to negotiate between different formats in order to obtain the most optimal combination of requested quality and available resources.

However, in some cases there are intrinsic needs for transcoding. For instance, when terminals are used that have only support for limited signalling or codec support. To give such systems access to streaming content, transcoders can be applied. Typically, these transcoders are special kinds of the more general 'media gateways'. A media gateway is currently typically used between traditional GSTN ('switched telephony') networks and packet networks (voice over IP solutions), or as H.320/H.323 gateways. There is an ongoing standardisation effort for these systems. This is a joint effort between IETF-Megaco workgroup and ITU study group 16, resulting in the (quite complex) Megaco protocol (also referred to as H.248). The requirements can be found in RFC2805: http://www.ietf.org/rfc/rfc2805.txt . The protocol actually describes the interactions that are possible between a media gateway (that manages the different connection-endpoints of the various streams) and a media gateway controller. The media-gateway controller is responsible for associating different endpoints, and allocating possible media-translation

functions between endpoints. Endpoints can be physical endpoints (e.g. ISDN channels or leased lines) or logical endpoints (e.g. RTP stream endpoints). A lot of information about media gateway controls can be found on the Megaco Web site: http://www.ietf.org/html.charters/megaco.html. For value-added RTSP streaming proxies, the media gateway controller is typically co-located with the RTSP proxy, and the media gateway itself resides in the RTP datapath between origin server and client. This is, however, not typical behaviour of a media gateway, and it is a research question whether such a configuration should work.

6.4 Transcoding and XML/HTML

Transcoding of HTML or XML content to specialised layouts is quite common these days. This is often performed at the origin server, using XSL stylesheets, but stylesheets can in principle also be applied on the client or on a proxy (although the latter has some security implications!) All current Web-environments support this; it allows one to separate the content from the presentation. For instance, XML content can easily be transcoded into HTML4 or WML content using the proper stylesheets. A lot of information about XSL, and the usage of XSL for transcoding is available on the Web, see e.g. http://www.w3.org/Style/XSL/ for an overview of the latest standards, ongoing standardisation efforts and existing implementations.

7 Authorisation, authentication, and accounting

Another aspect of Content Distribution networks deals with *AAA*: authorisation, authentication and accounting. It is clear that an increasing part of the content of the Internet is access-controlled. Therefore, proper authentication, accounting, and access control is necessary, certainly for third-party content-delivery service providers. There is currently not a single, or standardised, way of doing this.

7.1 What is AAA?

Today we live in a world where almost everything must be protected from misuse and where nothing is free. An increasing part of the content of the Internet is access-controlled. When providing commercial network services and content to public, there are three things that are commonly needed. These are authentication, authorisation and accounting. Authentication is needed to make sure that the user of the service is who he claims to be. This is quite important, because you don't want that someone else is using the service or content you have paid for. Usually authentication is provided by using a shared secret or a trusted third party. Related to authentication is authorisation. After the user has been authenticated we need a way to ensure that the user is authorised to do the things he is requesting. For example, if you are a normal user you don't have the permissions to access all the files in a file system. Usually authorisation is provided by using access control lists or policies. Accounting is the process in which the network service provider collects information of the network usage for billing, capacity planning and other purposes. This is important for the service provider, because there is no such thing as a free lunch.

7.2 AAA definitions

AAA stands for *Authentication, Authorisation and Accounting*. These are the three basic issues that are encountered frequently in many network services. Examples of these services are dial in access to Internet, electronic commerce, Internet printing, and Mobile IP.

But what do we exactly mean by these terms Authentication, Authorisation and Accounting? These are quite broadly used, but their meanings can be mixed up. The following list defines how they are used in this document.

- Authentication is the act of verifying a claimed identity, in the form of a pre-existing label from a mutually known namespace, as the originator of the message (message authentication) or as the channel end point[27](see also http://www.ietf.org/internet-drafts/draft-ietf-mobileip-aaa-reqs-01.txt).

[27] Glass, S. & Hiller, T. & Jacobs, S. & Perkins, C. , Mobile IP Authentication, Authorisation, and Accounting Requirements, Internet draft (work in progress), 11.2.2000.

- Authorisation is the act of determining whether a particular right can be granted to the presenter of a particular credential. This particular right can be, for example, an access to a resource[28] [see also http://www.ietf.org/internet-drafts/draft-ietf-mobileip-aaa-reqs-01.txt].
- Accounting can be defined as the functionality concerned with linking the usage of services, content and resources to an identified, authenticated and authorised person responsible for the usage, and a context in which these are used (e.g. time-of-day or physical location). Accounting includes the actual data logging, as well as management functionality to make logging possible. Accounting gathers collected resource consumption data for the purpose of capacity and trend analysis, auditing and billing. The information is ordered into user and session records and stored for later use for any of these three purposes[29].

One thing worth noticing is that the term authentication is used to denote the act of verifying an identity. This is noteworthy since the term has also other meanings, for example, the act of proving the authenticity of any object or piece of information.

Accounting has a strong relationship with authentication of users and authorisation of service usage. For example, to allocate service usage to the right user, it is necessary to have proof of the identity of the user, i.e., to authenticate the user. Because of the sensitive nature of accounting data, it may only be made available to authorised users. In the case of a prepaid service, the use of a service may only be authorised when the balance on the user's prepaid account is sufficient. Therefore, authentication, authorisation and accounting (AAA) are usually considered in combination.

7.3 AAA standardisation

Proposals for the AAA protocols and systems are currently being developed in the *AAA working group* (IETF's AAA working group: http://www.ietf.org/html.charters/aaa-charter.html) of the IETF. The goal of the AAA working group is to define one protocol that implements authentication, authorisation and accounting and is general enough to be used in a variety of applications. Currently only separate protocols are available to implement authentication, authorisation and accounting functionality. This is not desirable, because there are a lot of applications where they are needed together. There is also another group in the IETF called the *AAA Architecture Research Group* (AAAARCH, http://www.phys.uu.nl/~wwwfi/aaaarch/charter.html), which is responsible for developing a generic AAA architecture[30].

---

[28] Glass, S. & Hiller, T. & Jacobs, S. & Perkins, C. , Mobile IP Authentication, Authorisation, and Accounting Requirements, Internet draft (work in progress), 11.2.2000.
[29] GigaABP/D2.1, Jonkers, H. (ed.), Hille, S.C., Tokmakoff, A.& Wibbels, M., A functional architecture for the financial exploitation of network-based services, Enschede, Telematica Instituut, 2000.
[30] de Laat, C. & Gross, G. & Gommans, L. & Vollbrecht, J. & Spence, C., Generic AAA architecture, Internet Draft (work in progress), January 2000. < http://www.ietf.org/internet-drafts/draft-irtf-aaaarch-generic-00.txt .

7.4 AAA in a CDN

Source: Extensible Proxy Services Framework -IETF draft-tomlinson-epsfw-00.txt (http://www.ietf.org/internet-drafts/draft-tomlinson-epsfw-01.txt).

The AAA requirements for a CDN service environment are driven by the need to ensure authorisation of the client, publishing server or administrative server attempting to inject proxylet functionality, to authenticate injected proxylets, and to perform accounting on proxylet functions so the client or publishing server can be billed for the services. In addition, AAA is also required for a host willing to act as a remote callout server.

A typical CDN has relationships with publishers and provides them with accounting and access-related information. This information is typically provided in the form of aggregate or detailed log files.

In addition, these CDNs typically collect accounting information to aid in operation, billing and SLA verification. Since all accounting data is collected within the CDN's administrative domain there is no need for generalised systems or protocols.

Figure 23 contains a diagram of the trust relationships between the different entities in the service environment caching proxy architecture. These trust relationships govern the communication channels between entities, not necessarily the objects upon which the entities are allowed to operate (source: "Extensible Proxy Services Framework", IETF Internet draft, http://www.ietf.org/internet-drafts/draft-tomlinson-epsfw-01.txt) .

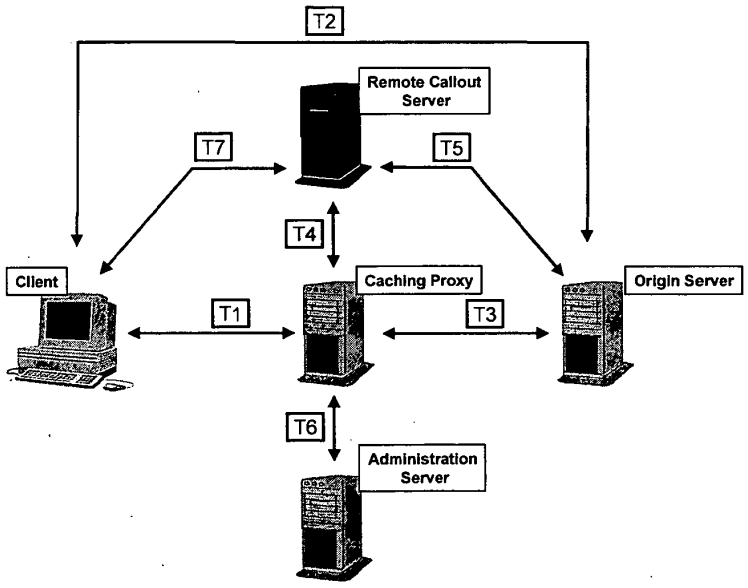


Figure 23: AAA trust relationships.

7.4.1 AAA in the existing Web system model

In the traditional client/server Web model, only T2 (end-to-end) and T1/T3 (hop-by-hop) are present. For T2, HTTP1.1 contains the WWW-Authenticate header for a server to indicate to the client what authentication scheme to use and the authorisation header for the client to present credentials to a server. The client presents these credentials if it receives a 401 (Unauthorised) response. HTTP authentication mechanisms that do not involve clear text transmittal of a password are detailed. At the user level, the mechanism used by the server to authorise and authenticate a client is challenge/response with some kind of login box, but there is no requirement for AAA in general. Access control lists can be used to fine tune control. In this case, the server could deny a client access to a particular object. In addition, if the server uses SSL[31], the client is assured of privacy in its transactions and can send a clear text password. In the other direction, there is no support for a client to authenticate a server. Since the client must discover the server's URL somehow, authentication of the source of the URL can provide some assurance that the URL is trusted. Typically, a person obtains the URL through some non-computational means and the client initiates the connection, so the client must know through some non-computational means that the URL is trusted. Examples of where a client can obtain a URL are through an e-mail message from a friend or co-worker, from a print or TV advertisement, or as a link form another Web page. However, unless the client is running secure DNS, the client can't determine whether the server's DNS entry has been hijacked. If SSL is used, then bi-directional authentication is possible. However, SSL primarily performs encryption, which might be unnecessary for a particular application, and additionally requires a different URL scheme (HTTPS instead of HTTP).

The addition of a proxy without a service environment (except perhaps for caching) changes the trust model to split T2 into T1 and T3 (although this does not mean that T2 is equivalent to T1 and T3). To the server, the proxy acts as a client, while to the client, it acts as the server. HTTP 1.1 contains a header, Proxy-Authenticate, that the proxy sends back to the client along with a 407 (Proxy Authentication Required) if the client must authenticate itself with the proxy. The client then sends back the Proxy-Authorisation header with credentials. This addresses the T1 relationship in the client to proxy direction. The T3 relationship in the proxy to server direction is addressed by having the server respond with a 407 (Proxy Authentication Required) and the Proxy-Authenticate header. Since Proxy-Authenticate is a hop-by-hop header, it can be used to authenticate the proxy to server connection just as it is used for the client to proxy connection. But there is still a lack of authorisation and authentication in the proxy to client and server to proxy direction, just as for end-to-end security. For a proxy acting as an avatar, the client is likely to have obtained the URL from a system administrator or other trusted source. Similarly, for a proxy acting as a surrogate, the publishing server typically has a business relationship with the surrogate provider, and the surrogate's URL or address is obtained by the server through some undefined, but necessarily secure means, because the surrogate provider wants to charge the publisher and prohibit unauthorised discovery.

[31] SSL (Secure Sockets Layer) is a commonly-used protocol for managing the security of a message transmission on the Internet.

7.4.2 AAA in the service environment caching proxy model

The lack of a mechanism whereby a client can authorise a proxy and a proxy can authorise a server means that the reverse directions of T1 and T3 are not addressed by HTTP/1.1. In the service environment caching proxy architecture, servers provide the caching proxy with computational objects (rule modules and proxylets) and therefore must be authorised to do so.

Therefore a service environment caching proxy acting as a surrogate must be able to demand authentication information from a server and a server must be able to respond with authentication information appropriate to the request, to authorise the server to provide computational objects. Moreover, a mechanism must be provided whereby a service environment caching proxy acting as a surrogate can authenticate individual proxylets and rule modules provided by an authorised server, if necessary.

For T1, the existing HTTP Proxy-Authenticate mechanism allows the service environment caching proxy acting as an avatar to authorise the client, but there is no mechanism for authentication of individual proxylets and rule modules, generating the requirement: This means that a mechanism must be present whereby a service environment caching proxy acting as an avatar can authenticate individual proxylets and rule modules provided by an authorised client, if necessary.

The proxy to client direction of T1 requires authentication, even though none is supplied in standard HTTP/1.1. Because a client will be providing computational objects to an avatar, it is essential that the client knows it can trust a service environment caching proxy acting as an avatar; otherwise, the computational objects may be provided to an unauthorised or hostile proxy, much to the client's detriment.

Finally, services run on the service environment caching proxy need to be paid. In other words, the service environment caching proxy server must be able to deliver secure, non-repudiable accounting information to a billing entity.

7.4.3 AAA in the Remote Callout Server model

In addition to the injection of proxylet functionality on the caching proxy, the caching proxy can also make use of a remote callout engine to modify particular objects. This architectural piece gives rise to the trust relationship T4, between the caching proxy and the remote callout engine, T5, between the remote callout engine and the server, and T6, between the client and the remote callout engine.

Existing remote callout protocols leverage off of HTTP authentication for the remote callout server. The ICAP specification explicitly states that an ICAP server acts as a proxy for purposes of authentication so a proxy client can send any Proxy-Authenticate and Proxy-Authorisation headers, although other hop-by-hop headers are not forwarded. However, this has little use for purposes of authenticating trust relationships T7 and T5. The remote callout server may require that the client or publishing server authenticate separately from the proxy, if the remote callout server is owned and administered by a separate entity from the proxy. In addition, a message from the caching proxy to a server that generates a 407 (Proxy Authentication Required) may or may not have been processed by the ICAP server, but in any event, the server won't know that the message

was so processed. The server responds to the sender of the message, namely the caching proxy. The caching proxy must respond with its credentials, the ICAP server is essentially invisible as far as the server is concerned.

Trust relationships T7 and T5 could derive transitively from T1/T4 and T3/T4. In that case, authorisation granted by/to the caching proxy is considered to be authorisation granted by/to the remote callout server. If the remote callout server is in the same administrative domain as the caching proxy, as is assumed in the ICAP specification, this is likely to be the case. However, in the general case, where the remote callout server resides outside the domain of the service environment caching proxy, authorisation by/of the caching proxy server is insufficient. A mechanism is required whereby, when the remote callout server is outside the administrative domain of the caching proxy, the remote callout server can directly authenticate with the publishing server and/or with the client, and the client or publishing server can directly authorise a remote callout server independent of the proxy. This requirement, if imposed on the HTTP stream between the client and server, would remove the invisibility of the remote callout server. However, this requirement could be met by an out-of-band authentication procedure, for example, using Diameter, in which case the remote callout server would remain invisible during HTTP transactions. ACLs could be established on the server allowing or denying access to the particular data objects for the remote callout server, at the expense of making the remote callout server visible to HTTP streams. Note that there is no need to authenticate computational objects because the remote callout server, by definition, does not receive computational objects from the client and/or publishing server.

The trust relationship T4 is on the remote callout to proxy connection. If the remote callout server is in a separate domain, authentication is required between the remote callout server and the caching proxy. Again, proxy authentication can be used in the remote callout to proxy direction, but there is no way for the caching proxy to authenticate the remote callout server. When the remote callout server is outside the administrative domain of the caching proxy, some means of authenticating the remote callout server with the caching proxy is required.

We also require uniform mechanisms on both the forward and reverse directions of T4, and T7 and T5 as well: The new authentication mechanism for the relationship T4 in the proxy to remote callout direction should be uniform with the mechanism in the opposite direction, either by implementing the new mechanisms in a manner similar to the old or by supplementing the old mechanisms with new. Authentication mechanisms for T7 and T5 may be uniform with other authentication mechanisms.

The requirement on T7 and T5 is looser in order to avoid overly constraining the mechanisms for verifying the other trust relationships, in which backward compatibility considerations may play a large role.

Finally, services run on the remote callout server need to be paid. The remote callout server must therefore be able to deliver secure, non-repudiable accounting information to a billing entity. Most likely, the billing entity will be the administrative server, but it may be another. If the billing entity is the administrative server, and the remote callout server is outside the domain of the caching proxy, the method whereby the accounting information is delivered must be secure and allow non-repudiation, so that the owners of the remote callout server can be assured of proper billing and payment.

7.4.4 AAA in the Administrative Server model

The administrative server is responsible for injecting proxylets into the service environment caching proxy, and for collecting accounting information from the service environment caching proxy and, transitively, from the remote callout server. The proxylets injected by the administrative server may run at an additional level of trust from those introduced by clients and publishing servers, since they may be involved in collecting accounting information or in other sensitive tasks.

From a practical standpoint, the administrative server is highly likely to be within the same administrative domain as the caching proxy, but as with the remote callout server, the case where it is not may also occur. This requires that trust relationship T6 be verified. Therefore, a mechanism must be present whereby, when the administrative server is outside the domain of the caching proxy, mutual authentication between the caching proxy and administrative server is possible.

The administrative server also requires some means of obtaining accounting information from the caching proxy and remote callout server: The administrative server must obtain accounting information that is secure and non-repudiable from the caching proxy and remote callout server.

Finally, if the administrative server is allowed to inject proxylets at an additional trust level, an additional authentication mechanism may be required: If the administrative server can inject proxylets at a higher trust level into the service environment proxy, a mechanism must be present whereby the additional trust level can be verified (possibly with human involvement).

7.5 Accounting in peered CDNs

Peering or interconnecting CDNs introduces the need to obtain accounting data from a foreign domain. This requirement means that customers of a peered CDN service (publishers, clients, and CDNs) must now have a generalised or standard means of obtaining accounting information to support current as well as planned business models. For example, the desire to implement business models such as "Pay-per-View" may require that there exist a mechanism for authenticating and authorising clients at a delivery point that lies in a foreign domain/CDN. See also section 3.3.

CDN peering must provide the ability for the content provider to collect accounting data regarding the delivery of their content by the peered CDNs. Accounting CDN Peering Gateways (CPGs) exchange the data collected by the interior accounting systems. This interior data may be collected, via, e.g., FTP, from the surrogates by the Accounting CPGs. Accounting CPGs may transfer the data to exterior neighbouring Accounting CPGs on request (push), in an asynchronous manner (push), or a combination of both. Accounting data may also be aggregated before it is transferred. The ability to aggregate statistical and access related information is essential to allow for scalability within the proposed solution. Figure 24 shows a diagram of the entities involved in the accounting peering system.

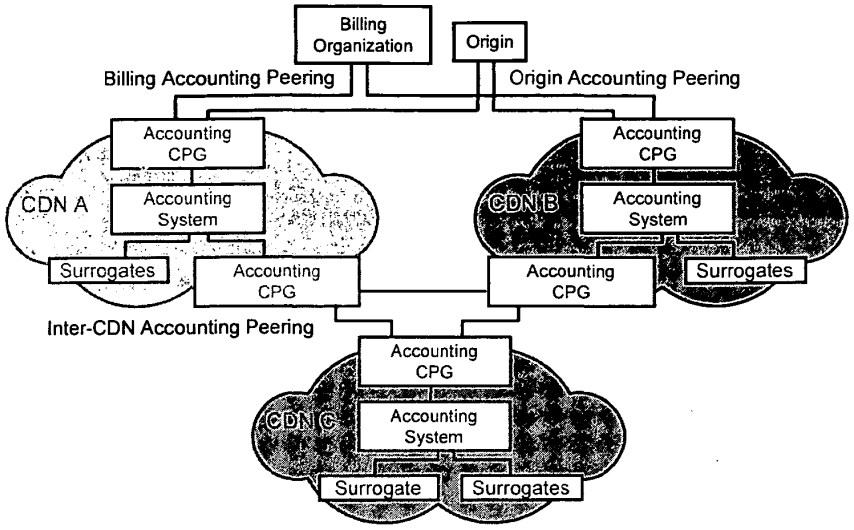


Figure 24: Accounting peering system architecture (source: Content Internetworking Architectural Overview, IETF Internet draft, http://www.ietf.org/internet-drafts/draft-green-cdnp-gen-arch-03.txt).

Three CDN accounting peering relationships are expected to be common in the near future:

- Inter-CDN accounting peering,
- Billing organisation accounting peering, and
- Origin accounting peering.

Inter-CDN accounting peering involves exchanging accounting information between individual CDNs in an inter-network of peered CDNs. Billing organisation peering involves exchanging of accounting information between CDNs and billing organisations. Origin accounting peering involves the exchange of accounting information between CDNs and the owner of the original content.

It is not necessary for an Origin to peer directly with multiple CDNs in order to participate in CDN peering. Origins participating in a single home CDN will be indirectly peered by their home CDN with the inter-network of CDNs the home CDN is member of. Nor is it necessary to have a Billing Organisation peer, since this function may also be provided by the home CDN. However, Origins that directly peer for accounting may have access to greater accounting detail. Also, through the use of accounting peering, third party billing can be provided.

7.6 DRM

Digital Rights Management (DRM) is the process of protecting and managing the rights of all participants engaged in the electronic commerce and digital distribution of content.

DRM technologies are being developed as a means of protection against online piracy of commercially marketed material. DRM has proliferated through the widespread use of

Napster and other peer-to-peer file exchange programs. It will become an important issue in CDNs as well, since original content will be distributed over the network.

DRM tools allow content providers to deliver songs, videos, books, and other media over the Internet in a protected, encrypted file format. Media files will be packaged, encrypted and locked with a key. This key is stored in an encrypted license, which is usually distributed separately but could also be transported with the media in some case. Other information may be added to the media file, such as the URL where the license can be acquired. A clearinghouse can be used to store the specific rights or rules of the license and implement the media rights manager license services. The role of the clearinghouse is to authenticate the consumer's request for a license. The protected file can be easily distributed over the Internet, placed on media servers for streaming, or placed on a Web site for download since only licensed customers are allowed to actually view the content.

DRM helps enable:

- Protection of digital content. By scrambling or encrypting content DRM enables authors and publishers to send digital content over an unsecured network so that content can be read only by the intended recipients (key owners).
- Secure content distribution. Once the digital content is protected via DRM encryption, the proper key is needed to decrypt the content and render it readable. Without the key, the file is unintelligible. Anyone can have access to the encrypted content, but it will be of no use without the decryption key.
- Content authenticity. A message digest is created from a one-way hash function when the original, authentic content is published.
- Transaction non-repudiation. Digital signatures are used.
- Market participant identification. A digital certificate is created using a cryptographic technique that binds a person's identity with his or her public cryptographic key. A digital certificate combines an individual's public key, other identity information and one or more digital signatures. The digital signatures belong to certificate authorities trusted to attest that the public key, in fact, belongs to the person named in the certificate.

For CDNs, this means that the issue of distributing content from the origin server over the CDN network to local servers can be solved with DRM. After all, DRM facilitates controlled distribution of content over an insecure network, like the Internet. Figure 25 shows a possible implementation of DRM functionality in a CDN network.

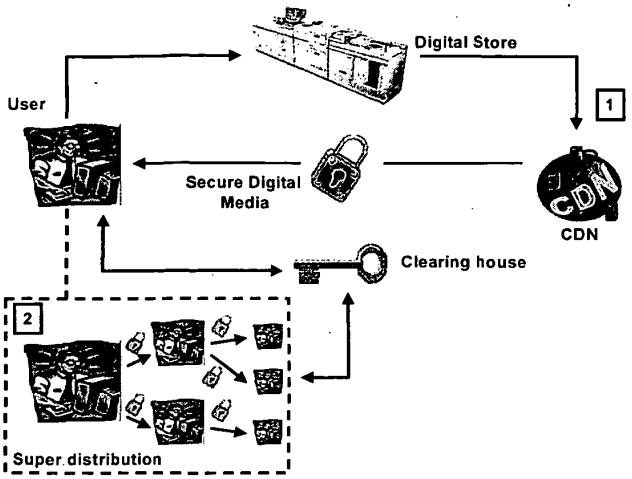


Figure 25: Digital Rights Management in a CDN. [1] The file is optimally delivered world-wide via a CDN, resulting in optimal end-user experience and expanded customer reach. [2] file sharing has been transformed into a new revenue channel.

The areas of DRM and accounting and billing are closely tied. People need to be paid royalties for their intellectual property. DRM allows one to protect that content in a controlled way and it can easily be used, when tied to mediation systems, to support various business models. For instance, the customer may order three movies and view them at leisure over 24 hours, order a single movie for one viewing, or pay a month's subscription for unlimited access to a movie library.



The Moving Picture Experts Group (MPEG) develops standards for digital video and digital audio compression. MPEG is working on an elaborate DRM scheme for all sort of multimedia in the MPEG-21 framework[32]. Both MPEG-2 and MPEG-4 allow a DRM scheme. In MPEG-2 content can both be uniquely identified as containing copyrighted material as well as a being protected. The identification part contains a unique 32-bit number and a number pointing to a registration authority. To protect an MPEG-2 stream

[32] http://www.cselt.it/mpeg/standards/ipmp/index.htm

provisions are made to signal that the streams are scrambled and to signal the authority. The problem with this scheme, however, is that each authority can use its own scheme, which makes it hard to make interoperable hardware.

In MPEG-4 there are more elaborate hooks and identification and protection are tightly build into the system layer, which makes it possible to build secure MPEG-4 delivery chains in efficient ways. The bitstream contains information that enables the terminal to select a particular DRM scheme for processing the bitstream. Extensions of the basic MPEG-4 scheme provide means to describe how streams can be decoded and encoded. An independent registration office is used for any party to register its DRM scheme.

DRM may also pose a problem for CDNs going further than just caching. In models like TVanytime[33] or MPEG-21, content never leaves the DRM environment. In these schemes the key needed to unlock the scrambled stream guarantees this. As long as a CDN just caches the stream this is unproblematic although there may be some legal repercussions. However, if the stream has to be modified, for example to down sample it for use on a mobile device, the stream has to be descrambled by the CDN. This means that the CDN must be authorised either by a key related to the key of the end user, or by the content owner. In the first case, the owner must be sure that the transcoding proxy will rescramble the content otherwise an unencoded stream will get out. Either way, the content owner must trust the CDN.

7.7 Lack of AAA in current CDNs

Because AAA functionality and AAA models are quite new and rapidly changing, there are not any standardised methods available.

Moreover, during our research survey for state-of-the-art content delivery networks, little information about AAA or security in general was found. A main source of information about these subjects logically is provided by the IETF working groups and drafts related to CDN aspects. However, many of the drafts do admit that the services they provide will only work for insecure content. Others make no mention of that.

Clearly the aspect of security is considered important but recommendations are barely given. Maybe the aspect of security is too elaborate to discuss and one desires to focus on the key problem. The discussions about AAA functionality in these CDN related working groups stick to a similar level.

The CDN providers of today are also not involved in much AAA and security activity. The reasons for neglecting these aspects could be:

- The CDN providers are still busy trying to build up their CDN network. This includes for instance the installation of hardware devices. Security is of later concern.
- The need for detailed AAA information is not present. One is, at most, interested in high-level Web server statistic at the moment. As for accounting, CDN providers such as Akamai Technologies Inc. and Digital Island Inc. list flat fees per megabit per second of usage; complex accounting strategies are not used.

[33] www.tv-anytime.org

- The implementation of a balanced end-to-end security architecture is difficult, expensive, and time consuming.
- Most current Internet customers don't want to experience any constraint on the use of the content he/she wants to acquire. AAA for content provision is not done. The customer expects to be king.

7.8 Accounting revenue sources

Several possible accounting revenue sources may become important for future CDN exploiters. The table below gives a short list of today's situation and the future possibilities.

Revenue source	Today	5-10 years
Pay-per-view	Sex only	e-learning, video-on-demand, high profile events
Subscription	Non-existing	Important model for quality content
Syndication	Taking off	Important model
Format licensing	Works for top brands	Growing
Merchandise	Few examples	Works for top brands
Advertising	Works for top brands	Works better due to better ratings
Sponsoring	First initiatives	Important model
SMS	Many examples	Replaced by more advanced services
Telephone voting	Well established	Well established
Product placement	Established	Bigger

8 Other platforms and system architectures

In a way, CDN providers offer a (middleware) platform for a wide range of interactive functions, from searching to user profiling to order processing. Middleware platforms, like CORBA and DCOM, enable CDN providers to cost effectively and transparently provide services, content management, and accounting and billing. These and other middleware platforms are described in the Telematica Instituut Middleware state of the art deliverable[34].

The areas of distributed operating systems and parallel computing and middleware platforms seem to come closer. They might even benefit from each other. This section will discuss several other platform technologies and system architectures that deal with the aspect of distribution information at a similar level as CDNs do.

8.1 The Globe middleware, GlobeDoc and the GDN.

Globe[35] is a middleware platform to help design wide area distributed applications. It is a research project of the computer systems group of Maarten van Steen and Andrew Tanenbaum at the Vrije Universiteit of Amsterdam. It has three principal design objectives: support a uniform model of distributed computing, support a flexible implementation framework, and ensure world-wide scalability. To test their ideas they have designed two major applications on top of Globe. GlobeDoc is a Globe based scalable, http-interoperable implementation of the Web, and the Globe Distribution Network (GDN), a content distribution network for freely available software.

8.1.1 The Globe system

The Globe system[36] is a wide area distributed system that is constructed as a middleware layer on top of Unix and Windows NT (although the latter seems not to be available yet). Globe has an object model and a collection of basic support services. Globe objects can be shared by many distributed processes, which can be distributed over the planet. Support services include naming and locating objects.

For world-wide scalability, objects have to provide support for partitioning and replication. Such support is not provided by "standard" middleware systems like CORBA or DCOM. Support for replication is provided by distributed file systems like AFS or CODA and optionally by the Web when using various complex caching strategies. However, for each of these systems, this strategy is fixed, whereas Globe's policy is very flexible and on a per object basis.

---

[34] Hulsebosch, B., Teeuw, W., and Poortinga, R., "Middleware", Tintel state-of-the-Art deliverable, Telematica Instituut, 1999, Enschede.
[35] http://www.cs.vu.nl/~steen/globe/
[36] M. Steen, P. Homburg, A.S. Tanenbaum, Globe a Wide-Area Distributed System. IEEE concurency Jan. March 1999
http://globe.cs.vu.nl:23003/nl/vu/cs/globe/proj/papers:/ftp/ieeeconc.99.org.pdf

The fundamental abstraction of the Globe system is that of a distributed shared object (DSO). Each DSO offers one or more interfaces, with each interface consisting of a set of methods. The objects interact and communicate only through this interface. A DSO has state that can be physically distributed over many different address spaces, which means that its state can be partitioned and replicated on many different machines. Processes are unaware of this because all non functional aspects like transport of method invocations, location, migration, replication of its state and security are hidden by the interface and are handled by the object itself using only a minimum of supporting services. A distributed object is built from local objects that reside in different address spaces, and communicate with each other. The Globe model thus abstracts not a master object, which is somehow magically cached when accessed over a network, but all replicas of one semantically defined object together. The local object gives a local "view" on the object in the way that it finds most useful or is cheapest. All replicas are equal, although some may, if convenient, be more equal than others.

To invoke a method of an object, the Globe system must first bind to this object. To do this, it contacts a contact address, that describes its network address, and the protocol through which the binding takes place. Binding then results in the interface of the object being placed in the clients address space, together with the implementation of that interface. This is called the *local representative* or *local object*.

A local object consists itself out of sub objects.

- A semantics object that is user defined and implements the functionality of the DSO.
- A communication sub object that is responsible for sending and receiving messages from other local objects.
- A replication object that implements the replication strategy that is appropriate for this particular object.
- A control object that handles the control flow within the local object.
- A security sub object that handles security.
- A persistence object that handles persistence.

This modular architecture allows for objects to be implemented in different implementation languages, to run them run on a variety of platforms and enables them to communicate using different protocols.

8.1.2 The GlobeDoc System

The GlobeDoc system[37] is a scalable implementation of the Web, based on Globe instead of HTTP. It is designed to be very scalable as the number of users increases. This scalability is mostly simply inherited from Globe by wrapping each document in a DSO. Replication and mirroring is done automatically by the system. In this way it is very similar to a Content Distribution Network.

---

[37] See. I. Kuz, P. Verkaik I. van der Wijk, M. van Steen, and A. S. Tanenbaum. Beyond HTTP: an implementation of the Web in Globe. Technical Report IR 465 dept. Mathematics and Computer Science, VU Amsterdam http://globe.cs.vu.nl:23003/nl/vu/cs/globe/proj/papers:/http/IR-465.99.pdf.

The distinguishing feature of GlobeDoc is that every document has its own replication strategy, rather than a one size fits all one. It gives near optimal performance, under performance metrics depending on the situation and the whim of the document provider.

GlobeDoc supports location independent names called HFN's (human friendly names) which allows documents to be produced and maintained on different sites all over the world whilst presenting them as a single Web site. The system will serve up the replica closest to the client.

A new experimental system is Globule[38], a platform which automates all aspects of replicating Web documents at a world-wide scale: server-to-server peering negotiation, creation and destruction of replicas, selection of the most appropriate replication strategies on a per-document basis, consistency management and transparent redirection of clients to replicas. To facilitate the transition from a non-replicated server to a replicated one, Globule is implemented as a module for the Apache Web server.

8.1.3 The Globe Distribution Network (GDN).

The GDN is a content distribution network build on top of Globe[39]. In its initial implementation it is aimed at the distribution of freely distributed software, because it is a good testbed: many files, many potential users (which are more likely to be beta software hardened) and a rapidly changing use pattern of files. There are also interesting copyright issues which have to be dealt with.

Like in GlobeDoc, every software package is wrapped in a Globe DSO. To make the threshold for use as low as possible, the GDN is accessible through a standard Web browser thus integrated with the Web.

The GDN itself consists of a number of modified HTTPD's running on machines all over the world. The HTTPD interprets a HTTP request, calls the corresponding DSO and sends the HTML formatted result back to the browser.

Users must choose a GDN-HTTPD preferably the closest one. Once connected to the GDN however, the storage location becomes transparent, and the GDN will find the nearest replica using the Globe location service. In particular, if a client has a local HTTPD it will be the closest, and the local representative built in the HTTPD will act as a replica for the DSO, which means that downloading is fast. This is called a GDN-proxy.

The GDN is similar to Napster and Gnutella in that it allows unreliable servers to become part of the network. The GDN is protected against unauthorised use and failures. It is designed to protect against (at least) two versions of unauthorised use: violating the integrity of the software, and distribution of commercial software or copyrighted music or

[38] G. Pierre, M. van Steen. "Globule: a Platform for Self-Replicating Web Documents." Technical Report IR-483, January 2001. http://globe.cs.vu.nl:23003/nl/vu/cs/globe/proj/papers:/http/IR-483.01.pdf.

[39] A. Bakker, E. Amade, G. Ballintijn, I. Kuz, P. Verkaik, I. van der Wijk, M. van Steen, A.S. Tanenbaum. "The Globe Distribution Network". Proc. 2000 USENIX Annual Conf. (FREENIX Track), San Diego, June 18-23, 2000, pp. 141-152. http://globe.cs.vu.nl:23003/nl/vu/cs/globe/proj/papers:/http/freenix.00.pdf.

films. Unfortunately as of writing, the full security scheme has not been fully implemented yet.

8.1.4 Status

Since October 2000 the Globe system is available for download at the Globe site (under a BSD license), see http://www.cs.vu.nl/pub/globe/ and http://globe.cs.vu.nl:23003/nl/vu/cs/globe/proj/papers:/http/IR-476.00.pdf.

It is claimed to be ready to be exposed to the public. It is now in version 0.8.0. The installation contains a detailed installation manual (http://globe.cs.vu.nl:23003/nl/vu/cs/globe/proj/giddy/releases:/version-0.7.0/gog-0.7.ps.gz), a name server, a location server, an object server and a Globe http server. In addition the Globedoc system and the GDN are included as well as some supporting utilities. The GlobeDoc works at least on the VU site whereas the GDN has currently four different nodes (http://globe.cs.vu.nl:23003/nl/vu/cs/globe/proj/gindex:/).

8.2 Globus, a Grid middleware layer

The Grid is the emerging computational and networking infrastructure designed to provide uniform and access to data, computational resources and humans on an extremely heterogeneous wide area network. The Globus[40] project aims to provide middleware services to support Grid computing environments. Globus is currently used for large computational tasks requiring large amounts of parallel computing, manipulating large amounts of data, and tele-immersion applications requiring multiple synchronised audio and video streams.

Remark: We believe that the specialised user base of Globus and the Grid is a purely sociological phenomenon. The high performance computing community has simply been networked longer then every body else (the Internet started by connecting supercomputing centres) and has the motive and the means (in money and brain power) to tackle the problems of large scale distributed processing. Big science is traditionally an international affair where the sharing of computational and data resources between non-trusting organisations is a necessity. High performance computing aims to do parallel processing in an environment with 100 Gbit/sec internal shared memory connections with predictable microsecond latencies, and a lowly 100 Mbit/sec connection with latencies up a tenth of a second to a remote computer at the other end of the globe.

Globus has four main components:
1. The Grid security infra structure (GSI) provides authentication and authorisation services using public key infra structure or Kerberos.
2. The Globus Resource Management architecture provides a language for specifying application requirements and mechanisms for immediate and advance reservations of resources. It also has several mechanisms for submitting jobs to remote machines.

[40] http://www.globus.org/research/papers/anatomy.pdf.

3. The Globus Information Management architecture provides a distributed collection of information servers on which to publish and retrieve resource information. They are accessed by higher level services, which perform resource discovery, scheduling and configuration.
4. The Globus Data Management architecture provides two components: a universal data transfer protocol called GridFTP and a replica management infrastructure for managing multiple copies of shared data sets. GridFTP is a secure and efficient data transport protocol based on the FTP standard.

Globus is a set of C-libraries that run on top of Unix. Some of Globus services require a daemon to run on the machine. However, Globus is designed to be "a bag of tools" i.e. it is a design goal to make at least part parts usable independent of each other.

We now discuss a few components that may be useful in the context of a Content Distribution network.

8.2.1 Grid Security Infrastructure.

The Grid security Infrastructure is designed for inter-site security, that makes use of the best security infra structure that is locally available. It grew out of the political and practical need NOT to have every site in the PACI testbed run Kerberos[41]. Likewise it allows site managers to keep control over their own resources, which is a key for the acceptance of a Grid infra structure. GSI allows a single sign on to a Globus network.

GSI is based on Credentials representing the identity of each entity such as a user, resource or program. A certification authority ties an identity to a public key pair by signing a certificate. Each resource can specify its policy how to accept incoming requests. The GSI is then responsible for verifying the global identity but then maps this on a local sites subject name and leaves the rest of the access control there.

GSI can work with multiple Credential Authorities, and allows storing the users private key on a smart card. A scheme exists to securely interface GSI with standard Web browsers

8.2.2 Globus Resource Management

8.2.2.1 QoS Management

The Globus Architecture for Reservation and Allocation (GARA) provides QoS mechanisms for network applications that have strongly varying network flows with high and low latency, and flows that may change their requirements dynamically during their lifetime[42]. It has a policy driven framework that allows for example to respond to resource availability by reducing rates (for example for video or large transfers) by introducing data compression for non-critical users.

---

[41] Kerberos is a secure method for authenticating a request for a service in a computer network.
[42] http://www.globus.org/documentation/incoming/iwqos_adapt1.pdf and http://www.globus.org/documentation/incoming/iwqos.pdf

GARA provides advance reservation and end-to-end management for quality of service of different types of resources, including networks, CPU's and disks.

A GARA system consists of a number of resource managers that each implement reservation, control and monitoring operations for a specific resource. This provides a more uniform interface than a "bandwidth broker" favoured in the network literature, and simplifies end-to-end QoS management strategies. Security is provided by the Globus security infrastructure. The network QoS manager uses the expedited forwarding-per-hop behaviour specified by the IETF Differentiated Service Workgroup. With careful admission control, it allows to build a QoS system with reasonably strong bandwidth guarantees, even though traffic is treated as an aggregate in the core of the network. To do so, the resource manager enables reservation requests by configuring the routers that it controls. In particular it configures the ingress routers that it controls to classify, police, mark and potentially shape, all packets that belong to a flow for which the reservation has been authorised, as is normally done for differentiated services. The expedited forwarding per-hop behaviour drops packets that exceed the reservation, but allows small bursts of excess traffic using a token-bucket mechanism.

8.2.2.2 The GRAM resource manager

The Grid Resource Allocation Manager provides an interface to scheduling and allocation primitives as they are found on the supporting OS and various distributed allocation programs like CONDOR[43].

8.2.3 Globus Data Management

See http://www.globus.org/research/papers/msc01.pdf : Secure, Efficient Data Transport and Replica Management for High-Performance Data-Intensive Computing. B. Allcock, J. Bester, J. Bresnahan, A. Chervenak, I. Foster, C. Kesselman, S. Meder, V. Nefedova, D. Quesnel, S. Tuecke. (Submitted to IEEE Mass Storage Conference, April 2001).

8.2.3.1 Replica Management

The Globus Replica management system provides the following services:

- Creating copies of a partial or complete file,
- Registering these copies in a Replica Catalogue,
- Allowing users and applications to query the catalogue to find all existing copies of a particular file or collection of files,
- Selecting the "best" replica for access based on storage and network performance predictions made by a Grid information service, There is work in progress to build higher level services for automatic creation of new replicas at desirable locations.

At the lowest level lies a Replica Catalogue, which allows users to register a set of files, which may be distributed over a WAN and which may contain replicas, as a single collection.

[43] http://www.globus.org/documentation/incoming/iwqos.pdf

8.3 Parlay

The Parlay Group was formed as a non-profit entity in 1998 in order to create open, technology independent Application Programming Interfaces (APIs), which enable Network Operators, Independent Software Vendors, and Service providers to generate and provide products and/or services that use the functionality resident in networks and that are suitable for and operate across multiple networks and to accelerate the adoption of these APIs through the sponsorship of developer education programs, certification efforts, and initiative promotion. The Parlay Group aims to create an explosion in the number of communication applications by specifying and promotion open Application Programming interfaces (APIs) that intimately link IT applications with the capabilities of the communications world. They enable carriers and independent software vendors to create applications (using existing network resources) that cross the traditional boundaries of technology, location and business. Members of the group are most major telecom- and IT companies like Alcatel, Cisco, Compaq, Ericsson, HP, IBM, Intel, Siemens, Lucent, SUN etc.

The purposes for which the Parlay Group[44] is organised are:

- To define, establish and support a common specification for industry standard Application Programming Interfaces (APIs), and to facilitate the production of test suites and applicable reference code in multiple technologies which provide a common foundation for the introduction of related products and services by developers across Wireless, Internet Protocol and Public Switched Networks;
- To provide a forum and environment whereby the Corporation's Members may meet to approve suggested revisions and enhancements that evolve the initial specifications; to make appropriate submissions to established agencies and bodies with the purpose of ratifying these specifications as an international standard; and, to provide a forum whereby users may meet with developers and providers of products and services to identify requirements for interoperability and general usability;
- To educate the business and consumer communities as of the value, benefits and applications for the Parlay APIs through publicity, publications, trade show demonstrations, seminars and other programs established by the Corporation;
- To support the creation and implementation of uniform conformance test procedures and processes which seek to assure the compliance of Parlay API implementations with the specifications;
- To maintain relationships and liaison with educational institutions, government research institutes, other technology consortia, and other organisations that support and contribute to the development of the specification.
- To foster competition in the development of new products and services based on specifications developed by the Corporation in conformance with all applicable antitrust laws and regulations.

Business View

The Parlay APIs (see Figure 26) exposes basic capabilities of the network provider in a secure and manageable way to a wide variety of application developers. Parlay-based services can be widely deployed in a variety of domains:

[44] http://www.parlay.org/about/index.asp

- Network Provider equipment
- Application Service Provider
- Service Bureau
- Enterprise
- Desktops
- Information Appliances
- Intelligent Handsets

Technology View

The following interfaces are defined:

Framework Interface Set

These provide the supporting capabilities necessary for the Service Interfaces access in a secure and manageable manner.

Service Interface Set

These offer applications access to a range of network capabilities and information. Functions provided by the service interfaces allow access to traditional network capabilities such as call management, messaging, and user interaction. The service interfaces also include generic application interfaces to ease the deployment of communications applications.

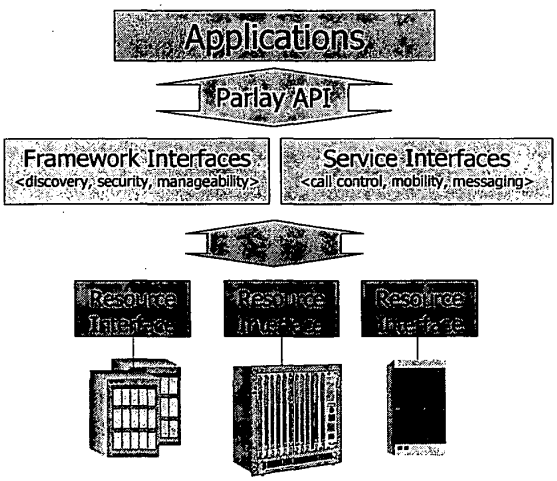


Figure 26: Parlay API.

Relationship with Other Standards

There is a proposed alignment between Parlay 2.1, JAIN SPA 2.0 (JAIN-Parlay), ETSI SPAN 3, and 3GPP OSA Call Control (see following paragraphs). People from these standardisation groups discuss issues in joint meetings and produce results that are commonly agreed upon.

8.4 3GPP-OSA

3GPP

3GPP stands for Third Generation Partnership Project, a co-operation of standardisation bodies (among which ETSI) called partners. The partners have agreed to co-operate in the production of globally applicable Technical Specifications and Technical Reports for a 3rd Generation Mobile System based on evolved GSM core networks and the radio access technologies that they support (i.e., Universal Terrestrial Radio Access (UTRA) both Frequency Division Duplex (FDD) and Time Division Duplex (TDD) modes). The partners have further agreed to co-operate in the maintenance and development of the Global System for Mobile communication (GSM) Technical Specifications and Technical Reports including evolved radio access technologies (e.g. General Packet Radio Service (GPRS) and Enhanced Data rates for GSM Evolution (EDGE)). More information can be found at: http://www.3gpp.org .

OSA

The 3GPP Technical Specification Group Core Network Workgroup 5 defines the Open Service Architecture (OSA). OSA defines an architecture that enables operator and third party applications to make use of network functionality through an open standardised interface (the OSA Interface). OSA provides the glue between applications and service capabilities provided by the network. In this way applications become independent from the underlying network technology. The applications constitute the top level of the Open Service Architecture (OSA). This level is connected to the Service Capability Servers (SCSs) via the OSA interface. The SCSs map the OSA interface onto the underlying telecommunications specific protocols and are therefore hiding the network complexity from the applications. More information about the Core Network WG 5 can be found at http://www.3gpp.org/TSG/CN5.htm.

8.5 JAIN

The JAIN initiative, organised by Sun in 1998, addresses the needs of next-generation telecom networks by developing a set of industry-defined APIs for Integrated Networks. Network services today are typically built using proprietary interfaces that inhibit the marketplace for new services. Members of the JAIN community have joined forces to define open APIs based on Sun's Java platform, thus allowing service providers to rapidly create and deploy new flexible, revenue-generating services. Information about the JAIN program can be found at http://java.sun.com/products/jain/.

The objective of the JAIN initiative is to create an open value chain from 3rd-party service providers, facility-based service providers, telecom providers, and network equipment providers to telecom, consumer and computer equipment manufacturers.

The JAIN APIs are a set of Java technology based APIs which enable the rapid development of Next Generation telecom products and services on the Java platform. The JAIN APIs bring service portability, convergence, and secure network access to telephony and data networks.

By providing a new level of abstraction and associated Java interfaces for service creation across Public Switched Telephone Network (PSTN), packet (e.g. Internet Protocol (IP) or Asynchronous Transfer Mode (ATM)) and wireless networks, JAIN technology enables the integration of Internet (IP) and Intelligent Network (IN) protocols. This is referred to as Integrated Networks. Furthermore, by allowing Java applications to have secure access to resources inside the network, the opportunity is created to deliver thousands of services rather than the dozens currently available. Thus, JAIN technology is changing the telecommunications market from many proprietary closed systems to a single network architecture where services can be rapidly created and deployed.

JAIN technology is being specified as a community extension to the Java Platform. It consists of two API Specification areas of development:

- The *Protocol API Specifications* specify interfaces to wireline, wireless and IP signalling protocols
- The *Application API Specifications* address the APIs required for service creation within a Java framework spanning across all protocols covered by the Protocol API Specifications

9 Conclusions

In this section we will end this state of the art survey with several evaluating remarks concerning the CDN matters described. We will do this CDN evaluation by means of a SWOT analysis. A SWOT analysis is an effective method of identifying the Strengths and Weaknesses, and to examine the Opportunities and Threats of CDNs. After the SWOT analysis, we will end with some research opportunities for CDN.

9.1 Strength of current CDN approaches

The strength of a CDN lies in the fact that it adds intelligence to network infrastructure. This intelligence can be leveraged as a platform to host value-added services within the network infrastructure. Such value-added services include the proper distribution and storage of content. As a result, the consumer network edge can be leveraged for strategically placed value-added services in the data plane. Examples of value-added services in the data plane are services for personalisation, ad insertion, content adaptation and virus filtering. Furthermore, CDN services provide performance (8 seconds rule, quality of service) and content (dynamic, streaming) differentiation. Bringing content closer to its receivers results in faster download times. As a result, it preserves the existing customer relationship, generates a higher margin revenue stream, and a reduction of the server load (due to a reduction of the processing time).

9.2 Weakness of current CDN approaches

On the other hand, there are several drawbacks for using CDNs. The costs for exploitation of a CDN are relatively high. We have observed a weak community on accounting and billing models. Service providers express a strong interest in accounting issues but don't actually contribute to a solution. Rules for proxy functionality are barely defined and the protocol modules are unclear. The current delivery of content is mainly based on unicast. Multicast functionality is desired considering the increased demand for audio and video content (live sporting events and fashion shows, for example).

9.3 Opportunities for future CDNs

The latter issue brings us immediately to an opportunity of a CDN: its inherent architecture is well suited for such multicast events. Combined with replication technologies a CDN has potential to offer efficient multicast delivery of especially rich content. The CDN market will be driven by the proliferation of streaming media. As streaming gains widespread adoption, CDN market growth will accelerate. As a result, the cost of CDN products and services will decrease over time, driving adoption rates up. This is stimulated by an increasing Web site traffic demand for bandwidth. Moreover, CDN peering allows for broader reach, scale, and enhanced performance across global networks. Proper and reliable distribution and management of content becomes very important; content must be distributed and stored in advance of demand. New value-added services for content distribution, adaptation, or negotiation can easily be implemented and offer large opportunities for a successful future of CDN.

9.4 Threats for future CDNs

Several threats, however, could spoil this prosperous CDN future. Legal issues form an important threat. How far can we go in adapting original content for network distribution and delivery? Content caching and replication is an important functionality in a CDN, what if content owners don't agree with such distribution of their content. Digital rights management solutions are currently not advanced enough to tackle these problems. The ability to intelligently link and monitor centralised content with edge delivery systems is critical to the deployment of content delivery networks. Questions can be set by this ability. Without scalable and reliable distributed storage and edge servers, CDNs are vulnerable. Security aspects (authentication, authorisation and denial of service protection) are hardly spoken of. The current CDN business models are changing rapidly. How do you make money in the CDN value chain as the business models evolve? Finally, there is the promise of the "infinite bandwidth future"; an illusion or reality?

9.5 CDN research opportunities

Based on the above SWOT analysis, we observe the following research opportunities for future CDN developments.

- *ASP and CDN synergy*: The distribution and delivery of content using a CDN and the ASP of databases (Storage Service Provision) come very close. Anyway, with Internet becoming a large archive, indexing and tagging data becomes important. This relates to data management issues (MPEG7, datawarehousing solutions, etc.). Creating synergy may bust developments.
- *GRID and CDN synergy*: The GRID infrastructure is emerging and peer-to-peer computing is reviving, if not a hype. Anyway, the areas of distributed operating systems and parallel computing on the one hand (from which GRID comes) and middleware platforms on the other hand (from which CDN comes) come closer. Integrating their strong points may create new opportunities.
- *Broadcasting*: Video is 'hot'. Traditionally, CDNs focus on the delivery of streamed data (video) in particular. That is why satellite and broadcasting companies show up in this CDN area. Broadcasting companies more and more deliver their TV programmes via Internet as well (video on demand, e.g., www.omroep.nl, www.bbc.co.uk, or www.cnn.com). Noting that telecom companies and content creators are integrating (e.g., AOL and Time Warner, www.aoltimewarner.com; or Telefonica and Endemol Entertainment, www.endemol.com), this is an interesting area to create win-wins.
- *Personalisation and localisation*: Mobility is a trend, no discussion about that. Traditional CDNs mainly focus on the *content*: organising and delivering it. With mobility showing up, not only the content, but also the *context* becomes important. Content adaptation for mobile and wireless devices, adapting content to personal preferences and location-based services show up. This is a major research area for CDNs.
- *Globalisation*: From a business point of view, CDNs mean globalisation. The Internet bridges distance while guaranteeing what is delivered (Service Level Agreements) and how it is delivered (Quality of Services). Interesting research issues include the authentication and authorisation of users; accounting, payment and billing issues; and digital rights management.

Appendix A - CDN glossary of terms

This section consists of the definitions of a number of terms used to refer to roles, participants, and objects involved in CDNs.

These terms are mainly obtained from IETF drafts:

- http://www.ietf.org/internet-drafts/draft-day-cdnp-model-04.txt
- http://www.ietf.org/internet-drafts/draft-tomlinson-epsfw-00.txt
- http://events.stardust.com/cdn/documents/CDN_whitepaper_v3CH.PDF

AAA	Accounting, Authorisation, and Authentication
accounting	Measurement and recording of Distribution and Delivery activities, especially when the information recorded is ultimately used as a basis for the subsequent transfer of money, goods, or obligations. Accounting can be defined as the functionality concerned with linking the usage of services, content and resources to an identified, authenticated and authorised person responsible for the usage, and a context in which these are used (e.g. time-of-day or physical location). Accounting includes the actual data logging, as well as management functionality to make logging possible. Accounting gathers collected resource consumption data for the purpose of capacity and trend analysis, auditing and billing. The information is ordered into user and session records and stored for later use for any of these three purposes.
accounting system	A collection of Network Elements that supports Accounting for a single CDN.
aggregator	A distributed or multi-network CDN service provider that places its CDN services in the PoPs of as many facilities-based providers as possible, creating an internetwork of content servers that cross multiple ISP backbones.
authorisation	Authorisation is the act of determining whether a particular right can be granted to the presenter of a particular credential. This particular right can be, for example, an access to a resource.
authentication	Authentication is the verification of a claimed identity, in the form of a pre-existing label from a mutually known namespace, as the originator of the message (message authentication) or as the channel end point.

authoritative request-routing system	The Request-Routing System that is the correct/final authority for a particular item of Content.
avatar	A caching proxy located at the network access point of the user agent, delegated the authority to operate on behalf of, and typically working in close co-operation with a group of user agents.
cache	A program's local store of response messages and the subsystem that controls its message storage, retrieval, and deletion. A cache stores cacheable responses in order to reduce the response time and network bandwidth consumption on future, equivalent requests. Any client or server may include a cache, though a cache cannot be used by a server that is acting as a tunnel.
caching proxy	A proxy with a cache, acting as a server to clients, and a client to servers. A caching proxy is situated near the clients to improve Internet performance problems related to congestion. Caching proxies cache objects based on client demand, so they may not help the distribution of load of a given origin server. Caching proxies are often referred to as "proxy caches" or simply "caches". The term "proxy" is also frequently misused when referring to caching proxies.
CDN	"Content Delivery Network" or "Content Distribution Network". A collection of Network Elements arranged for more effective delivery of Content to Clients. Typically a CDN consists of a Request-Routing System, Surrogates, a Distribution System, and an Accounting System. [Editor note: we need to clarify what is the "minimum" CDN. One possibility is that a collection of Surrogates is the minimum. Another possibility is that Surrogates and a Request-Routing System is the minimum.].
CDN peering	CDN peering allows multiple CDN resources to combined so as to provide larger scale and/or reach to participants than any single CDN could achieve by itself.
CDN peering gateway	The interconnection of CDNs occurs through network elements called CDN Peering Gateways (CPGs).
client	The origin of a Request and the destination of the corresponding delivered Content.
content	Digital data resources. One important form of Content with additional constraints on Distribution and Delivery is Continuous Media.

content-delivery network	See: CDN.
content-distribution network	See: CDN.
content peering	A function by which operators of two different CDNs can share content, maintain consistent content delivery levels across their infrastructures, and bill one another for services rendered.
content provider	Provider of original content
content signal	A message delivered through a Distribution System that specifies information about an item of Content. For example, a Content Signal can indicate that the Origin has a new version of some piece of Content.
content server	The server on which content is delivered from. It may be an origin server, replica server, surrogate, or parent proxy.
continuous media	Content where there is a timing relationship between source and sink; that is, the sink must reproduce the timing relationship that existed at the source. The most common examples of Continuous Media are audio and motion video. Continuous Media can be real-time (interactive), where there is a "tight" timing relationship between source and sink, or streaming (playback), where the relationship is less strict.
CPG	See: CDN-Peering Gateway
delivery	The activity of presenting a Publisher's Content for consumption by a CLIENT. Contrast with Distribution and Request-Routing.
distribution	The activity of moving a Publisher's Content from its Origin to one or more Surrogates. Distribution can happen either in anticipation of a Surrogate receiving a Request (pre-positioning) or in response to a Surrogate receiving a Request (fetching on demand). Contrast with Delivery and Request-Routing.
distribution system	A collection of Network Elements that support Distribution for a single CDN. The Distribution System also propagates Content Signals.
edge services	The delivery of content from a surrogate to an end user across a single last-mile hop. Requires caching at the edge

	of a service provider's network.
inbound / outbound	Inbound and outbound refer to the request and response paths for messages: "inbound" means "travelling toward the origin server", and "outbound" means "travelling toward the user agent".
interception proxy (a.k.a. "transparent proxy" or "transparent cache")	The term "transparent proxy" has been used within the caching community to describe proxies used with zero configuration within the user agent. Such use is somewhat transparent to user agents. Due to discrepancies (see definition of "proxy" above), and objections to the use of the word "transparent", we introduce the term "interception proxy" to describe proxies that receive redirected traffic flows from network elements performing traffic interception. Interception proxies receive inbound traffic flows through the process of traffic redirection (such proxies are deployed by network administrators to facilitate or require the use of appropriate services offered by the proxy). Problems associated with the deployment of interception proxies are described in the companion document "Known HTTP Proxy/Caching Problems"[19]. The use of interception proxies requires zero configuration of the user agent, which act as though communicating directly with an origin server.
load balancing	Intelligent functions in IP networks – either bundled into routers or run as separate appliances – that determine which servers are least loaded and balance requests among server clusters accordingly.
mapping	See "Request-Routing".
multicast	Multicast is communication between a single sender and multiple receivers on a network. Within the streaming context this means that only one media stream has to be set up at the server side that can be viewed or listened to by a potentially unlimited number of clients making Multicast content delivery an extremely bandwidth efficient method.
network element	A device or system that affects the processing of network messages.
non-transparent proxy	See "Proxy".
origin	The point at which Content first enters a Distribution System. The Origin for any item of Content is the server or set of servers at the "core" of the distribution, holding the "master" or "authoritative" copy of that Content.

origin server	The server on which a given resource resides or is to be created. The origin server is the one that is refreshed by the content provider. The origin server communicates updates to the many distributed surrogate servers, often via IP Multicast technology.
peering	See: CDN Peering, Content Peering
PoP	Points of Presence. An IP network service provider's central office, which connects an end user, such as a customer, to the Internet over a last-mile access link.
proxy	An intermediary program, which acts as both a server and a client for the purpose of making requests on behalf of other clients. Requests are serviced internally or by passing them on, with possible translation, to other servers. A proxy MUST implement both the client and server requirements of this specification. A "transparent proxy" is a proxy that does not modify the request or response beyond what is required for proxy authentication and identification. A "non-transparent proxy" is a proxy that modifies the request or response in order to provide some added service to the user agent, such as group annotation services, media type transformation, protocol reduction, or anonymity filtering. Except where either transparent or non-transparent behaviour is explicitly stated, the HTTP proxy requirements apply to both types of proxies.
proxylet	Executable code modules that have a procedural interface to the caching proxy's core services. Proxylets may be either downloaded from content servers or user agents, or they may be preinstalled on the caching proxy.
proxylet library	A language binding dependent API on the service environment caching proxy platform with which proxylets link. This provides a standardised and strictly controlled interface to the service execution environment on the proxy.
publisher	The party that ultimately controls the content and its distribution.
reachable surrogates	The collection of Surrogates that can be contacted via a particular Distribution System or Request-Routing System.
redirector	A tool that enables content providers to redirect requests to their own DNS servers to the DNS server of their CDN provider. Also, a lookup service that uses metrics such as user proximity and server load to determine which

	surrogate delivers content to the requesting user.
remote callout server	A co-operating server, which runs services as the result of network protocol messaging interactions to/from a service environment caching proxy.
request	A message identifying a particular item of Content to be delivered.
request-routing	The activity of steering or directing a Request from a Client to a suitable Surrogate, which is able to service a Client request.
request-routing system	A collection of Network Elements that support Request-Routing for a single CDN. A Request-Routing Peering System represents the request-routing function of the CDN peering system. It is responsible for routing client requests to an appropriate peered CDN for the delivery of content.
reverse proxy caching	Use of surrogates or cache servers to extend a publisher's origin point to distributed points of presence (PoPs) that are physically closer to end-users.
RTP	RTP (Real Time Protocol) [RFC-1889] is the protocol that runs on top of UDP used for transport of real-time data, including audio and video. RTP consists of a data- and a control part called RTCP. The data part of RTP is a thin protocol providing support for applications with real-time properties such as continuous media (e.g., audio and video), including timing reconstruction, loss detection, security and content identification (the 'payload').
RTSP	RTSP (Real Time Streaming Protocol) [RFC-2326] is a communications protocol for control of the delivery of real-time media. It defines the connection between streaming media client and server software, and provides a standard way for clients and servers from a number of vendors to stream multimedia content. It can be seen as the "Internet VCR remote control protocol". RTSP is an application-level protocol designed to work with lower-level protocols like RTP to provide a complete streaming service over Internet.
rule module	A collection of message pattern descriptions and consequent actions that are used to match incoming protocol messages and process their contents if a match occurs.
service	Work performed (or offered) by a server. This may mean

simply serving simple requests for data to be sent or stored (as with file servers, gopher or http servers, e-mail servers, finger servers, SQL servers, etc.); or it may be more complex work, such as that of IRC servers, print servers, X Windows servers, or process servers.

service environment caching proxy

A caching proxy which has functionality beyond the basic short-circuit request fulfilment, making it capable of executing extensible (programmable) services, including network transactions with other hosts for purposes of modifying message traffic.

service execution environment

The environment on the caching proxy that allows new services to be defined and executed.

Surrogate

A gateway co-located with an origin server, or at a different point in the network, delegated the authority to operate on behalf of, and typically working in close co-operation with, one or more origin servers. Responses are typically delivered from an internal cache.

Or: A delivery server, other than the Origin. Receives a mapped Request and delivers the corresponding Content.

Surrogates may derive cache entries from the origin server or from another of the origin server's delegates. In some cases a surrogate may tunnel such requests.

Where close co-operation between origin servers and surrogates exists, this enables modifications of some protocol requirements, including the Cache-Control directives in [4]. Such modifications have yet to be fully specified.

Devices commonly known as "reverse proxies" and "(origin) server accelerators" are both more properly defined as surrogates.

Syndication

The supply of material for reuse and integration with other material, often through a paid service subscription. The most common example of syndication is in newspapers, where such content as wire-service news, comics, columns, horoscopes, and crossword puzzles are usually syndicated content. Newspapers receive the content from the content providers, reformat it as required, integrate it with other copy, print it, and publish it. For many years mainly a feature of print media, today content syndication is the way a great deal of information is disseminated across the Web.

Syndicator	Content assembler.
Transparent proxy	See "proxy".
Trigger	A rule that matches a network protocol message, causing a proxylet to execute or other action to occur on the matched message segment.
Unicast	Unicast is communication between a single sender and a single receiver on a network. Within the streaming context this means that for every client requesting a certain audio and/or video asset a new media stream has to be set up between server and client making Unicast content delivery extremely bandwidth intensive.
User agent	The client which initiates a request. These are often browsers, editors, spiders (Web-traversing robots), or other end user tools.

Appendix B - Overview of CDN organisations

See also http://www.webreference.com/internet/software/site_management/cdns.html.

Table 3: CDN organisation types and their products.

Organisation	Web site	Organisation Type	Product Name	Product Type
Activate	www.activate.com	streaming-media caching		
Adero	www.adero.com	CDN service provider	GlobalWise Network	
Aerocast	www.aerocast.com	broadband streaming video distribution		
Akamai	www.akamai.com	CDN service provider (incl. streaming)		
AppStream	www.appstream.com	CDN service provider		
AT&T	www.att.com	CDN service provider	Intelligent Content Distribution Service	
Axient	www.axient.com	CDN service provider		
BackStream	www.backstream.com	CDN service provider		
CacheFlow	www.cacheflow.com	caching hardware	clQ content delivery architecture:	
			– Cacheflow	edge caching
			– clQ Director	content management
			– clQ Sever Accelerator	server-side caching
			– clQ Streaming Services	streaming-media solutions
CacheWare	www.cacheware.com	CDN software		
Caspian Networks	www.caspian.com			
Cereva Networks	www.cereva.com			
Cidera	www.cidera.com			
Cisco	www.cisco.com	caching hardware, network	Content Distribution Manager	
			Content Engine	
			Content Router	
			CSS Switch	load balancing
Clearway	www.clearway.com			
ClickArray Networks	www.clickarray.com			
Digital Fountain	www.digitalfountain.com			
Digital Island (note 1)	www.digitalisland.com	CDN service provider	Custom Host	hosting
		(incl. authentication, streaming)	Footprint Streaming Solutions:	
			– Footprint Live	CDN service for live broadcasting events

			– Footprint On-Demand	streaming on demand
			– Footprint Media Services	syndication, sponsorship, DRM
Digital Pipe	www.digitalpipe.net			
Dynamai		CDN service provider (satellite-based)		
Edgix	www.edgix.com			
e-Media	www.e-media.com	streaming-media solutions		
Enron	www.enron.net	streaming-media solutions		
epicRealm	www.epicrealm.com	CDN service provider		
eScene Networks	www.escene.com		Content Delivery	suite of applications
			Streamline	streaming-media solutions
Exodus	www.exodus.net		Datavault Service	backup & storage
			Managed Services	systems management Web site
			Professional Services	consultancy
			Security Service Pack	security
			Streaming Media Monitoring Service	streaming-media monitoring
F5 Networks	www.f5.com			
Genuity	www.genuity.com	CDN service provider		
Globix	www.globix.com	CDN service provider (incl. streaming)	EarthCache CDN	
HTRC Group	www.htrcgroup.com	? Market Analysts		
iBEAM	www.ibeam.com	CDN service provider (incl. streaming)		
iKnowledge	www.iknowledgeinc.com			
Imminet	See: Lucent Technologies			
InfoLibria	www.infolibria.com	caching hardware		
Inktomi	www.inktomi.com	caching software	Traffic Server	network caching platform
Intel	www.intel.com	CDN service provider (incl. streaming)		
Into	www.intonetworks.com	CDN service provider (incl. streaming)		
iSyndicate	www.isyndicate.com	syndicator		
Jupiter Research	www.jup.com			
Keynote Systems	www.keynote.com			
Kinecta	www.kinecta.com		Kinecta Syndication Server	
			Kinecta Content Directory	

			Kinecta Content Metrics	
Lucent Technologies	www.lucent.com	network	Imminet:	
			– Imminet WebCache	caching
			– Imminet WebDirector	load balancing
			– Imminet WebStream	streaming-media service
			– Imminet WebDNS	redirection service
Madge.web	www.madgeweb.com	CDN service provider		
Microspace Communications Corp.	www.microspace.com			
Minerva Networks	www.minervanetworks.com	IP television		
Mirror Image	www.mirror-image.com	CDN service provider	instaDelivery Internet services:	
		(incl. streaming)	– instaContent	content distribution
			– instaSpeed	caching
			– instaStream	streaming-media service
Net 36	www.net-36.com			
NetActive	www.netactive.com			
NetworkAppliance	www.netapp.com	caching hardware		
Nextpage	www.nextpage.com		NXT 3	CDN service
NLANR	www.squid-cache.org	caching software		
Nortel	www.nortel.com	network	Alteon:	
			– Alteon Content-Intelligent Web Switches	load balancing
			– Alteon Integrated Service Director	traffic offloading service
			– Alteon Personal Content Director	redirection service
			Shasta 5000 Broadband Service Node	broadband server
			Shasta Personal Content Portal	personalisation services
Novell	www.novell.com	caching software		
Orblynx	www.orblynx.com			
Predictive Networks	www.predictivenetworks.com			
Reliacast	www.reliacast.com			
Sandpiper Networks	See: Digital Island			
SkyStream Networks	www.skystream.com			
SolidSpeed Networks	www.solidspeed.com	CDN service provider		
Sonicity	www.sonicity.com	CDN service provider		
SpectraRep	www.spectrarep.com			
Speedera	www.speedera.com	CDN service	Speedera Content	CDN service

		provider	Delivery Network	
			Speedera Download Service	CDN service
			Speedera Live Streaming	live streaming
			Speedera Failover	fall-back service
			Speedera SSL Service	CDN service for e-business
			Speedera Streaming Service	streaming-media service
			Speedera Traffic Balancer	load balancing
			Speedeye	content management
Talarian	www.talarian.com			
Tanto	www.tanto.de	syndicator		
Tier 1 Research	www.Tier1Research.com			
TV Files	www.tvfiles.com			
UCSB	www.cs.ucsb.edu	research		
Unitech Networks	www.unitechnetworks.com	caching	Netplicator	edge server
Volera	www.volera.com		Volera Excelerator	caching
			Content Exchange	caching
WebEver	www.webever.com	manufacturer		
XOsoft	www.xosoft.com	manufacturer		
Yahoo!	www.yahoo.com			

Legenda:

Organisation Type:	CDN	CDN service provider
	caching hardware	vendor of caching hardware for CDNs
	caching software	vendor of caching software for CDNs
	manufacturer	CDN product manufacturer
	network	vendor of network infrastructure for CDNs
	research	Research institute
	syndicator	

Organisation types state of the art deliverable:	content provider
	syndicator
	distributor
	connectivity provider
	server-capacity provider
	product

Index

3
3GPP 72

A
AAA 52, 77
  standardisation 53
Accounting 53, 77
Adero 5
Aggregator 77
Akamai 5
Authentication 52, 77
Authorisation 53, 77
Avatar 78

B
BCDF 8

C
Cache 78
Cache Digests 28
Cached Delivery 33
CacheWare 5
Caching proxies 26
CARP 30
CC/PP 41, 44
CDI 7
CDN 1, 78
  AAA 54
  architectures 24
  business models 12, 14
  business roles 15
  components 24
  functionality 14
  future scenarios 21
  market forecasts 6
  opportunities 74
  peering 19, 78
  peering gateway 19, 78
  product manufacturers 18
  protocols 24
  research opportunities 75
  sevice providers 4
  standardisation 7
  strength 74
  threads 75
  weakness 74
Cidera 5
Clearway 5
Client 78
ConNeg 39
Content 78
Content adaptation 22, 44
  techniques 45
Content Alliance 8
Content Bridge 8
Content consumer 18
Content distribution service provider 16
Content negotiation 36
Content Negotiation
  transparent 40
Content provider 16
CPG 79

D
DIAMETER 41
Digital Island 5
Digital Rights Management 59
DRM 59

E
EpicRealm 5

G
GARA 68
Globe 64
  Distribution Network 66
  status 67
  system 64
GlobeDoc 65
Globus 67
  Data Management 69
  GRAM 69
  Replica Management 69
  Resource Management 68
Gnutella 66
Grid 67
  security 68

H
HTCP 30

I
IBeam 6
ICAP 8, 31, 44, 46
  architecture 46
  benefits 46
  forum 8
  limitations 50
  opportunities 49
ICP 28
Insertion 21
  ad banners 21
  regional data 21
Internet 12
  business models 13
  trends 12
Internet Engineering Task Force (IETF) 7
ISMA 8
ISP 18

J
JAIN 72

K
Kerberos 68

L
Load balancing 80

M
Media gateways 50
MIDCOM 7, 50
Middle boxes 50

MIDTAX 50
MIME 36
Mirror Image 6
MMUSIC 42
Monetising content distribution services 17
Multicast 80
Multicast Split 34

N
Napster 66

O
OPES 7, 22, 31, 44
  architecture 31
Origin 80
Origin server 81
OSA 72

P
PAM 41
Parlay 70
  Business View 70
  Technology View 71
Pass-Through Delivery 35
Peering 81
PoP 81
Proxies 25
  Filtering Requests 25
  Performance 26
  Sharing Connections 25
  streaming 32
Proxy 81
Proxylet 81
  library 81
Publisher 81
Pushcache 6

Q
QoS 68

R
RADIUS 41
Redirector 81
Remote callout server 82
Replication 33
Request 82
Request-routing 78, 82
RMRG 8
RMT 7
RTP 82
RTSP 82

S
SDP 42
Server capacity provider 18
Service 83
SSL 55
Streaming 8
  challenges 32
  media adaptation 22
Surrogate 83
Syndication 13, 83
Syndicator 16, 84

T
Transcoding 50
  XML/HTML 51
Trigger 84

U
Unicast 84
Unicast Split 34
User (agent) profiles 40

V
Virus scanning 21

W
W3C 8
WEBI 7
WMF 8
WREC 7

C2

Brian Davison's Web-Caching Bibliography

Brian D. Davison
Department of Computer Science
Rutgers, The State University of New Jersey
New Brunswick, NJ 08903 USA
davison@cs.rutgers.edu

June 28, 2000

1 Introduction

This document is an attempt at building a comprehensive bibliography of web-caching publications. At present, it reflects my bias of what documents are interesting, so I'm sure there are more that can be included. Additions and corrections are welcome and should be sent to *davison@cs.rutgers.edu*. I prefer *complete* citations, in bibtex format. Links to on-line copies of documents are also appreciated.

This bibliography is one of the resources available in Brian Davison's Web-Caching pages (http://www.web-caching.com/).

You may also want to take a look at the Bibliographies on Distributed Computing, Networking and Telecommunications at http://wheat.uwaterloo.ca/bibliography/Distributed/index.html and the WEBBIB project which provides an online bibliogrpahy of distributed systems topics as they relate to the World Wide Web at http://www.cs.wpi.edu/~webbib/.

2 Citations

[AFA97] [AFAW97] [ANF97] [ALF98] [Abd98] [ASA+95] [AW96a] [ASF98] [ABCdO96] [dAMdO+96] [AAY97] [AC98a] [ACF+98a] [ACF+98b] [AW96b] [AW97] [AC98b] [BMS96] [BBM+97] [BDR97] [BC98] [BO00] [BS96] [BLFF96] [FGM+99] [Bes95b] [BCC+95a] [Bes95a] [BCC+95b] [BC95] [Bes95c] [Bes96b] [BC96] [Bes96a] [BVG96] [BH96] [BDH+94] [Bra96] [wBC95] [BCFP98] [BCF+99] [BGM98] [BD96] [CDF+98] [CFKL95] [CI97] [CZB98] [CGL+98] [CC96] [Cat92] [CP95] [CDN+96] [CY97] [CMPS95] [CLM99] [CLLM00] [Cor96] [Cor98] [CB95] [CB96] [CB97] [CB98] [CBC95] [CJ97] [Cun97] [DHS93] [Dan98] [Dav99b] [Dav99a] [Dav99c] [Den96] [DCW96] [Din96] [DP96] [DMP95] [DFKM97a] [DFKM97b] [DBCK98] [Duc99] [DMF97] [Dyk98] [DJD99] [CKR99] [CKO99] [CKZ99] [CK99] [CK00] [EWCS96] [FCAB98] [FJCL99] [FRC98] [FRCar] [Fel98] [FCD+99] [FVYI00] [FB96] [GRC97] [GCR97] [GCR98] [GC97] [Gla94] [GPB98] [Gri97] [GB97] [GPV98] [GNPV98] [GS96a] [GS96b] [GS97] [HMY97] [Hed98] [HN96] [HWMS98] [HSY98] [HK97] [HJWC98] [IKY97] [IST98] [JC98] [JDB96] [JK97] [JK98] [Kah97] [KKO98] [WMS98a] [KS98a] [KW97a] [KW97b] [KW98] [KMK99] [KR99] [KW99] [KA99] [KS98b] [KLM97] [KSW98b] [KSW98a] [LG98] [LHC+98] [LSCH98] [LWS+99] [LD99] [LAJF98] [Liu98] [LC97] [LC98] [TB97] [LOG96] [LA94] [Luo98] [Mah99] [MWE00] [MEW00] [MLB95] [MR97] [MS97] [MSC98] [Mar96] [MC98] [Mar99] [MR98] [WMFMA98] [Mel96] [MBV97] [MA98] [Mog95] [Mog96] [MDFK97a] [MDFK97b] [MJ98] [MAWM98] [Nau98] [NLN98] [Pad95] [PM96] [Par96] [Pet98] [PR94] [PK96] [Pit97] [Pit98] [RCG98] [RSGR00] [RF98b] [RF98a] [RV98] [RBR98] [RS98] [RW98] [RS99] [RID00] [RH98] [SKS98] [SSV96] [SSV97] [Sch96] [SSV98] [SSV99] [SAYZ99] [SW97] [TM00] [TG97a] [TG97b] [TVDS98] [THVK98] [Tew98] [GR96] [GR98] [THO96] [Tou98] [Tur98] [VR98] [VL98] [VdJM98] [Vix98] [VAK+98] [Wac96] [WC96] [Wan99] [WMS98b] [Wes95a] [Wes95b] [WC97b] [WC97a] [WAS+96] [Wil98] [WS97] [WM99b] [WM99a] [WVS+99] [WAB+96] [Woo96] [WWB96] [WA97] [Wor94] [YFVI00] [YS98] [ZFJ97] [ZMN+98]

References

[AAY97] Jussara Almeida, Virgílio A. F. Almeida, and David J. Yates. Measuring the behavior of a world wide web server. In *Proceedings of the Seventh Conference on High Performance Networking (HPN)*, pages 57–72. IFIP, April 1997.

[ABCdO96] Virgílio A. F. Almeida, Azer Bestavros, Mark Crovella, and Adriana de Oliveira. Characterizing Reference Locality in the WWW. In *Proceedings of the IEEE-ACM Conference on Parallel and Distributed Information Systems (PDIS'96)*, Miami Beach, FL, December 1996.

[Abd98] Ghaleb Abdulla. *Analysis and Modeling of World Wide Web Traffic.* PhD thesis, Computer Science Department, Virginia Tech., 1998.

[AC98a] Jussara Almeida and Pei Cao. Measuring Proxy Performance with the Wisconsin Proxy Benchmark. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998. Also available as Technical Report 1373, Computer Sciences Dept., Univ. of Wisconsin-Madison, April, 1998.

[AC98b] Jussara Almeida and Pei Cao. Wisconsin proxy benchmark 1.0. Available from http://www.cs.wisc.edu/~cao/wpb1.0.html, 1998.

[ACF+98a] Virgílio Augusto F. Almeida, Márcio G. Cesário, Rodrigo C. Fonseca, Jr. Wagner Meira, and Cristina D. Murta. Analyzing the behavior of a proxy server in light of regional and cultural issues. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[ACF+98b] Virgílio Augusto F. Almeida, Márcio G. Cesário, Rodrigo C. Fonseca, Jr. Wagner Meira, and Cristina D. Murta. The influence of geographical and cultural issues on the cache proxy server workload. In *Poster proceedings of the Seventh World Wide Web Conference*, Brisbane, Australia, April 1998.

[AFA97] Ghaleb Abdulla, Edward A. Fox, and Marc Abrams. Shared User Behavior on the World Wide Web. In *Proceedings of WebNet97*, Toronto, Canada, October 1997.

[AFAW97] Ghaleb Abdulla, Edward A. Fox, Marc Abrams, and Stephen Williams. WWW Proxy Traffic Characterization with Application to Caching. Technical Report TR-97-03, Computer Science Department, Virginia Tech, March 1997.

[ALF98] Ghaleb Abdulla, Binzhang Liu, and Edward A. Fox. Searching the World-Wide Web: Implications From Studying Different User Behavior. In *Proceedings of WebNet98*, Orlando, FL, November 1998.

[ANF97] Ghaleb Abdulla, A. H. Nayfeh, and Edward A. Fox. A Realistic Model of Request Arrival Rate to Caching Proxies. Submitted for publication., 1997.

[ASA+95] Marc Abrams, Charles R. Standridge, Ghaleb Abdulla, Stephen Williams, and Edward A. Fox. Caching proxies: Limitations and potentials. In *Proceedings of the Fourth International World-Wide Web Conference*, pages 119–133, Boston, MA, December 1995. Available from http://www.w3j.com/1/abrams.155/paper/155.html.

[ASF98] Manuel Afonso, Alexandre Santos, and Vasco Freitas. QoS in web caching. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[AW96a] Marc Abrams and Stephen Williams. Complementing surveying and demographics with automated network monitoring. *World Wide Web Journal*, 1(3), June 1996.

[AW96b] Martin F. Arlitt and Carey L. Williamson. Web server workload characterization: The search for invariants. In *Proceedings of the ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems*, May 1996.

[AW97] Martin F. Arlitt and Carey L. Williamson. Internet Web servers: Workload characterization and implications ftp://ftp.cs.usask.ca/pub/discus/paper.96-3.ps.Z. *IEEE/ACM Transactions on Networking*, 5:631–644, October 1997.

[BBM+97] Michael Baentsch, L. Baum, Georg Molter, S. Rothkugel, and P. Sturm. World-Wide Web caching: The application-level view of the Internet. *IEEE Communications*, June 1997.

[BC95] Azer Bestavros and Carlos Cunha. A Prefetching Protocol Using Client Speculation for the WWW. Technical Report TR-95-011, Boston University, CS Dept, Boston, MA, April 1995.

[BC96] Azer Bestavros and Carlos Cunha. Server-initiated Document Dissemination for the WWW. *IEEE Data Engineering Bulletin*, 19(3):3–11, September 1996.

[BC98] Paul Barford and Mark Crovella. Generating representative web workloads for network and server performance evaluation. In *Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '98/PERFORMANCE '98)*, pages 151–160, Madison, WI, June 1998.

[BCC+95a] Azer Bestavros, Robert Carter, Mark Crovella, Carlos Cunha, Abdelsalam Heddaya, and Suliman Mirdad. Application Level Document Caching in the Internet. Technical Report TR-95-002, Boston University, CS Dept, Boston, MA, January 1995.

[BCC+95b] Azer Bestavros, Robert Carter, Mark Crovella, Carlos Cunha, Abdelsalam Heddaya, and Sulaiman Mirdad. Application-Level Document Caching in the Internet. In *Proceedings of the Second International Workshop on Services in Distributed and Networked Environments (SDNE'95)*, Whistler, British Columbia, June 1995.

[BCF+99] Lee Breslau, Pei Cao, Li Fan, Graham Phillips, and Scott Shenker. Web Caching and Zipf-like Distributions: Evidence and Implications. In *Proceedings of INFOCOM'99*, 1999. An earlier version of this paper can be found as Technical Report 1371, Dept. of Computer Science, University of Wisconsin-Madison, April 1998.

[BCFP98] Lee Breslau, Pei Cao, Li Fan, and Graham Phillips. On the Implications of Zipf's Law for Web Caching. Technical Report 1371, Computer Science Department, University of Wisconsin-Madison, April 1998.

[BD96] Alan R. Bugos and Shuang Deng. Access traffic model inferred from WWW servers. In *Proceedings of the Fourth International Conference on Telecommunication Systems Modeling and Analysis*, March 1996.

[BDH+94] C. Mic Bowman, Peter B. Danzig, Darren R. Hardy, Udi Manber, and Michael F. Schwartz. The Harvest information discovery and access system. In *Proceedings of the Second International World Wide Web Conference*, pages 763–771, October 1994.

[BDR97] Gaurav Banga, Fred Douglis, and Michael Rabinovich. Optimistic Deltas for WWW Latency Reduction. In *Proceedings of the USENIX Technical Conference*, 1997.

[Bes95a] Azer Bestavros. Demand-based Document Dissemination for the World Wide Web. Technical Report TR-95-003, Boston University, CS Dept, Boston, MA, February 1995.

[Bes95b] Azer Bestavros. Demand-based Document Dissemination to Reduce Traffic and Balance Load in Distributed Information Systems. In *Proceedings of the Seventh IEEE Symposium on Parallel and Distributed Processing (SPDP'95)*, San Antonio, TX, October 1995.

[Bes95c] Azer Bestavros. Using Speculation to Reduce Server Load and Service Time on the WWW. In *Proceedings of CIKM'95: The Fourth ACM International Conference on Information and Knowledge Management*, Baltimore, MD, November 1995. Also available as Technical Report TR-95-006, Computer Science Department, Boston University.

[Bes96a] Azer Bestavros. Middleware Support for Data Mining and Knowledge Discovery in Large-scale Distributed Information Systems. In *SIGMOD'96 Data Mining Workshop*, Montreal, Canada, June 1996.

[Bes96b] Azer Bestavros. Speculative Data Dissemination and Service to Reduce Server Load, Network Traffic and Service Time for Distributed Information Systems. In *Proceedings of the International Conference on Data Engineering (ICDE'96)*, New Orleans, LA, March 1996.

[BGM98] Eric A. Brewer, Paul Gauthier, and Dennis McEvoy. The long-term viability of large-scale caching. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[BH96] Jean-Chrysostome Bolot and Philipp Hoschka. Performance engineering of the world wide web: Application to dimensioning and cache design. In *Proceedings of the Fifth International World Wide Web Conference*, Paris, France, May 1996.

[BLFF96] Tim Berners-Lee, Roy T. Fielding, and Henrik Frystyk Nielson. Hypertext Transfer Protocol — HTTP/1.0. RFC 1945, http://ftp.isi.edu/in-notes/rf1945.txt, May 1996.

[BMS96] Michael Baentsch, Georg Molter, and Peter Sturm. Introducing application-level replication and naming into today's web. In *Proceedings of the Fifth International World Wide Web Conference*, Paris, France, May 1996.

[BO00] Greg Barish and Katia Obraczka. World Wide Web Caching: Trends and Techniques. *IEEE Communications Magazine Internet Technology Series*, May 2000. To appear.

[Bra96] Tim Bray. Measuring the web. In *Proceedings of the Fifth International World Wide Web Conference*, Paris, France, May 1996.

[BS96] David Barnes and Neil G. Smith. An Analysis of World-Wide Web Proxy Cache Performance and its Application to the Modelling and Simulation of Network Traffic. In *Proceedings of the Fourth International Conference on Telecommunication Systems Modeling and Analysis*, March 1996.

[BVG96] Jean-Chrysostome Bolot and Andres Vega-Garcia. Performance engineering of the world wide web: Application to dimensioning and cache design. In *Proceedings of IEEE Infocom*, San Francisco, CA, March 1996.

[Cat92] V. Cate. Alex— A global filesystem. In *Proceedings of the USENIX File System Workshop*, pages 1–12, Ann Arbor, MI, May 1992.

[CB95] Mark E. Crovella and Azer Bestavros. Explaining World Wide Web Traffic Self-Similarity. Technical Report TR-95-015, Boston University, CS Dept, Boston, MA, October 1995.

[CB96] Mark Crovella and Azer Bestavros. Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes. In *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '96)*, Philadelphia, PA, May 1996.

[CB97] Mark E. Crovella and Azer Bestavros. Self-similarity in World Wide Web traffic: evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6):835–846, December 1997.

[CB98] Mark Crovella and Paul Barford. The Network Effects of Prefetching. In *Proceedings of IEEE Infocom '98*, San Francisco, CA, 1998. More detailed version available as Boston University Computer Science Department Technical Report, TR-97-002, February 1997.

[CBC95] Carlos R. Cunha, Azer Bestavros, and Mark E. Crovella. Characteristics of WWW Client-based Traces. Technical Report TR-95-010, Computer Science Department, Boston University, July 1995.

[CC96] Robert L. Carter and Mark E. Crovella. Dynamic Server Selection Using Bandwidth Probing in Wide-Area Networks. Technical Report TR-96-007, Computer Science Department, Boston University, March 1996.

[CDF+98] Ramón Cáceres, Fred Douglis, Anja Feldmann, Gideon Glass, and Michael Rabinovich. Web proxy caching: The devil is in the details. In *Workshop on Internet Server Performance (WISP'98)*, Madison, WI, June 1998.

[CDN+96] Anawat Chankhunthod, Peter B. Danzig, Chuck Neerdaels, Michael F. Schwartz, and Kurt J. Worrell. A Hierarchical Internet Object Cache. In *Proceedings of the USENIX Technical Conference*, San Diego, CA, January 1996.

[CFKL95] Pei Cao, Edward W. Felten, Anna R. Karlin, and Kai Li. A Study of Integrated Prefetching and Caching Strategies. In *Proceedings of the ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems*, May 1995.

[CGL+98] José Cardoza, Pedro F. Goncalves, Alessandro Lima, Lucianna Valadares, Cynthia Tercero, Silvio L. Meira, Ana Carolina Salgado, and Fabio Q. B. da Silva. A framework for developing information indexing, filtering, and searching tools on the internet. In *Proceedings of the Eighth Annual Conference of the Internet Society (INET'98)*, Geneva, Switzerland, July 1998.

[CI97] Pei Cao and Sandy Irani. Cost-Aware WWW Proxy Caching Algorithms. In *Proceedings of the USENIX Symposium on Internet Technoloy and Systems*, pages 193–206, December 1997.

[CJ97] Carlos R. Cunha and Carlos F. B. Jaccoud. Determining WWW User's Next Access and Its Application to Prefetching. In *Proceedings of Second IEEE Symposium on Computers and Communications (ISCC'97)*, Alexandria, Egypt, July 1997.

[CK99] Edith Cohen and Haim Kaplan. Exploiting Regularities in Web Traffic Patterns for Cache Replacement. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing.* ACM, 1999.

[CK00] Edith Cohen and Haim Kaplan. Prefetching the means for document transfer: A new approach for reducing Web latency. In *Proceedings of IEEE INFOCOM*, 2000. To appear.

[CKO99] Edith Cohen, Haim Kaplan, and Jeffrey D. Oldham. Policies for Managing TCP Connections under Persistent HTTP. In *Proceedings of the Eighth International World Wide Web Conference*, Toronto, Canada, May 1999.

[CKR99] Edith Cohen, Balachander Krishnamurthy, and Jennifer Rexford. Efficient algorithms for predicting requests to Web servers. In *IEEE INFOCOM '99*, March 1999. An extended version is available as AT&T Research Technical Report 981221-01, December 1998.

[CKZ99] Edith Cohen, Haim Kaplan, and Uri Zwick. Connection Caching. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing.* ACM, 1999.

[CLLM00] Cho-Yu Chiang, Yingjie Li, Ming T. Liu, and Mervin E. Muller. On Request Forwarding for Dynamic Web Caching Hierarchies. In *Proceedings of the 20th International Conference on Distributed Computing Systems (ICDCS'00)*, Taipei, Taiwan, April 2000.

[CLM99] Cho-Yu Chiang, Ming T. Liu, and Mervin E. Muller. Caching Neighborhood Protocol: a Foundation for Building Dynamic Caching Hierarchies with WWW Proxy Servers. In *Proceedings of the 29th International Conference on Parallel Processing (ICPP'99)*, pages 516–523, Aizu, Japan, September 1999.

[CMPS95] Abhijit K. Choudhury, Nicholas F. Maxemchuk, Sanjoy Paul, and Henning G. Schulzrinne. Copyright protection for electronic publishing over computer networks. *IEEE Network*, 9(3):12–20, May/Jun 1995.

[Cor96] Andrew Cormack. Web caching. Available from http://www.jisc.ac.uk/acn/caching.html, 1996.

[Cor98] Andrew Cormack. Experiences with installing and running and institutional cache. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[CP95] Lara D. Catledge and James E. Pitkow. Characterizing Browsing Strategies in the World Wide Web. *Computer Networks and ISDN Systems*, 26(6):1065–1073, 1995.

[Cun97] Carlos R. Cunha. *Trace Analysis and Its Applications to Performance Enhancements of Distributed Information Systems.* PhD thesis, Computer Science Department, Boston University, 1997.

[CY97] Ken-ichi Chinen and Suguru Yamaguchi. An interactive prefetching proxy server for improvement of WWW latency. In *Proceedings of the Seventh Annual Conference of the Internet Society (INET'97)*, Kuala Lumpur, June 1997.

[CZB98] Pei Cao, Jin Zhang, and Kevin Beach. Active Cache: Caching Dynamic Contents on the Web. In *Proceedings of IFIP International Conference on Distributed Systems Platforms and Open Distributed Processing (Middleware '98)*, 1998.

[dAMdO+96] J.M. de Almeida, Cristina D. Murta, Adriana de Oliveira, M. A. de Souza Mendes, and Virgílio A.F. Almeida. Performance analysis and modeling of a WWW Internet server. In *Proceedings of the Fourth International Conference on Telecommunication Systems Modeling and Analysis*, March 1996.

[Dan98] Peter B. Danzig. NetCache architecture and deployment. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[Dav99a] Brian D. Davison. A Survey of Proxy Cache Evaluation Techniques. In *Proceedings of the Fourth International WWW Caching Workshop (WCW99)*, pages 67–77, San Diego, CA, March 1999.

[Dav99b] Brian D. Davison. Simultaneous Proxy Evaluation. In *Proceedings of the Fourth International WWW Caching Workshop (WCW99)*, pages 170–178, San Diego, CA, March 1999.

[Dav99c] Brian D. Davison. Web traffic logs: An imperfect resource for evaluation. In *Proceedings of the Ninth Annual Conference of the Internet Society (INET'99)*, June 1999.

[DBCK98] Fred Douglis, Thomas Ball, Yih-Farn Chen, and Eleftherios Koutsofios. The AT&T Internet Difference Engine: Tracking and Viewing Changes on the Web. *World Wide Web*, 1(1):27–44, January 1998. Also available as AT&T Labs - Research Technical Report 97.23.1.

[DCW96] Gihan V. Dias, Graham Cope, and Ravi Wijayaratne. A smart Internet caching system. In *Proceedings of the Sixth Annual Conference of the Internet Society (INET'96)*, 1996.

[Den96] Shuang Deng. Empirical model of WWW document arivals at access links. In *Proceedings of the 1996 IEEE International Conference on Communication*, June 1996.

[DFKM97a] Fred Douglis, Anja Feldmann, Balachander Krishnamurthy, and Jeffrey C. Mogul. Rate of Change and other Metrics: a Live Study of the World Wide Web. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems (USITS '97)*, December 1997. An extended version is available as AT&T Labs - Research Technical Report 97.24.2.

[DFKM97b] Fred Douglis, Anja Feldmann, Balachander Krishnamurthy, and Jeffrey C. Mogul. Rate of Change and other Metrics: a Live Study of the World Wide Web. Technical Report 97.24.2, AT&T Labs - Research, December 1997. This is an extended version of a paper that appeared in USITS'97 [DFKM97a].

[DHS93] Peter B. Danzig, Richard S. Hall, and Michael F. Schwartz. A case for caching file objects inside internetworks. *Computer and Communications Reviews*, 23(4):239–248, October 1993.

[Din96] Adam Dingle. Cache consistency in the HTTP 1.1 draft. In *ICM Workshop on Web Caching*, Warsaw, Poland, 1996.

[DJD99] Sandra G. Dykes, Clinton. L. Jeffery, , and Samir Das. Taxonomy and Design Analysis for Distributed Web Caching. In *Proceedings of the 32nd Hawaii International Conference on System Sciences (HICSS '99)*, January to appear, 1999. Distributed Caching and Replication Minitrack.

[DMF97] Bradley M. Duska, David Marwood, and Michael J. Feely. The Measured Access Characteristics of World-Wide-Web Client Proxy Caches. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems (USITS '97)*, December 1997.

[DMP95] Chris Dodge, Beate Marx, and Hans Pfeiffenberger. Web cataloging through cache exploitation and steps toward consistency maintenance. *Computer Networks and ISDN Systems*, 27(6):1003–1008, April 1995.

[DP96] Adam Dingle and Tomas Partl. Web cache coherence. *Computer Networks and ISDN Systems*, 28(7-11):907–920, May 1996.

[Duc99] Dan Duchamp. Prefetching hyperlinks. In *Proceedings of the Second USENIX Symposium on Internet Technologies and Systems (USITS '99)*, Boulder, CO, October 1999.

[Dyk98] Sandra G. Dykes. Cooperative Web Caching Using Server-Directed Proxy Sharing. Technical Report CS-98-01, Division of Computer Science, University of Texas at San Antonio, San Antonio, TX, April 1998.

[EWCS96] Yasuhiro Endo, Zheng Wang, J. Bradley Chen, and Margo Seltzer. Using Latency to Evaluate Interactive System Performance. In *Proceedings of the Symposium on Operating System Design and Implementation*, Seattle, WA, October 1996.

[FB96] Armando Fox and Eric A. Brewer. Reducing WWW Latency and Bandwidth Requirements by Real-Time Distillation. In *Proceedings of the Fifth International World Wide Web Conference*, Paris, France, May 1996.

[FCAB98] Li Fan, Pei Cao, Jussara Almeida, and Andrei Z. Broder. Summary Cache: A Scalable Wide-Area Web Cache Sharing Protocol. In *Proceedings of ACM SIGCOMM*, September 1998.

[FCD+99] Anja Feldmann, Ramón Cáceres, Fred Douglis, Gideon Glass, and Michael Rabinovich. Performance of Web Proxy Caching in Heterogeneous Bandwidth Environments http://www.research.att.com/~douglis/papers/. In *Proceedings of IEEE INFOCOM*, pages 106–116, March 1999.

[Fel98] Anja Feldmann. Continuous online extraction of HTTP traces from packet traces. In *World Wide Web Consortium Workshop on Web Characterization*, Cambridge, MA, November 1998. Position paper.

[FGM+99] Roy T. Fielding, Jim Gettys, Jeffrey C. Mogul, Henrik Frystyk, L. Masinter, P. Leach, and Tim Berners-Lee. Hypertext Transfer Protocol — HTTP/1.1. RFC 2616, http://ftp.isi.edu/in-notes/rfc2616.txt, June 1999.

[FJCL99] Li Fan, Quinn Jacobson, Pei Cao, and Wei Lin. Web Prefetching Between Low-Bandwidth Clients and Proxies: Potential and Performance. In *Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '99)*, Atlanta, GA, May 1999.

[FRC98] Anja Feldmann, Jennifer Rexford, and Ramon Caceres. Reducing overhead in flow-switched networks: An empirical study of web traffic. In *Proceedings of IEEE INFOCOM*, pages 1205–1213, April 1998.

[FRCar] Anja Feldmann, Jennifer Rexford, and Ramon Caceres. Efficient Policies for Carrying Web Traffic over Flow-Switched Networks. *IEEE/ACM Transactions on Networking*, To appear.

[FVYI00] Daniela Florescu, Patrick Valduriez, Khaled Yagoub, and Valerie Issarny. Caching strategies for data-intensive web sites. In *Proceedings of the Ninth International World Wide Web Conference*, Amsterdam, May 2000.

[GB97] Steven D. Gribble and Eric A. Brewer. System Design Issues for Internet Middleware Services: Deductions from a Large Client Trace. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems (USITS '97)*, December 1997.

[GC97] Gideon Glass and Pei Cao. Adaptive Page Replacement Based on Memory Reference Behavior. In *Proceedings of the ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems*, pages 115–126, Seattle, WA, June 1997. ACM Press.

[GCR97] Syam Gadde, Jeff Chase, and Michael Rabinovich. Directory Structures for Scalable Internet Caches. Technical Report CS-1997-18, Department of Computer Science, Duke University, November 1997.

[GCR98] Syam Gadde, Jeff Chase, and Michael Rabinovich. A taste of crispy squid. In *Workshop on Internet Server Performance (WISP'98)*, Madison, WI, June 1998.

[Gla94] Steven Glassman. A Caching Relay for the World Wide Web. In *Proceedings of the First International Conference on the WWW*, 1994. Also published in "Computer Networks and ISDN Systems", volume 27 (1994), Number 2, November 1994.

[GNPV98] Christian Grimm, Marc Neitzner, Helmut Pralle, and Jens-S. Vöckler. Request routing in cache meshes. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[GPB98] Arthur Goldberg, Ilya Pevzner, and Robert Buff. Characteristics of Internet and Intranet Web Proxy Traces. In *Computer Measurement Group Conference CMG98*, Anaheim, CA, December 1998.

[GPV98] Christian Grimm, Helmut Pralle, and Jens-S. Vöckler. Load and traffic balancing in large scale cache meshes. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[GR96] David G.Thaler and Chinya V. Ravishankar. A Name-Based Mapping Scheme for Rendezvous. Technical Report CSE-TR-316-96, Dept. of Electrical Engineering and Computer Science, University of Michigan, June 1996.

[GR98] David G.Thaler and Chinya V. Ravishankar. Using name-based mappings to increase hit rates. *IEEE/ACM Transactions on Networking*, 6(1):1–14, February 1998.

[GRC97] Syam Gadde, Michael Rabinovich, and Jeff Chase. Reduce, Reuse, Recycle: An Approach to Building Large Internet Caches. In *The Sixth Workshop on Hot Topics in Operating Systems (HotOS-VI)*, pages 93–98, May 1997.

[Gri97] Steven D. Gribble. UC Berkely home IP HTTP traces. Available at: http://www.acm.org/sigcomm/ITA/, July 1997.

[GS96a] James Gwertzman and Margo Seltzer. The Case for Geographical Pushcaching. In *Proceedings of the Workshop on Hot Operating Systems*, 1996.

[GS96b] James Gwertzman and Margo Seltzer. World-Wide Web Cache Consistency. In *Proceedings of the USENIX Technical Conference*, San Diego, CA, January 1996.

[GS97] James Gwertzman and Margo Seltzer. An analysis of geographical push-caching. Available at http://www.eecs.harvard.edu/vino/web/server.cache/icdcs.ps, 1997.

[Hed98] Abdelsalam Heddaya. Dynacache: Weaving caching into the internet. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[HJWC98] Bradley Huffaker, Jaeyeon Jung, Duane Wessels, and K. Claffy. Visualization of the growth and topology of the nlanr caching hierarchy. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[HK97] Saied Hosseini-Khayat. *Investigation of Generalized Caching.* PhD thesis, Dept. of Electrical Engineering, Washington University, St. Louis, MO, August 1997.

[HMY97] Abdelsalam Heddayn, Sulaiman Mirdad, and David Yates. Diffusion-based caching along routing paths. In *Proceedings of the NLANR Web Cache Workshop*, June 1997.

[HN96] John Heaton and George Neisser. The WWW infrastructure at manchester and future developments. In *ICM Workshop on Web Caching*, Warsaw, Poland, 1996.

[HSY98] Vegard Holmedahl, Ben Smith, and Tao Yang. Cooperative Caching of Dynamic Content on a Distributed Web Server. In *Proceedings of the 7th IEEE International Symposium on High Performance Distributed Computing (HPDC-7)*, pages 243–250, Chicago, IL, July 1998.

[HWMS98] John H. Hine, Craig E. Wills, Anja Martel, and Joel Sommers. Combining client knowledge and resource dependencies for improved world wide web performance. In *Proceedings of the Eighth Annual Conference of the Internet Society (INET'98)*, Geneva, Switzerland, July 1998.

[IKY97] Hiroyuki Inoue, Kanchana Kanchanasut, and Suguru Yamaguchi. An adaptive WWW cache mechanism in the AI3 network. In *Proceedings of the Seventh Annual Conference of the Internet Society (INET'97)*, Kuala Lumpur, June 1997.

[IST98] Hiroyuki Inoue, Takeshi Sakamoto, and Suguru Tamaguchi. Webhint: An automatic configuration mechanism for optimizing world wide web cache system utilization. In *Proceedings of the Eighth Annual Conference of the Internet Society (INET'98)*, Geneva, Switzerland, July 1998.

[JC98] Quinn Jacobson and Pei Cao. Potential and Limits of Web Prefetching Between Low-Bandwidth Clients and Proxies. Technical Report 1372, Computer Science Dept., Univ. of Wisconsin-Madison, February 1998. An extended version is available [FJCL99].

[JDB96] Clinton L. Jeffery, Samir Das, and Gary Bernal. Proxy Sharing Proxy Servers. In *Proceedings of the IEEE etaCOM Conference*, Portland, OR, May 1996.

[JK97] Zhimei Jiang and Leonard Keinrock. Prefetching Links on the WWW. In *ICC '97*, pages 483–489, Montreal, Canada, June 1997.

[JK98] Zhimei Jiang and Leonard Kleinrock. An adaptive network prefetch scheme. *IEEE Journal on Selected Areas in Communications*, 16(3):358–368, April 1998.

[KA99] Balachander Krishnamurthy and Martin Arlitt. PRO-COW: Protocol Compliance on the Web. Technical Memorandum 990803-05-TM, AT&T Labs Research, August 1999.

[Kah97] Brewster Kahle. Preserving the Internet. *Scientific American*, 276(3):82–83, March 1997.

[KKO98] Jussi Kangasharju, Young Gap Kwon, and Antonio Ortega. Design and implementation of a soft caching proxy. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[KLM97] Thomas M. Kroeger, Darrell D. E. Long, and Jeffrey C. Mogul. Exploring the bounds of web latency reduction from caching and prefetching. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems (USITS '97)*, December 1997.

[KMK99] Balachander Krishnamurthy, Jeffrey C. Mogul, and David M. Kristol. Key differences between HTTP/1.0 and HTTP/1.1. In *Proceedings of the Eighth International World Wide Web Conference*, pages 659–673, Toronto, Canada, May 1999.

[KR99] Balachander Krishnamurthy and Jennifer Rexford. En Passant: Predicting HTTP/1.1 traffic. In *Proc. Global Internet Symposium*, December 1999.

[KS98a] Serge Krashakov and Lev Shchur. WWW caching in Russia — current state and future development. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[KS98b] P. Krishnan and Binay Sugla. Utility of co-operating web proxy caches. In *Proceedings of the Seventh International World Wide Web Conference*, Brisbane, Australia, April 1998.

[KSW98a] Michal Kurcewicz, Wojtek Sylwestrzak, and Adam Wierzbicki. A distributed WWW cache. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[KSW98b] Michal Kurcewicz, Wojtek Sylwestrzak, and Adam Wierzbicki. A filtering algorithm for proxy caches. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[KW97a] Balachander Krishnamurthy and Craig E. Wills. Study of piggyback cache validation for proxy caches in the world wide web. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems (USITS '97)*. USENIX Association, December 1997.

[KW97b] Balachander Krishnamurthy and Craig E. Wills. Piggyback cache validation for proxy caches in the world wide web. In *Proceedings of the NLANR Web Cache Workshop*, Boulder, CO, June 1997.

[KW98] Balachander Krishnamurthy and Craig E. Wills. Piggyback server invalidation for proxy cache coherency. *Computer Networks and ISDN Systems*, 30, 1998. Also in *Proceedings of the Seventh International World Wide Web Conference*, pages 185-193, Brisbane, Australia, April 1998.

[KW99] Balachander Krishnamurthy and Craig E. Wills. Analyzing factors that influence end-to-end Web performance. Technical Memorandum 991129-07-TM, AT&T Labs Research, November 1999.

[LA94] Ari Luotonen and Kevin Altis. World-wide web proxies. *Computer Networks and ISDN Systems*, 27(2):147–154, 1994.

[LAJF98] Binzhang Lui, Ghaleb Abdulla, Tommy Johnson, and Edward A. Fox. Web Response Time and Proxy Caching. In *Proceedings of WebNet98*, Orlando, FL, November 1998.

[LC97] Chengjie Liu and Pei Cao. Maintaining Strong Cache Consistency for the World-Wide Web. In *The 17th International Conference on Distributed Computing Systems*, pages 12–21, May 1997. An extended version appears in IEEE Transactions on Computers, April 1998.

[LC98] Chengjie Liu and Pei Cao. Maintaining strong cache consistency for the world-wide web. *IEEE Transactions on Computers*, April 1998. A condensed version appears in the 17th Int'l Conf. on Distr. Systems, May 1997.

[LD99] Vincenzo Liberatore and Brian D. Davison. Data Dissemination on the Web: Speculative and Unobtrusive. Technical Report 99-23, UMIACS, University of Maryland, College Park, College Park, MD, May 1999.

[LG98] Ulana Legedza and John Guttag. Using network-level support to improve cache routing. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[LHC+98] Jieun Lee, Hyojung Hwang, Youngmin Chin, Hyunchul Kim, and Kilnam Chon. Report on the costs and benefits of cache hierarchy in Korea. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[Liu98] Binzhang Liu. Characterizing Web Response Time. Master's thesis, Virginia Tech., Blacksburg, VA, April 1998.

[LOG96] A. Lopez-Ortiz and Daniel M. German. A multicollaborative push-caching HTTP protocol for the WWW. In *Proceedings of the Fifth International World Wide Web Conference*, Paris, France, May 1996. Also available as Technical Report CS-96-12, Dept. of Computer Science, Universtiy of Waterloo.

[LSCH98] Wen-Syan Li, Junho Shim, K. Selcuk Candan, and Yoshinori Hara. Webdb : A web query system and its modeling, language, and implementation. In *Proceedings of the IEEE Advances in Digital Libraries Conference (ADL'98)*, 1998.

[Luo98] Ari Luotonen. *Web Proxy Servers.* Prentice Hall, Upper Saddle River, NJ, 1998.

[LWS+99] Wen-Syan Li, Yi-Leh Wu, Junho Shim, Kyoji Hirata, Sougata Mukherjea, Divy Agrawal, Yoshinori Hara, Reiko Ito, Yutaka Kimura, Kazuyuki Shimazu, and Yukiyoshi Saito. WebDB hypermedia database system. *Institute of Electronics, Information and Communication Engineers (IEICE) Transactions on Information Systems*, 1999. To appear.

[MA98] Daniel A. Menasce and Virgilio A. F. Almeida. *Capacity Planning for Web Performance: Metrics, Models, and Methods.* Prentice Hall, Upper Saddle River, NJ, 1998.

[Mah99] Anirban Mahanti. Web Proxy Workload Characterisation and Modelling. Master's thesis, Department of Computer Science, University of Saskatchewan, September 1999.

[Mar96] Evangelos P. Markatos. Main memory caching of web documents. *Computer Networks and ISDN Systems*, 28(7-11):893–905, May 1996.

[Mar99] Evangelos P. Markatos. On Caching Search Engine Results. Technical Report 241, Institute of Computer Science, Foundation for Research & Technology - Hellas (FORTH), January 1999.

[MAWM98] Cristina Duarte Murta, Virgilio A. F. Almeida, and Jr. Wagner Meira. Analyzing performance of partitioned caches for the WWW. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[MBV97] Ingrid Melva, Hanny Bekker, and Ton Verschunren. Building a web caching system — architecture considerations. In *Proceedings of the Eighth Joint European Networking Conference (JENC8)*, Edinburgh, Scotland, May 1997. http://www.uninett.no/prosjekt/desire/jenc8/webcache-jenc8.ps.

[MC98] Evangelos P. Markatos and Catherine E. Chronaki. A top–10 approach for prefetching the web. In *Proceedings of the Eighth Annual Conference of the Internet Society (INET'98)*, Geneva, Switzerland, July 1998. Also available as ICS-FORTH Technical Report 173.

[MDFK97a] Jeffrey Mogul, Fred Douglis, Anja Feldmann, and Balachander Krishnamurthy. Potential Benefits of Delta-encoding and Data Compression for HTTP. In *Proceedings of ACM SIGCOMM*, pages 181–194, September 1997. An extended and corrected version appears as Research Report 97/4a, Digital Equipment Corporation Western Research Laboratory, December, 1997.

[MDFK97b] Jeffrey Mogul, Fred Douglis, Anja Feldmann, and Balachander Krishnamurthy. Potential Benefits of Delta-encoding and Data Compression for HTTP. Research Report 97/4a, Digital Equipment Corporation, Western Research Laboratory, December 1997.

[Mel96] Ingrid Melva. Why Internet Service Providers should integrate web caches into their networks. In *ICM Workshop on Web Caching*, Warsaw, Poland, 1996.

[MEW00] Anirban Mahanti, Derek Eager, and Carey Williamson. Temporal locality and its impact on web proxy cache performance. *Performance Evaluation*, September 2000.

[MJ98] David Mosberger and Tai Jin. httperf—A tool for measuring web server performance. In *Workshop on Internet Server Performance (WISP'98)*, Madison, WI, June 1998.

[MLB95] Radhika Malpani, Jacob Lorch, and David Berger. Making World Wide Web Caching Servers Cooperate. In *Proceedings of the Fourth International World Wide Web Conference*, Boston, MA, December 1995.

[Mog95] Jeffrey C. Mogul. The case for persistent-connection HTTP. In *Proceedings of the ACM SIGCOMM'95 Conference*, pages 299–313, October 1995.

[Mog96] Jeffrey C. Mogul. Hinted caching in the web. In *Proceedings of the Seventh ACM SIGOPS European Workshop*, Connemara, Ireland, September 1996.

[MR97] Carlos Maltzahn and Kathy J. Richardson. Performance issues of enterprise level web proxies. In *Proceedings of the ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems*, pages 13–23, Seattle, WA, June 1997. ACM Press.

[MR98] Ian Marshall and Chris Roadknight. Linking cache performance to user behaviour. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[MS97] Stephen Manley and Margo Seltzer. Web Facts and Fantasy. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems (USITS '97)*, December 1997.

[MSC98] Stephen Manley, Margo Seltzer, and Michael Courage. A self-scaling and self-configuring benchmark for web servers. In *Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '98/PERFORMANCE '98)*, pages 270–271, Madison, WI, June 1998.

[MWE00] Anirban Mahanti, Carey Williamson, and Derek Eager. Traffic analysis of a web proxy caching hierarchy. *IEEE Network*, 14(3):16–23, May/June 2000.

[Nau98] Andrey Naumenko. Some Recommendations on Building Proxy Caching Service. In *AACE WebNet 98*, Orlando, FL, November 1998.

[NLN98] Nicolas Niclausse, Zhen Liu, and Philippe Nain. A new and efficient caching policy for the world wide web. In *Workshop on Internet Server Performance (WISP'98)*, Madison, WI, June 1998.

[Pad95] Venkata N. Padmanabhan. Improving World Wide Web Latency. Technical Report UCB/CSD-95-875, UC Berkeley, May 1995.

[Par96] Tomas Partl. A comparison of WWW caching algorithm efficiency. In *ICM Workshop on Web Caching*, Warsaw, Poland, 1996.

[Pet98] Arthur Petrosyan. Development of the Armenian regional academical cache service. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[Pit97] James E. Pitkow. In search of reliable usage data on the WWW. In *Proceedings of the Sixth International World Wide Web Conference*, Santa Clara, CA, April 1997.

[Pit98] James E. Pitkow. Summary of WWW characterizations. In *Proceedings of the Seventh International World Wide Web Conference*, Brisbane, Australia, April 1998.

[PK96] James E. Pitkow and C. M. Kehoe. Emerging trends in the WWW user population. *Communications of the ACM*, 36(6):106–108, 1996.

[PM96] Venkata N. Padmanabhan and Jeffrey C. Mogul. Using predictive prefetching to improve world wide web latency. *Computer Communication Review*, 26(3):22–36, July 1996.

[PR94] James E. Pitkow and Mimi Recker. A Simple Yet Robust Caching Algorithm Based on Dynamic Access Patterns. In *Electronic Proceedings of the Second World Wide Web Conference '94: Mosaic and the Web*, 1994.

[RBR98] Pablo Rodriguez, Ernst W. Biersack, and Keith W. Ross. Improving the WWW: Caching or multicast. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[RCG98] Michael Rabinovich, Jeff Chase, and Syan Gadde. Not all hits are created equal: Cooperative proxy caching over a wide area network. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[RF98a] Mike Reddy and Graham P. Fletcher. An Adaptive Mechanishm for Web Browser Cache Management. *IEEE Internet Computing*, 2(1), January 1998.

[RF98b] Mike Reddy and Graham P. Fletcher. Intelligent web caching using document life histories: A comparison with existing cache management techniques. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[RH98] Mark Russell and Tim Hopkins. CFTP — A caching FTP server. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[RID00] Daniela Rosu, Arun Iyengar, and Daniel Dias. Hint-based Acceleration of Web Proxy Caches. In *Proceedings of the 19th IEEE International Performance, Computing, and Communications Conference (IPCCC 2000)*, Phoenix, AZ, February 2000. To appear.

[RS98] Alex Rousskov and Valery Soloviev. On performance of caching proxies. In *Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '98/PERFORMANCE '98)*, pages 272–273, Madison, WI, June 1998.

[RS99] Alex Rousskov and Valery Soloviev. A performance study of the squid proxy on http/1.0. *World Wide Web*, 2(1-2):47–67, January 1999.

[RSGR00] Mohammad S. Raunak, Prashant Shenoy, Pawan Goyal, and Krithi Ramamritham. Implications of proxy caching for provisioning networks and servers. In *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 2000)*, pages 66–77, Santa Clara, CA, June 2000.

[RV98] Luigi Rizzo and Lorenzo Vicisano. Replacement policies for a proxy cache. Research Note RN/98/13, Department of Computer Science, University College London, 1998.

[RW98] Alex Rousskov and Duane Wessels. Cache digests. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[SAYZ99] Ben Smith, Anurag Acharya, Tao Yang, and Huican Zhu. Exploiting Result Equivalence in Caching Dynamic Web Content. In *Proceedings of Second USENIX Symposium on Internet Technologies and Systems*, pages 209–220, October 1999.

[Sch96] Henning G. Schulzrinne. World wide web: Whence, whither, what next? *IEEE Network*, 10(2):10–17, Mar/Apr 1996.

[SKS98] Stuart Schechter, Murali Krishnan, and Michael D. Smith. Using path profiles to predict HTTP requests. In *Proceedings of the Seventh International World Wide Web Conference*, Brisbane, Australia, April 1998.

[SSV96] Peter Scheuermann, Junho Shim, and Radek Vingralek. WATCHMAN: A Data Warehouse Intelligent Cache Manager. In *Proceedings of the 22nd International Conference on Very Large Databases*, 1996.

[SSV97] Peter Scheuermann, Junho Shim, and Radek Vingralek. A case for delay-conscious caching of web documents. In *Proceedings of the Sixth International World Wide Web Conference*, Santa Clara, CA, April 1997.

[SSV98] Junho Shim, Peter Scheuermann, and Radek Vingralek. A Unified Algorithm for Cache Raplacement and Consistency in Web Proxy Servers. In *Proceedings of the Workshop on the Web and Data Bases (WebDB98)*, 1998. Extended version will appear in Lecture Notes in Computer Science, Springer Verlag, 1998.

[SSV99] Junho Shim, Peter Scheuermann, and Radek Vingralek. Proxy Cache Design: Algorithms, Implementation and Performance. *IEEE Transactions on Knowledge and Data Engineering*, 1999. To appear.

[SW97] Joel Sommers and Craig E. Wills. Prefetching on the web using client and server profiles. Technical Report WPI-CS-TR-97-2, Computer Science Department, Worcester Polytechnic Institute, June 1997.

[TB97] Sau Loon Tong and Vaduvur Bharghavan. Alleviating the Latency and Bandwidth Problems in WWW Browsing. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems (USITS '97)*, December 1997.

[Tew98] Russell Baird Tewksbury. Is the internet heading for a cache crunch? In *Proceedings of the Eighth Annual Conference of the Internet Society (INET'98)*, Geneva, Switzerland, July 1998.

[TG97a] Linda Tauscher and Saul Greenberg. How People Revisit Web Pages: Empirical Findings and Implications for the Design of History Systems. *International Journal of Human Computer Studies*, 47(1):97–138, 1997.

[TG97b] Linda Tauscher and Saul Greenberg. Revisitation Patterns in World Wide Web Navigation. In *Proceedings of the ACM SIGCHI'97 Conference on Human Factors in Computing Systems*, Atlanta, GA, March 1997. ACM Press.

[THO96] Joe Touch, John Heidemann, and Katia Obraczka. Analysis of HTTP performance. Available at http://www.isi.edu/lsam/publications/http-perf/, August 1996.

[THVK98] Renu Tewari, Michael Hahlin, Harrick M. Vin, and Jonathan S. Kay. Beyond Hierarchies: Design Considerations for Distributed Caching on the Internet. Technical Report TR98-04, Department of Computer Science, University of Texas at Austin, February 1998.

[TM00] Wenting Tang and Matt W. Mutka. Intelligent browser initiated server pushing. In *Proceedings of the IEEE International Performance, Computing, and Communications Conference*, Phoenix, AZ USA, February 2000.

[Tou98] Joe Touch. The LSAM proxy cache: a multicast distributed virtual cache. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[Tur98] Volker Turau. What practices are being adopted on the web? *Computer*, 31(5):106–108, May 1998.

[TVDS98] Renu Tewari, Harrick M. Vin, Asit Dan, and Dinkar Sitaram. Resource-based Caching for Web Servers. In *Proceedings of the SPIE/ACM Conference on Multimedia Computing and Networking (MMCN)*, San Jose, CA, January 1998.

[VAK+98] Geoffrey M. Voelker, Eric J. Anderson, Tracy Kimbrel, Michael J. Feeley, Jeffrey S. Chase, Anna R. Karlin, and Henry M. Levy. Implementing cooperative prefetching and caching in a globally-managed memory system. In *Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '98/PERFORMANCE '98)*, pages 33–43, Madison, WI, June 1998.

[VdJM98] Ton Verschuren, André de Jong, and Ingrid Melve. Web caching meshes: Hit or miss? In *Proceedings of the Eighth Annual Conference of the Internet Society (INET'98)*, Geneva, Switzerland, July 1998.

[Vix98] Paul Vixie. Hyper Text Caching Protocol — HTCP/0.0. Internet Draft, March 1998.

[VL98] Juan Ramon Velasco and Luis Alberto Velasco Lucianez. Benefits of compression in HTTP applied to caching architectures. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[VR98] Vinod Valloppillil and Keith W. Ross. Cache Array Routing Protocol v1.0. Internet Draft, February 1998.

[WA97] Roland P. Wooster and Marc Abrams. Proxy caching that estimates page load delays. In *Proceedings of the Sixth International World Wide Web Conference*, pages 325–334, Santa Clara, CA, April 1997.

[WAB+96] Allison Woodruff, Paul M. Aoki, Eric Brewer, Paul Gauthier, and Lawrence A. Rowe. An investigation of documents from the world wide web. In *Proceedings of the Fifth International World Wide Web Conference*, Paris, France, May 1996.

[Wac96] Stuart B. Wachsberg. Efficient information access for wireless computers. Master's thesis, University of Waterloo, Waterloo, Ontario, CA, September 1996. Available from http://www.shoshin.uwaterloo.ca/papers.html.

[Wan99] Jia Wang. A survey of web caching schemes for the Internet. *ACM Computer Communication Review*, 29(5):36–46, October 1999.

[WAS+96] Stephen Williams, Marc Abrams, Charles R. Standridge, Ghaleb Abdulla, and Edward A. Fox. Removal Policies in Network Caches for World-Wide Web Documents. In *Proceedings of ACM SIGCOMM*, pages 293–305, Stanford, CA, 1996. Revised March 1997.

[wBC95] Hans werner Braun and Kimberly C. Claffy. Web Traffic Characterization: an Assessment of the Impact of Caching Documents from NCSA's Web Server. *Computer Networks and ISDN Systems*, 28(1&2):119–125, December 1995.

[WC96] Zheng Wang and Jon Crowcroft. Prefetching in World-Wide Web. In *Proceedings of IEEE Globecom*, London, December 1996.

[WC97a] Duane Wessels and Kimberly Claffy. Application of Internet Cache Protocol (ICP), version 2. RFC 2187, http://nic.ddn.mil/ftp/rfc/rfc2187.txt, September 1997.

[WC97b] Duane Wessels and Kimberly Claffy. Internet Cache Protocol (ICP), version 2. RFC 2186, http://nic.ddn.mil/ftp/rfc/rfc2186.txt, September 1997.

[Wes95a] Duane Wessels. Intelligent caching for world-wide web objects. In *Proceedings of the Fifth Annual Conference of the Internet Society (INET'95)*, Honolulu, HI, June 1995.

[Wes95b] Duane Wessels. Intelligent caching for world-wide web objects. Master's thesis, University of Colorado, Boulder, CO, January 1995. Available from http://itp-www.colorado.edu/~wessels/Proxy/wessels-thesis.ps.

[Wil98] Bert Williams. Transparent web caching solutions. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[WM99a] Craig E. Wills and Mikhail Mikhailov. Examining the cacheability of user-requested web resources. In *Proceedings of the Fourth International WWW Caching Workshop (WCW99)*, San Diego, CA, 1999.

[WM99b] Craig E. Wills and Mikhail Mikhailov. Towards a better understanding of web resources and server responses for improved caching. In *Proceedings of the Eighth International World Wide Web Conference*, Toronto, Canada, May 1999.

[WMFMA98] Jr. Wagner Meira, Erik Fonseca, Cristina Murta, and Virgílio Almeida. Evaluating and configuring WWW cache hierarchies. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[WMS98a] Bob Warfield, Terry Mueller, and Peter Sember. Monitoring the performance of a cache for broadband customers. In *Proceedings of the Eighth Annual Conference of the Internet Society (INET'98)*, Geneva, Switzerland, July 1998.

[WMS98b] Bob Warfield, Terry Mueller, and Peter Sember. Monitoring the performance of a cache for broadband customers. In *Proceedings of the Eighth Annual Conference of the Internet Society (INET'98)*, Geneva, Switzerland, July 1998.

[Woo96] Roland Peter Wooster. Optimizing response time, rather than hit rates, of WWW proxy caches. Master's thesis, Virginia Polytechnic Institute, December 1996. http://scholar.lib.vt.edu/theses/materials/public/etd-34131420119653540/etd-title.html.

[Wor94] Kurt Worrell. Invalidation in large scale network object caches. Master's thesis, University of Colorado, Boulder, CO, 1994.

[WS97] Craig E. Wills and Joel Sommers. Prefetching on the web through merger of client and server profiles. Available at http://www.cs.wpi.edu/~cew/papers/webprofile.ps.gz, June 1997.

[WVS+99] Alec Wolman, Geoffrey M. Voelker, Nitin Sharma, Neal Cardwell, Anna Karlin, and Henry M. Levy. On the scale and performance of cooperative web proxy caching. In *Proceedings of the 17th Symposium on Operating Systems Principles*, December 1999.

[WWB96] Roland O. Wooster, Stephen Williams, and Patrick Brooks. HTTPDUMP: Network HTTP packet snooper. Working paper available at http://www.cs.vt.edu/~chitra/work.html, April 1996.

[YFVI00] Khaled Yagoub, Daniela Florescu, Patrick Valduriez, and Valerie Issarny. Caching strategies for data-intensive web sites. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, Cairo, Egypt, September 2000.

[YS98] Yuping Yang and Mukesh Singhal. A tree dictionary web cache. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

[ZFJ97] Lixia Zhang, Sally Floyd, and Van Jacobson. Adaptive web caching. In *Proceedings of the NLANR Web Cache Workshop*, Boulder, CO, June 1997.

[ZMN+98] Lixia Zhang, Scott Michel, Khoi Nguyen, Sally Floyd, and Van Jacobson. Adaptive web caching: Towards a new global caching architecture. In *Proceedings of the Third International WWW Caching Workshop*, Manchester, England, June 1998.

C4

The Case for Geographical Push-Caching

James S. Gwertzman

Division of Applied Sciences
Harvard University
Cambridge, MA 02138

Margo Seltzer

Division of Applied Sciences
Harvard University
Cambridge, MA 02138

Abstract

Most wide-area caching schemes are client initiated. *Decisions on when and where to cache information are made without the benefit of the server's global knowledge of the situation. We believe that the server should play a role in making these caching decisions; and we propose* geographical push-caching *as a way of bringing the server back into the loop. The World Wide Web is an excellent example of a wide-area system that will benefit from geographical push-caching, and we present an architecture that allows a Web server to autonomously replicate HTML pages.*

1 Introduction

The World-Wide Web [1] operates for the most part as a cache-less distributed system. When two neighboring clients retrieve a document from the same server, the document is sent twice. This is inefficient, especially considering the ease with which Web browsers allow users to transfer large multimedia documents.

To combat this problem, some Web browsers have begun to add local client caches. These prevent one client from transferring the same document twice. Some networks are also beginning to add Web proxies [8, 9] that prevent two clients on the same campus network from transferring the same document twice.

The problem with both these schemes is that they are myopic. A client cache does not help a neighboring computer, and a campus proxy does not help a neighboring campus. Furthermore, these caches are usually limited in size. Disk might be cheap, but as the size of multimedia files increases it will be impossible to cache everything. As a result these caches will only be able to store the most popular items even though there is still some demand for other, less popular items.

The solution is for clients to share each other's cache space. The degree to which a file is replicated should be proportional to that file's global popularity, and clients should retrieve files from the the nearest cache to minimize network traffic. The server can satisfy both goals by deciding when and where to cache files. Furthermore, when the server decides where to cache a file, it can make this decision using its knowledge of network topology and the file's access history for even greater network bandwidth savings.

2 Motivation

Preliminary analysis of Web server access logs show that a few files on each server are responsible for most traffic from that server. Servers can therefore save a great deal of bandwidth and decrease their load by making access to those files more efficient. The slight load increase due to replicating and distributing files will be offset by this savings, and the client in turn will see increased performance.

Analysis also shows that the access pattern for each file is not uniform. File requests are often clustered geographically which implies that judicious selection of cache sites will provide excellent bandwidth and latency savings. Furthermore, we have also found strong correlations between geographical distance and Internet metrics like latency and network hops over short distances. This implies that geographical distance, which is easy to figure out using the `nets.unl.now` file at `nic.merit.edu` [10], can be used to predict network topology, which can be hard to figure out.

3 Architecture

There are two components to our proposed Web system: a modified HTTP server and a replication service. The modified server is responsible for tracking geographical access information for its files and for accepting and offering cached replicas of other files. This can be done by modifying a proxy server, such as the CERN proxy server [8], to accept files for replication using a modified POST request.

The replication service keeps track of modified HTTP servers that are willing to serve replicated files, the amount of available free space on each server, and each server's average load. The replication service

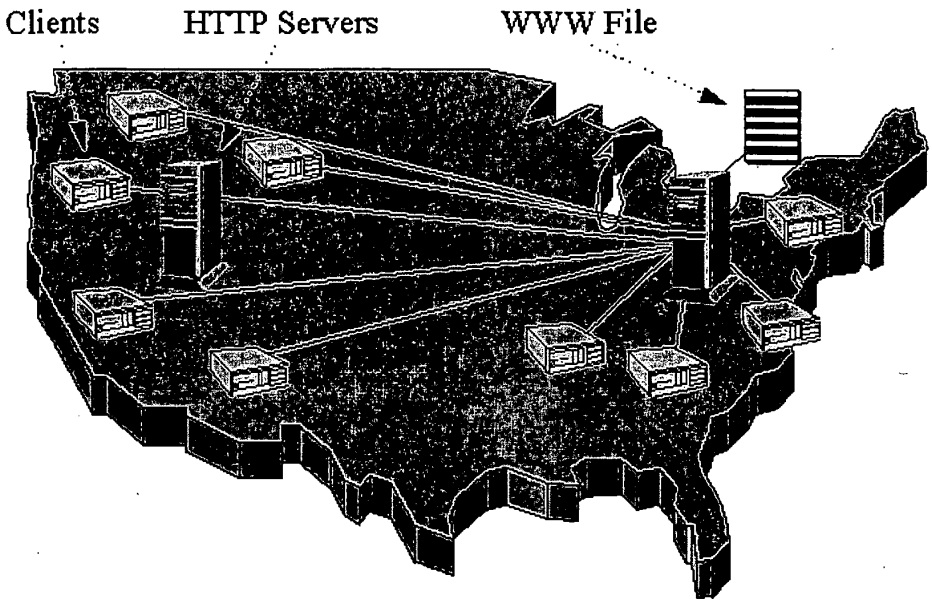


Figure 1: Before file replication takes place: several clients accessing a World Wide Web file on the east coast.

works with the modified server to decide where a given file should be cached.

We must minimize the amount of state that each Web server stores for its files, or else we will face scalability problems. We therefore track geographical access information in a coarse manner. We are currently using states and countries since these can easily be obtained from network addresses.

When the demand for a file exceeds a replication threshold, the server replicates it. We are using trace-driven simulation to determine reasonable values for the replication threshold, and we expect it to be dynamic.

The replication service decides where to replicate the file given its access history. The goal is to pick a server to cache the file that will minimize the amount of bandwidth and/or expected latency used in the future. We predict this by using the file's access history.

Figure 1 and figure 2 illustrate the replication process in action. Several clients from across the United States are accessing a file on an east coast server. The east coast server replicates the file such that network bandwidth is minimized, and the file ends up on a west coast server.

The replication service maintains a list of all HTTP servers willing to replicate files; it must choose one of them to cache the file. We are considering several algorithms to determine the optimal cache location. If the service knows the Internet's topology it can solve the problem by finding a good solution to the corresponding graph partitioning problem.

If only coarse grained information is available about the Internet, such as average latency between servers (available from `traceroute`), a better solution is to iterate over a representative sample of the available servers, calculating bandwidth savings for each. The service would then replicate the file on the server that would have reduced network bandwidth the most.

Once the primary server gives a file to another server for caching, the primary server can forget about the other server. The primary server's load will drop as clients begin to access the file from the new server. Should the primary server's load climb high enough that it must replicate the file again, the primary server will choose a different server to cache it on since the access patterns will have changed. Likewise, if the new server's load climbs high enough such that it must replicate the file, it will be cached in yet another place,

BEST AVAILABLE COPY

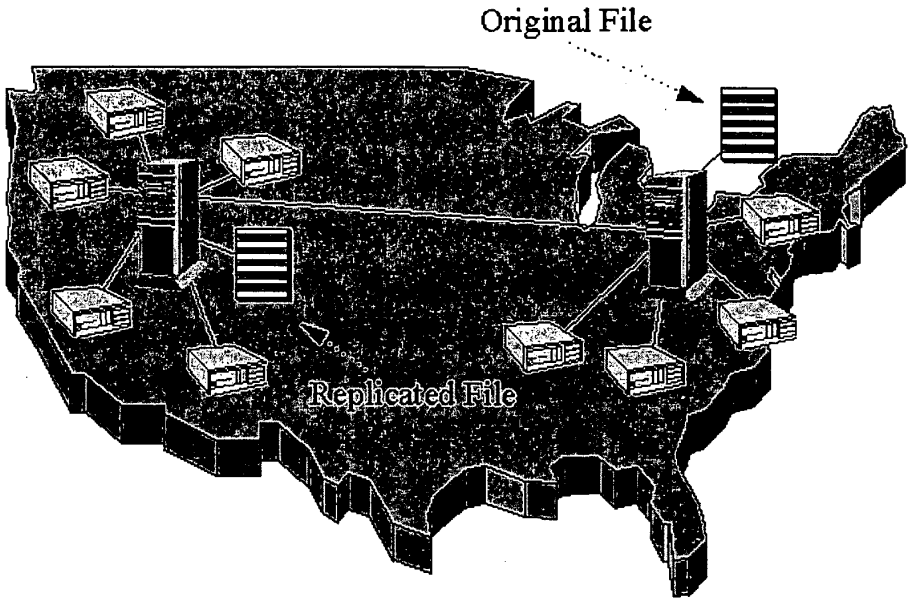


Figure 2: After file replication has taken place: the file has been replicated onto a west coast server so as to minimize network bandwidth.

because the access patterns for the new server will be very different than those for the old server.

There are two issues that must still be addressed for this scheme: file consistency and resource discovery. A server may determine that its copy of a file is out of date by using the *get-if-modified-since* HTTP request. This is an efficient way to both check consistency and to request the new file in the event it has been modified, but it is too expensive to use every time a file is requested.

Since weak-consistency should be acceptable for the Web, we are using a scheme developed for the Alex [5] file system. With the exception of dynamic pages which can not be cached we expect the Web to obey the same principle as FTP: the older a file is, the less likely it is to be modified. Therefore, the older the file that an HTTP server is caching, the less frequently the HTTP server must poll to check if its copy is still up-to-date. This is very efficient compared to checking for every request, and the client will be able to force a poll if it is essential to use the latest file.

As for resource location, there are several groups working on this problem. Until this problem is solved we are using a technique proposed by Blaze [3], which we call the "1-800 technique". Clients "call" the primary server to ask for the "server nearest you." This is not elegant, but it works because latency is currently more critical than bandwidth. The expense of querying a distant server once is amortized over the many local requests that are thereby made possible. Eventually there will need to be a cleaner solution so that all dependence on the original server can be removed. Otherwise we will face scalability and reliability problems.

4 Other Applications

Throughout this paper we have referred to HTTP servers and files. This was for the purposes of clarity, as well as to provide a focus for our research. We expect our results to be applicable to any wide-area distributed system, however; not just the World Wide Web. One application for geographical push-caching that we have in mind is to replicate not only data files but also services themselves.

A good example would be Archie [6], whose load problems are notorious. If Archie were to be written in a machine-independent network-service scripting lan-

BEST AVAILABLE COPY

guage (e.g. Tcl [11]), its code could be replicated and cached just like a Web file. This might also be the answer to how to cache dynamic pages, such as those generated by cgi-bin scripts that are used to create Web pages on the fly.

5 Related Work

There is little work on caching in large-scale distributed systems outside of distributed file systems, since only in the past few years has the attention of the distributed systems community turned toward globally distributed systems such as the World-Wide Web and FTP. Several groups are working on similar problems, but most are focused on client-side caching. Only one group that we know of, at Boston University, is also working on server-initiated caching.[2]

One such client-side system is included in The Harvest system [4]. Their *object caching subsystem* provides a hierarchically organized means for efficiently retrieving Internet objects such as FTP and HTML files. The group at the University of Colorado at Boulder has also examined the resource location problem [7] and the consistency problem [13].

Blaze [3] has addressed caching in a large-scale system. His research focused on distributed file systems, but can be applied to FTP or the Web. Finally, the Alex system [5] was designed to provide a means of caching FTP files. Of these three systems, Blaze's design comes closest to our own since it supports replication when demand becomes too high, and because it lets clients use any nearby cache. It does not, however, provide the server with control over where replicas are placed.

6 Conclusion

We do not believe that geographical push-caching should replace client-initiated caching. These two techniques address the same problem on different time-scales and are therefore complimentary. Client-initiated caching responds quickly to local changes in a file's popularity, but can not alleviate a global rise in demand. Likewise, server-initiated caching can not cope very well with sudden, localized jumps in popularity, but is best suited to handling long-term file request trends.

We close with this reminder of why server-initiated caching is necessary, taken from the Web home page of the WebLouvre [12].

> Note: Starting end of October 1994, we are currently experiencing severe network problems on our 256 Kb school Internet connection. Please be understanding! I am still looking for a site willing to mirror the WebLouvre exhibit (30 Mb in all), preferably in the USA.

References

[1] T. Berners-Lee, R. Cailliau, J-F. Groff, and B. Pollermann. World-wide web: The information universe. *Electronic Networking Research, Applications and Policy*, 2(1):52–58, 1992.

[2] Azer Bestavros. Demand-based document dissemination for the world-wide web. Technical Report 95-003, Boston University, 1995.

[3] Matthew A. Blaze. Caching in large-scale distributed file systems. Technical Report TR-397-92, Princeton University, January 1993.

[4] C. Mic Bowman, Peter B. Danzig, Darren R. Hardy, Udi Manber, and Mich ael F. Schwartz. Harvest: A scalable, customizable discovery and access system. Technical Report CU-CS-732-94, University of Colorado, Boulder, 1994.

[5] Vincent Cate. Alex - A global filesystem. In *USENIX File Systems Workshop Proceedings*, pages 1–12, Ann Arbor, MI, May 21 - 22 1992. USENIX.

[6] Alan Emtage and Peter Deutsch. Archie - an electronic directory service for the internet. In *Proceedings of the USENIX Winter Conference.* USENIX, January 1992.

[7] James D. Guyton and Michael F. Schwartz. Locating nearby copies of replicated internet servers. Technical Report CU-CS-762-95, University of Colorado at Boulder, 1995.

[8] Ari Luotonen and Kevin Altis. World-wide web proxies. In *Computer Networks and ISDN systems.* First International Conference on the World-Wide Web, Elsevier Science BV, 1994. available from 'http://www.cern.ch/ PapersWWW94/ luotonen.ps'.

[9] Mosaic-x@ncsa.uiuc.edu. Using proxy gateways. World-Wide Web. available from 'http://www.ncsa.uiuc.edu/SDG/Software/Mosaic/Docs/proxy-gateways.html'.

[10] nic.merit.edu. ftp://nic.merit.edu/nsfnet/announced.networks/nets.unl.now.

[11] John K. Ousterhout. Tcl: An embeddable command language. In *USENIX Conference Proceedings*, pages 133–146, Washington, D.C., January 22-26 1990. USENIX.

[12] Nicolas Pioch. Le weblouvre. World-Wide Web. http: //mistral.enst.fr/ pioch/louvre/louvre.shtml.

[13] Kurt Jeffery Worrell. *Invalidation in Large Scale Network Object Caches.* PhD thesis, University of Colorado–Boulder, 1994.

Digital Forensic Analysis of E-Mails: A Trusted E-Mail Protocol

Gaurav Gupta*
Senior Research Fellow
Bureau of Police Research and Development

Chandan Mazumdar
Professor
Jadaypur University

M. S. Rao
Director /Chief Forensic Scientist
Ministry of Home Affairs

Abstract

E-mail has revolutionized business, academic, and personal communication The advantages of e-mail include speedy delivery, ease of communication, cost effectiveness, geographical independence, and the portability of mailboxes. The last two are the biggest advantages over snail mail. However, with e-mail comes the threat of a genuine user being compromised through key loggers, social engineering, shoulder surfing, password guessing and other similar, though less technical, methods. This passive espionage can have a direct impact on the genuine user in terms of denial of information, loss of money, loss of time, mental harassment and an attack of personal privacy. To enable digital forensic analysis of e-mails, we propose behavioral biometric based authentication, which is analogous to a signature in paper documents. In the proposed system, if someone other than a genuine user tries to authenticate himself, then detection and fixing is possible.

Introduction

Most countries recognize e-mail as legitimate document evidence. E-mails have been used as substantial sources of evidence in cases of homicide, cyber stalking, harassment, spoofed identity and espionage. The digital forensic aspect of e-mails (e-mail forensics) requires urgent attention, due to its impact in solving most of the cases of Computer Frauds and Cyber Crimes (CFCC). To make things worse, investigative and law enforcement agencies are under-prepared to tackle the explosion of this new unseen, unheard, and innovative way of committing crime. Technologies such as quantum computing, DNA computing, and "Adaptive or Reconfigurable Computing," [1], [16] make hardware behave flexibly and can be tailored to imitate various stipulations. The latest Wi-Fi technology and migration of wireless standards 802.11b to five times faster 802.11g [1], [16] has forced rethinking about security and authentication systems.

The evolution of sophisticated and powerful digital technological solutions needs to be matched by development of tamper proof security solutions. The existing protocols, such as kerberos, one time passwords, and methods such as encryption and steganography provide only limited security from direct and active attacks such as sniffing, analyzing traffic, breaking into servers, breaking encryption and exploiting existing protocols for replay attacks. Passive espionage attack methods, i.e.. use of key loggers, password guessing techniques, password crackers, shoulder surfing, social engineering, and other similar less technical methods for compromising the authentication token, pose a very serious threat to the integrity of the genuine user account. The existing biometrics systems, in spite of being highly reliable, lack portability and cost effectiveness and are statically bound to a fixed location. Thus, biometrics systems are not suitable for email applications where portability of mailboxes over geographical boundaries along with efficiency of cost are major driving factors. Our heavy dependency on user name and password combination to authenticate provides a window of opportunity for criminals to acquire the authentication token for malicious and unlawful gain. The direct impact of passive espionage of emails includes denial of information, loss of money, loss of valuable time, mental harassment, and an attack on personal privacy. In this paper an attempt has been made to identify the peculiar characteristics which can form a basis for the development of trusted email protocol for authentication. The major issues addressed by the study are:

- Problems associated with geographical independence, portability of mailboxes, simultaneous multiple logins, and authentic date and time stamps.
- Problems arising due to espionage on a genuine user account, i.e. violation of privacy and denial of access leads to loss of valuable information, money, time and reputation.
- Embarrassment and harassment caused by illegal use of an account, requiring the user to prove his innocence.

We propose a trusted e-mail protocol, which can provide the information of when, where and how many times the e-mail has been accessed. Simultaneously, multiple logins can be prevented. We will put forward the solution from the perspective of a digital forensic expert making use of identified behavioral biometrics characteristics of the genuine user, i.e. the keystroke dynamics and the audio-visual speech recognition **(AVSR)** for the purpose of authentication through the software layer. This will also help immensely in the cost and time effective analysis of digital forensic cases.

Previous Work

RSA [2], [3], SSL [4], PGP [5], one-time passwords [6], the Kerberos [7], [8], SSID [9], WEP [10], Open Authentication, Shared Key Authentication, and MAC Address Authentication for wireless networks [10] for secure authentication are vulnerable to passive espionage attacks. If an attacker is smart enough to break any one password of a user, then he is certain to break others, too, as they will be similar to the one the attacker cracked.

The impact of a passive attack is far more serious, as any computer literate with little knowledge of the Internet can use a script to launch an attack. A lot of work has been done to address direct attack, and many secure protocols and encryption techniques have been developed. The smart card based authentication has also been used, but it is vulnerable to tampering, and can be acquired by illegal means, such as stealing or using force.

All these problems can be countered by the use of a behavioral biometrics based authentication system. Such systems are relatively economical and if implemented properly, offer a very high level of security. The limited success of initially deployed keyboard dynamics techniques in practical implementation was mainly due to non-consideration of conditions such as illness, drunkenness, and age-related limitations. Hence we propose the following identified peculiar characteristics with associated relative weights, which will be derived considering these factors, to achieve reliable solutions with minimization of false acceptance rate (FAR), false rejection rate (FRR). We propose to broaden the points of calculation of the identified peculiar characteristics for generation of a reference template, by taking into consideration all the above-discussed factors to make the proposed system more reliable.

Basic Characteristics of Behavioral Biometric Based Authentication

Universality: The universality of characteristics means that every person should have the characteristics irrespective of geographical, religion, or any other constraining boundary. The proposed behavioral biometrics based authentication system will use keyboard dynamics and audio-visual speech recognition (AVSR). These characteristics are universal and hence, suitable for the authentication system.

Uniqueness: It is essential that no two people are the same in terms of characteristics, i.e. there should be a significant scope for differentiating one person from another. This property is achieved by keyboard dynamics and corroborative audio-visual speech recognition (AVSR)..

Permanence: Permanence requires that the characteristics are invariant with time i.e. the degree of variation should be in a range where natural variations should not result in false rejection or false acceptance.

Quantifiable: This property requires that the characteristics can be measured quantitatively.

Proposed Protocol

We propose the use of a software layer that can also be used as a plug-in, based on the identified peculiar characteristics and their weights, to generate a contemporary reference template for authenticating the genuine user. That layer would also monitor patterns of the user's work to broaden the base of characteristics and determine the possible natural variations and their permitted levels specific to each user. The proposed system has three phases. First is the enrollment of the genuine user.. The second phase is authentication of the user, using BBBA, when he or she accesses the system, Third is key generated signing of email, using keyboard dynamics and AVSR.

Keyboard Dynamics

In this paper we will put forward the mechanism for efficient and tamper proof authentication, using identified peculiar characteristics of the genuine user through keyboard dynamics and audio-visual speech recognition, to generate a "reference template" i.e. behavioral biometrics based authentication system. The proposed system will minimize the False Acceptance Rate (FAR) and False Rejection Rate (FRR), due to its ability to learn throughout its life cycle, therefore overcoming the high rate of rejection of the genuine user when he is ill, drunken, tired, injured or aging. The proposed system will also help in both detection and digital forensic analysis, as any attempt by an attacker to compromise the patterns of AVSR and keyboard dynamics will be logged with date and time stamps.

Keystroke dynamics, also referred to as typing rhythms, is considered one of the most unusual and innovative biometric technologies. It is a fairly new biometric technology and is still underdeveloped and underutilized [11], [12], and [13]. Keystroke dynamics looks at the way a person types on a keyboard. Specifically, keyboard dynamics measures two distinct variables (the identified peculiar characteristics): "dwell time," which is the amount of time a person holds down a particular key, and "flight time," which is the amount of time a person takes between keys.

Also, additional variables can be used for more reliability. These include the time taken in between every key as distinguished from every other key and the time taken between combinations of keys. Keyboard dynamics systems can measure one's keyboard input up to 1000 times per second. Keystroke dynamics requires, as with most biometrics technologies, a "reference template" [14]. This involves an initial session with a person using a keystroke dynamic system, so that the system can construct or build the "reference template" by detecting the person's typing rhythms.

Keystroke dynamics is behavioral in nature, hence if developed and implemented properly will offer a maximum level of tamper proof secure authentication. Enrollment, as well as identification, goes undetected by the user; that is, it is passive, occurring without user knowledge. Another inherent benefit to using keystroke dynamics as an identification device is that the hardware (i.e. keyboard) is inexpensive and non-intrusive. Also even if someone has physical access to system and boots through floppy and CD drives, it is extremely difficult to remove the files and folders of the proposed system, as it is coupled with keyboard drivers. Hence if tampering occurs, the system will not work. This is not the case with other biometrics systems.

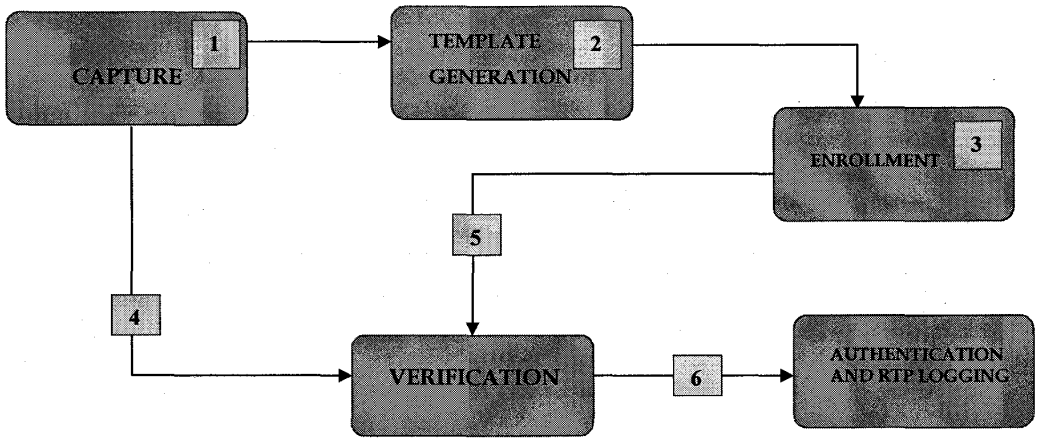


Figure 1. Four Step Mechanism of Behavioral Biometrics Based Authentication

Visual Interactivity: Audio-Visual Speech Recognition [17], [18]

Hindrances such as background noises, sore throat, and other illnesses of the genuine user posed a threat to robust speech recognition systems. These implicit problems can be over come by using the visual features of genuine user to make reliable audio-visual speech recognition systems. The use of visual features in AVSR is justified by both the audio and visual modality of the speech generation and the need for features that are invariant to acoustic noise perturbation. The speaker independent audio-visual continuous speech recognition system relies on a robust set of visual features obtained from the accurate detection and tracking of the mouth region.

Working

The proposed BBBA will be an added software layer over the existing email system and will be compatible to all types of existing email systems. The BBBA, using Keyboard Dynamics and audio-visual speech recognition (AVSR), will generate a Reference Template, which will later be used for authentication purposes, along with the user name/password combination. The reference template will be logged and will be useful for digital forensic analysis in the following scenarios:

- When a genuine person denies that he has accessed the system, i.e. to prove his innocence or otherwise.
- When somebody fraudulently uses or tries to use the system. Here two cases arise: one, when the user name/password combination is correct, but the reference template does not match, and the other, when both the user name/password combination and reference templates do not match.

We can also generate a unique key based on Keyboard Dynamics and AVSR, which could be used to sign the emails. This will help in linking the body of the email/text to the person and to tackle "man in the middle" attacks. The BBBA will help immensely in

cases where a genuine user is implicated, i.e. when a criminal uses a passive espionage attack to gain an authentication token and misuses the account of the genuine user by threatening someone, or by sending pornographic and obscene material to harm the genuine user's reputation. The proposed system can prove the innocence of the genuine user and also any intentional disguise. The RTP will log the attempted authentication template information, which can be retrieved from the server by competent authorities in order to detect and fix the espionage and to establish authentic date and time stamps (ADTS).

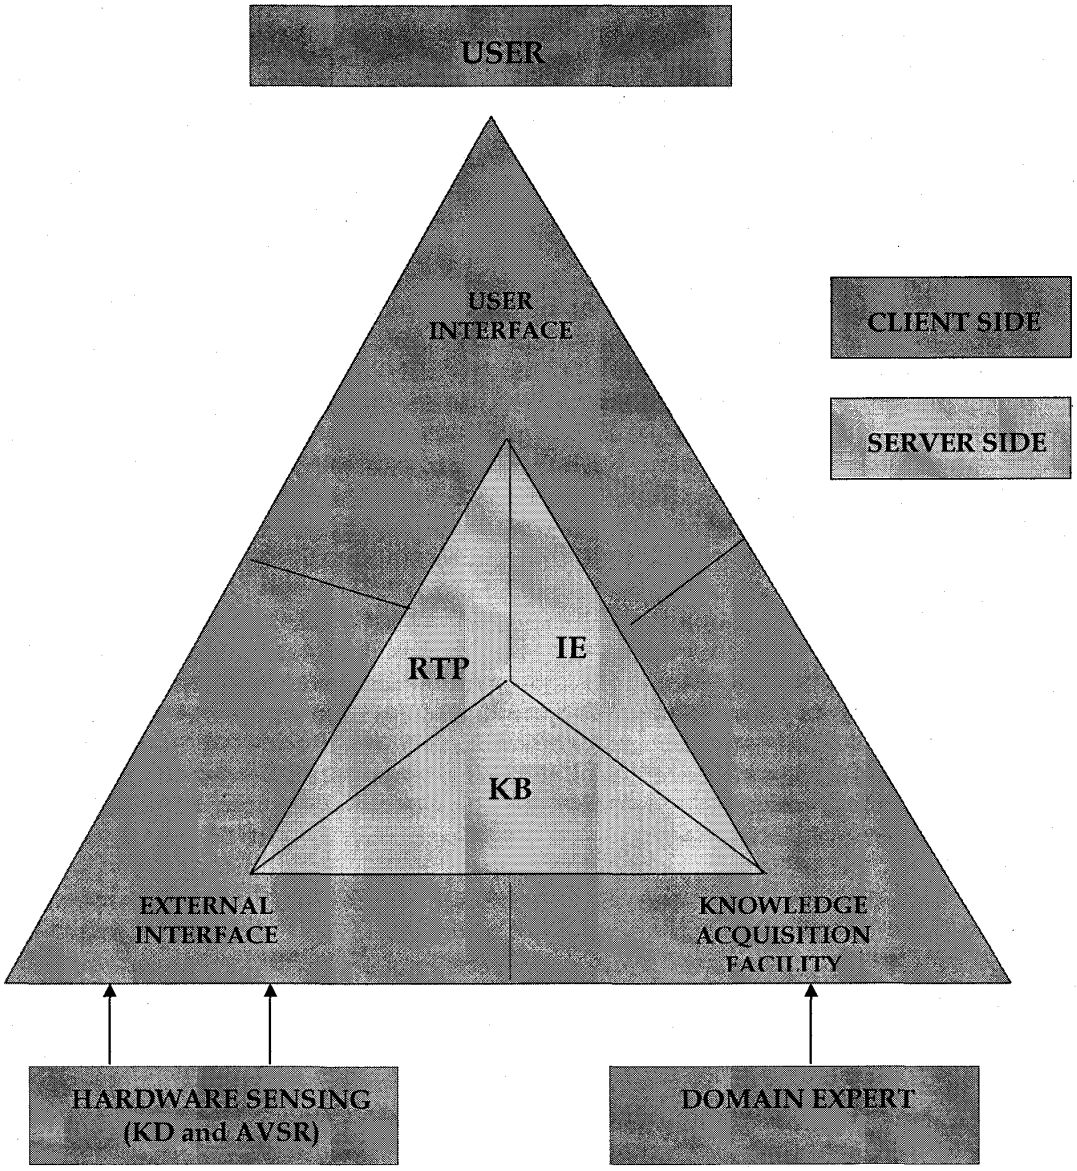


Figure2. Architecture of Behavioral Biometrics Based Authentication System

RTP: Reference template profiling for digital forensic analysis
IE: Inference engine authenticating through reference template for genuineness
KB: Knowledge base of reference templates

False Positives and False Negatives

Behavioral based biometrics systems are analogous to signatures on paper documents, as they may be known to anybody, but still extremely difficult to forge or self-disguise. With behavioral based biometrics systems, no one except the genuine user can gain entry. The identified peculiarities of keyboard dynamics and AVSR, specific to a genuine user, can form the basis for development of an expert system, which is able to monitor these patterns to establish the genuineness of the user. Because the chances of false positives and negatives are crucial to establishing the acceptance of such a system, there must be a learning capability which can adapt to gradually changing unique peculiar characteristics. The proposed system will also consider the impact of a change of systems, malfunctioning equipment, illness of user, drunken user, and effect of age, i.e. time on reliability in terms of false positives and negatives for rejection or acceptance.

In order to be implemented, the proposed system combines learning and updating according to user habits, and reflects even gradual change in a contemporary reference template, hence minimizing the impact of age, fatigue, illness, and time. Also the proposed system will authenticate the genuineness of the user, based on results deducted from a combined analysis of peculiar identities and characteristics, allowable natural variations, their respective assigned weights to minimize chances of false positives and negatives, and pattern recognition using keyboard dynamics and audio-visual speech recognition (AVSR). The proposed system will require a user to feed their characteristics through an initial interactive session. Experiments conducted so far, show that it is usable and provides industry-acceptable results. This inexpensive, scaleable, easy to deploy, and proactive concept adds a secure layer for raising the threshold of strong authentication.

Passwords are the most popular and firmly entrenched form of computer security and access used today. Passwords are also the most vulnerable security method, due to the ease with which they can be cracked and carelessly shared or posted. The concept presented here virtually eliminates this problem by providing an additional layer of strong user authentication to existing protocol systems. This new layer is based on the science of behavioral biometrics based authentication BBBA and can accurately determine whether the person typing and speaking is authorized to have access to the network or resource they are requesting. This method is unobtrusive as it allows the genuine user to log on in a manner with which they are familiar. The only new step is the initial enrollment process, where the legitimate user provides a series of typing and oral samples to train the proposed system to recognize their unique rhythm. The advantage of this system is that even if the user's authentication token is compromised, the user's unique pattern makes it next to impossible for criminals to get access to his/her account, much like your signature, which everyone knows, but can not execute.

The combination of keyboard dynamics and the characteristics identified using AVSR provides unique, measurable characteristics for a human being that can be used to authenticate the person. It eliminates entry gained by spuriously generated passwords through direct attacks, using mechanized methods. The proposed system, in conjunction with existing protocols, makes compromising a genuine user account next to impossible. The success rate, efficiency, implementation, effect of natural variations, and the chances of false negatives and positives have been considered and an acceptable solution level has been achieved in experimental protocol [11], [12], [13], [14] and [15].

For more efficient and user-friendly implementation, a layered approach has been devised. The specific steps necessary to establish the objectives of the protocol include:

- Identifying and defining unique peculiar characteristics of the individual user.
- Making an artificially intelligent system with a capability to learn.
- Developing an inference engine to authenticate a user with permissible natural variation.
- Collaborative authentication using AVSR capabilities to guarantee the genuineness of the user.
- Methods to deal with False Acceptance Rate (FAR) and False Rejection Rate (FRR).
- Automated filtering of peculiar characteristics and enhancement of the database depending on the changing environment .
- Establishment of standards, principles, quality, and admissibility according to law.
- Creation of appropriate tests to check the reliability of each step of the process.
- Creation of appropriate tests to determine the effect of human interaction and involvement on each step of the process.
- Measurement of the scope of error (mainly human) and ways to minimize them.
- Ways to minimize human interaction and maximize automatic detection, initialization, and control of the knowledge base involved in a digital forensic examination.

Conclusion

The proposed behavioral biometrics based authentication is based on universal characteristics, making it suitable to use for authentication. The BBBA can counter passive espionage attacks of:

- Key logger
- Social engineering
- Shoulder surfing
- Password guessing
- Password cracking tools
- Internal security breaches, including negligence
- Casual sharing of accounts.

Hence, BBBA is a less expensive and passive method which provides high security. At the same time, its use allows resources to be conserved, as there is no need to rotate and reassigns password on regular basis. The global reference template will log the reference template and location coordinates, along with defined authenticated date and

time stamps of any false attempt. Thus, the system is useful in digital forensic analysis, as it helps to locate suspects and match their samples to recorded reference templates. Because only a genuine user's reference template can be authenticated, it tackles the problem of integrity of a genuine user account, along with the question of simultaneous multiple logins.. The captured reference template of a genuine user will be transmitted in encrypted form for authentication purposes, based on a session to avoid sniffing and reply attacks. Innocence can also be proved by comparing the logged global reference template. Because behavioral biometrics based characteristics are universal in nature, they can be used for tracking and monitoring purposes and for reliable and tamper proof authentication and identification.

The key advantages are:

- The portability of implementation scores over other conventional biometrics methods such as fingerprints, electronic signatures [19], facial, voice, lip movement recognition [20], and iris recognition systems as a biometrics system requires costly hardware, which is generally not available or portable. Contrary to this, BBBA requires only a keyboard.
- Economical and efficient security without costly hardware, which can be customized.
- Its passive in nature, requiring no special skills and offering forensic aspect implementation.
- Capability for global monitoring and tracking.

The constraints are:

- Only genuine user will be allowed.
- Slight computational overhead.
- Person dependent and will fail if person is not in normal state.
- Need for improving ERR, FAR, FRR.

Future work

- Universal implementation can be used to have interconnected reference templates for multiple applications, to avoid several different points of authentication, thereby saving resources and improving efficiency.
- A global scope will help to counter the misuse of digital technologies by criminals. The Scope of Global Monitoring through the Reference Template's universality will be explored.
- A feasibility analysis and testing of similar systems for other commonly used applications, such as PIN numbers for smart cards, will be explored, as well as the quantification of the role of the proposed system in tracking criminals and digital forensic analysis.

© 2004 International Journal of Digital Evidence

About the Authors

Gaurav Gupta * is a Senior Research Fellow in the Computer Forensic Division of the Bureau of Police Research and Development in Ramanthapur, Hyderabad, India. He can be reached through e-mail at gaurav_jindalin@rediffmail.com.

Chandan Mazumdar is a professor in the Department of Computer Science and Engineering at Jadaypur University in Kolkata, India. He can be reached through e-mail at chandanm@vsnl.com.

Dr. M. S. Rao is Director-cum-Chief Forensic Scientist at the Ministry of Home Affairs in the Government of India. He can be reached through e-mail at msrnd@indiatimes.com.

References

[1] The Hindu, 26th June 2003 India.

[2] Planet, Introduction to Public-Key Infrastructure, 2001, URL:

http://www.iplanet.com/developer/docs/articles/security/pki.html.

[3] Waters J. Digital Certificates, 2000. URL:

http://oden.csom.umn.edu/idsc6452/Papers/JWaters/Paper.htm.

[4] Netscape, Secure Sockets Layer, 2000 URL:

http://home.netscape.com/security/techbriefs/ssl.html.

[5] Computer Networks by Andrew S. Tanenbaum

[6] GSP Services, One-time Passwords In Everything, 2001. URL:

http://www.gsp.com/cgi-bin/man.cgi?section=4&topic=opie.

[7] Duke University, Kerberos: Strengths and Weaknesses, August 18, 1997

URL: http://www.oit.duke.edu/~rob/kerberos/kerbasnds.html.

[8] ACIS 5584, Kerberos Authentication in Windows 2000, February 4, 2001.

URL: http://filebox.vt.edu/users/hmao/project1/win2k.htm.

[9] Internet Security Systems, http://www.iss.net/wireless/

[10] RSA Security, http://www.rsasecurity.com/rsalabs/technotes/

[11] Monrose, F. Rubin, D. Keystroke dynamics as a biometric for authentication. March 3,1999,

[12] URL: http://www.cs.columbia.edu/~hgs/teaching/security/hw/keystroke.pdf.

[13] Automated Identification and Data Capture Biometrics Web Site, 2001 URL:http://www.tech.purdue.edu/it/resources/aidc/BioWebPages/Biometrics_Keystroke.html/(2001-05-02).

[14] BioPassword, 2001. URL: http://www.biopassword.com/.

[15] Government Computer News, BioPassword Security Checks User's Typing Pattern, April5, 2001 URL: http://www.washtech.com/news/software/8838-1.html.

[16] Quick Silver Technology, URL:http:// www.qstech.com

[17] Visual Interactivity: Audio Visual Speech Recognition by Ara V. Nefian, Lu Hong Liang, Xiao Xing Liu, Xiaobo Pi from http://www.intel.com/research/mrl/research/avcsr.htm#.

[18] Deravi, F. Audio-Visual Person Recognition for Security and Access Control , January 9, 1999, URL: http://www.jtap.ac.uk/reports/htm/jtap-038.html.

[19] CyberSign, 2001. URL: http://www.cybersign.com/.

[20] BioID, BioID Enterprise 3.0, 2001. URL: http://www.bioid.com

Proxy Signatures for Delegating Signing Operation

Masahiro MAMBO Keisuke USUDA Eiji OKAMOTO

School of Information Science
Japan Advanced Institute of Science and Technology
15 Asahidai Tatsunokuchi Nomi
Ishikawa, 923-12, Japan
Email: mambo@jaist.ac.jp

Abstract

In this paper a new type of digital proxy signature is proposed. The proxy signature allows a designated person, called a proxy signer, to sign on behalf of an original signer. Classification of the proxy signatures is shown from the point of view of the degree of delegation, and conditions of a proposed proxy signature for partial delegation are clarified. The proposed proxy signature scheme is based on the discrete logarithm problem. Compared to the consecutive execution of the ordinary digital signature schemes, it has a direct form, and a verifier does not need a public key of a user other than the original signer in the verification stage. Moreover, it requires less amount of computational work than the consecutive execution of the signature schemes. Due to this efficiency together with the delegation property, an organization, e.g. a software company, can very efficiently create many signatures of its own by delegating its signing operations to multiple employees. Another attractive feature of the proposed schemes is their high-applicability to other ordinary signature schemes based on the discrete logarithm problem. For instance, designated confirmer proxy signatures can be constructed. Furthermore, using a proposed on-line proxy updating protocol, the original signer can revoke proxies of dishonest proxy signers.

1 Introduction

A hand-written signature, or a seal, is effective for the authentication of both the content of a document and the identity of its owner or creator. An electronic alternative to them is considered to be a digital signature, where only a person possessing a secret s can produce numbers satisfying predetermined equations which include a public value v compatible with s. Strictly speaking, there is one difference between the hand-written signature and the seal. The former is created only by the true signer, but the latter can be created by any person that is given the seal from its possessor. Considering a relation such that a digital signature can be created by any person that is given a secret from its original possessor, the digital signature might be regarded as an electronic alternative only to the seal. However more detail observation is needed.

In the following discussion we take account of the degree of the importance of a document and the number of created seals or the number of pairs of secret and public keys. If a user uses only one seal that is officially certified, s/he should not hand it over even to a reliable person. This is because an important document, e.g. a document for user's asset, is sealed exactly by the single seal. Similarly a secret should not be given to others if only one pair of secret and public keys is used. Handing over a secret is physically possible, but it breaches a fundamental security assumption of the digital signature mechanism. From this viewpoint, the seal and the digital signature function as the hand-written signature does.

Even today, there are countries in the world where seal is customarily used in daily life. One good point of the seal is that one can leave it with other reliable persons, e.g. a secretary, and the secretary uses it for less important and time-consuming procedural documents. As mentioned in the above paragraph, a single seal should not be given to others. However, if one creates multiple seals, keeps at least one seal at hand for important documents, and hands over extra seals to a secretary for less important documents, such an implementation is very useful for the possessor of the seals.

In a similar way multiple pairs of secret and public keys are created and used by each user in the digital signature. Although multiple pairs of keys are very useful

Permission to make digital/hard copies of all or part of this material for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copyright is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires specific permission and/or fee.
CCS '96, New Delhi, India
© 1996 ACM 0-89791-829-0/96/03..$3.50

as in the seal, it is not always implementable in terms of increased burden of key management. On the other hand, in some cases, whether for multiple pairs of keys or for a single key, a pair of keys prepared for important documents should also be used for delegating signing operation without breaching the fundamental security assumption. Therefore, a digital signature corresponding to the transferable seal among the multiple seals should be invented. It is not unimaginable that such a digital signature becomes popular in the future cyberspace.

No straightforward method corresponding to the transferable seal has been proposed so far, at least not at the application level. In the strict sense, a physical device, e.g. an IC card, provides a functionality of the transferable seal by the following implementation. A secret is hidden inside the device and protected by a tamper-proof mechanism. The device is designed in such a way that any person given it can create a signature without inputting any additional information, like a password. Although such a device operates as a transferable seal, a scheme using a physical device is essentially not cost effective. From this view, it is important to create a new type of cost effective digital signature corresponding to the transferable seal without relying on any physical device.

Suppose an employee in a company needs to go on a business trip, and he has no special portable computer available for the access to the company. It is highly probable that he will not be able to gain access to any terminal of the world-wide computer network. During the trip he will receive many emails, electronic mails. Some may need to be responded to quickly. For instance, the employee is involved in several different projects, and he has to give comments on plans of projects.

Before going on a trip, he plans to forward his email to his secretary. The forwarded email is carefully prepared, and it can be read by the secretary [MO95]. Then the secretary responds to the email in place of the employee according to his decision policy told beforehand, or if the employee can make a call to the office, the secretary follows instructions given by him. Since his comment is influential in the decision making, the message should be signed. Then *how can the secretary digitally sign the message under the employee's name instead of the employee?*

More meaningful example is the following. Suppose a software company digitally signs all of its programs under its secret s in order to certify the correctness of their content. Attached digital signatures offer a functionality of identifying the creator of the programs and, more importantly, of detecting any kind of alteration to the programs. The most conceivable threat is the infection with computer viruses. The president of the company knows that the fraction of programs infected with viruses before put on the market is not negligible. She does not want to give s to programmers. Her first idea is to ask all programmers to submit their programs. After checking the content of the programs, she signs them under s by herself. But this idea is not good enough since she cannot deal with enormous amount of programs. Thus a new method should be sought where the president gives each programmer a secret value which is distinct from s, but a signature created from which shows an agreement of the company. With the use of such a signature, the company can simultaneously produce a lot of products accompanied by its signatures, which are generated by multiple employees, and the signing operation is performed in an efficient way.

In this paper a digital signature, a digital *proxy signature*, which is well suited to the situations above, is studied. The proxy signature allows a user to delegate his signing operation to a designated signer, a *proxy signer*. Naturally speaking, it is better to construct an efficient signature scheme in terms of computational work and signature length. Hence, a new type of digital proxy signature, which is more efficient than previous ones, is proposed.

Related work

The problem of delegation of rights, called a proxy problem, itself is not new. If we do not limit our concern on constructing a digital signature, which is not the consecutive execution of several digital signatures described later on, there is previous work on this problem, especially in distributed systems. Please refer to Section 5 of [Neu93] and the references of [VAB91] for these contributions. Large amounts of the previous work deals with a mechanism for proxy-based user authentication. If an authorized user, i.e. a grantee or a delegate, can show s/he is the right person to be authorized, then any of the authorized user's action is considered to be an action of the user who has authorized the right, i.e. a grantor or a delegator.

A digital signature scheme is a method to realize the user authentication because the digital signature scheme ensures both message authenticity and identity of its creator. By the consecutive execution of ordinary digital signature schemes to be described in delegation by warrant in Section 2, one can establish a proxy-based user authentication, e.g. [VAB91, Neu93]. In this paper a new type of digital proxy signature scheme is shown, which is not constructed by the consecutive execution but rather created in a direct form. That is, we do not need to execute two component signature schemes, one for warrant creation by a delegator and one for signature creation by a delegate, successively. The proposed signature scheme bears the proxy mechanism by its own. Consequently the proposed scheme requires only one verification in contrast to two verifications in the consecutive scheme, and it is more efficient.

Meanwhile, there are other digital signatures related to the delegation of right. In undeniable signatures [CA89, Cha90], the confirmation operation of a verifier is partially delegated to a signer. This confirmation operation can be delegated to a single person [BCD91] or to multiple persons [Per91] after a preliminary negotiation. In designated confirmer signatures [Cha94, Oka94], an original signer can delegate the confirmation capability to another person without a preliminary negotiation.

The delegations described above are all for the signature verification. In contrast, our concern is on the delegation of signing capability. [UMUO96], which is highly devoted to computer-virus protection, uses a proxy signature as a tool to automatically generate a signature right after compiling a source program. There is no time for an executable program to be infected with viruses before signature creation. The signature scheme in [UMUO96] is different from a scheme proposed in this paper.

Main subjects of this paper are to clarify conditions of a new type of digital proxy signature(Section 2) and to propose a straightforward and concrete proxy signature scheme(Section 3) . The proposed scheme is not constructed by combining component public key signature schemes. It is widely applicable to schemes based on the discrete logarithm problem(Subsection 3.3). Moreover, it should be stressed that proposed schemes have a revocation capability of authorized proxies(Section 4).

2 Classification and conditions

It is assumed that a signer Shigeo asks a proxy signer Peggy to carry out signing instead of him, and a verifier Veronica checks the validity of created signatures.

There are different types of delegation, *full delegation*, *partial delegation* and *delegation by warrant*.

(Full delegation) In the full delegation, a proxy signer is given the same secret s that an original signer has, so that she can create the same signature he creates. Obviously, when the proxy signer deliberately signs a document unfavorable for the original signer, her mischievous action is not detected because the signature created by the proxy signer is indistinguishable from the signatures created by the original signer.

(Partial delegation) In the partial delegation, a new secret σ is created from s, which follows the modification of a verification equation, and σ is given to a proxy signer in a secure way. The created signature is checked by the modified equation, but not by the original equation. That implies a signature created by the proxy signer is distinguishable from a signature created by the original signer, and the original signer, who has found a signed document with the content unfavorable for him, can distinguish his ordinary signature from a proxy signature for partial delegation. More notably, a proxy signature for each proxy is distinct. This property corresponds to a fact that seals with different rings or marks leave different images on a sheet. In this delegation, only the public key of the original signer is required for the verification. As far as the author's knowledge, this type of delegation has not appeared in the literature.

(Delegation by warrant) The last delegation is implemented by using a warrant, which certifies that Peggy is exactly the signer to be entrusted. Delegation by warrant is performed by the consecutive execution of signing of the public key signature scheme, and this type of delegation has appeared in the literature, e.g. [VAB91, Neu93]. There are two types of signature schemes for this approach.

In the first approach, a warrant is composed of a message part and an original signer's signature for a public key of Peggy. Or the warrant is just composed of only a message declaring Peggy is designated as a proxy signer. Given the warrant, Peggy signs a document under her secret by an ordinary signature scheme, and a valid proxy signature consists of a created signature y_w^p together with the warrant. It should be remarked that y_w^p is not related to the original signer's public key v in this case.

In the second approach, a warrant is composed of a message part and an original signer's signature for a newly generated public key. The secret key compatible with this generated public key is given to Peggy in a secure way. Given the warrant, Peggy signs a document under the given secret by an appropriate signature scheme. A created message and its verification is similar to that in the first approach except that only the original signer's public key is required.

The former and the latter approaches correspond to two classes of proxies in [Neu93], called a delegate proxy and a bearer proxy, respectively.

Proxy signature schemes can be constructed for each of these delegation types, and the most adequate scheme should be selected depending on the user's need for security, message length and signer's and/or verifier's computational ability. The partial delegation, and the delegation by warrant, are more secure than the full delegation in a sense that created proxy signatures are distinguishable from ordinary signatures. Peggy for partial delegation can never create an original signer's signature for the ordinary signature scheme. The last delegation can be implemented by ordinary signature schemes without any modification.

For proxy problems mentioned in the introduction, especially in the software company example, a fast signature scheme is demanded. The advantages of the partial delegation, including the merit on the processing speed,

are shown in the subsection 3.4.

As described in [Neu93], the delegation by warrant is appropriate for restricting documents to be signed, e.g. a warrant states that applicable documents are only non-commercial ones. The partial delegation does not have such a property.

Meanwhile, since a proxy in the partial delegation does not specify its valid period, a proxy signer can create a proxy signature for ever. If a warrant does not show the period of its validity, the delegation by warrant has the same problem. In these cases the original signer should revoke a created signature or the signing capability of the proxy signer. In Section 4, *proxy revocation* is discussed.

Hereafter the second type of proxy signatures are simply called proxy signatures if it does not lead any confusion. The following conditions are imposed on proxy signatures.

[Conditions of proxy signatures(partial delegation)]

(i) **(Unforgeability)** Besides an original signer only a designated signer, called a proxy signer, can create a valid proxy signature for the original signer.

(ii) **(Proxy signer's deviation)** A proxy signer cannot create a valid proxy signature not detected as her signature.

(iii) **(Secret-keys' dependence)** A new secret σ is computed from a secret s of an original signer.

(iv) **(Verifiability)** From proxy signatures a verifier can be convinced of the original singer's agreement on the signed message either by a self-authenticating form or by an interactive form.

(v) **(Distinguishability)** Valid proxy signatures are distinguishable from valid self-signing signatures in polynomial time or size computation. Here, the self-signing signature means an ordinary signature created by the original signer.

(vi) **(Identifiability)** An original signer can determine from a proxy signature the identity of the corresponding proxy signer.

(vii) **(Undeniability)** Once a proxy signer creates a valid proxy signature for an original signer, it is not disavowed even by the proxy signer.

Proxy signer's unforgeability of other proxy signers' signatures is included in the second condition. The seventh condition simply means the proxy signer cannot take back what she claimed, and it does not demand existence of a disavow protocol shown in [CA89, Cha90].

3 Proposed proxy signature

3.1 A proxy signature scheme widely applicable to schemes based on the discrete logarithm problem

A verification equation using a public value v can be modified in a way such that a signature generated from a new secret σ satisfies a modified equation. In our scheme only the public key is modified and an original signature scheme, to which the modified public key is applied, itself is kept unchanged.

Denote by v a public value of an original signer in an implemented public-key signature scheme. Our scheme is based on the discrete logarithm problem. $v \equiv g^s \bmod p$, where $s \in_R Z_{p-1}\backslash\{0\}$ is a secret of the original signer and p is a prime number whose length is taken greater than 512 bits. g is a generator for Z_p^*. The fundamental modification is described as follows.

[Basic Protocol]

Step 1. (Proxy generation) An original signer generates a random number $k \in_R Z_{p-1}\backslash\{0\}$ and computes $K = g^k \bmod p$. After that, he calculates $\sigma = s + kK \bmod p - 1$.

Step 2. (Proxy delivery) The original signer gives (σ, K) to a proxy signer in a secure way.

Step 3. (Proxy verification) The proxy signer checks a congruence such that

$$g^\sigma \equiv vK^K \bmod p. \qquad (1)$$

If (σ, K) passes this congruence, she accepts it as a valid proxy. Otherwise, she rejects it and requests him a valid one, or she stops this protocol.

Step 4. (Signing by the proxy signer) When the proxy signer signs a document m_p for the sake of the original signer, she uses the σ as an alternative to s, and executes the ordinary signing operation. Then, $(m_p, (\text{Signature of the original scheme}), K)$ serves as a created proxy signature.

Step 5. (Verification of the proxy signature) The verification of the proxy signature is carried out by the same checking operation as in the original signature scheme except for the extra computation $vK^K \bmod p$. The computed value $v'(\equiv vK^K \bmod p)$ is dealt with as a new public value.

Since the congruence for verification at the last step uses a public value v of Shigeo, Veronica can confirm his involvement in the proxy signature.

3.2 Security considerations

A crucial point of the proposed proxy construction is that σ is a part of a signature (σ, K) of an original signer that possesses a secret s compatible with a public value v and that the proxy signer uses σ as a secret for the main signature scheme. A signature scheme (congruence (1)) for the creation of σ corresponds to the modified schemes [YL93, Shi94] of the ElGamal signature scheme [ElG85]. The former modified scheme is based on a congruence $g^{\sigma} \equiv vK^{\{Km \bmod q\}} \bmod p$, where $q|p-1$. The latter modified scheme is based on a congruence $g^{\sigma} \equiv v^m K^K \bmod p$. The proxy signature implicitly includes the original signer's signature for a predetermined constant number, 1, in either [YL93] or [Shi94].

In the similar context, the following congruence can be used as an alternative to the congruence (1) among various signature schemes [Shi94, NR94, HPM94] based on the discrete logarithm problem.

$$g^{\sigma} \equiv v^K K \bmod p \tag{2}$$

By replacing the exponent K with $h(K, C)$, where $h(\cdot, \cdot)$ is a one-way hash function and C is a constant, one can easily observe the congruence (2) has a form similar to the checking congruence of the Schnorr signature scheme [Schn91]. Note that the scheme in [UMUO96] is based on the Schnorr signature scheme.

The conditions mentioned in Section 2 are discussed in the followings.

(On proxy signer's deviation) Concerning the second condition of the proxy signature, the proxy signer faces the following problem.

> Given a proxy (σ, K) and a public value v of the original signer, i.e. given (σ, K, v), succeed in
>
> (a) computing either s or k, or
>
> (b) generating a valid new proxy $(\tilde{\sigma}, \tilde{K})$ satisfying $\tilde{\sigma} \equiv s + \tilde{k}\tilde{K} \bmod p-1$ and $\tilde{K} \equiv g^{\tilde{k}} \bmod p$.

(b) is more generally written as

> (b') computing $(\tilde{\sigma}, \tilde{K}, \tilde{v})$ satisfying $\tilde{\sigma} \equiv \tilde{v}\tilde{K}^{\tilde{K}} \bmod p$ for some user's $\tilde{v}$.

Since (σ, K, v) itself is a signature, the success in either (a) or (b) means breaking the signature schemes [YL93, Shi94].

Concerning (b'), a proxy signer could happen to know a proxy of some user with a public key $\tilde{v}$. Following the security analysis [Nyb94] on [YL93], the proxy signer selects a number $u \in Z_{p-1} \backslash \{0\}$ and computes $U = g^u \bmod p$. Then he computes

$$\begin{cases} \tilde{v} & \equiv & v^U \bmod p, \\ \tilde{K} & \equiv & KU \bmod p \text{ and} \\ \tilde{\sigma} & \equiv & (\sigma + uK)U \bmod p-1. \end{cases} \tag{3}$$

Since $\tilde{K}^{\tilde{K}} \equiv g^{(k+u)KU} \bmod p$ and $g^{\tilde{\sigma}} \equiv g^{sU+(k+u)KU}$, $(\tilde{\sigma}, \tilde{K}, \tilde{v})$ satisfies $\tilde{\sigma} \equiv \tilde{v}\tilde{K}^{\tilde{K}} \bmod p$. For instance, by selecting $u = (p-1)/2$, i.e. $U \equiv -1 \bmod p$, the proxy signer for a public key v can forge a proxy of the public key $v^{-1} \bmod p$.

This type of forgery is very hard to avoid in a public-key cryptography based on the discrete logarithm problem. If a user with a pair (s, v) of secret and public keys happens to find in a public-key list $\tilde{v}$ s.t. $\tilde{v} \equiv v^U \bmod p$ for some U, he can obtain the secret $\tilde{s}$ corresponding to $\tilde{v}$ by computing $sU \bmod p-1$. Then he can disguise the user possessing the public key $\tilde{v}$. The difference between this type of attack in the public-key cryptography and that in the proxy signature is whether an attacker can find the secret $\tilde{s}$. In the proxy signature scheme secrets of original signers are never revealed to proxy signers. As in the public-key cryptography, when $\tilde{v}$ is given at first, the probability of successfully computing a valid proxy of a target $\tilde{v}$ in polynomial time or space might be negligible in the proxy signature.

So far, the similar vulnerability has not been found in the congruence (2).

(On unforgeability) As for the first condition, forgery by the third parties including a verifier and excluding a proxy signer for other original signer is considered to be more difficult than the proxy signer's deviation because a forger has to solve the problem (b) only from the knowledge of K, v and created proxy signatures. As mentioned above, a proxy signer for other original signer possibly derives a valid proxy for the original signer. Such a fraud is a potential danger in the public-key cryptography based on the discrete logarithm problem.

(On secret-keys' dependence) This condition is satisfied since a new secret σ of a proxy (σ, K) is computed from s by the congruence (1).

(On verifiability) New public keys are computed only from the original public key v and released K, and the use of v is an evidence for the original signer's agreement.

(On distinguishability) An original signature and proxy signatures utilize distinct congruences for verification so that these verification congruences distinguish self-signing signatures and proxy signatures.

(On identifiability) As to the sixth condition, K is the key to identify the proxy signer. Since no one except the original signer can generate with high probability a legal pair of σ and K for a target public key, only the person that has given the (σ, K) is the signer capable to create the corresponding proxy signatures. The original signer can identify her from K. Even so, the proxy signer is anonymous to the third parties because K tell them nothing on the identity of the proxy singer.

(On undeniability) As in ordinary digital signatures, a proxy signer cannot disavow an proxy signature

she has created.

Next we discuss the similarity of our scheme to the seal. Like a seal, any person that has been handed out a valid secret can generate a proxy signature, and different secrets provide different proxy signatures. As one can specify a seal from an image put on a piece of paper, an original signer can determine the proxy from a created proxy signature, and he can identify who has signed. This property is ensured by the unforgeability, the condition on proxy signer's deviation and the identifiability of the proposed signature scheme. Thus, if a proxy signer create a proxy signature without following instructions of an original signer, then the original signer finds out the proxy signer and takes a proper treatment against her. For example he could revoke the proxy of the dishonest proxy signer by the method described in Section 4.

In schemes based on a key certificate, i.e. delegation by warrant, the validity of a public key is explicitly verified with the use of the key certificate [Ken93]. Conversely, the validity of a secret key, in other words the validity of a new public key, is implicitly verified with the use of K and information hidden in a created signature.

3.3 Examples of proxy signatures

In this subsection the high applicability of the proposed signature schemes is exemplified in two implementations, one for ElGamal signature scheme and the other for Okamoto's provably secure practical signature scheme [Oka92]. In some of signature schemes based on the discrete logarithm problem, like in [ElG85], a plaintext m should have a certain structure, or a one-way function should be applied to m before signing operation [ElG85], but these processes are omitted in this paper.

[Proxy signature for ElGamal scheme]

A public key v of an original signer is computed by $v = g^s \bmod p$. p satisfies $|p| \geq 512$ for preserving scheme's security. s_p is a secret of a proxy signer, and v_p is a proxy signer's public key satisfying $v_p \equiv g^{s_p} \bmod p$. After proxy submission process, proxy (σ_p, K) satisfying the following congruences is obtained. $\sigma_p \equiv s + kK \bmod p-1$, where $K \equiv g^k \bmod p$ and $k \in_R Z_{p-1}\backslash\{0\}$.

Step 1. (Signing) A proxy signer generates a random number $r \in_R Z^*_{p-1}$, and computes $x = g^r \bmod p$. For a message m_p to be signed the proxy signer computes $y = r^{-1}(m_p - x\sigma_p) \bmod p-1$. $(m_p, (x, y, K))$ is sent to a verifier.

Step 2. (Verification) A verifier checks $(m_p, (x, y, K))$ by the congruence

$$g^{m_p} \equiv (vK^K)^x x^y \bmod p.$$

If the check has succeeded, $(m_p, (x, y, K))$ is properly signed. Otherwise, signature is not valid, or the message might be manipulated.

[Proxy signature for Okamoto scheme]

$(p, q, g_1, g_2, \eta, v)$ are public keys of an original singer, and they satisfy the following conditions. p and q are primes and satisfy $|p| \geq 512$, $|q| \geq 140$ and $q|p-1$. g_1 and g_2 are order q generators in Z^*_p. The original signer's public key v is computed by $v = g_1^{-s_1} g_2^{-s_2} \bmod p$, where $s_1, s_2 \in_R Z_q\backslash\{0\}$. There exists a relation $g_2 \equiv g_1^{\alpha} \bmod p$, $\alpha \neq 0, 1$. η is a security parameter and is set to $\eta \geq 128$. This length comes from the output length of a one-way hash function $h(\cdot,\cdot)$, for example MD5. s_{p1} and s_{p2} are secrets of a proxy signer and v_p is a proxy signer's public key satisfying $v_p = g_1^{-s_{p1}} g_2^{-s_{p2}} \bmod p$, where $s_{p1}, s_{p2} \in_R Z_q\backslash\{0\}$. After proxy delivery process, the proxy $(\sigma_{p1}, \sigma_{p2}, K)$ satisfying the following congruences is obtained by the proxy signer.

$$\begin{aligned} \sigma_{p1} &\equiv s_1 + k_1 K \bmod q \text{ and} \\ \sigma_{p2} &\equiv s_2 + k_2 K \bmod q, \end{aligned}$$

where $K \equiv g_1^{-k_1} g_2^{-k_2} \bmod p$ and $k_1, k_2 \in_R Z_q\backslash\{0\}$. k_1 and k_2 are kept secret by the original signer.

Step 1 (Signing)

Step 1-1. A proxy signer generates random numbers $r_1, r_2 \in_R Z_q\backslash\{0\}$, and computes $x = g_1^{r_1} g_2^{r_2} \bmod p$.

Step 1-2. For a message m_p, she calculates $e = h(x, m_p)$, whose length $|e|$ is equal to $\eta(= 128)$ bits.

Step 1-3. She computes (y_1, y_2) s.t.

$$\begin{aligned} y_1 &= r_1 + e\sigma_{p1} \bmod q \\ &\equiv r_1 + e(s_1 + k_1 K) \bmod q, \text{ and} \\ y_2 &= r_2 + e\sigma_{p2} \bmod q \\ &\equiv r_2 + e(s_2 + k_2 K) \bmod q. \end{aligned}$$

Then $(m_p, (e, y_1, y_2, K))$ is sent to a verifier.

Step 2 (Verification) In order to check the validity of the delivered proxy signature, Veronica computes at first $x' = g_1^{y_1} g_2^{y_2} (vK^K)^e \bmod p$, and examines if the equation $e = h(x', m_p)$ holds. $(m_p, (e, y_1, y_2, K))$ is considered properly signed only when the check succeeds.

Similar to these examples, *designated confirmer proxy signatures* and *blind proxy signatures* are constructed by encapsulating our proxy into the designated confirmer signature [Oka94] and the blind signature [Oka92], based on the discrete logarithm problem.

3.4 Comparison among different types of proxy signatures

Proxy signatures for ElGamal scheme, Schnorr scheme and Okamoto scheme are evaluated in terms of message length and amount of computational work. The results are shown in Tables 1,2 and 3 at the appendix. For the evaluation of delegation by warrant we have used the relation such that a warrant consists of a message part m_w and its signature part.

In the ElGamal scheme, the amount of computational work is smaller in the proxy signature for the partial delegation than that with a warrant. The computational work is divided into three stages: proxy creation, signature creation and signature verification. As a total, proxy signature for the partial delegation cuts one inversion operation and 25 % of modular multiplications of proxy signature with a warrant. One proxy produces many proxy signatures, which indicates that the computational work of signature creation and signature verification should be compared. Signature creation spends the same amount of computational work in both delegations. The partial delegation is 33 % faster in verification than the delegation by warrant. Proxy signatures for partial delegation have shorter signature length than proxy signatures with a warrant.

In the Schnorr scheme, the computational work is also smaller in the proxy signatures for partial delegation than that with a warrant. As a total, two hashing operations and about 16 % of modular multiplications are reduced in the partial delegation. Although the signature creation requires the same amount of computation in both delegations, the partial delegation lessens in verification one hashing operation and about 24 % of modular multiplications compared with the delegation by warrant. If $|m_w| > 244$ (bits) is satisfied, the message length always becomes shorter in the partial one. 244 bits is almost 30 English letters, and a warrant would not become so short in the real world.

In the Okamoto scheme, again the partial delegation is faster than that with a warrant. As a total, two hashing operations and about 19 % of modular multiplications are reduced in the partial delegation. The partial delegation lessens in verification one hashing operation and about 33 % of modular multiplications compared with the delegation by warrant. The gross difference in the message length is only 104 bits, and if $|m_w| > 104$ (bits) is satisfied, the message length always becomes shorter in the partial one.

In the above three examples the computational advantage mainly stems from the reduction of computational work in verification stage. The signature creation requires the same amount of computation in both partial delegation and delegation by warrant.

In Tables 1,2 and 3, the length of public keys is not counted in the total length. When a public key is sent together with a signature as a public key certificate, the actual signature length becomes longer. The public keys might be preliminary delivered to each user and each user stores them. In this case, the message length does not become longer, but a user needs to have a memory resource for keeping public keys. A user possibly receive a public key from a key certificate center every time a signer uses it. Such an approach consumes the corresponding amount of channel capacity. In all cases, public keys consume some computer resource, and the consumed resource can be estimated by the length of information, e.g. the length of a public key. The proxy signature with a warrant uses two public keys, one for an original signer and one for a proxy signer, but the proxy signature for the partial delegation uses only one public key of an original signer. The public key is longer than 512 bits as described in the subsection 3.1. That means even in the Schnorr and Okamoto schemes, the net message length is indeed shorter in the partial delegation.

Overall, the partial delegation is faster than the delegation by warrant, and the bit length needed for verification is considered to be shorter in the partial one.

4 Proxy revocation

In proxy signature schemes a proxy signer might give her proxy to others. She seems to receive no harm from such a deviation because her proxy does not produce a signature corresponding to her public key v_p. However, an original signer can tell the identity of the proxy signer from created proxy signatures according to the identifiability condition. Even if the proxy is given to other party and a proxy signature is created by the party, the original signer considers the designated proxy signer has created the proxy signature.

If a party to whom a proxy has passed or a true proxy signer frequently signs an malicious message, the original signer should revoke proxies of dishonest proxy signers in order to stop further abuses.

There are following two simple ways to revoke the proxy signer's signing capability.

> (I) To make a revocation list publicly seen. When a problem occurs, the original signer puts K or v_p to the list. Every verifier checks this list at first, and if no corresponding K or v_p is found in the list, she starts her verification.
>
> (II) To change the public key of the original signer, and accordingly update all proxies of honest proxy signers.

The above revocation mechanisms (I) and (II) require that revocation information, i.e. updated revocation list

in (I) and changed public keys in (II), be swiftly and widely propagated to all potential verifiers of a proxy signature.

For the second type of revocation a honest proxy signer needs to have a new proxy after the renewal of the public key. As described in the subsection 3.1, a designated proxy signer has already been given a proxy of the old public key in a secure way. By the following protocol an honest proxy signer can update her proxy through an insecure channel, thereby avoiding to go on a visit to the original signer again.

[On-line proxy updating protocol]

Step 1. (New public-key creation) An original signer selects his new secret $s' \in_R Z_{p-1}\backslash\{0\}$ and computes his new public key v' by $v' = g^{s'} \bmod p$. Then the original signer announces his revocation to all of honest proxy signers of him.

Step 2. (Identification) After the announcement, a proxy signer who wants to update her proxy asks him to send a new proxy. To this end, the proxy signer proves her identity by some identification protocol.

Step 3. (New proxy creation and implicit delivery) If the original signer is convinced of her identity, he looks for her old secret proxy variable k in his secret proxy variable list. He calculates her old proxy $\sigma = s + kK \bmod p-1$ and her new proxy $\sigma' = s' + k'K' \bmod p-1$, where $k' \in_R Z_{p-1}\backslash\{0\}$, $K' = g^{k'} \bmod p$. Then he computes $\overline{\sigma} = \sigma' - \sigma \bmod p-1$, and returns $(\overline{\sigma}, K')$ to her.

Step 4. (New proxy construction) Using the received information the proxy signer calculates a part of her new proxy (σ', K') simply by $\sigma' = \sigma + \overline{\sigma} \bmod p-1$, and checks its validity by $g^{\sigma'} = v'K'^{K'} \bmod p$.

(General attack without relevant knowledge) An attacking third party wiretaps transmitted information $\overline{\sigma}$ and K'. K is a known value because K is a part of a signature. Therefore the attacker tries to find s' of a new public key v' satisfying $\overline{\sigma} + s + kK \equiv s' + k'K' \bmod p-1$, given $\overline{\sigma}$, K' and K. At the same time the attacker tries to generate a valid new proxy $(\tilde{\sigma}', \tilde{K}')$ s.t. $g^{\tilde{\sigma}'} \equiv v'\tilde{K}'^{\tilde{K}'} \bmod p$. More generally, the attacker collects information for proxy up-dating. Let γ be the number of collected information. Then the attacker tries to find s' or to generate a valid new proxy $(\tilde{\sigma}', \tilde{K}')$ from the congruences for $0 \le j \le \gamma$,

$$\begin{aligned}\overline{\sigma}^{(j)} &\equiv \sigma'^{(j)} - \sigma^{(j)} \bmod p-1 \\ &\equiv s' - s + k'^{(j)}K'^{(j)} - k^{(j)}K^{(j)} \bmod p-1, \quad (4)\end{aligned}$$

given $(\overline{\sigma}^{(j)}, K'^{(j)}, K^{(j)})$ for all j.

This third party's attack is not successful because there are two unknown values $k^{(j)}$ and $k'^{(j)}$ for each congruence. Even though s and s' are the same for all j, there are $2\gamma + 2$ independent variables in γ congruences, and it cannot be solved under the discrete logarithm assumption. It is also hard to generate a valid new proxy $(\tilde{\sigma}', \tilde{K}')$. The attacker knows nothing on a part of the proxy $\sigma^{(j)}$ for any j.

(With proxy signer's knowledge) A honest proxy signer might turn an attacker after her proxy updating. By the on-line proxy updating protocol the proxy signer receives information $(\overline{\sigma}^{(0)}, K'^{(0)})$. Then she wiretaps updating information for other proxy signers. What she has to do is to find s' or $k'^{(0)}$ or to generate a valid proxy $(\tilde{\sigma}', \tilde{K}')$ from the congruences (4) with additional knowledge $(\sigma'^{(0)}, \sigma^{(0)})$.

If she can cooperate with a group G_{proxy} of proxy signers who have succeeded in proxy updating. Then the attacker can use more information $\sigma'^{(l)}$ and $\sigma^{(l)}$ for all l of members of G_{proxy}.

With the knowledge $(\sigma'^{(l)}, \sigma^{(l)})$ for $l = 0$ and all l of members of G_{proxy}, the congruences (4) for such l are alternatively expressed by

$$\begin{aligned}\sigma'^{(l)} &\equiv s' + k'^{(l)}K'^{(l)} \bmod p-1, \text{ and} \quad (5)\\ \sigma^{(l)} &\equiv s + k^{(l)}K^{(l)} \bmod p-1.\end{aligned}$$

Solving these congruences for s' or $k'^{(l)}$ given the knowledge $(K'^{(l)}, \sigma'^{(l)}, K^{(l)}, \sigma^{(l)})$ is exactly the problem (a) of the proxy signer's deviation explained in the subsection 3.2. This problem is believed to be hard to solve.

The values $K'^{(j)}$ and $K^{(j)}$ for $j \in Z_{\gamma+1}\backslash\{0\}$ and $j \neq l$ are selected independent of $K'^{(l)}$ and $K^{(l)}$, where l represents the members of G_{proxy}. Hence, the knowledge of $(\overline{\sigma}^{(j)}, K'^{(j)}, K^{(j)})$ for all $j \in Z_{\gamma+1}\backslash\{0\}$ and $j \neq l$ is not useful for solving the congruence (5). The other direction is also true. The knowledge $(\sigma'^{(l)}, \sigma^{(l)}, K'^{(l)}, K^{(l)})$ for all l is no help for solving the congruence (4) for $j \neq l$.

Because of the original signer's coin flip for every proxy generation, a proxy can be securely updated through an insecure channel with the use of the proposed on-line proxy updating protocol.

5 Conclusion

Proxy signatures are discussed and a straightforward and concrete proxy signature schemes has been shown. The proxy signature allows a designated party to sign on behalf of an original party.

In the up-coming highly computer-oriented society, some organizations, e.g. a municipal office and a software company, have to deal with a large amount of signing operations. The proxy signature enables us to cope with such cumbersome signing operations by delegating

them to many co-operators. Then the signing operations are executed in a very efficient way.

Different types of delegation, full delegation, partial delegation and delegation by warrant, are classified. The proposed proxy signature for partial delegation has a computational advantage over the proxy signature with a warrant.

When an abuse of a proxy is conducted in the proposed proxy signature, an original signer can identify the deviating proxy signer and consequently revoke the abused proxies.

Acknowledgement

The authors would like to thank the anonymous referees for their comments on our paper. The authors also would like to thank Dr. Rene Peralta for his advice on English expressions. Finally, the authors appreciate a comment of Dr. Jonathan Brandt on the name of our signature scheme.

References

[BCD91] J. Boyar, D. Chaum and I. Damgård: *"Convertible Undeniable Signatures,"* Lecture Notes in Computer Science 537, Advances in Cryptology -Crypto '90, Spring-Verlag, pp. 189-205 (1991).

[Cha90] D. Chaum: *"Zero-Knowledge Undeniable Signatures,"* Lecture Notes in Computer Science 473, Advances in Cryptology -Eurocrypt '90, Spring-Verlag, pp. 458-464 (1991).

[Cha94] D. Chaum: *"Designated Confirmer Signatures,"* Lecture Notes in Computer Science 950, Advances in Cryptology -Eurocrypt '94, Spring-Verlag, pp. 86-91 (1995).

[CA89] D. Chaum and H. van Antwerpen: *"Undeniable Signatures,"* Lecture Notes in Computer Science 435, Advances in Cryptology -Crypto '89, Spring-Verlag, pp. 212-216 (1990).

[ElG85] T. ElGamal: *"A Public-Key Cryptosystem and a Signature Scheme Based on Discrete Logarithms,"* IEEE Trans. on Information Theory, Vol. IT-31, No. 4, pp. 469-472 (Jul. 1985).

[HPM94] P. Horster, H. Petersen and M. Michels: *"Meta-ElGamal Signature Schemes,"* 2nd ACM Conf. on Computer and Communications Security, pp. 96-107 (Nov. 1994).

[Kali91] B. S. Kaliski: *"A Response to DSS,"* (Nov. 1991).

[Ken93] S. Kent: *"Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management,"* Request for Comments 1422 (Feb. 1993).

[MO95] M. Mambo and E. Okamoto: *"Proxy Cryptosystems,"* IEICE Technical Report Vol. 95, No. 240, ISEC95-18, pp. 31-40 (Sept. 1995).

[Neu93] B. C. Neuman: *"Proxy-Based Authorization and Accounting for Distributed Systems,"* Proc. of the 13th International Conference on Distributed Computing Systems, pp. 283-291 (May 1993).

[Nyb94] K. Nyberg: *"Comment New Digital Signature Scheme Based on Discrete Logarithm,"* Electronics Letters, Vol. 30, No. 6, page 481 (Mar. 1994).

[NR94] K. Nyberg and R. A. Rueppel: *"Message Recovery for Signature Schemes Based on the Discrete Logarithm Problem,"* Lecture Notes in Computer Science 950, Advances in Cryptology -Eurocrypt '94, Spring-Verlag, pp. 182-193 (1995).

[Oka92] T. Okamoto: *"Provably Secure and Practical Identification Schemes and Corresponding Signature Schemes,"* Lecture Notes in Computer Science 740, Advances in Cryptology -Crypto '92, Spring-Verlag, pp. 31-53 (1993).

[Oka94] T. Okamoto: *"Designated Confirmer Signatures and Public-Key Encryption are Equivalent,"* Lecture Notes in Computer Science 839, Advances in Cryptology -Crypto '94, Spring-Verlag, pp. 61-74 (1994).

[Per91] T. P. Perdersen: *"Distributed Provers with Applications to Undeniable Signatures,"* Lecture Notes in Computer Science 547, Advances in Cryptology -Eurocrypt '91, Spring-Verlag, pp. 221-242 (1991).

[Schn91] C. P. Schnorr: *"Efficient Signature Generation by Smart Cards,"* Journal of Cryptology, Vol. 4, No. 3, pp. 161-174 (1991).

[Shi94] A. Shimbo: *"Multisignature Schemes Based on the ElGamal Scheme,"* Proc. of the 1994 Symposium on Cryptography and Information Security, SCIS94-2C (Jan. 1994). [In Japanese]

[UMUO96] K. Usuda, M. Mambo, T. Uyematsu and E. Okamoto: *"Proposal of an Automatic Signature Scheme Using a Compiler,"* IEICE Trans. Fundamentals, Vol. E79-A, No. 1 (Jan. 1996).

[VAB91] V. Varadharajan, P. Allen and S. Black: *"An Analysis of the Proxy Problem in Distributed Systems,"* Proc. 1991 IEEE Computer Society Symposium on Research in Security and Privacy, pp. 255-275 (May 1991).

[YL93] S. M. Yen and C. S. Laih: *"New Digital Signature Scheme Based on Discrete Logarithm,"* Electronics Letters, Vol. 29, No. 12, pp. 1120-1121 (1993).

Appendix

A method used by Kaliski [Kali91] is adopted to asses the amount of computational work. The following is basic rules of his method.

Let $WM(b)$ and $WS(b)$ be amount of work to perform b-bit modular multiplication and squaring, respectively. Then basic rules of the computation are;

1. $WS(b) = 0.75WM(b)$, and
2. $WM(b_1)/b_1^2 = WM(b_2)/b_2^2$.

The first condition is a rule of thumb. In our estimation, amount of work, say W, of raising s exponentiations, e.g. $g^s \bmod p$, is $|s|WS(|p|) + 0.5|s|WM(|p|)$. By applying the basic rules above, $W = 640WM(512)$ is derived.

There are more sophisticated techniques which effectively reduce the amount of computational work, but any one of processing techniques is enough to compare our targets. The above W is computed based on the well-known square and multiply approach.

Numbers: Amount of work to perform modular multiplication.
The length of modulus is 512 bits.

$WI(b)$: Amount of work to perform b-bit modular inversion .

$WH(b)$: Amount of work to compute a hash function with a b-bit long input.

m_w: Message part of a warrant, created by original singer

m_p: Message signed by a proxy singer
m_p is also used for the evaluation of original schemes since a proxy signature scheme for full delegation is exactly the same as the original scheme.

(Abbreviations)

Com.:	Computation
cre.:	creation
P.C.:	Proxy Creation
Sig. cre., S.C.:	Signature creation
Ver.:	Verification

Table 1: Performance of proxy signatures for the ElGamal scheme

		ElGamal		
		Original (or Proxy, Full)	Proxy	
			Partial	Warrant
Total length		$1024+\|m_p\|$	$1536+\|m_p\|$	$2048+\|m_p\|+\|m_w\|$
Com.	Proxy cre.	–	641	$642+WI(512)$
	Sig. cre.	$642+WI(512)$	$642+WI(512)$	$642+WI(512)$
	Ver.	1921	2562	3842
	Total	$2563+WI(512)$	$3845+WI(512)$	$5126+2WI(512)$

Table 2: Performance of proxy signatures for the Schnorr scheme

		Schnorr		
		Original (or Proxy, Full)	Proxy	
			Partial	Warrant
Total length		$268+\|m_p\|$	$780+\|m_p\|$	$536 + \|m_p\| + \|m_w\|$
Com.	P.C.	–	175	$175+WH(\|m_w\|)$
	S.C.	$175+WH(\|m_p\|)$	$175+WH(\|m_p\|)$	$175+WH(\|m_p\|)$
	Ver.	$336+WH(\|m_p\|)$	$512+WH(\|m_p\|)$	$672+WH(\|m_p\|)+WH(\|m_w\|)$
	Total	$511+2WH(\|m_p\|)$	$862+2WH(\|m_p\|)$	$1022+2WH(\|m_p\|)+2WH(\|m_w\|)$

Table 3: Performance of proxy signatures for the Okamoto scheme

		Okamoto		
		Original (or Proxy, Full)	Proxy	
			Partial	Warrant
Total length		$408+\|m_p\|$	$920+\|m_p\|$	$816+\|m_p\|+\|m_w\|$
Com.	P.C.	–	351	$351+WH(\|m_w\|)$
	S.C.	$351+WH(\|m_p\|)$	$351+WH(\|m_p\|)$	$351+WH(\|m_p\|)$
	Ver.	$512+WH(\|m_p\|)$	$688+WH(\|m_p\|)$	$1024+WH(\|m_p\|)+WH(\|m_w\|)$
	Total	$863+2WH(\|m_p\|)$	$1390+2WH(\|m_p\|)$	$1726+2WH(\|m_p\|)+2WH(\|m_w\|)$

BEST AVAILABLE COPY

-C5

Atlanta Olympics WOMplex



By Andy Stanford-Clark

This article describes the major components of the Atlanta Olympics WOMplex and attempts to give a flavor of some cutting-edge technology developed for the project.

For the summer of '96, a team of IBM programmers and engineers, based in Southbury, Connecticut, built the official Internet Web site for the Atlanta Olympic Games. The project was immense in every respect—from the amount of content in the site, to the volume of traffic through the site, to the amount of equipment deployed, and the number of hours worked by the team.

WOMplex is an umbrella term that describes the combination of the Web Object Manager (WOM) Internet content and application server, the scalable architecture, the requisite scalable parallel processing hardware, the content authoring environment, the load-balancing systems, the user-level applications on the Web site, and the subsystems that support those applications. In other words, a WOMplex is a scalable, manageable, globally distributable Internet content and application server comprising both hardware and software.

Statistics

The URLs http://www.atlanta.olympic.org and http://results.atlanta.olympic.org together constituted the largest Web server in the world. During the 17 days of the Olympics, the site received just over 192 million hits, with peak daily hits of 16.9 million. Drilling down to the "Sneak Peek" application, which provided a constantly changing montage of captured images from 38 video feeds from the Games in Atlanta, the servers handled up to 21,000 images per hour. What makes these figures especially impressive is that every page served was built dynamically from its component objects at the instant the page was requested.

To provide this level of service 24 hours a day required non-trivial amounts of equipment. The Atlanta site was served by 50 nodes of a distributed IBM RS/6000™ Scalable POWERparallel™ system, spread across five locations in four countries. Most of the processors were at the primary WOMplex site at one of Integrated Systems Solutions Corporation's (ISSC's) major Facilities Management Centers in Southbury. This site provided the required infrastructure for a project of this scale—Uninterruptable Power Supplies(UPS), good physical security, high-bandwidth connectivity into major U.S. Network Access Points (NAPs), and raised floor.

Several nodes—typically some portion of an RS/6000 SP frame—joined the WOMplex at four mirror sites:

- Cornell Theory Center at Ithaca, New York
- IBM U.K. Laboratories at Hursley, U.K.
- University of Karlsruhe near Mainz, Germany
- Keio University near Tokyo, Japan

Each mirror site replicated most of the content of the Olympic site. Notable exceptions were the results system and the rapidly changing Sneak Peek images, which were served only from the primary site at Southbury.


Andy Stanford-Clark

Presented at Get Connected Technical Interchange '96, IBM Hursley U.K., October 1996

System Architecture for Atlanta WOMplex

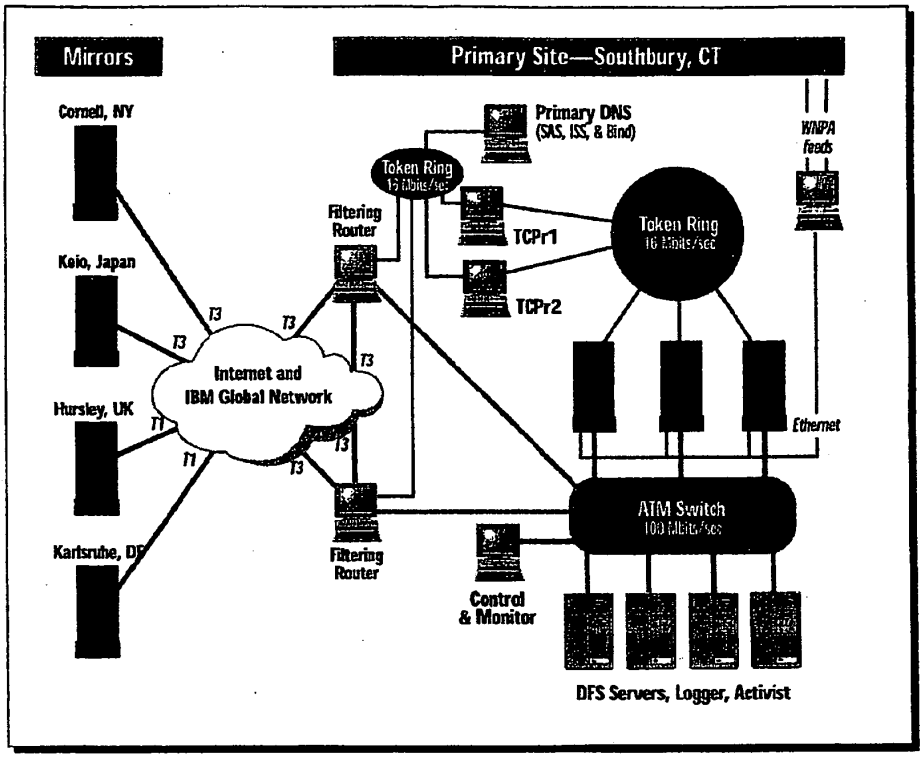


Figure 1. System architecture used for the Olympic games in Atlanta

During the last few months before the games, tickets for seats to events at the games were sold over the Internet by a server based on the IBM Net.Commerce server. In all, this generated over $5 million in ticket sales.

System Architecture

Figure 1 shows the overall system architecture for the Atlanta WOMplex. The "cloud" combines the Internet and the IBM Global Network. T3 connections are DS3 links providing 45 Mbits per second. Filtering routers are Cisco® 7500 Series routers with sets of carefully crafted filter rules to allow only the desired packets in and out from desired sources and destinations. TCPr1 and TCPr2 are the TCP routers, or WOMsprayers—these will be discussed later. Together SAS, ISS, and Bind form the name resolution service. The WNPA feed is a 56 Kbits per second leased-line connection to Atlanta, which provided the feed of results to the databases that were used to drive the Internet Results System. Several leased lines were installed for redundancy. The roles of the machines labeled Logger and Activist are discussed below.

Infrastructure

The IBM RS/6000 SP system provided the WOMplex infrastructure. At the main Southbury site, three frames of the RS/6000 SP were filled with mainly "thin" nodes (equivalent to the RS/6000 Model 390). We did not use the RS/6000 SP high-speed switch in this application—all communication between processors was done over the Asynchronous Transfer Mode (ATM) network.

Control workstations for the main SP frames were in the machine room near the RS/6000 SP. Access to the SP control workstations at the remote mirror sites was handled through X-Windows sessions running across the Internet, using

IBM 000027

Secure SHell (SSH) (described below) for secure connections. The Operations Center in Southbury could control every processor in the WOMplex—right down to power cycling.

ATM Network

The outbound network for data from the Web site was routed over a 100 Mbits per second ATM network. ATM uses centralized switching technology to provide a high-bandwidth network. The site networking was highly complex, with multiple Token-Ring, Ethernet™, and ATM networks connected to the Internet through two filtering routers. IBM 8260 switching hubs enabled this level of connectivity. A 16 Mbits per second Token Ring handled the inbound networking. *Note:* Network requirements for a Web site are highly asymmetrical—one packet that contains an HTTP request can provoke the transfer of potentially many megabytes of data from the servers.

Transarc DFS Cell

Transarc's Distributed File System (DFS) ensured that all WOMplex Web servers had access to up-to-date information. (For those familiar with Andrew File System (AFS®), think of DFS as the "grown up" AFS.) DFS is based on the fileset concept, a collection of files handled in the same manner as a single group. Each fileset in a DFS cell resides on a nominated server; it may have mirrors on other nominated servers. Fileset replication happens in two ways: on a timer interval or (often more useful) on a manual "push" from the primary server, when the application or user knows that a batch of updates has been completed.

Each WOM server in the WOMplex was a DFS client configured to request files from the nearest DFS server. At remote mirror locations, the clients retrieved their data from a local server that was kept in sync with the main servers in Southbury. One big advantage of DFS in the Web server environment is client caching. During the Olympics, most clients achieved a 98% cache hit rate.

The primary DFS server in Southbury was distributed across three RS/6000 Model 570 servers and two RS/6000 SP wide nodes. It used 128 GB of IBM 7137 RAID storage, configured with replication and mirroring to give 32 GB of effective storage with three-way redundancy in addition to the RAID failsafe.

The entire cell was administered centrally from Southbury. Problems with filter rules on routers and TCP/IP routing across the Internet make setting up a very large distributed DFS cell a complex process; although when it is running, the system works extremely well. The single DFS cell spanned six locations: Southbury (the main server), Cornell, Keio, Hursley, Karlsruhe (the mirror sites), and Atlanta, where the content was being generated for the Web site.

Internet Connectivity

To ensure the availability of enough bandwidth for the Olympics, four DS3 circuits (T3—45 Mbits per second) were connected into four major U.S. Network Access Points (NAPs): Mae-East, Mae-West, Sprint NJ, and Ameritech™ Chicago. Regardless of the origin of a request, by using a routing algorithm that minimized hop counts, the response was sent to the major NAP closest to the source of the request. Most of the world's major non-U.S. service providers connect directly into one of these four NAPs, so the site was well connected. Multiple circuits also protected the site against any link failure. Packets were automatically routed over other links or could be manually routed around trouble spots, which occurred from time to time at various NAPs.

Security

The entire Atlanta Internet site, the software development environment, and all DFS infrastructure were located outside the IBM firewalls. Consequently, site security was always a high priority for everyone because such a high visibility site was an obvious target for hackers.

High security must be built into a site from the ground up.

High security must be built into a site from the ground up; it cannot be retrofitted onto an existing site. The filtering routers provided the main protection against invasion, and the filter rules for these machines were a major design and maintenance effort. Interworking with the remote mirror sites involved setting filter rules at the remote sites to allow only User Datagram Protocol (UDP) and Transmission Control Protocol (TCP) packets on specific port numbers in and out from specified Internet Protocol (IP) addresses.

A major problem for the primary site at Southbury was the need for extensive access to systems inside the DMZ ("De-Militarized" Zone) at the remote sites; the rest of the world could only have port 80 World Wide Web access. Heterogeneous router technology across the world ensured that many people learned a lot about writing filter rules in a short space of time!

Kerberos and authentication servers were used to secure DFS. Sophisticated password schemes for

the machines were generated and then frequently changed. The apparent inconvenience of complex password schemes was easy to justify in the high profile environment of the Olympics.

Transmission of system control information, including machine root passwords, across the Internet was a potential major exposure. Early adoption of Secure SHell (SSH) for all terminal communication between machines solved this problem. AIX® and Windows® 95 clients for SSH allowed developers to use SSH as a direct replacement for the telnet they would have used otherwise. A version of the SSH client that worked through a SOCKS firewall allowed developers inside the IBM firewall to access the servers and development lab machines.

ISS provided another level of security with the "caching, routing name server," which was deployed as the primary name server for the Olympics site. In this Domain Name Server (DNS) emulation mode, ISS had authority over one or more specified subdomains. If ISS received a name resolution request for a subdomain over which it did not have authority or a name that was not being used for load balancing, then ISS passed the request to one of two backend, ordinary, name servers—typically knowledgeable DNSs that could resolve nearly any request.

The WOM system has been iteratively developed to support high-volume Web sites for sporting events.

The IP address of the original request determines the backend name server that receives the request. ISS configuration has a table of subnet mask specifications. The client IP address, or name server, requesting the name resolution is matched against this table. If a match occurs, the client is considered internal, and passed to one backend DNS; otherwise, it is considered external and sent to the other one.

From the practical standpoint, even with a distributed, multi-site operation, the appropriate client machines can access all of the main site's internal server names and addresses, whereas the rest of the world can only access the limited set of names and addresses (possibly only the Internet address) that the DNS administrator made available. The ISS name server also blocks the DNS Zone Transfer operation, a popular method to access lists of internal machines that are considered hacking targets.

The WOM Web Server

The Atlanta Web site used the latest version of IBM's Web Object Manager system, known as the *WOMserver*. The WOM system has been iteratively developed to support high-volume Web sites for sporting events and IBM external Web sites since 1994. Notable events have included Wimbledon® Tennis, Masters Golf, U.S. Open™ Tennis, and the ACM Chess Challenge (Kasparov versus Deep Blue).

Object-oriented Perl was the language used for the Atlanta server. Supporting applications and subsystems on the site were written in C, Perl, and Java™. Several prototype versions of the WOMserver written entirely in Java were also deployed in the WOMplex for experimental, demonstration, and testing purposes. A specialized binary server, based on the Apache Web server and run on a separate port number, handled binary objects—GIF, JPEG, AVI, MPEG, WAV, Java classes, and Bamba audio and video. All references to binary objects in the HTML source dynamically served as pointers to this port (for example, http://www.atlanta.olympic.org:8080/).

The data flow for entering authored content into the WOMserver had several steps. Content was authored by several types of systems, including traditional HTML editors, image editing tools, and a new generation of WOMtools for directly authoring into the WOM internal object format. This content was then "checked in" to the WOM database. Other systems, based on IBM Digital Libraries technology, generated Web pages from dynamically changing data, such as sports results.

WOMwire was the protocol used for communication among the various tools, including several batch upload tools. This protocol defines the way in which objects and their metadata are transmitted to the WOM system. The WOM database used the DB2/6000™ relational model to store the data and metadata object.

The WOMserver can access data and metadata stored in several formats. For the Atlanta WOMplex, the data and metadata were exported as flat files, called *servables*, from the DB2® database into the DFS filesystem. The DB2 export mechanism also triggered the DFS fileset replication system to push the updated files out to all servers at each WOMplex site.

Web-based applications for the Atlanta site would traditionally have been written as Common Gateway Interface (CGI) programs. However, they were written for the WOMserver in an object-oriented Perl or Java environment. They were run inside the server to avoid the overheads associated with executing CGI applications on traditional Web servers.

WOM architecture allows an executable program to generate any part of a served page;

the applications also fit within this paradigm. Detailed discussion of the WOM Network Application Platform is outside the scope of this article.

The WOMserver is *multi-homed*—it can support several distinct Web sites concurrently in the same server. This makes the WOMplex, with global load-balancing systems, a potential home for many Web sites at the same time—ideal for providing content-hosting services.

Centralized Logging

An important aspect of WOMplex operation is collecting HTTP access and error logs from the servers distributed across the world and bringing them into a central location for monitoring and analysis. The "Grim Reaper" performs this function. The Grim Reaper, a tree-structured plumbing of sockets and UNIX® pipes, streamed log records from each server to a confluence point at each site, then across the Internet to a central collection point on a machine called the *logger* (see Figure 1).

When everything is streaming nicely, the task of the Grim Reaper is straightforward. The problems occur when one of three events occurs:

- The servers produce log data faster than the local Grim Reaper can take it from them.
- Network delays or failures cause holdups in the Grim Reaper plumbing, causing a back-log of data.
- The final receiver of the logs cannot cope with the rate at which they are arriving from the Grim Reaper.

Any of these scenarios, coupled with the simple mandate that no log records can be discarded, make the Grim Reaper's job very difficult. As a result, the Reaper is a sophisticated technology. At any stage, log streams can be diverted into a file on a local disk to hold data, which for some reason, cannot be sent on. Later, when connections are restored or congestion eases, files are spooled back into the network. The Grim Reaper also can deal with requirements such as off-line transfer of batched log data (for example, late at night with less network traffic) and on-the-fly data compression for transmission across slower network connections.

After arriving at the logger machine, the log streams were transferred across a high-speed protocol Common Link Access to Workstation (CLAW) link to an 8-way CMOS 390 Squadron mainframe. The logs were then pre-processed by a series of pipes and filters, then implemented using MVS® Pipes. Individual log records were added to a table in a DB2/MVS database. From here, SQL queries could extract useful information from the logs, including hit counts and rates, online data mining using Activist, and real-time "click tracking." The raw log data was spooled to tape for later analysis.

Content Authoring Environment

The WOMwire protocol defines how authoring and site management tools communicate across the wire with the WOM database. The client part of the WOMwire protocol was implemented as an ActiveX control. The authoring tools, developed for the Windows 95 and Windows NT™ environments in Visual Basic®, C, or C++, could trivially access the WOM database by including the WOMwire control.

WOM objects can contain text, graphics, sound, video, page layout information, or executable code. These objects and their associated metadata can be checked in and out of the system, similar to a programmer's source code control system. The WOMbler created and edited content and managed metadata. WOMexplorer was used to navigate and reorganize a WOMplex site.

WOM objects can contain text, graphics, sound, video, page layout information, or executable code.

Each page on the site fits into an information taxonomy, described by a set of taxonomy tags. These tags position the page somewhere in an *information space*, which has a different hierarchy than the *data space* hierarchy imposed by the HTML links between a set of pages. This exciting new approach to Web site design is outside the scope of this article.

A Visual Basic application called the *Wu-tool classifier* was used to classify existing pages on the site. Users could classify each page into the information taxonomy by choosing from sets of predefined categories that applied to the page in question.

With the WOMpix tool, users could process images by cropping, resizing, reformatting, and labeling, before checking them into the WOM database. WOMpix dealt with many thousands of images that appeared on the site during the Olympics. It also allowed metadata, such as the source agency and photographer, to be recorded with the image. WOMpix could also generate thumb-nail images and image water-marking.

Load Distribution

Load distribution and balancing within the WOMplex uses a two-tier architecture. At each site, local load balancing distributes incoming traffic toward a single IP address across multiple backend WOMservers. Between sites, a global load-balancing system allows logical Web sites to be moved between physical server locations.

Interactive Session Support

Both tiers of the load-balancing system use software called Interactive Session Support (ISS). This is the interactive component of the RS/6000 SP LoadLeveler™ package for batch and interactive load balancing across a set of RS/6000 SP nodes or across other UNIX workstations, notably RS/6000, HP™, Sun®, and Silicon Graphics.

ISSmonitor collates agent responses, starts and stops remote agents (using inetd), and reacts appropriately if agents do not report in for an extended period of time.

Each server runs the ISSagent, which runs a user-defined metric that determines how busy the server is. The metric can be any executable program or pipeline of UNIX commands, providing it returns an integer at the end. For example, a popular choice is `ps -ef | wc -l`, which reports the number of processes currently running on the system and provides a good measure of machine activity. Far more complex metrics can be contrived, including an executable program to generate the metric value. The metric used for the Olympics counted the number of established socket connections on port 80 of each server.

ISSagents periodically report to a central ISSmonitor, which collates agent responses, starts and stops remote agents (using `inetd`), and reacts appropriately if agents do not report in for an extended period of time. The ISSmonitor also uses Internet Control Message Protocol (ICMP) ping packets to ensure that servers are still "alive" before recommending them for use. The ISSmonitor transfers a *load profile* of server activity to the WOMsprayer via a socket. The WOMsprayer uses the load profile to compile a weighted round-robin schedule for incoming traffic. Any server whose ISSagent does not respond is eliminated immediately from the load-balancing schedule to prevent incoming traffic from hitting a non-responsive server and receiving `Connection Refused` error message.

In another operational mode, ISS works as a load-balancing Domain Name Server (DNS). A generic Internet name, such as www.atlanta.olympic.org, is assigned to the cluster of servers. When a client machine or another name server in the hierarchical DNS system requests the IP address for that name, ISS ensures that the name server returns the IP address of the least heavily loaded server to provide the required service. ISS can work in DNS mode in one of two ways.

ISS can cooperate with an existing DNS, by regularly updating the "name to address mapping" database of the name server. Although this works well for low volume, long-lived connections, it is not responsive enough for high-volume, short-lived transactions usually experienced by a Web server.

ISS also operates as a high-performance, load-sensitive DNS, totally replacing the existing name server. This mode of operation can handle Web server traffic and provide a weighted round-robin mechanism for distributing high-volume load evenly across the backend servers.

ISS also offers an enhanced security, caching name server facility, described in the Security section above.

Local Load Balancing—WOMsprayer

Load-balancing software technology from IBM Watson Research has several names: TCP router, Encapsulated Cluster, WOMsprayer, ShockAbsorber, and Network Dispatcher. The *WOMsprayer*, a network router between two physical networks, offers a single IP address to the world. When a TCP connection is made to that address on the inbound network, the WOMsprayer software modifies the Medium Access Control (MAC) address field of the incoming connection request packet and transfers the packet to its outbound network. The MAC address change redefines the final packet destination, and the WOMsprayer arranges for it to be delivered to one of several backend servers.

The backend machines have unique IP addresses and also have the IP address of the WOMsprayer aliased onto their inbound network interface. Consequently, when this modified connection request packet arrives at a backend machine, it appears that the packet was destined for this server; the server does the required "accept" to create the TCP session. The WOMsprayer maintains a table of currently active connections, so any further incoming packets for that TCP session, such as packet acknowledgments, are routed to the correct backend server.

A weighted round-robin schedule determines which backend server receives an incoming session request. This schedule is constructed from the server activity profile that the ISSmonitor periodically supplies to the WOMsprayer. If ISS has not heard from a server or has reason to believe that a server is having problems, the

server is quickly excluded from the WOMsprayer schedule. No new connections are routed to that server until it rejoins the cluster.

The backend server's default route serves as the network path from the server to the client browser. This route generally does not use the same machine as the WOMsprayer. Even if the same server is used, the WOMsprayer does not intercept the returning packets.

Global Load Balancing

Two mechanisms allow Web traffic to be spread effectively across the multiple sites comprising the WOMplex: ISS-DNS and the WOMbat.

ISS-DNS

In its DNS emulation mode, ISS allows the DNS administrator to configure a table of Internet domain names to IP address mappings. This differs from an ordinary DNS in that more than one IP address may be associated with any given Internet name. An ISSagent runs on each server that is designated to serve any given name. According to the health, load, and the configured balancing policy for that name, incoming resolution requests for that name receive the IP address of the most appropriate server.

In the WOMplex, the ISSagents run on the machine that hosts the WOMsprayer for each site. Two balancing policies can be configured: Round and Fail.

Round. This policy does a weighted round-robin of the IP addresses in the list associated with that name. It is distributed according to two factors:

- The server activity profile, as determined from the input provided by the ISSagents
- A server weighting factor, allowing for different server capacity at each site—different numbers of WOMserver nodes in the RS/6000 SPs at each site, in the case of the WOMplex

Fail. This attempts to use the first IP address in the list—the "preferred" site—until that site becomes unavailable. Then, ISS switches to a weighted round-robin of the remaining IP addresses in the list, the same as in the Round policy. This system allows a single, logical Web site (defined by a specific Internet name) to be distributed across multiple physical locations, which could be a cluster of servers controlled by a WOMsprayer. Adjustments to the system configuration can be made dynamically, allowing extra servers to be switched-in, if necessary, to cope with excess load. The system enables automatic fail-over to other servers if a server experiences downtime, which can be fail-over to other servers from a preferred or primary server or omission of unresponsive servers from the round-robin schedule.

WOMbat

WOMbat uses ping to "echo-locate" a client machine and determine the nearest WOMplex site in network round-trip terms. With this knowledge, WOMserver can then divert the client to its nearest server site, and hopefully receive better service, thanks to the better network connectivity to that site.

When a client browser visits a WOMplex site for the first time, the server knows it has not seen that client before, because the client does not have a browser *Cookie* for that site. The server generates a Cookie, a unique key to identify that client in the future, and sends it to the browser to use whenever it visits that site. The Cookie is also stored in the Global Profile Daemon (GPD), a database of profile information, keyed by browser Cookie. The server also invokes— asynchronously in the background—the WOMbat Driver before composing the HTML page requested by the client and returning it to the client.

In its DNS emulation mode, ISS associates more than one IP address with any given Internet name.

The WOMbat Driver sends details of the client's IP address and its browser Cookie to a WOMbat daemon at each site in the WOMplex. The WOMbat daemon sends a volley of ICMP echo request packets (ping) to the client IP address and measures the round-trip time of any responses that return within a pre-specified time-out period.

WOMbat daemons send back to the WOMbat Driver the average network response time taken to ping the client machine. The WOMbat Driver waits for a specified length of time for the WOMbat daemons to respond, then determines which server is closest to the client by comparing the responses. Server availability information also can be factored into this calculation, so that the closest server might be rejected in favor of a more distant, but less busy server. The user's profile in the GPD contains the identity of the preferred server.

The next time the user requests a page from the Web server, the server daemon interrogates the GPD using the browser Cookie as a key, and discovers that the preferred server for that client is not this server. The WOMserver then builds the requested page for the client, but rewrites all embedded links to follow-on pages and to

embedded images as URLs on the preferred server rather than on local links. For example, if the preferred server was www.hursley.atlanta.olympic.org and the local link would normally have been the following:

```
<IMG SRC="/data/image1.gif">
```

then, the link sent back to the client is rewritten as follows:

```
<IMG SRC="http://www.hursley.atlanta.
olympic.org/data/image1.gif">
```

When the client browser follows any of these links, they are retrieved from the preferred server. From that point on, all links will point back to that server, so the client has effectively been "moved" to that server.

WOMplex separates logical services from physical servers, which allows a Web site to be hosted from any site in the WOMplex.

Virtual Addressing

AIX and other UNIX operating systems can alias multiple IP addresses onto a single network interface; that is, the server can "pretend" to have several different IP addresses on the same network at the same time. This was used two ways in the Atlanta WOMplex: multi-homing and service virtualization.

Multi-Homing. A single WOMserver instance can serve several Web sites. The IP site addresses are aliased onto the server's Internet network interface. The WOMserver configuration has separate "trees" of URL data for each hosted address. When a Web page request arrives, the WOMserver interrogates the socket connection to identify which IP address it was sent to, then serves the appropriate page from the correct data tree. Many "standard" Web servers consider this a standard feature, such as Virtual Hosting in Apache.

Service Virtualization. WOMplex separates logical services from physical servers, which allows a Web site to be hosted from any site in the WOMplex. IP aliasing within a site separates logical from physical and frees any dependence on specific hardware. All externally visible IP addresses are aliases, not the "real" base IP address of any machine. Therefore, if a server fails, the service IP addresses supported by the failed server can be transferred readily to another machine. One additional requirement, apart from adding the IP address aliases, is to flush the arp cache on the nearest network router, which forces it to reestablish the new IP to MAC address mappings. The System Network Management Protocol (SNMP) interface to the router can often automate this operation.

Global Profile Daemon

The Global Profile Daemon stores and caches client profile information, using a client browser Cookie as a retrieval key. Any WOM application or the WOMserver itself can store profile information in GPD. Profile data is held as a file of lines containing `name = value` pairs. Keys used for the Atlanta WOMplex were constructed from the concatenation of the following:

- Three randomly chosen letters
- Cookie creation timestamp (seconds since Epoch, plus an offset)
- Cyclic Redundancy Check (CRC) checksum digits

GPD uses a multi-level cache to balance the global distribution requirements of profile data and fast access to user profiles by the WOMserver. Closest to the WOMserver, GPD maintains an in-memory cache with a Least Recently Used (LRU) policy; therefore, people who are known to be currently moving around the Web site are most likely to be in memory.

To locate profiles efficiently, the directory structure has directory names 'A' through 'Z'—three levels deep. For example, the profile with the key ADF1234567123 would be found in the file with the following full path: /dfs/profiles/A/D/F/3217654321FDA.

Note that the name given to the profile file was the reverse of the key string, that is, 3217654321FDA. This is because DFS uses its own hashing algorithm for efficient searching of directory contents—based on the first character of the filename. Consequently, the first character of the filename should be a random character. This avoids having all files in a directory starting with the same first character, and thus all falling into the same hash bucket. We chose three directory levels and 26 entries per level, knowing that DFS caching is optimized for a certain maximum directory content size.

At the next level, DFS caches recently retrieved profiles on the local disk of the server machine. At a site level, recently accessed profiles are stored in a DFS fileset that physically resides on disks at that site. This avoids trans-Atlantic connections for many profiles. Finally, the main DFS server at the primary WOMplex site in Southbury stores the global GPD repository.

When a client comes back to the Web site (at any physical location) with a previously assigned Cookie, the WOMserver requests information on that client's profile from its local (on the same machine) GPD instance. GPD checks the in-memory LRU cache and returns the information if it is there. If not, GPD must fetch the profile from the global repository because that client has not been to this physical site recently.

GPD reads the profile from the global DFS directory and copies the file to a local DFS fileset located physically on the local site. Then GPD sends a UDP broadcast to other servers at its local site informing them that the client with a specified Cookie has come to the site, and may soon request pages from them. Other GPD instances at that site receive the broadcast and go to the local DFS fileset to read the profile information into the memory cache. Thus, only the first GPD at a site has to do the trans-Atlantic journey to get the profile data of a returning client. The other servers (who are likely to receive requests from that client in the near future, thanks to the WOMsprayer) pick it up from a fileset that is local to their physical site.

The CRC digits are built into the Cookie. This makes it more difficult for anyone to tamper with a Cookie while trying to obtain profile information about another user by reverse engineering their Cookie. Cookies with invalid checksums are not accepted by the WOM system.

Applications

This section describes applications that were designed and built for the Atlanta Web site. For various reasons, not all of these went into full production on the site. They were all working in the lab, so we have included them here for the sake of completeness.

Sneak Peek

Sneak Peek was the largest "Net Cam" ever. Live video feeds from 38 locations at the Atlanta Games were fed to "The Fishbowl" in Atlanta. Here they went through a bank of "frame grabbers" and were checked into the WOM system. Images were captured approximately every five seconds from each feed. Because of the high volatility of this data, the images were not stored into the WOM DB2 database, but went directly into DFS. HTML pages to render the latest set of images and to navigate between Olympic venues were generated on-the-fly to pick up the most recent image file names.

The human operator in the Fishbowl received all feeds concurrently and could select channels that currently had activity on them. The navigation options reflected the choices presented by the dynamic HTML pages in the Sneak Peek section of the site. The operator could choose the most stunning images and build a daily "scrapbook" of images worth keeping. These images were titled, cropped, assigned metadata using the WOMpix tool, and checked into the main WOM database.

Sneak Peek images were not mirrored across the world because of their highly volatile nature—they were all served from the primary WOMplex site at Southbury. This became an extremely popular application, and many hundreds of thousands of images were downloaded.

Sneak Peek images were all served from the primary WOMplex site at Southbury.

Bamba

Bamba is IBM's streaming audio and video technology for the Web based on G.723 low bit-rate streaming standards. For more information about the technology and to download the decoders and encoders, visit the Bamba home page on the IBM AlphaWorks™ Web site (http://www.alphaworks.ibm.com/).

Bamba has three advantages over other Internet audio and video formats:

- Data is streamed, so you can view and download it simultaneously, avoiding huge waits while the whole file downloads.
- The content is encoded at a low bit rate, making it suitable for transfer over a standard 28.8 Kbit modem.
- G.723 is a standard.

Three forms of Bamba content were used at the Atlanta Olympics:

- **Audio clips:** Recordings of many interviews with athletes, Olympics organizers, and news reports. A regular Web server handles Bamba audio clips in the same way as GIF images. On the Atlanta site, they were one data type served by the binary server.
- **Video clips:** Encoded highlights of the sports events. Bamba video clips are also regular binary content that can be served by an ordinary Web server. They were served from the binary server during the Olympics.
- **Live audio:** Atlanta local radio station, WGST Atlanta, was broadcast over the Internet 24 hours a day during the Olympics. Bamba

Reflector allowed many people to listen to the radio station concurrently.

Bamba Reflector

In typical use, a small Pentium™ PC would do real-time Bamba audio encoding. It is not realistic to have hundreds or thousands of clients connecting to this machine to listen to the audio stream, so Bamba Reflector was created. The Reflector connects to the transmitter as the only client to that machine, it then accepts connections from Bamba clients. Whatever the Reflector receives from the transmitter is transferred to the clients. There must be a certain amount of buffering and flow-control to deal with clients connected over different speed links. The Reflector intelligently drops Bamba packets without losing stream synchronization for any clients who fall too far behind in receiving Bamba packets. Each Reflector can support many hundreds of client connections concurrently. One Reflector can act as a source to other reflectors, so a whole world-wide tree of Reflectors can potentially serve many thousands of clients across the globe.

Users could specify to the WOMserver the specific sports, countries, and athletes of interest to follow as the Games progressed.

The Reflector configuration can limit connections to a specified upper limit or it can be off-line—not broadcasting at the moment. When a connection is not possible, it can send out an information message in one of three forms: a pre-recorded audio message built into the client, a text message sent from the Reflector, or an audio message sent from the Reflector. Various combinations are also possible.

ISS load balancing can be used in front of a bank of Bamba Reflectors to balance incoming connection requests to a single advertised Internet name across the available servers to ensure an even load across the Reflectors.

WOMplex Boiler Room

The Boiler Room monitoring system allowed statistics to be gathered from various parts of the WOMplex, its applications, and subsystems. The system structure is described below.

Each application to be monitored had a `funneld` (`funnel` daemon) client. To make this a simple task, developers had client code for Java, Perl, and C available. The protocol used throughout this system was Internet Relay Chat (IRC), transmitting data in a textual format with each monitored system using a different IRC channel number. The client code connected to a `funneld` instance ran on each server. The `funneld` instances connected to other `funnel` daemons in a tree structure. An IRC server was at the root of the tree of `funnel` daemons.

The Boiler Room monitors were Java applets embedded in HTML pages on the WOMplex Web site. These applets used IRC client classes written in Java that connected back to the IRC server. The monitor applications then tuned in to the IRC channels required to construct their displays and waited for information to arrive. A variety of dials, counters, and indicators were implemented as Java classes that could be combined in various ways to display the statistics for any system.

Since the Boiler Room system could also tail log files and receive data through sockets and Interprocess Communications (IPC) message queues, it was easy to hook up monitors for almost anything. A special *chart recorder* client logged IRC channel data for subsequent analysis.

Information Taxonomy

Early in this project we realized that a very large site needs some kind of information taxonomy, or hierarchy, mapped to the site; otherwise, it becomes completely unmanageable and almost impossible to navigate. This important issue required the skills of an information architect.

The metadata for every page, image, and multimedia clip checked into the WOM database included a set of category tags. These tags position the object in the *information space* of the Web site. This information space is dimensionally exclusive from the *data space* imposed by the HTML hyperlinks built into the documents. Tools, such as the Wu-tool, categorize pages on the site. Once the whole site has been categorized, facilities can be provided to navigate the site in information space, as well as the traditional HTML navigation. For example, someone looking at pages about Men's Canoeing might also be interested in other pages that are "close" in the information space. For example, in this case, Kayak events might be of interest.

As an interesting new area for Web site navigation and management, this will undoubtedly play an important role in the future.

Personal Home Page

The combination of generating dynamic pages in the WOMserver and identifying individual clients from their browser Cookie allowed a personal home page to be created for each visitor to the site.

Users could specify to the WOMserver the specific sports, countries, and athletes of interest to follow as the Games progressed. If they provided

IBM 000035

this information, they could then request their personal home page (/myhome.html) from the WOMserver. The Home Page application would build a personalized page for that user, including highlights of the day for the sports, countries, and athletes of interest. This was coupled with Activist, which could suggest links of interest for the user, based on previous browsing history.

Activist

Activist is an online Web log data mining tool. Activist was a DB2 client application connected to the DB2 server running on the CMOS 390 parallel sysplex. This machine was amassing HTTP log data from all the servers in the WOMplex via the Grim Reaper. By following the click-track of each unique client (based on their browser Cookie), Activist would compile a "persona" for that client. Once this was established, it was possible to ask questions of Activist, such as "Based on my browsing patterns in the site so far, show me some links that I might find of interest that I haven't seen yet," or "Show me some links that I might like to see, based on where others like me have been."

The phrase "like me" in this context is derived from a data mining technique called *segmentation*. Segmentation allows a data set to be partitioned into several clusters of individuals who show a statistically significant similarity in certain areas of the data. In this application, if you had looked at Canoeing and Cycling, and many others had looked at Canoeing, Cycling, and Basketball, then Activist would probably offer you some Basketball links to view.

This powerful technology will prove useful in making sense of the vast amounts of log data accumulated by every Web server in the world.

Country and Sport Pages

As information such as results, news reports, news articles, news wire feeds, and photographs flowed into the site during the Games and checked into the WOM database, each page was categorized into its appropriate position in the site's information taxonomy. Every few hours, a batch job was run to compile a home page for every country participating in the games and for each sport. This home page contained links to everything recently added to the site relating to that country or sport. This provided an easy way to follow the progress of a particular country or sport and to catch up on recent happenings in that area.

Guestbook, Search, Feedback

For the Atlanta site, these standard applications were built in the WOMserver application environment and run as dynamic extensions to the server. The Guestbook and Feedback applications made a client/server connection to a remote machine running DB2 that stored the information submitted by the user. Other applications provided a password-protected Web interface to these databases to allow the information to be extracted and analyzed.

Postcards

The Postcards application allowed users to "send a postcard" of any image on the Atlanta site to someone else on the Internet using E-mail. When a full-size version of an image was requested, usually by clicking on its thumb-nail image, one option on the navigation bar of the page containing the full-size image was to "Send a Postcard."

With this option, the user could specify the recipient's E-mail address, a message to go with the Postcard, and the sender's identity. The recipient would then receive an E-mail message indicating that a "Postcard from Atlanta" was waiting. By connecting to the URL given in the E-mail message, the recipient would go to a page that was generated dynamically, based on a unique identifier that was part of the Postcard URL. The page would contain the selected image, the message from the sender, and the sender's name.

Eye on the Olympics

"Eye on the Olympics," a Java applet, provided a way to "keep an eye" on what was happening at the Olympics, without having to continuously surf around the site. The applet used IRC to connect to a server at the WOMplex site. It presented a set of icons: one for each sporting discipline. The client displayed a moving ticker-tape of news flashes. This ticker-tape had an area that could be reached by clicking on a sports icon, which contained results and news stories relating to that sport.

The server used the metaphor of a television station with several "programs"—one per discipline—each created by an automated producer. The producer's job was to choose pieces of information being fed into the Web site and assemble them in a form for presentation in the Eye on the Olympics. The server sent out a *play list* of current articles, then sent new articles that had been added since the last play list was

Segmentation allows a data set to be partitioned into several clusters of individuals who show a statistically significant similarity in certain areas of the data.

broadcast. The client applet received the play list and modified its display behavior accordingly, adding new items as they were transmitted and deleting old items that were no longer in the play list. The producers relied heavily on the information taxonomy to determine which new items being checked into the site were of interest to their programs.

The client applet used HTTP to retrieve non-text items, such as images, from the main Web site, using a URL sent over IRC. The applet also used HTTP to retrieve a file containing the initial set of current content when the applet was first invoked.

Conclusion

The Atlanta Olympics Web site was the largest Web site in the world and handled an immense amount of traffic during the 1996 Summer Games. The WOMplex technology was pioneered to support large sporting events on the Internet. The Olympics gave the WOMble Team a chance to push forward the technology frontiers, and has paved the way for products using WOM technology in the future.

The WOMserver was designed as an Internet application platform, of which a sporting event server was just an application. The suitability of this architecture as an application environment has been proven in the "trial by fire" experienced during the Olympics.

This experience greatly advanced load balancing and global scalability. The systems deployed for the Olympics showed that a Global WOMplex is a viable concept.

The author gratefully acknowledges the contributions of Sean Martin and David Grossman to this article, and also wishes to thank the Team at the IBM Southbury lab.



***Andy Stanford-Clark**, IBM Corporation, 150 Kettletown Road, Southbury, CT 06488. Internet: andysc@vnet.ibm.com. Dr. Stanford-Clark is an Internet advanced technology consultant in IBM's Internet Division. His responsibilities include exploring emerging Internet technologies and working with customers to implement Internet-based projects. Dr. Stanford-Clark is currently in the U.S. on international assignment from IBM Hursley, U.K. He has a BS in Computing and Mathematics and a PhD in Parallel Computing from the University of East Anglia in the U.K.*

GO TO NEXT STORY | BACK TO CONTENTS

C6

Locating Objects in Wide-Area Systems

Maarten van Steen (contact)

Vrije Universiteit, Department of Mathematics & Computer Science
De Boelelaan 1081a, 1081 HV, Amsterdam, The Netherlands
tel: +31 (0)20 444 7784, fax: +31 (0)20 444 7653
e-mail: steen@cs.vu.nl

Franz J. Hauck

University of Erlangen–Nürnberg, IMMD4
Martensstr. 1, D–91058, Erlangen, Germany
e-mail: hauck@informatik.uni-erlangen.de

Philip Homburg
Andrew S. Tanenbaum

Vrije Universiteit, Department of Mathematics & Computer Science
De Boelelaan 1081a, 1081 HV, Amsterdam, The Netherlands
e-mail: {philip,ast}@cs.vu.nl

Locating Objects in Wide-Area Systems

Maarten van Steen, Franz J. Hauck,
Philip Homburg, Andrew S. Tanenbaum

Abstract

Locating mobile objects in a worldwide system requires a scalable location service. An object can be a telephone or a notebook computer, but also a software or data object, such as a file or an electronic document. Our service strictly separates an object's name from the addresses where it can be contacted. This is done by introducing a location-independent object handle. An object's name is bound to its unique object handle, which, in turn, is mapped to the addresses where the object can be contacted. To locate an object, we need only its object handle. We present a scalable location service based on a worldwide distributed search tree that adapts dynamically to an object's migration pattern to optimize lookups and updates.

1 Introduction

In the near future we can expect hundreds of millions of users to have access to a global Information Superhighway. A large part of that information network will be mobile: telephones, faxes, notebook computers, personal assistants, etc. But we can also expect software and data to be mobile. For example, a Web page may move as its owner changes computers; likewise, a shared electronic document may travel between its users. Another example is a mobile agent that moves through the network in search of specific resources for its owner. Components in a network capable of changing locations, and which may be implemented in software, hardware, or a combination thereof, are collectively referred to as **mobile objects**.

Supporting mobile objects means that a client should be able to contact an object even if he does not know its current location. Moreover, locating the object should be completely hidden from the client. For example, in Personal Communications Systems (PCS), a user should only have to dial a telephone number to contact the callee. It should not be necessary to know the callee's present location or how the callee is tracked. But mobile objects also need to contact other (possibly nonmobile) objects. When a mobile object moves to a new location, the object will have to find out which facilities it can use there. For example, a mobile computer may need to use the local printer. Likewise, it may want to contact the local Web server instead of having to use the server at its home location.

Being able to contact objects, whether they are mobile or not, is traditionally supported by a **naming service** which maintains a **binding** between an object's name and one or more addresses where the object can be contacted. As an analogy, a naming service is like a telephone book; a

binding corresponds to one of its entries. With mobile objects, names should always be resolved to a current address. To illustrate, a name such as *pcs://dept.univ.edu/Mary* may be dynamically bound to the network address of Mary's mobile computer. No matter where in the world that name is used, it should always be resolved to her computer's *current* address, which changes as she moves. In addition, applications on her mobile computer may use the name *local://usr/addr/lpr* for the local printer. In this case, as Mary travels around the world, the printer's name on her computer needs to be dynamically rebound to the address of the nearest printer server. Thus, unlike the world of cellular telephony with their fixed bindings of device to telephone number, in the computer world, the addresses used to reach objects change as an object moves, and the mapping of names to addresses must also change.

Changing the address of an object affects the name–to–address binding. If such changes hardly ever occur, then constructing a worldwide scalable naming service is feasible, as demonstrated by the Internet's Domain Name System [8] and the X.500 Directory Service [10]. However, if bindings change frequently, as in the case of mobile objects, we have a much more difficult problem.

In this paper, we focus on a wide-area naming service that provides flexible and easily adaptable name–to–address bindings. The service is currently being developed as part of Globe, an object-based worldwide distributed system aimed to support a billion users each having thousands of objects [3]. The paper is organized as follows. In Section 2 we explain and motivate the basic architecture of the Globe naming service. The main goal of this paper is to explain how objects are located, which is described in Section 3. Related work is discussed in Section 4. We give our conclusions in Section 5.

2 Binding Names to Addresses

To discuss name–to–address binding in wide-area systems, we assume that all (hardware and software) objects have symbolic ASCII names, such as *pcs://dept.univ.edu/Mary*. Also, each object is assumed to have one or more addresses where a client can contact it. By way of analogy, an owner of a cellular telephone is also assumed to have a name which can be registered in a telephone directory. The owner's telephone number corresponds to the address where he can be reached. Unlike cellular telephones, computer objects often have two or more addresses. For example, a replicated file will be known to its users under one name, which is mapped to several addresses, one for each copy. A user asking for the file generally does not care which copy is selected.

To make name resolution efficient for wide-area systems, names often contain location information. For example, the Uniform Resource Locator (URL) *ftp://ds.internic.net/nic/rfc/rfc1737.txt* is the name of a Web page containing the text of RFC 1737. The name reflects where the page is stored (*ds.internic.net*), allowing part of the name resolution process to take place at that location. Similarly, telephone numbers also contain location information: a worldwide number like *+31 20 444 7784* gives the country (31 is The Netherlands), the city (20 is Amsterdam), and a specific telephone exchange (the 444 exchange of the Vrije Universiteit in Amsterdam–Buitenveldert).

However, using location information in names can make it difficult to handle migration. If an object moves, we may have to change its name, or otherwise make that name become a forwarding reference. In wide-area systems, the latter can lead to long chains of references, which

are inefficient and susceptible to network failures. What is needed is a naming facility that hides all aspects of an object's location. Users should not be concerned where an object is located or whether it can move.

These requirements can be met by introducing a two-level naming hierarchy as shown in Figure 1. The first level deals with hierarchically-organized, user-defined name spaces. These name spaces are handled by what we call a distributed **object naming service**. A name is bound to an **object handle**, which is a globally unique and location-independent object identifier.

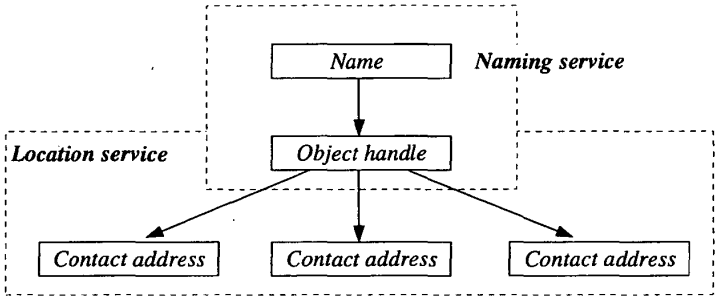


Figure 1: A two-level naming hierarchy that allows an object's name and contact addresses to be independently changed.

The second level deals with mapping each object handle to a set of contact addresses, and is handled by a distributed **object location service**. In contrast to traditional naming services, a location service is designed to support frequent updates and lookups of contact addresses such as needed for mobile objects. It is not concerned with naming.

As an illustration, consider an officeless company whose employees are located across the country, normally working at home, or visiting customers. Using our approach, we assign a lifetime and location-independent telephone number to the company. This number corresponds to an object handle. A naming or directory service would maintain a mapping between the company's name and it's lifetime telephone number. The telephone number is used by a location service to redirect incoming calls to, for example, the nearest employee currently working. An employee's own telephone number is registered at the location service when he or she starts work, and is unregistered again when he or she finishes. An employee's telephone number corresponds to a contact address. Note that how the company is named, and in which directories its name is registered, is no longer important. Naming has been fully separated from how and where we contact the company.

In distributed systems, this way of locating a service is also known as **anycasting**: a client requires a particular service, but is really not interested which server will handle the request. Using our approach, the service is assigned a unique object handle, and each server registers its network address under that object handle. A client has the service's name resolved to the object handle, which is then subsequently resolved to the address of any server that can handle the request

An object handle is designed specifically for the location service. It contains a globally unique *service-independent object identifier* which is very similar to a UUID in DCE [11]. Additionally,

an object handle may contain information that can be used to assist in locating the object. For example, it may be known that an object will move only within a certain region. Instead of conducting a global search for such an object, it is more efficient to start the search process in that specific region. Therefore, it makes sense to encode this information into the object handle.

Requirements for a location service

Clearly, our two-level approach makes sense only if we can indeed provide a scalable and efficient naming and location service. The feasibility of developing scalable naming services has been demonstrated by systems such as DNS. This is not yet the case for location services, for which the following requirements will have to be met:

Scalability. The service should allow clients and objects to be located anywhere in the world, and be able to support a huge number of objects. We anticipate that eventually, one billion users with 1000 objects each will be registered with the location service, adding up to 10^{12} objects.

Locality. Assuming that the cost for looking up an address generally increases with the length of the route to that address, we require that the location service exploits locality. This means, for example, that if an object has its address near to the client, finding the object should be fairly cheap.

Stability. Objects may differ with respect to their migration patterns. For example, a Web page may possibly move through the entire network in a seemingly random way, whereas a mobile computer may possibly move only within a city. An object is said to be **stable** with respect to a region, if its addresses most often are in that region. If an object is stable with respect to a region R, we require that searching or updating one of its addresses in R is cheaper than when the object is not stable with respect to R.

Fault tolerance. The location service should be resilient to node and link failures and should continue to operate in the presence of network partitions. The service should, at the least, degrade gracefully in terms of performance and functionality.

Location services are not new and have shown to be relatively easy to implement in local distributed systems. However, they become much more complicated when scalability is taken into account as we discuss next.

3 Tracking Distributed Objects in Globe

In this section we discuss the architecture of Globe's scalable location service. We shall provide only an outline of the architecture, further information can be found in [13].

Basic operations

In our model for tracking objects, we assume a hierarchical decomposition of a (worldwide) network into regions. This decomposition is relevant to only the location service. With each region we associate a **directory node**, capable of storing addresses that lie within that region. This leads to a logical tree-based organization as shown in Figure 2. Addresses are assumed to be location dependent: the region in which an address lies is encoded in the address itself.

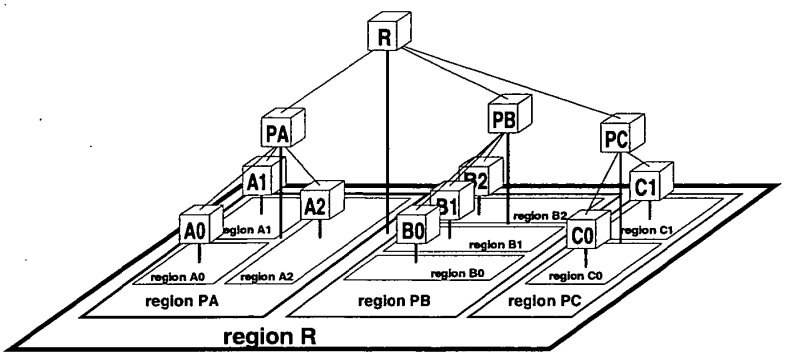


Figure 2: The logical organization of the location service as a virtual search tree.

The location service normally stores new addresses at the leaf node representing the region in which the address lies. For each new object, it constructs a path of forwarding pointers from the root to each leaf node where an address is stored. Addresses and forwarding pointers are stored in **contact records**. An implication of this design is that in the worst case, it is always possible to locate every object by following the chain of pointers from the root node. In practice, we can do much better than this, as described later.

In principle, a request for insertion of a previously unregistered object begins at a leaf node and is propagated up the tree to the root. Then, a path of forwarding pointers is established from the root to the leaf node where the insertion takes place. A contact record containing a forwarding pointer is created at each intermediate node. The address itself is finally stored only in the leaf node. When a part of the path already exists, for example, when inserting a second address in a different region, only the missing pointers are established. This is shown in Figure 3. In the case that there is already a contact record for the object at the leaf node, the new address is simply added to that record.

Deleting a contact address is straightforward and is done as follows. First, the address is found through a search path up the tree, starting at the leaf node representing the region in which the address lies (note that this information is encoded in the address). Once the contact record in which the address is stored, has been found, it is removed from that record. If a contact record no longer contains contact addresses or forwarding pointers, it is deleted. The parent directory node is informed that it should delete its forwarding pointer to that record, possibly leading to the (recursive) deletion of the object's contact record at the parent node.

Looking up a contact address is done as follows. A client process passes an object handle to

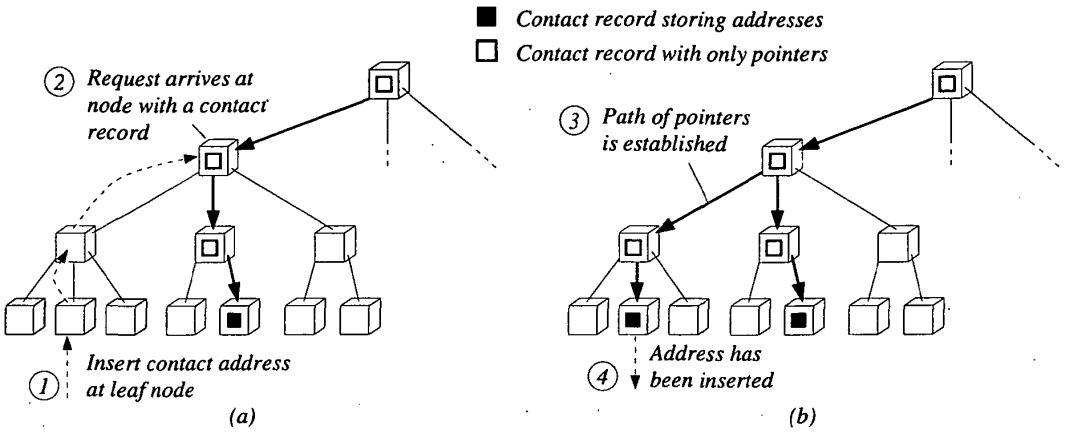


Figure 3: Inserting a contact address when the object is already known. Only the missing pointers are established.

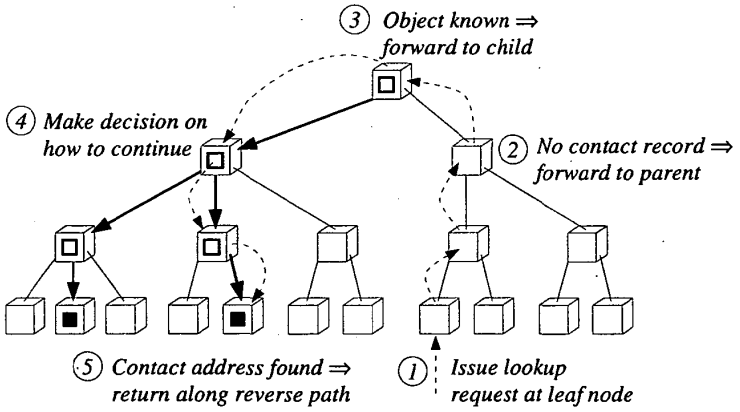


Figure 4: The default approach for looking up a contact address.

the leaf node of the region where that process resides. (We require that there is exactly one such leaf node.) As shown in Figure 4, a search path is established starting at the client's leaf node, and going upwards to the first directory node where the object is already known. In the worst case, this means propagating the request up to the root. The path then continues downwards to a leaf node, whose addresses are then returned to the requester.

Concurrent update and lookup operations are allowed, although no ordering is guaranteed when requests are issued at different nodes. In particular, we have the following consistency rule:

> **Request Consistency**: Update requests issued at the same leaf node are completed in the order they were issued. Update operations issued at different leaf nodes are completed in an arbitrary order.

Dynamic optimizations

The location service has full control over the placement of addresses in contact records. Consequently, if we can place addresses in stable locations, we can make effective use of pointer caches during lookup operations. By default, an object's address is stored in its contact record at the leaf node where it was initially inserted. Now, consider some region R as shown in Figure 5, and assume that an object O is changing its addresses regularly between the subregions S_1, S_2, and S_3. For simplicity, assume that there is always at least one address somewhere in R, so that there will always be a nonempty contact record for O at directory node $dir(R)$.

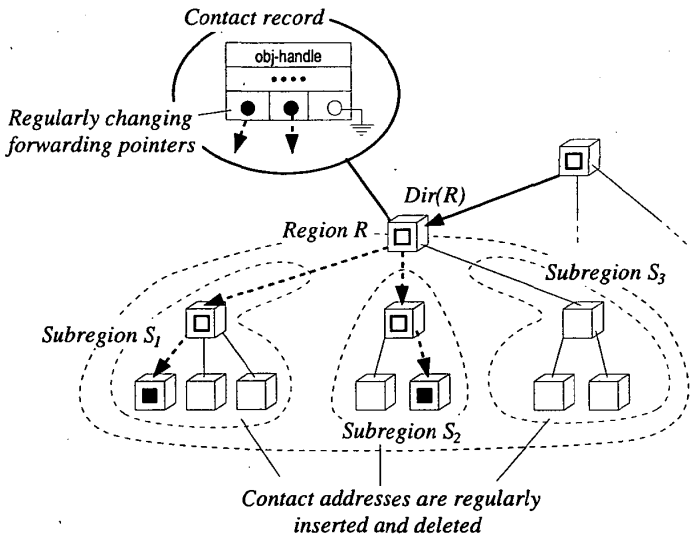


Figure 5: The situation of an object regularly moving between subregions. The solution is to store its changing address in $dir(R)$.

Each time the object moves to a new subregion S_k, the location service creates a path of forwarding pointers from $dir(R)$ to a leaf node in S_k. Likewise, when moving out of S_k the path has

to be deleted. If migration occurs regularly, it makes sense to store the address in the object's contact record at $dir(R)$. This not only saves the cost of path maintenance, but more important is that addresses from any of the subregions are now stored at a stable place, namely at the directory node $dir(R)$.

By storing addresses in stable contact records, our model leads to the construction of a search tree *per object*, in which contact records tend to remain in place, even for mobile objects. This permits us to effectively shorten search paths by caching *pointers* to contact records, since they are stable, even if the corresponding objects are not. Specifically, a pointer to the directory node with a contact record containing addresses, is cached at each node of the search path when returning the answer to the leaf node where a lookup request originated.

The combined effect of pointer caches and stable contact records should not be underestimated. An object that moves primarily within a region R can be tracked by just two successive lookup operations: the first one at the leaf node servicing the requesting process, and the second one at the directory node for region R. Moreover, our solution forwards a request in the direction of an address. This is a considerable improvement over existing approaches.

Of course, the migration behavior of an object may change. For example, assume the contact record at $dir(R)$ has contained an address for subregion S_k for quite some time. In that case, the address will be propagated to a directory node in S_k, because apparently, stability occurs in a smaller region than R. Stability is measured by timestamping addresses and forwarding pointers, as well as recording how long an object has not been in a specific region. In all cases, history is taken into account by weighted accumulation of old and new timing information.

Fault tolerance

Any service available in a wide-area distributed system should mask failures of the underlying network to its clients. This is not always possible, because, for example, network partitions may last for hours. Fault tolerant behavior for our location service can be partly expressed in terms of the following informal progress rules:

> **Global progress:** The effect of an update operation initiated at an arbitrary leaf node is eventually visible to a lookup operation initiated at any of the leaf nodes of the search tree.
>
> **Local progress:** The effect of an update operation at a directory node D should be immediately visible at each directory node in the (connected) subtree rooted at D.

The second rule states that an update operation should come into effect at a particular directory node as soon as it is issued at that node, and perhaps before the operation completes. The completion of an operation may depend on the response of the parent node, which may be temporarily unreachable. In other words, update operations are to be handled in an asynchronous fashion. Satisfying the progress rules effectively means that our service is resilient to nodes being unreachable, either caused by network partitioning or because a node has crashed.

To handle operations asynchronously while satisfying the rule for request consistency, incoming requests are first appended to a *queue* and subsequently forwarded to the parent node. Queued

operations at a particular node are evaluated in the order they have been appended, and yield a *tentative* result. As soon as the parent agrees with the update, the operation is completed by removing it from the queue and making its result authoritative. If the parent disagrees with the update, for example, when a new address should be stored at a higher-level directory node, the operation is only removed from the queue.

There are several benefits to this approach. First, queuing operations allows us to maintain a consistent, albeit tentative view of the data maintained at a node without having to block any requests until the associated operation is fully completed. This makes the location service resilient to network partitions.

Recovering from node crashes is harder, and is subject to further research. However, a crashed node can easily re-invoke incomplete operations by having its children re-issue their requests, although this is clearly not enough to restore the node's original data set. This approach is very similar to sender-based message logging as discussed in [4], or using queued RPCs as in the Rover toolkit [5]. By replicating authoritative data among the directory nodes in each subtree, the location service can be made resilient against a single node failure in each path originating at the root.

Scalability

Clearly, to construct a worldwide scalable location service it is necessary to adopt hierarchical solutions. For example, nonhierarchical solutions such as the read/write sets proposed in [9] will not do, as it is much harder to exploit locality. However, our search tree described so far obviously does not yet scale. In particular, higher-level directory nodes not only have to handle a relatively large number of requests, they also have enormous storage demands. Our solution is to partition a directory node into one or more **directory subnodes**, such that each subnode is responsible for a subset of the records originally stored at the directory node.

As an example, we can use the first n bits of an object's handle to identify the subnode responsible for that object. Subnodes of a particular directory node need not communicate with each other since they maintain different subsets of objects, and all operations are performed on a per-object basis. Communication between directory nodes in the original search tree takes place only between their respective subnodes. To illustrate, Figure 6 shows a search tree in which the root node has been partitioned into four subnodes based on the first two bits of the object handle ($n = 2$), and each of the leaf nodes into two subnodes ($n = 1$). (We note that we have developed more realistic hashing-based methods than explained here. For example, we also have to take into account that the total number of links between parent and children subnodes remains manageable.)

As communication between directory nodes in the original search tree now takes place between their respective subnodes each subnode should be aware of how the directory node with which it communicates is actually partitioned. This information is contained in a separate tree management service. This service also maintains the mapping of subnodes to physical nodes. Partitioning and mapping information is assumed to be relatively stable, so that it can be easily cached by subnodes. This assumption is necessary to avoid having to query the management service each time a subnode needs to communicate with its parent or children, which would turn the

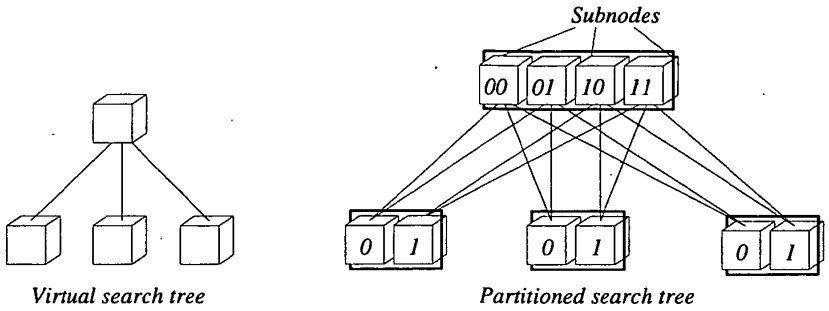


Figure 6: A search tree and a corresponding logical tree after partitioning the directory nodes into subnodes.

management service into a potential communication bottleneck.

4 Related work

Location services are becoming increasingly important as mobile telecommunication and computing facilities become more widespread. So far, mobility is almost invariably connected to *mobile hosts*. A characteristic feature of these hosts is that their mobility is directly coupled to that of their user. This has two important consequences that do not apply to our location service. First, the speed of migration is limited to the maximum speed at which a person can move (typically 1000 km/hour in an airplane), making it possible to adopt a strategy in which data structures gradually adapt as the object moves. This has been successfully applied to several models for location services (see e.g. [1, 7]), but these techniques cannot be used in our case. Second, a host is always at precisely one location. There is no notion of multiple addresses per object as we have introduced in our model. In contrast to current approaches, we can effectively deal with replication.

The situation becomes entirely different when dealing with *mobile software objects*. An important distinction with keeping track of mobile hosts, is that there are many more objects than hosts, immediately leading to a scalability problem. Also, to ensure scalability, it is necessary to take mobility patterns of an individual object into account as well. In Emerald [6], mobile objects are tracked through chains of forwarding pointers, combined with techniques for shortening long chains, and a broadcast facility when all else fails. Such an approach does not scale to worldwide networks. An alternative approach to handle worldwide distributed systems is the Location Independent Invocation (LII) described in [2]. However, LII uses a global naming service as a fallback mechanism, where it assumes that the update–to–lookup ratio is small. Designing a global location service that is not based on such an assumption is an important goal of our research.

A seemingly promising approach that has been advocated for large-scale systems are SSP chains [12]. SSP chains allow object references to be transparently handed over between processes, at the expense of gradually constructing a chain of forwarding references to the object. A serious drawback of this approach is that exploiting locality is completely neglected. In particu-

lar, a *home* must keep track where the object is during its entire lifetime, possibly years after the object last lived there. Also, it is unclear how fault tolerance can be efficiently dealt with. SSP chains therefore do not seem to scale to worldwide systems.

Contributions of our approach

One of the main advantages of our approach is that our location service can handle objects that have several contact addresses and that show arbitrary migration patterns. We do not adapt *update and search strategies* to migration patterns, but adapt *the search tree on a per-object basis* instead. By registering contact addresses in the smallest region in which (part of) the object is moving we can make effective use of pointer caches. The combined effect is an extremely short search path, in the optimal case of only length two, from a client to the object. In just two hops it is possible to locate even seemingly randomly migrating objects. This is a considerable improvement over existing approaches.

The use of pointer caches instead of data caches has also been proposed for Personal Communications Systems (PCS). The main reason to apply caching in those cases is to avoid excessive network traffic to the *home location* of a host, which forms the root of a two-level search tree such as in GSM. Caching is done at the second level, by pointing to locations where the host is expected to be found. Cache consistency is achieved either by invalidation on demand, or through active updates. However, caches in PCSs do not account for update patterns. A distinctive feature of our approach compared to PCSs, is that we have several levels allowing us to exploit locality more effectively by inspecting succeeding expanded regions at linearly incrementing costs. On the other hand, locality is also exploited in location updates, making our pointer caches highly effective.

5 Conclusions

The Globe location service offers a novel approach to locating objects in a worldwide context, using location-independent object identifiers instead of user-defined names. Our approach allows an object to update its contact addresses independent of how users have named the object. Locating an object proceeds through a worldwide search tree which dynamically adapts itself on a per-object basis. By storing addresses at stable locations and subsequently caching pointers to those locations, it is possible to contact an object in just two steps, irrespective of the object's migration pattern. An important distinction with current approaches, is that the search path is directed toward the object's present location.

Presently, our research is continuing in two directions. First, we are building a prototype implementation for experimentation and validation of our approach. Second, we are enhancing the basic algorithms described in this paper to include fault tolerance. Also, algorithms for selecting the most appropriate nodes for storing an object's contact addresses need to be further improved.

References

[1] B. Awerbuch and D. Peleg. "Online Tracking of Mobile Users." *J. ACM*, 42(5):1021–1058, Sept. 1995.

[2] A. Black and Y. Artsy. "Implementing Location Independent Invocation." *IEEE Trans. Par. Distr. Syst.*, 1(1):107–119, Jan. 1990.

[3] P. Homburg, M. van Steen, and A. Tanenbaum. "An Architecture for A Scalable Wide Area Distributed System." In *Proc. Seventh SIGOPS European Workshop*, pp. 75–82, Connemara, Ireland, Sept. 1996. ACM.

[4] D. Johnson and W. Zwaenepoel. "Sender-Based Message Logging." In *Proc. 17th Annual International Symposium on Fault-Tolerant Computing*, pp. 14–19, Pittsburgh, PA, July 1987. IEEE.

[5] A. D. Joseph, J. A. Tauber, and M. F. Kaashoek. "Mobile Computing with the Rover Toolkit." *IEEE Trans. Comput.*, 46(3):337–352, Mar. 1997.

[6] E. Jul, H. Levy, N. Hutchinson, and A. Black. "Fine-Grained Mobility in the Emerald System." *ACM Trans. Comp. Syst.*, 6(1):109–133, Feb. 1988.

[7] P. Krishna, N. Vaidya, and D. Pradhan. "Location Management in Distributed Mobile Environments." In *Proc. Parallel and Distributed Information Systems*, pp. 81–88. IEEE, 1994.

[8] P. Mockapetris. "Domain Names - Concepts and Facilities." RFC 1034, Nov. 1987.

[9] S. Mullender and P. Vitányi. "Distributed Match-Making." *Algorithmica*, 3:367–391, 1988.

[10] S. Radicati. *X.500 Directory Services: Technology and Deployment.* International Thomson Computer Press, London, 1994.

[11] W. Rosenberry, D. Kenney, and G. Fisher. *Understanding DCE.* O'Reilly & Associates, Sebastopol, CA., 1992.

[12] M. Shapiro, P. Dickman, and D. Plainfossé. "SSP Chains: Robust, Distributed References Supporting Acyclic Garbage Collection." Technical Report 1799, INRIA, Rocquencourt, France, Nov. 1992.

[13] M. van Steen, F. Hauck, and A. Tanenbaum. "A Model for Worldwide Tracking of Distributed Objects." In *Proc. TINA '96*, pp. 203–212, Heidelberg, Germany, Sept. 1996. Eurescom.

Multilevel Filesystems in Solaris Trusted Extensions

Glenn Faden
Sun Microsystems
17 Network Circle
Menlo Park, CA, USA
1-650-786-4003
glenn.faden@sun.com

ABSTRACT

Multilevel security is typically implemented by assigning fine-grained security contexts, such as sensitivity labels to all subjects and objects. These extended security contexts require modifications to standard filesystems, and interfaces that affect system throughput and application compatibility. This trade-off between policy enforcement and performance tends to marginalize these systems to special-purpose environments. This paper describes a light-weight approach which avoids the requirement for customized filesystems or modified applications. Instead, the system is partitioned into labeled zones. Subjects and objects are associated with these zones from which they inherit their sensitivity labels. This structured approach to data separation makes it possible to implement mandatory access control on a mainstream operating system.

Categories and Subject Descriptors

D.4.3[**Operating Systems**]: File Systems Management, - *file organization*

D.4.4[**Operating Systems**]: Security and Protection – *access controls*

General Terms

Design, Security

Keywords

multilevel security, mandatory access control, file labeling

1. INTRODUCTION

In both corporate and government organizations, policies exist to limit access to information. A first step in establishing any access control policy is to evaluate the risks that are inherent in unauthorized disclosure of various types of data. Data labeling has been used for decades[12] to specify the sensitivity of information, and the risk of its disclosure. Access control policies can be expressed using these sensitivity labels to specify valid data flows. For example, a policy may state that data marked *Confidential* should not be distributed outside the organization.

Copyright © 2007 Sun Microsystems, Inc..
SACMAT'07, June 20–22, 2007, Sophia Antipolis, France.
ACM 978-1-59593-745-2/07/0006.

When these policies are advisory in nature, they are referred to as Discretion Access Control (DAC). However, when such policies are enforced automatically and cannot be circumvented, they represent Mandatory Access Control (MAC). Two examples of MAC policies are Multilevel Security (MLS) and Domain Type Enforcement (TE)[5]. Both approaches use labels to partition objects and to define the rules for the flow of data.

Systems that implement MLS and TE have historically done so by attaching security contexts to all of the resources which they protect (objects) and to the users or processes which act on them (subjects). These security contexts are used by policy managers to determine whether a given subject has access to the objects. For example, when one label dominates another, data at the lower level can be read by subjects at the higher-level.

In order to uniquely label each object, storage must be allocated to retain the labels, and interfaces must be developed to maintain the label associations. Since these labels are used for mandatory access control, they must be rigorously protected from unauthorized modification. As required by the mandatory access policy, even the owner of the file is typically not authorized to change these labels. Only a special administrative role, such as an information security officer, is authorized to reclassify labeled data. Protecting a file's security context is at least as important as protecting the actual data in the file. This approach tends to require customized filesystems since standard filesystems don't provide the necessary protection model.

Modern commercial filesystems implement some kind of DAC policy which allows the subject who created the file (the owner) to restrict access to the file. But these filesystems don't have a notion of security beyond that associated with ownership. Therefore, operating systems that implement MAC have generally provided customized filesystems to support mandatory security contexts. For example, the Trusted Solaris™ OS extended the inodes of its UFS filesystem to support labels and other attributes[6, 11]. SELinux uses Extended Attributes (EA) for security contexts. According to J. Morris[7], for a filesystem to support SELinux security context labels, it needs EA support and a handler for the EA security namespace. Such filesystems currently include ext3, ext2, XFS and ReiserFS.

This paper describes a new approach in which a mandatory MLS policy is applied automatically, so that it can be used with any existing filesystem. This transparent approach has recently been integrated into the Solaris 10 OS and its open source counterpart, OpenSolaris[10]. The new MLS policy is enforced by enabling

Solaris Trusted Extensions[4]. The motivation for this transparent approach was to provide a single commercial OS with an optional MLS policy, that would support the full range of applications, patches, and hardware platforms.

2. GRANULARITY OF THE POLICY

A successful security policy should do more than prevent unauthorized access. It should be unobtrusive and straight-forward to minimize user frustration and confusion. Excessive complexity should be avoided in favor of strategies that are intuitive to users. Previous MLS systems have had limited success because their policies were invasive and incompatible with existing applications and procedures. Therefore it is worth reconsidering the benefits of fine-grained versus coarse-grained MLS policies.

Operating systems usually protect filesystem objects at the granularity of a single file. More fine-grained controls have been implemented in services such as multilevel versions of the X11 Window System [2] and per-row protections in Oracle Label Security [8], but these extensions are not interpreted by the OS kernel.

On the other hand, a less fine-grained approach makes sense for structured multilevel data. Although it may seem obvious that a MLS policy should be implemented at the same granularity as the traditional DAC policy, the protection policies for multilevel data involve special procedures and conventions that are distinct from unclassified data policies. For example, a user should not be able to determine that higher-level information even exists unless the user is cleared to read that data.

Just as in the physical world, sensitive digital information requires special handling procedures. For example, a folder marked Secret should not contain Top Secret data. The hierarchy of data classifications should be reflected in the organization of the data. Therefore, a more appropriate granularity for an OS-enforced MLS policy is at a directory (folder) level. But this model can be further optimized to handle just those directories whose labels differ from their parent directory. We can distinguish these special directories by making them the root nodes of mounted filesystems. In this way, only directories corresponding to mounted filesystems need to be explicitly validated by the MLS policy. All other directories and files can be labeled implicitly as an attribute of the mounted filesystem in which they exist.

3. FILESYSTEMS AND NAMESPACES

In the Solaris OS, new filesystems can be mounted over any arbitrary directory. In fact, filesystems can be mounted recursively on top of other filesystems with arbitrary complexity. However, there are a variety of restrictions which apply to the mount points. For example, only privileged users may mount (or unmount) filesystems. In this way the set of available pathnames is strictly controlled. We refer to the set of available pathnames as the file namespace, or simply *namespace*.

The Solaris kernel maintains an in-core table of mounted filesystem that is used to lookup pathname components which correspond to mounted filesystems. Some mounted filesystems correspond to actual mass storage, and some are simply aliases to directories in other filesystems. The term *filesystem* refers to both on-disk formats and virtual types. For example, Personal Computer File System (PCFS), UNIX File System (UFS), and Zettabyte File System (ZFS) are all on-disk formats, whereas the Loopback File System (LOFS), Network File System (NFS), and Common Internet File System (CIFS), are mappings to other filesystems.

3.1 The Zettabyte File System

The Zettabyte File System [14] provides several new abilities which facilitate filesystem labeling. It provides Copy-on-Write semantics (COW) which enables almost instantaneous cloning of filesystem snapshots. It also allows an arbitrary number of filesystems to be created from an arbitrary number of pooled disks. The filesystems, known as *data sets* in ZFS parlance, include an attribute identifying their corresponding directories, and are automatically mounted as needed. This makes it convenient to use large number of data sets, such as associating a unique data set with each user's home directory.

3.2 The Loopback File System

The Loopback File System (called *bindfs* in Linux) is a special case in which a portion of an existing filesystem can be exposed with a new name and protection context within a different filesystem. For example, a portions of a filesystem can be exposed as read-only in a new context, even though it is writable in the original context. This concept will be explored in section 7, Labeled File Sharing.

3.3 The Network File System

Finally, the Network File System makes it possible to share or export arbitrary directories from local filesystems, such as ZFS data sets, so that they are remotely accessible. By enforcing the MLS policy on both trusted NFS clients and trusted NFS servers the MLS policy can be extended to apply to heterogeneous environments of trusted and untrusted clients and servers.

Together, these three filesystems facilitate the implementation of labeled filesystems without explicit labeled attributes. Although a wide variety of filesystems are supported, we will focus on these three filesystems in the remainder of the paper.

4. LABELS AND PATHNAMES

In order to implement a file MLS policy, the kernel must be able to unambiguously determine the sensitivity label of any file from its pathname. Since the Solaris kernel can identify the mounted filesystem corresponding to any pathname, this can be extended to determine the sensitivity label of any mounted filesystem.

However, this is not as simple as it may seem. A large Solaris system using network storage may have thousands of mount points. New filesystems, such as ZFS, reduce the distinction between directories and mount points. So a reliable and automated mechanism is essential to ensure that the mounted filesystems are labeled correctly and consistently.

Solaris Trusted Extensions eliminates the need to explicitly associate labels with mounted filesystems. Instead, the kernel derives the label of each mount point implicitly from other attributes. The system does not provide any interface for directly associating a label with a mount point, thereby eliminating an administrative burden and a source of human error.

5. LABELED ZONES

A unique namespace is established for each sensitivity label that is defined on the system. The namespace and its default label are tied together in a containment mechanism known as a *zone* [16]. A zone is a virtualized environment in which related applications run, and appear to those processes to be a complete operating system. However, each of these zones is just an isolated labeled execution environment sharing a single Solaris kernel. The isolation between labeled zones applies to all system services including networking, the window system, and interprocess communication. By default, no communication is possible between processes in differently labeled zones. However, the implementation is transparent to applications since they appear to be running in separate OS instances.

Although any two zones have a label relationship, the zones are more isolated than is normally required by an MLS policy. For example, the ability for a higher-level subject to read lower-level objects is not provided by default. Processes in higher-level zones can't observe lower-level processes, and lower-level files are not visible unless explicitly shared.

5.1 The Global Zone

A special global zone is used to administer these labeled zones, and to establish the data flow policies between them. From an MLS perspective, the global zone is a multilevel environment. Its processes run at the highest label and dominate all other labels. The global zone is inaccessible to normal users; restricted access is provided to administrative roles via the Role-Based Access Control [3]Faden, Glenn Faden, Glenn Faden, Glenn Faden, Glenn . Users and roles are assigned a range of sensitivity labels with which they may operate. For each such label, a corresponding labeled zone is instantiated, and users are constrained to run their processes in the corresponding zones.

5.2 The Trusted Path

Although the user can't start processes in the global zone, a *trusted path* menu is provided with which they may interact. The trusted path provides a reliable communication mechanism that can never be obscured or interfered with by untrusted processes. It provides a set of user interfaces and services to access labeled zones, allocatable devices, and administrative roles.

Users are initially authenticated using the window system in the global zone. Then they are presented with a set of trusted path interfaces to select an initial label or clearance. They may specify their labeling preferences and assign labels to individual desktop workspaces. The window system is able to initiate processes on behalf of users in the appropriately labeled zones.

For processes running in labeled zones, there are no interfaces to change their label or transition to another zone. However, an all-privileged global zone process can fork off a process in a labeled zone. To do so it must close any open files and relinquish any privileges beyond the limit established for the labeled zone. Once the transition occurs, the process and all of its descendents are contained within their labeled zone and cannot return to the global zone.

5.3 Namespace Hierarchy

Each zone has a unique root directory, which is a mounted ZFS data set. The mount point and all the files within it are labeled with the zone's label; they are only writable within the zone. When a zone is configured it is assigned a set of filesystems. These filesystems either belong to the zone, or are imported from another zone. Directories imported from other zones retain the label of their owning zone For example, most of the global zone's executable files and libraries are imported through a special read-only mount option called a loopback mount. Such files are protected at the lowest label on the system and are immutable in labeled zones.

Before any filesystem can be imported into a zone, the kernel performs a label comparison between the owning zone and the importing zone. For a read-only mount to succeed, the importing zone's label must dominate (have a higher sensitivity label than) the owning zone. An attempt to start a zone with an incorrect mount configuration will be logged and terminated. The MLS policy is always enforced and is not compromised by administrative errors. Once the mount configuration is corrected, the administrator can boot the zone.

A read-write mount is only permitted when the label of the zone equals the label of the filesystem to be mounted. Since each zone within a single system has a unique sensitivity label, the label of any locally exported filesystem will never equal the label of the importing zone and will never be mounted as writable. Therefore, no file is writable by more than one zone on any system. Zones with disjoint labels can never share their files.

However, remote filesystems accessed via the NFS protocol may be mounted for both read and write modes, if their remote label matches that of the importing zone. Otherwise, the MLS mount policy is identical for both remote and local filesystem imports.

The implementation imposes a natural hierarchical structure to the labels in the file namespace. The root of the filesystem has the highest label in the zone and sensitivity labels decrease monotonically as we traverse pathnames. Subjects cannot determine the existence of any higher-level filesystem objects since there are none present in their namespace.

5.4 Automatic File Labeling

One of the benefits of this implementation is that the MLS policy for read and write access to a file does not require explicit label comparisons. Instead of performing a MLS check every time a file is opened, the MLS check occurs when the filesystem containing the file is mounted. In general, the only time that a file's label needs to be determined is for human consumption, such as for auditing or for data archival. When a user queries the system to display the file's properties, or requests a printed copy of the file, the label is computed and displayed appropriately. Since looking up a file's label is rarely required, it doesn't need to be particularly efficient. So file labels are neither stored in the filesystem nor cached in the Solaris kernel.

The label of a file is invariant with respect to the label of the requestor. Since each zone has its own file namespace, the kernel first resolves all pathnames into a global zone representation before resolving any labels. In doing so it must deal with hard and soft links. Hard linking is a mechanism in the Solaris OS that provides the same underlying file node with multiple independent names in different directories. The system prevents hard links from spanning filesystems, ensuring that all hard links to a single file must have the same label. Soft links, also known as symbolic links,

may span filesystems so their label is resolved by following the links and determining the label of the final target.

The algorithm[1] for computing a file's label involves several steps. The global-zone relative pathname is computed and matched against the closest match in the mount table. If the entry corresponds to a local filesystem, the label of the zone owning the mount point is returned. For NFS mounts, if the remote server is a trusted host, the label is retrieved from the server. If the server is a single-level host, the label is determined from its network address.

The MLS policy prevents returning labels of files that are not dominated by the user's label. Since a user can only ask for a label of a file in its own namespace, the returned label will always be dominated by the requestor.

6. POLYINSTANTIATED DIRECTORIES

A traditional challenge in implementing MLS systems is the need to provide writeable instances of well known directories, such as `/tmp`, or user home directories. This is called *polyinstantiation*, since each label appears to have a unique instance of a directory with the same pathname. This problem has been solved in the past by customizing existing filesystem to provide special multilevel directories [11].

Since each zone has its own file namespace, multiple instances of the same pathname may exist in each labeled zone, but correspond to distinct directories. For example each zone will typically have a unique instance of a directory called `/export/home`. These directories are actually contained in different filesystems and are only writeable in the zone owning that filesystem.

Polyinstantiation of entire filesystems can also be accomplished by cloning ZFS snapshots. For example a ZFS data set (which is a type of filesystem) can be populated with lower-level information, and then preserved as a read-only snapshot. A unique instance of the snapshot can then be quickly cloned for each higher-level zone without copying the underlying data. Even though each zone initially shares a reference to the same underlying storage, the cloned instances are implicitly labeled by their owning zone. When a process in a labeled zone writes to any cloned filesystem object a unique instance of the file or directory is created for that zone.

7. LABELED FILE SHARING

A convention exists for exporting such directories in read-only mode so that they can be accessed by processes in higher-level zones. The convention is to mount such lower-level shared directories using the prefix /zone/*zonename* where *zonename* is the name of the lower-level zone. In this way the higher-level zone can read the shared portion of the lower-level zone's named space.

Figure 1 shows three labeled zones: `need-to-know`, `internal`, and `public`. The labels for these zones have a dominance relationship, where `need-to-know` strictly dominates `internal`, and `internal` strictly dominates `public`. All three of these zones have been configured by the policy set in the global zone to permit sharing of their instance of the `/export` directory. Since the `internal` zone's label strictly dominates the `public` zone, a process in the `internal` zone can read the `public` zone's `/public` directory using the pathname `/zone/public/export`.

Meanwhile, a process in the `need-to-know` zone can also read the `public` zone's `/export` directory use the `/zone/public/export` pathname, as well as the `internal` zone's `/export` directory using the `/zone/internal/export` pathname. Neither the `internal` nor `public` zones can read from the `need-to-know` zone, even though its `/export` directory is shared, because their labels do not dominate the `need-to-know` zone's label. The global zone's `/usr` directory is shared for reading by all labeled zones.

The enumeration of shared filesystems is specified by a global zone administrative role. Since zones have separate namespaces, it is not possible for a process in one zone to directly mount any files owned by another zone. But the global zone includes the root directories of both namespaces, so a privileged global zone role can create a read only loopback mount from a lower-level zone to a higher-level zone. Such mounts are always consistent with the MLS policy. If the global zone administrator mistakenly attempts to create a loopback mount from a higher-level zone to a lower-level zone, the operation will be denied.

A second mechanism for file sharing is the use of NFS exports. Higher-level labeled zone processes may mount lower-level shared directories for read-only access via standard NFS protocols. An `automount` daemon in each zone does this dynamically, or each zone can be configured with its own instance of `/etc/vfstab`

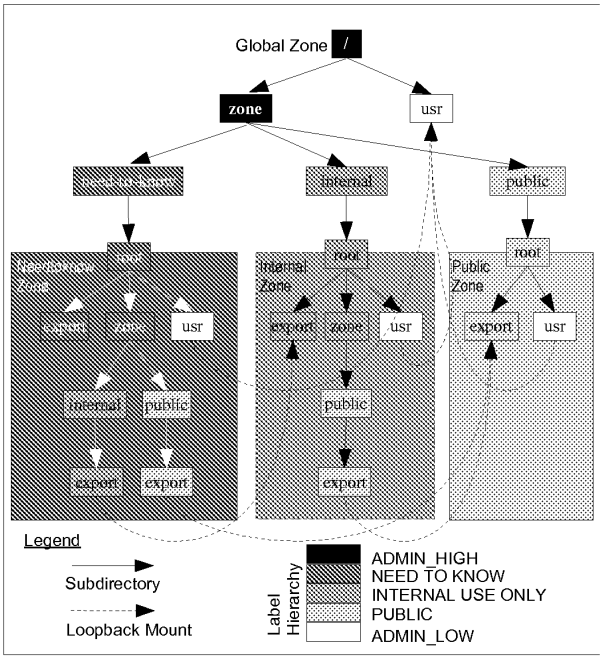


Figure 1. File sharing via loopback mounts.

[1]The complete implementation can be viewed at http://src.opensolaris.org/source/xref/onnv/onnv-gate/usr/src/uts/common/os/tlabel.c

to specify NFS mounts at boot time. Again, the MLS policy is enforced by the kernel to prevent inappropriate data flows.

Note that sharing is not transitive. A filesystem which is imported into any zone via an NFS mount cannot be exported from that zone. Each zone must import any shared filesystems from the entity that exported it. Therefore, the global zone's namespace includes only those portions of a labeled zone's namespace that were mounted from the global zone. Special privileges are required for global zone processes to observe the files or processes belonging to labeled zones.

A user in a zone can create symbolic links to files in these lower-level named spaces, so that they may appear at any point in the current zone's namespace to which the user has write access. The user may query the system to return the labels of any such files. Regardless of the pathanme used to refer to the file, the label will always reflect that of the zone from which it is imported.

Directories exported from the global zone are always assigned the label ADMIN_LOW. Note that any such lower-level mounts are always read-only. Higher-level mounts are prohibited.

8. UNIDIRECTIONAL DATA FLOWS

The ability to read lower-level files from a higher-level zone is called *reading down*. The lower-level process is only able to write at its own label, and cannot *write up*. However, two processes in separate zones can communicate using a FIFO conduit. A special rendezvous file, known as a *named pipe* can be created by one process, which then writes to the file. Another process reading from the file will receive the data written by the first process. The OS synchronizes these operations so that the reader will wait for data and terminate when the writer closes the file. This mechanism can be used by labeled zones to achieve a unidirectional data flow. If the writer creates the named pipe in a directory which is exported from a lower-level zone and imported into a higher-level zone, the named pipe will appear in the namespace of the higher level zone as a read-only object. It retains the label of the lower-level zone.

The apparent write up in this context is consistent with the MLS policy because the writer is is simply writing to an object at its own label, and cannot identify the process at at the other end of the FIFO. The reader in the higher-level zone cannot reply to the lower-level process. Unidirection flows provide a mechanism for audit trails and similar logs that cannot be destroyed by processes within the zone. Even an all-privileged process within the zone cannot modify or erase the log once it has been recorded by a process in the higher-level zone.

9. LABEL MAINTENANCE

Although files are automatically labeled when mounted, the system provides for archiving labels, and for relabeling of existing files.

For offline storage a label-aware version of the `tar` program is provided to record the label of each file when it is archived. The labels are recorded in an ancillary file format which is backward compatible with Trusted Solaris labeled archives. The archived labels are used for validation when files are restored to ensure that the file labels are not changed. A file is not restored unless the archived file's label is equal to the label of the zone into which it is being restored.

Authorized users may request that a the file label be reclassified. Separate authorizations are required for users to either upgrade or downgrade a file. Such requests are handled by a global zone service that validates the request, including a check to ensure the file is not currently in use. If the user is authorized and the file is unused, the file is relabeled by removing it from its current filesystem and copying it into a filesystem corresponding to the requested label. The pathname of the original file is adjusted to the namespace of the new filesystem.

An interesting side-effect of this operation is that it cannot be undone by the original subject, since the file disappears from its current directory in the original namespace.

10. RELATED WORKS

The Trusted Linux system proposed by Dalton and Choo [1] makes use of separate filesystem roots to achieve containment of web services. Their compartments are similar to the labeled zones in Solaris Trusted Extensions, but their labeling strategy does not include a hierarchical component. Therefore, there is no implicit policy for information flow other than complete isolation between compartments; explicit cross-compartment data flows are enumerated in a policy file. There is no concept of privileged processes within a compartment; transitions to `root` are prevented.

SELinux can be configured to enforce an MLS policy in addition to the primary TE Policy. To gain access to an object a subject must pass the traditional DAC policy, the TE policy, and the MLS policy. The following discsussion is focused on the MLS filesystem policy. Both Trusted Solaris [11] and SELinux [9] provide extensions to the `mount` command to specify the default security context for a filesystem. This requires an administrator to determine the appropriate context, which is a source of human error. Futhermore, neither of these systems require that the label of the mounted fileystem be dominated by the label of the mount point. They permit higher-level filesystems or individual higher-level files to exist in directories that are accessable from lower-level subjects, and they permit authorized subject to upgrade a file so that its label dominates its directory.

Both Trusted Solaris and the SELinux MLS policy prevent normal subjects from reading higher-level files, but the file names and their existence represent sensitive information which should be hidden from the namespace. Filtering such multilevel information in a directory is problematic. Although the system call for reading directory entries, `getdents`, can be modified to hide upgraded file and directory names, a lower-level subject can infer their existence by attempting to create duplicate entries at its own label. SELinux does not provide such filtering so higher level file names are visible to lower level subjects. In Trusted Solaris, the policy for hiding upgraded file names is configurable because its enforcement introduces additional overhead affecting performance. This contrasts with the Solaris Trusted Extensions architecture, where upgrading a file is done by moving it to a higher-level directory in another filesystem. No filtering is required because directories contain no upgraded entries.

11. CONCLUSION

Much has been written about configuring flexible security policies and developing tools to validate them [5, 13]. This paper asserts that a useful security policy is best implemented when it is automatically consistent and easily understood. The use of labeled zones eliminates the need for explicitly labeled filesystems or for specifying labels as mount attributes. This approach provides a significant set of advantages:

- Since no modifications are required to any filesystem structures, all commercially available filesystems are compatible with the MLS policies enforced in the kernel.
- Unlike systems, where labeling attributes are specified as context mount [9] options, the mount point labels are automatically applied. This prevents a source of human error and inconsistency.
- Since labeled zones provide unique instances of directories and services at each label, unmodified applications can run at multiple labels concurrently.
- Since the zones are completely isolated by default and applications require no modification, it is unnecessary to maintain complex policy files.

This new MLS architecture is currently undergoing Common Criteria certification using the Labeled Security (LSPP) and Role-Based Protection Profiles (RBACPP). This approach is not a panacea, however. One of its weaknesses is that the underlying filesystems are only protected while the trusted OS is running. Of course, this is also true of systems which implement explicit security contexts. For example, an explicitly labeled filesystem can be mounted on another system which ignores the explicit labels.

Another tradeoff is the requirement to configure a unique zone for each label. Although zones can be provisioned efficiently using ZFS cloning, the administrator must determine the set of labels and zones that will be required by the user community.

A more robust solution to this problem requires labeled protection for data at rest. Cryptographic filesystems provide the missing functionality. Encrypted ZFS data sets are currently in development [15]. Each ZFS data set can be independently encrypted, so that the cryptographic keys can be associated with their corresponding labels. The MLS policy is applied to key management so that the filesystem cannot be mounted without providing the proper key for that label. As work is ongoing in this area, the details of this implementation will be the subject of another paper.

12. ACKNOWLEDGMENTS

The author is thankful for the the efforts of the Rampart project team whose efforts made this possible.

REFERENCES

[1] Dalton C. and Choo T. S. *Trusted Linux: An operating system approach to securing e-services* Communications of the ACM, Volume 44, Issue 2. 2001

[2] Faden, Glenn *Reconciling CMW Requirements with Those of X11 Applications*, Proceedings of the 14th National Computer Security Conference, Washington, D.C. 1991

[3] Faden, Glenn *RBAC in UNIX administration* Proceedings of the fourth ACM workshop on Role-based access control, Fairfax, Virginia, United States, 1999, Pages: 95 – 101

[4] Faden, Glenn, *Solaris Trusted Extensions: An Architectural Overview* 2006 http://www.opensolaris.org/os/community/security/projects/tx/TrustedExtensionsArch.pdf

[5] Jaeger T., Zhang, X, and Cacheda F. *Policy management using access control spaces* ACM Transactions on Information and System Security (TISSEC) Volume 6 , Issue 3 , page 327-364. 2003

[6] Mc Dougall, R. and Mauro J. *Solaris Internals*, Chapter 15 UFS File System, page 764, 15.4 Access Control in UFS, Sun Microsystems Press, 2005

[7] Morris, James *Filesystem Labeling in SELinux* Linux Journal Contents #126, October 2004 http://www.linuxjournal.com/article/7426

[8] Oracle Corporation *Data Classification with Oracle Label Security* 2005 http://www.oracle.com/technology/deploy/security/db_security/pdf/twp_security_db_ols_10gr2.pdf

[9] Red Hat *Taking Avantage of SELinux in Red Hat Enterprise Linux,* Red Hat Magazine, Issue #6, April 2005 http://www.redhat.com/magazine/006apr05/features/selinux

[10] Solaris Trusted Extensions http://www.opensolaris.org/os/community/security/projects/tx

[11] Sun Microsystems, *Trusted Solaris 8: A Technical Overivew*, page 15. 2000 http://www.sun.com/software/whitepapers/wp-ts8/ts8-wp.pdf

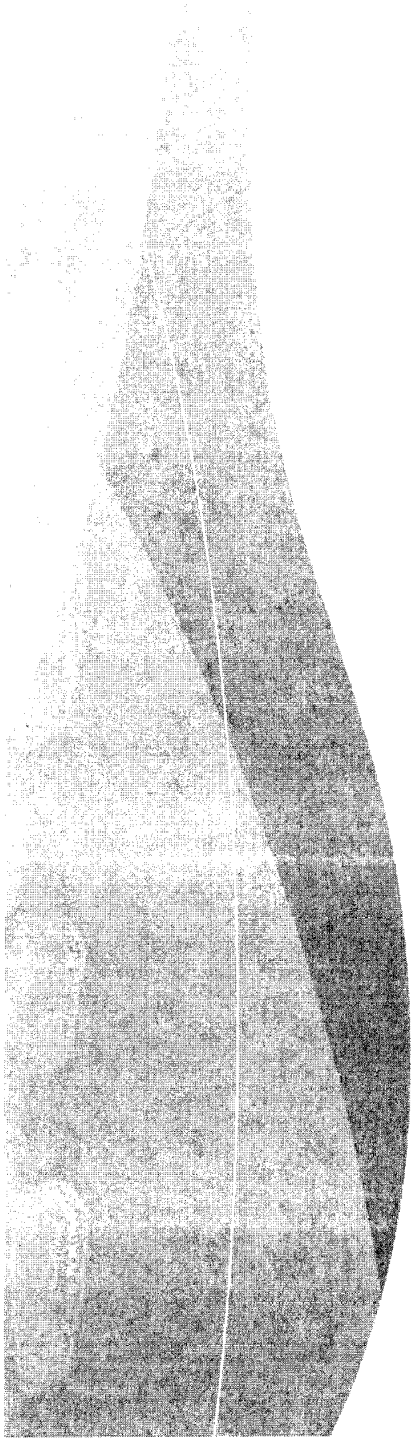
[12] US Department of Defense *Trusted Computer System Evaluation Criteria,* DoD 5200.28-STD. December 26, 1985.

[13] Zanin, G. and Mancini L. *Towards a formal model for security policies specification and validation in the selinux system,* Proceedings of the ninth ACM symposium on Access control models and technologies, pages 136-145, Yorktown Heights, New York, USA , 2004

[14] ZFS at OpenSolaris.org http://www.opensolaris.org/os/community/zfs

[15] ZFS on disk encryption support http://www.opensolaris.org/os/project/zfs-crypto/

[16] Zones at OpenSolaris.org http://www.opensolaris.org/os/community/zones



Solaris Trusted Extensions

Architectural Overview

Glenn Faden
April 2006

Table of Contents

Introduction....3
Compatibility....3
Efficiency....3
File System MAC Policy....4
Configurable Zone Privileges....5
Label Ranges and Sets....5
Network Labeling....6
Multilevel Services....6
Multilevel Desktop Sessions....8
Roles....8
Devices....8
Printing....8
Auditing....9
Conclusion....9

Introduction

Solaris™ Trusted Extensions is an optionally-enabled layer of secure labeling technology that allows data security policies to be separated from data ownership. While it has its roots in the multilevel **Trusted Solaris 8** OS, it has been integrated into the standard Solaris 10 Operating System. This new approach allows the Solaris operating system to support both traditional Discretionary Access Control (*DAC*) policies based on ownership, as well as label-based Mandatory Access Control (*MAC*) policies. The label-based policies for file systems and networks are light-weight and have been implemented within the standard Solaris 10 kernel, services and utilities. Unless the Trusted Extensions layer is enabled, all labels are equal so the kernel is not configured to enforce the MAC policies. For efficiency, a boolean value is maintained in the kernel to indicate whether labeling comparisons should be used in policy enforcement.

When the label-based MAC policies are enabled, *all* data flows are restricted based on a comparison of the labels associated with the subjects requesting access and the objects containing the data. Like other multilevel operating systems, Trusted Extensions meets the requirements of the **Common Criteria Labeled Security Protection Profile** (**LSPP**) and the **Role-Based Access Protection Profile (RBAC)**. However, the Trusted Extensions implementation is unique in its ability to provide high assurance, while maximizing compatibility and minimizing overhead.

Compatibility

Maximum compatibility with thousands of Solaris applications is achieved by building on existing Solaris features and using existing industry standards. No new protocols are required nor new file system attributes. Applications do not need to be modified nor profiled to bring them into conformance with the MAC policy. Instead, the entire application environment is virtualized for each label through the use of Solaris Containers (*zones*). This approach is referred to as *polyinstantiation* because there can be an instance of each resource and service available at each label. However, there is also a unification principle known as a *single system image* which is applied to the entire operating environment. All the zones are centrally administered from a special, protected global zone which manages the *Trusted Computing Base* (TCB) known as the *Trusted Path*. The zones share a single LDAP directory in which network-wide policy is defined, as well as a single name service cache daemon for synchronizing local databases. All labeling policy and account management is done from within the Trusted Path. MAC policy enforcement is automatic in labeled zones and applies to all their processes, even those running as `root`. Access to the Global Zone (and hence Trusted Path applications) is restricted to administrative roles.

Efficiency

Minimal overhead is achieved by moving access control checks to a higher level of enforcement. For example, instead of maintaining labels on fine-grained objects like files and directories, Trusted Extensions associates labels with Solaris Containers (zones), and network endpoints. Each zone is assigned a unique sensitivity label and can be customized with its own set of file systems and network resources. Each mounted file system is automatically labeled by the kernel when it is mounted. The file system label is derived from the label of the zone or host which is sharing it. All files and directories within the mounted file system have the same label as their mount point. Because no explicit extensions to the file or file system structure are required, essentially any file system that works on Solaris 10 will work when Trusted Extensions label enforcement is enabled. This

includes Sun issued file systems such as UFS, ZFS, SAM-FS and QFS as well as third-party file systems.

Processes are uniquely labeled according to the zone in which they are executing. All processes within a zone (and their descendants) must have the same label, and are completely isolated from processes in other zones. Unlike other virtualization technologies, there is no performance penalty for executing within a zoned environment as there is no emulation required for a Container. Labeled zones can be instantiated quickly by cloning a copy of a default zone. Disk usage is minimized by sharing immutable instances of most system files and by utilizing copy-on-write technology for the rest.

File System MAC Policy

A zone's local file systems are only writable at the zone's label, but can be shared with labeled zones via loopback or NFS mounts. Loopback mounts are used between zones running on the same host, and multi-level NFS is used for access between hosts. File systems that are shared by all zones on a system are always mounted read-only. This policy provides both confidentiality and integrity protection. Such file systems are assigned the lowest administrative label, `ADMIN_LOW`. Similarly, file systems imported from lower-level zones are assigned the label of the zone with which they are shared.

File sharing between sets of Trusted Extensions systems using NFS can be symmetric. Corresponding zones on each system (with matching labels) can have read-write access to each other's shared file systems. Zones which dominate (have higher labels) than the owning zone can be granted read-only access (depending on per-zone policy settings).

Writing up to higher-level regular files is not possible because such files are never visible within a labeled zone. However, writing up is possible using named pipes which are loopback mounted into higher-level zones. This unidirectional conduit is useful for implementing one-way guards and for tamper-proof logging.

The following is an example of the labels assigned to the mount points in a zone called `needtoknow`, whose label is `CONFIDENTIAL : NEED TO KNOW`. It dominates two zones, `internal` and `public`.

Mount Point	Access	Sensitivity Label
/	Read/Write	CONFIDENTIAL : NEED TO KNOW
/kernel	Read Only	ADMIN_LOW
/lib	Read Only	ADMIN_LOW
/opt	Read Only	ADMIN_LOW
/platform	Read Only	ADMIN_LOW
/sbin	Read Only	ADMIN_LOW
/usr	Read Only	ADMIN_LOW
/var/tsol/doors	Read Only	ADMIN_LOW
/tmp	Read/Write	CONFIDENTIAL : NEED TO KNOW
/var/run	Read/Write	CONFIDENTIAL : NEED TO KNOW

Mount Point	Access	Sensitivity Label
/home/gfaden	Read/Write	CONFIDENTIAL : NEED TO KNOW
/zone/public/export/home/gfaden	Read Only	PUBLIC
/zone/internal/export/home/gfaden	Read Only	CONFIDENTIAL : INTERNAL USE ONLY

Table 1: Labeled mount attributes

To prevent configuration errors and to simplify system administration, there are no interfaces for specifying the labels of mount points. Instead, the kernel determines the labels of all mount points based on host and zone labels, and ensures that the MAC policy is correctly implemented.

Configurable Zone Privileges

By default each labeled zone is completely isolated from all other labeled zones because their labels are required to be unique. No process in a zone can view or signal processes running in other zones. There are no privileges available for any process in a labeled zone to write to lower-level files. However, such policies as reading from files in lower-level zones, exporting directories to higher-level zones, and moving files into higher-level zones can be enabled by specifying the privileges available to each zone when it is booted[1] .

Privileges available to a zone can, in turn, be assigned to processes in the zone. However, a zone's privilege limit is an upper bound that applies to all processes (even root-owned) that are run in the zone. All policies which affect multiple zones, such as sharing of directories, are administered via the Trusted Path.

Label Ranges and Sets

Labels consist of hierarchical components called classifications (or levels) and a non-hierarchical components called compartments (or categories). The mapping of names to classifications and compartments is specified in a database which is private to the Trusted Path. The internal structure of labels is deliberately opaque to users and applications and might change in a future release. At least 256 classifications and 256 compartment bits are supported.

When two labels are compared, the first label can be greater than, less than, equal to, or disjoint from the second label. Classifications are compared as integers, and compartments are compared as bit masks. Labels are disjoint when each contains at least one compartment bit which is not present in the other.

A label range can be specified by an upper bound (called a clearance), and a lower bound. Administrative roles can use the Trusted Path to assign label ranges to users, network attributes, workstations, and allocatable devices.

1 The policy for reading down is configurable because it is not always appropriate. For example, it could result in processes at a higher level executing lower-level applications which manipulate the higher-level data. One way to mitigate that risk is to configure the zone without the privilege to do read down mounts. An additional way is to specify the `noexec` mount option for lower-level mounts.

Network Labeling

As in previous versions of Trusted Solaris software, remote hosts can be single level or multilevel. Single level hosts have an implicit label assigned to them based on their network or IP address. Non-label aware systems, such as workstations running Microsoft Windows™, are assigned a specific label for communications purposes. Multilevel hosts are trusted to operate at a range of labels, and explicitly specify the label of every network packet when communicating with other trusted systems. Packet labels are specified using the Commercial IP Security Option (CIPSO) which encapsulates a sensitivity label as an IP option. CIPSO is specified in the FIPS 188 Standard and is supported by Trusted Solaris 8 and other labeled systems.

When specifying the labeling policy for network attributes, both label ranges and sets of disjoint labels can be enumerated. This ability to precisely define the labeling policy is required to support various multilevel configurations including guards, NFS servers, Sun Ray servers, name servers, print servers, workstations, and high-assurance grid computing. An administrator can also assign a label range to a router even if the router does not interpret labels. Although zones have unique labels, specific multilevel services can be configured for each zone.

The network attributes database is maintained in an `LDAP` directory and shared by all trusted systems comprising a network of multilevel systems. `IPsec` can be used to authenticate the source IP addresses associated with incoming network packets. IPsec enforces integrity protection, and is used to encrypt data on multilevel networks.

Multilevel Networking

Zones can be configured to share a single IP address, or they can be assigned unique IP addresses. Similarly, they can share the same physical network interface, or can be configured to use separate network interfaces. Both shared and per-zone IP addresses can be used concurrently, with different labeling policies for each IP address. Solaris Zones technology allows multiple zones to share a single network interface through the use of virtual interfaces.

Sharing of IP addresses is possible in Trusted Extensions because each packet is implicitly or explicitly labeled. When a packet is received, the kernel uses the label of the packet to determine the appropriate zone to which it should be delivered. Sharing a single IP address for all zones is convenient for workstations and laptops, especially when DHCP is used. This simplifies deployment into infrastructures with limited IP addresses.

Using per-zone IP addresses is required when separate networks are in use, and might be appropriate when multilevel services are being provided. To enable multilevel services, a database of multilevel ports is maintained via the Trusted Path. A multilevel port is a special kind of reserved port whose multilevel semantics are administratively controlled. For each IP address, a range of labels, as well as explicit labels outside of that range can be configured for use by multilevel services. A privileged server can bind to a multilevel port using any IP addresses that are assigned to the server's zone. The server can receive requests at these labels and reply to any request. For multilevel TCP services, the reply is automatically sent using the label of the request without requiring any special programming in the server. For multilevel UDP services, the server must set a socket option to indicate the label of the reply. In either case, the server can query the kernel to determine the label of

each request and then restrict the reply accordingly.

The following block diagram shows some typical network paths that can be administratively configured.

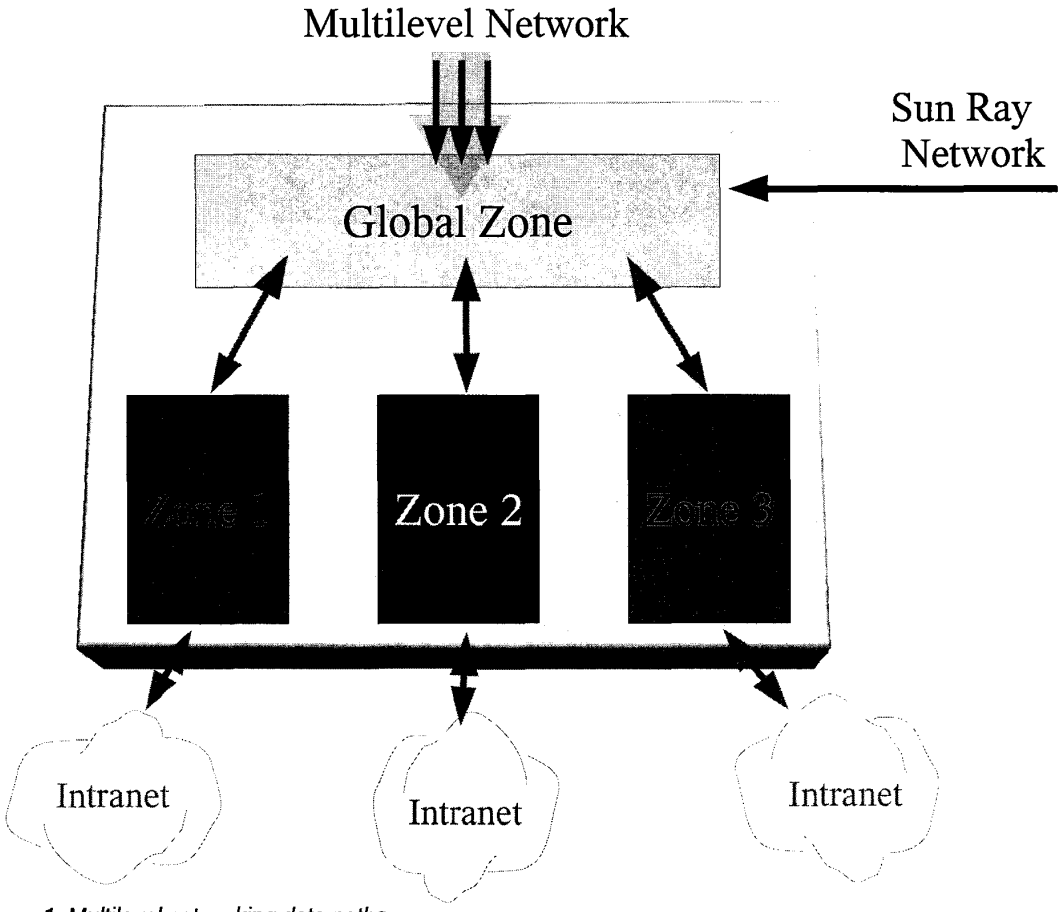


Figure 1: Multilevel networking data paths

Multilevel Services

By default Solaris 10 with Trusted Extensions enables the following multilevel services :

- X11 Window System
    - Common Desktop Environment
    - Java Desktop System
- Printing
    - Internet Protocol Printing

    - BSD Protocol Printing
- Network File System
- Sun Directory Server
- Label Translation Service
- Name Service Cache Daemon

All other services are polyinstantiated in each zone. However, additional multilevel services such as Web Servers and Secure Shell can be enabled administratively via the Trusted Path.

Multilevel Desktop Sessions

Users can log in via the Trusted Path and can be authorized to select their multilevel desktop preference (CDE or JDS). Once authenticated they are presented with an option to select an explicit label or a range of labels within their clearance and the label range of their workstation or desktop unit. The window system initiates a user session in the zone whose label corresponds to the user's default or minimum label.

The window system provides menus for interacting with the Trusted Path to change the label of the current workspace or to create additional labeled workspaces. For each selected label, the window system starts another user session in the corresponding zone. All of these user sessions run concurrently and are subjects of the user's identity that was established during the initial authentication. Each window is visibly labeled according to the zone or host with which it is associated. Although users can simultaneously interact with windows running in multiple zones, the applications themselves remain isolated.

Attempts to cut and paste data, or drag and drop files between clients running in different zone are mediated by the Trusted Path. Specific authorizations are required for upgrading or downgrading selections and files, and are prohibited by default. The following screen shot shows an authorized user interacting with the Trusted Path to upgrade a selection:

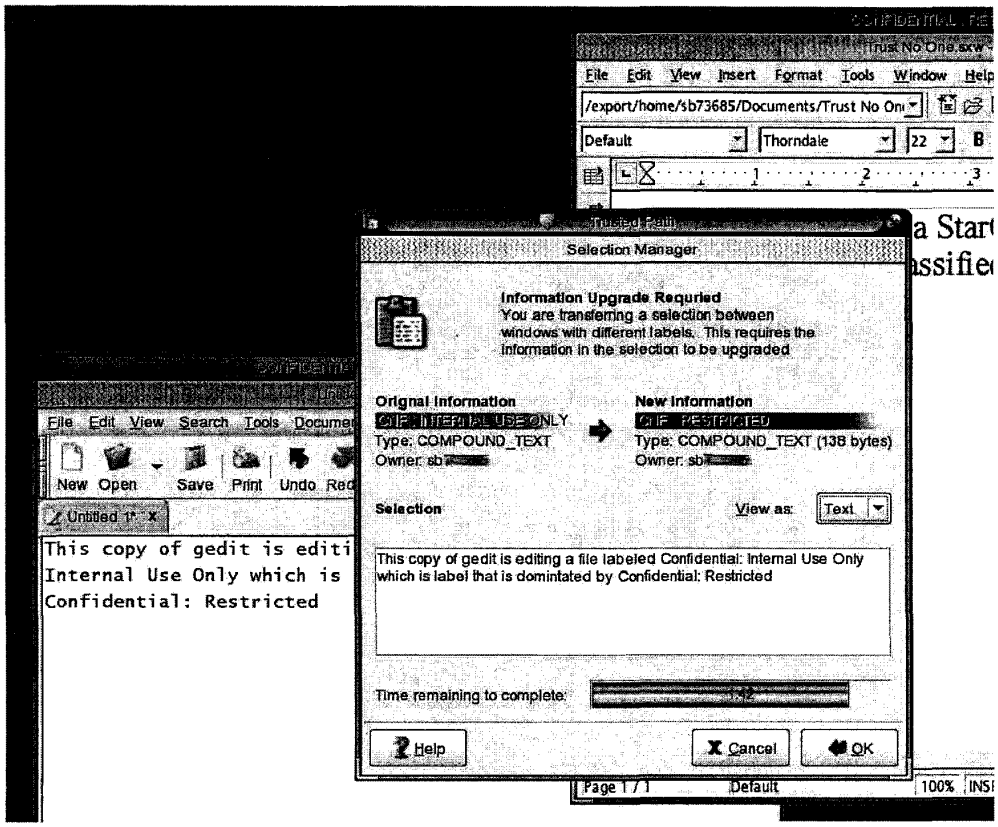


Figure 2: Multilevel Cut and Paste in Trusted JDS

Roles

By using the Trusted Path menu, authorized users can assume one or more roles which they have been assigned. For each assumed role a secondary authentication is required. Once authenticated, the window system creates a new administrative workspace for the role and starts another session. These administrative workspaces are protected from interference by untrusted X11 clients and non-role user logins. For roles who are cleared for all labels, their sessions are initiated as Trusted Path processes. Each role has a limited set of Role Based Access Control [RBAC] rights which restrict its access. Typically, two or more cooperating roles can be used to configure the system. For example, a system administrator role creates accounts and zones, while a security administrator assigns labels to them. Roles with sufficient rights can configure aspects of the MAC and DAC policies that apply to one or more zones

Devices

Devices represent a security threat because they can be used to import and export data from the system. In Trusted Extensions, removable media devices are administered through the Trusted Path menu. The window system provides a Trusted Path interface for device allocation which provides fine-grained access to specific

devices based on user authorizations and label ranges. For example, a user can be authorized to allocate the audio system (speaker and microphone) at a single level. Hot pluggable devices such as USB flash memory drives are also managed by the Trusted Path user interface. An authorized user can request to have such devices mounted into a zone whose label is within the user's label range and the device's label range. As an extra security measure, the raw device is not available within the labeled zone. This capability protects the integrity of the mounted file system and prevents unauthorized access.

Printing

Each printer is assigned a label range from which it will accept requests. Multilevel printers can accept jobs from labeled zones or remote hosts whose labels fall within their range. Each job can be encapsulated between reliably matching banner and trailer pages which indicate the label and handling caveats for the output. Each page can be automatically labeled with headers and footers corresponding to the sensitivity of the data.

Auditing

Trusted Extensions audit records are compatible with standard Solaris OS. They include the labels of subjects and objects, and additional label-related events. The auditing system is configured via the Trusted Path and is transparent to users and roles running in labeled zones. The auditing system is robust and cannot be tampered with by processes running in labeled zones. Even processes with all privileges cannot observe the audit trail nor tamper with any records.

Conclusion

Solaris Trusted Extensions builds on the security features in Solaris 10 in an upward compatible fashion. Users of Solaris 10 can enable Trusted Extensions by installing the Trusted Extensions administrative packages. Because it is built-into the Solaris OS, all of the latest Solaris functionality is supported by Trusted Extensions, and all hardware platforms are supported. The MAC policy is enforced automatically and transparently without requiring customized applications or complex application profiles that are difficult to deploy and maintain. This new approach makes it possible to deliver a cost-effective and interoperable system that simultaneously provides ease-of-use and a high level of information assurance.



Sun Microsystems, Inc. Network Circle, Santa Clara, CA 95054 USA Phone 1-650-960-1300 or 1-800-555-9SUN Web sun.com

©2006 Sun Microsystems, Inc. All right reserved Sun, Sun Microsystems, the Sun logo [insert all other Sun Product Names, Service Names, Program Names, Special Programs Logos and slogans listed in Tmark and referred to or displayed in the document] are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries."

PATENT COOPERATION TREATY

PCT

INTERNATIONAL PRELIMINARY REPORT ON PATENTABILITY
(Chapter I of the Patent Cooperation Treaty)

(PCT Rule 44*bis*)

Applicant's or agent's file reference GOS838	**FOR FURTHER ACTION**	See item 4 below
International application No. PCT/GB2005/001845	International filing date *(day/month/year)* 12 May 2005 (12.05.2005)	Priority date *(day/month/year)* 12 May 2004 (12.05.2004)
International Patent Classification (8th edition unless older edition indicated) See relevant information in Form PCT/ISA/237		
Applicant BLUESPACE GROUP LTD		

1. This international preliminary report on patentability (Chapter I) is issued by the International Bureau on behalf of the International Searching Authority under Rule 44 *bis*.1(a).

2. This REPORT consists of a total of 6 sheets, including this cover sheet.

   In the attached sheets, any reference to the written opinion of the International Searching Authority should be read as a reference to the international preliminary report on patentability (Chapter I) instead.

3. This report contains indications relating to the following items:

☒	Box No. I	Basis of the report
☒	Box No. II	Priority
☐	Box No. III	Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
☐	Box No. IV	Lack of unity of invention
☒	Box No. V	Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
☐	Box No. VI	Certain documents cited
☐	Box No. VII	Certain defects in the international application
☐	Box No. VIII	Certain observations on the international application

4. The International Bureau will communicate this report to designated Offices in accordance with Rules 44*bis*.3(c) and 93*bis*.1 but not, except where the applicant makes an express request under Article 23(2), before the expiration of 30 months from the priority date (Rule 44*bis* .2).

	Date of issuance of this report 14 November 2006 (14.11.2006)
The International Bureau of WIPO 34, chemin des Colombettes 1211 Geneva 20, Switzerland Facsimile No. +41 22 338 82 70	Authorized officer Dorothée Mülhausen e-mail: pt01@wipo.int

Form PCT/IB/373 (January 2004)

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

REC'D 23 AUG 2005
WIPO PCT

To:

see form PCT/ISA/220

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY
(PCT Rule 43*bis*.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/GB2005/001845	12.05.2005	28.10.2004

International Patent Classification (IPC) or both national classification and IPC
H04L12/58, H04L29/06

Applicant
BLUESPACE GROUP LTD.

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☒ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☐ Box No. VII Certain defects in the international application
- ☐ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA"). However, this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notifed the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of three months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

3. For further details, see notes to Form PCT/ISA/220.

Name and mailing address of the ISA:

European Patent Office - P.B. 5818 Patentlaan 2
NL-2280 HV Rijswijk - Pays Bas
Tel. +31 70 340 - 2040 Tx: 31 651 epo nl
Fax: +31 70 340 - 3016

Authorized Officer

Poppe, F

Telephone No. +31 70 340-4378

Form (PCT/ISA/237) (Cover Sheet) (January 2004)

Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of the international application in the language in which it was filed, unless otherwise indicated under this item.

   ☐ This opinion has been established on the basis of a translation from the original language into the following language , which is the language of a translation furnished for the purposes of international search (under Rules 12.3 and 23.1(b)).

2. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application and necessary to the claimed invention, this opinion has been established on the basis of:

   a. type of material:

   ☐ a sequence listing

   ☐ table(s) related to the sequence listing

   b. format of material:

   ☐ in written format

   ☐ in computer readable form

   c. time of filing/furnishing:

   ☐ contained in the international application as filed.

   ☐ filed together with the international application in computer readable form.

   ☐ furnished subsequently to this Authority for the purposes of search.

3. ☐ In addition, in the case that more than one version or copy of a sequence listing and/or table relating thereto has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

4. Additional comments:

Box No. II Priority

1. ☒ The validity of the priority claim has not been considered because the International Searching Authority does not have in its possession a copy of the earlier application whose priority has been claimed or, where required, a translation of that earlier application. This opinion has nevertheless been established on the assumption that the relevant date (Rules 43*bis*.1 and 64.1) is the claimed priority date.

2. ☐ This opinion has been established as if no priority had been claimed due to the fact that the priority claim has been found invalid (Rules 43*bis*.1 and 64.1). Thus for the purposes of this opinion, the international filing date indicated above is considered to be the relevant date.

3. Additional observations, if necessary:

Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	
	No:	Claims	1-26
Inventive step (IS)	Yes:	Claims	
	No:	Claims	1-26
Industrial applicability (IA)	Yes:	Claims	1-26
	No:	Claims	

2. Citations and explanations

see separate sheet

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

Reference is made to the following documents:

D1: EP 1 085 444 A
D2: EP 0 739 115 A

1. The present application does not meet the criteria of Article 33(1) PCT, because the subject-matter of **claims 1, 7 and 13** is not new in the sense of **Article 33(2) PCT**.

1.1. The document D1 discloses (the references in parentheses applying to this document):

a computerized messaging server in an electronic messaging system, comprising:

- a messaging module (par. 36) adapted to control a message database ("message store 130") storing messages exchanged in the messaging system, each message (par. 45) comprising one or more message components ("initial message 10 for context along with the text of replies 412, 414, 416 and 418"); and
- a governance module (par. 42) adapted to control a governance policy database storing a compliancy policy describing rules applicable to the message components ("rules for governing access to message 180"), the governance module further adapted to apply the rules described by the compliance policy to the message components ("control access").

The subject-matter of claim 1 is therefore not new.

1.2. The same reasoning applies, mutatis mutandis, to the subject-matter of the corresponding independent claims 7 and 13, which therefore is also considered not new.

2. Dependent claims 2-6, 8-12 and 14-26 do not contain any features which, in combination with the features of any claim to which they refer, meet the requirements of the PCT in respect of novelty and/or inventive step, see documents D1 and D2 and the corresponding passages cited in the search report.

INTERNATIONAL SEARCH REPORT

ational Application No
PCT/GB2005/001845

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/58 H04L29/06

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)
EPO-Internal, WPI Data, PAJ, INSPEC, COMPENDEX

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 1 085 444 A (MICROSOFT CORP) 21 March 2001 (2001-03-21) abstract paragraph '0036! paragraph '0041! - paragraph '0045! figure 4 figure 5	1-26
X	EP 0 739 115 A (AT&T IPM CORP) 23 October 1996 (1996-10-23)	1,7,13
A	column 5, line 51 - column 6, line 54 figure 3A figure 3B	2-6, 8-12, 14-26

-/--

[X] Further documents are listed in the continuation of box C.

[X] Patent family members are listed in annex.

° Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
11 August 2005	24/08/2005

Name and mailing address of the ISA	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Poppe, F

Form PCT/ISA/210 (second sheet) (January 2004)

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EDWARDS W K: "The design and implementation of the Montage multimedia mail system" COMMUNICATIONS FOR DISTRIBUTED APPLICATIONS AND SYSTEMS. CHAPEL HILL, APR. 18 - 19, 1991, PROCEEDINGS OF THE CONFERENCE ON COMMUNICATIONS SOFTWARE. (TRICOMM), NEW YORK, IEEE, US, vol. CONF. 4, 18 April 1991 (1991-04-18), pages 47-57, XP010039767 ISBN: 0-87942-649-7 abstract page 51, right-hand column, line 48 - page 52, left-hand column, line 2	1-26

INTERNATIONAL SEARCH REPORT

Information on patent family members

tional Application No
PCT/GB2005/001845

Patent document cited in search report		Publication date	Patent family member(s)		Publication date
EP 1085444	A	21-03-2001	US	6704772 B1	09-03-2004
			EP	1085444 A2	21-03-2001
			US	2004148359 A1	29-07-2004
EP 0739115	A	23-10-1996	CA	2172746 A1	21-10-1996
			EP	0739115 A2	23-10-1996
			JP	3470846 B2	25-11-2003
			JP	8307452 A	22-11-1996
			SG	93792 A1	21-01-2003
			US	5941946 A	24-08-1999

Form PCT/ISA/210 (patent family annex) (January 2004)

PATENT COOPERATION TREATY

REC'D 23 AUG 2005
WIPO PCT

From the
INTERNATIONAL SEARCHING AUTHORITY

To:

see form PCT/ISA/220

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY
(PCT Rule 43*bis*.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/GB2005/001845	12.05.2005	28.10.2004

International Patent Classification (IPC) or both national classification and IPC
H04L12/58, H04L29/06

Applicant
BLUESPACE GROUP LTD.

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☒ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☐ Box No. VII Certain defects in the international application
- ☐ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA"). However, this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notifed the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of three months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

3. For further details, see notes to Form PCT/ISA/220.

Name and mailing address of the ISA:

European Patent Office - P.B. 5818 Patentlaan 2
NL-2280 HV Rijswijk - Pays Bas
Tel. +31 70 340 - 2040 Tx: 31 651 epo nl
Fax: +31 70 340 - 3016

Authorized Officer

Poppe, F

Telephone No. +31 70 340-4378

Form (PCT/ISA/237) (Cover Sheet) (January 2004)

Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of the international application in the language in which it was filed, unless otherwise indicated under this item.

   ☐ This opinion has been established on the basis of a translation from the original language into the following language , which is the language of a translation furnished for the purposes of international search (under Rules 12.3 and 23.1(b)).

2. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application and necessary to the claimed invention, this opinion has been established on the basis of:

   a. type of material:

   ☐ a sequence listing

   ☐ table(s) related to the sequence listing

   b. format of material:

   ☐ in written format

   ☐ in computer readable form

   c. time of filing/furnishing:

   ☐ contained in the international application as filed.

   ☐ filed together with the international application in computer readable form.

   ☐ furnished subsequently to this Authority for the purposes of search.

3. ☐ In addition, in the case that more than one version or copy of a sequence listing and/or table relating thereto has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

4. Additional comments:

Box No. II Priority

1. ☒ The validity of the priority claim has not been considered because the International Searching Authority does not have in its possession a copy of the earlier application whose priority has been claimed or, where required, a translation of that earlier application. This opinion has nevertheless been established on the assumption that the relevant date (Rules 43*bis*.1 and 64.1) is the claimed priority date.

2. ☐ This opinion has been established as if no priority had been claimed due to the fact that the priority claim has been found invalid (Rules 43*bis*.1 and 64.1). Thus for the purposes of this opinion, the international filing date indicated above is considered to be the relevant date.

3. Additional observations, if necessary:

Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	
	No:	Claims	1-26
Inventive step (IS)	Yes:	Claims	
	No:	Claims	1-26
Industrial applicability (IA)	Yes:	Claims	1-26
	No:	Claims	

2. Citations and explanations

see separate sheet

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

Reference is made to the following documents:

D1: EP 1 085 444 A
D2: EP 0 739 115 A

1. The present application does not meet the criteria of Article 33(1) PCT, because the subject-matter of **claims 1, 7 and 13** is not new in the sense of **Article 33(2) PCT**.

1.1. The document D1 discloses (the references in parentheses applying to this document):

a computerized messaging server in an electronic messaging system, comprising:

- a messaging module (par. 36) adapted to control a message database ("message store 130") storing messages exchanged in the messaging system, each message (par. 45) comprising one or more message components ("initial message 10 for context along with the text of replies 412, 414, 416 and 418"); and
- a governance module (par. 42) adapted to control a governance policy database storing a compliancy policy describing rules applicable to the message components ("rules for governing access to message 180"), the governance module further adapted to apply the rules described by the compliance policy to the message components ("control access").

The subject-matter of claim 1 is therefore not new.

1.2. The same reasoning applies, mutatis mutandis, to the subject-matter of the corresponding independent claims 7 and 13, which therefore is also considered not new.

2. Dependent claims 2-6, 8-12 and 14-26 do not contain any features which, in combination with the features of any claim to which they refer, meet the requirements of the PCT in respect of novelty and/or inventive step, see documents D1 and D2 and the corresponding passages cited in the search report.

INTERNATIONAL SEARCH REPORT

Int[illegible]nal Application No
PCT/GB2004/004609

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/58 G06F17/60 H04L29/06

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)
EPO-Internal, WPI Data, PAJ, INSPEC, COMPENDEX

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 03/036492 A (IDETIC, INC; STIERS, TODD) 1 May 2003 (2003-05-01) abstract page 2, line 20 - page 3, line 10 page 4, line 10 - page 6, line 22 figure 1	1-35
A	EP 1 085 444 A (MICROSOFT CORP) 21 March 2001 (2001-03-21) abstract figure 5 paragraph '0012! paragraph '0015! - paragraph '0016! paragraph '0041! paragraph '0043! paragraph '0045!	1-35

-/--

[X] Further documents are listed in the continuation of box C.

[X] Patent family members are listed in annex.

° Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier document but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
15 February 2005	22/02/2005

Name and mailing address of the ISA	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL – 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Poppe, F

Form PCT/ISA/210 (second sheet) (January 2004)

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EDWARDS W K: "The design and implementation of the Montage multimedia mail system" COMMUNICATIONS FOR DISTRIBUTED APPLICATIONS AND SYSTEMS. CHAPEL HILL, APR. 18 - 19, 1991, PROCEEDINGS OF THE CONFERENCE ON COMMUNICATIONS SOFTWARE. (TRICOMM), NEW YORK, IEEE, US, vol. CONF. 4, 18 April 1991 (1991-04-18), pages 47-57, XP010039767 ISBN: 0-87942-649-7 abstract page 51, right-hand column, line 48 - page 52, left-hand column, line 2 -----	1-35

INTERNATIONAL SEARCH REPORT

formation on patent family members

Int ɔnal Application No
PCT/GB2004/004609

Patent document cited in search report		Publication date	Patent family member(s)		Publication date
WO 03036492	A	01-05-2003	WO	03036492 A1	01-05-2003
			US	2003055907 A1	20-03-2003
EP 1085444	A	21-03-2001	US	6704772 B1	09-03-2004
			EP	1085444 A2	21-03-2001
			US	2004148359 A1	29-07-2004

Form PCT/ISA/210 (patent family annex) (January 2004)

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

REC'D 18 FEB 2005
WIPO PCT

To:

see form PCT/ISA/220

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY
(PCT Rule 43*bis*.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/GB2004/004609	01.11.2004	31.10.2003

International Patent Classification (IPC) or both national classification and IPC
H04L12/58, G06F17/60, H04L29/06

Applicant
BLUESPACE GROUP LIMITED

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☐ Box No. VII Certain defects in the international application
- ☐ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA"). However, this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notifed the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of three months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

3. For further details, see notes to Form PCT/ISA/220.

Name and mailing address of the ISA:

European Patent Office - P.B. 5818 Patentlaan 2
NL-2280 HV Rijswijk - Pays Bas
Tel. +31 70 340 - 2040 Tx: 31 651 epo nl
Fax: +31 70 340 - 3016

Authorized Officer

Poppe, F

Telephone No. +31 70 340-4378

Form (PCT/ISA/237) (Cover Sheet) (January 2004)

Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of the international application in the language in which it was filed, unless otherwise indicated under this item.

   ☐ This opinion has been established on the basis of a translation from the original language into the following language , which is the language of a translation furnished for the purposes of international search (under Rules 12.3 and 23.1(b)).

2. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application and necessary to the claimed invention, this opinion has been established on the basis of:

   a. type of material:

   ☐ a sequence listing

   ☐ table(s) related to the sequence listing

   b. format of material:

   ☐ in written format

   ☐ in computer readable form

   c. time of filing/furnishing:

   ☐ contained in the international application as filed.

   ☐ filed together with the international application in computer readable form.

   ☐ furnished subsequently to this Authority for the purposes of search.

3. ☐ In addition, in the case that more than one version or copy of a sequence listing and/or table relating thereto has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

4. Additional comments:

Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	4,11,15,22,26,33
	No:	Claims	1-3,5-10,12-14,16-21,23-25,27-32,34,35
Inventive step (IS)	Yes:	Claims	
	No:	Claims	1-35
Industrial applicability (IA)	Yes:	Claims	1-35
	No:	Claims	

2. Citations and explanations

see separate sheet

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

Reference is made to the following documents:

D1: WO 03/036492 A
D2: EP 1 085 444 A

1. The present application does not meet the criteria of Article 33(1) PCT, because the subject-matter of **claims 1, 9, 14, 20, 25 and 31** is not new in the sense of **Article 33(2) PCT**.

1.1. Document D1 discloses (the references in parentheses applying to this document):

a computer-implemented method of delivering electronic messages, comprising:

- providing a messaging client with information identifying electronic messages available to the messaging client (p. 2, l. 20-26; p. 4, l. 10-16: "The end user receives the non-attachment portions of an email message (...).");
- caching copies of the electronic messages close to the messaging client (p. 5, l. 28-32: "The proxy server 100 stores end users' email messages in its cache 104."); and
- responsive to a request from the messaging client for the identified electronic messages, providing the cached copies of the electronic messages to the messaging client (p. 2, l. 20-26; p. 6, l. 19-22).

The subject-matter of claim 1 is therefore not new.

1.2. Furthermore document D1 discloses (the references in parentheses applying to this document):

a computer system adapted to be used in a messaging system (fig. 1), comprising:

- a cache management module adapted to manage a cache (fig. 1: "e-mail cache"), the cache storing messages exchanged by the messaging system (p. 5, l. 28-32); and

- an interface module adapted to interface with a messaging client (fig. 1: "e-mail client 202") and a messaging server (fig. 1: "source e-mail box 400"), the interface module adapted to receive messages from the server (p. 5, l. 16-22; p. 6, l. 1-6) and provide messages to the client (p. 6, l. 19-22), wherein messages received from the server are stored in the cache (p. 5, l. 28-32) and messages provided to the client are provided from the cache (p. 6, l. 19-22).

The subject-matter of claim 9 is therefore not new.

1.3. The same reasoning applies, mutatis mutandis, to the subject-matter of the corresponding independent computer program product claim 20, which therefore is also considered not new.

1.4. Furthermore document D1 discloses (the references in parentheses applying to this document):

a messaging server in a relational messaging system, comprising:

- a messaging module adapted to control a message database (p. 6, l. 1-6; fig. 1: "source e-mail box 400") storing messages exchanged in the relational messaging system, each message comprising one or more submessages (p. 6, l. 7-11: "MIME attachments");
- a security module adapted to control a security database storing security information for the submessages in the message database (implicit in p. 5, l. 10-15); and
- an interface module adapted to interact with a messaging client and a proxy server, the interface module adapted to provide the security information to the messaging client (p. 6, l. 12-16) and the submessages to the proxy server (p. 6, l. 1-6),

wherein the proxy server is adapted to cache the submessages (p. 5, l. 28-32) and, wherein the messaging client is adapted to retrieve the cached submessages from the proxy server and access the submessages using the security information (p. 6, l. 12-22).

The subject-matter of claim 14 is therefore not new.

1.5. The same reasoning applies, mutatis mutandis, to the subject-matter of the corresponding independent computer program product claim 25, which therefore is also considered not new.

1.6. Furthermore document D1 discloses (the references in parentheses applying to this document):

a relational messaging system adapted to exchange messages with a messaging client, comprising:

- a messaging server adapted to store messages exchanged in the relational messaging system (p. 6, l. 1-6; fig. 1: "source e-mail box 400"), each message comprising a container having links to one or more submessages and further adapted to provide the container to the messaging client (p. 2, l. 20-26; p. 6, l. 7-16); and
- a proxy server adapted to cache submessages received from the messaging server (p. 5, l. 28-32) and further adapted to provide the cached submessages to the messaging client (p. 2, l. 20-26; p. 6, l. 7-16).

The subject-matter of claim 31 is therefore not new.

2. Dependent claims 2-8, 10-13, 15-19, 21-24, 26-30 and 32-35 do not contain any features which, in combination with the features of any claim to which they refer, meet the requirements of the PCT in respect of novelty and/or inventive step, see documents D1 and D2 and the corresponding passages cited in the search report.

INTERNATIONAL SEARCH REPORT

[redacted]tional Application No
PCT/GB2004/000805

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/58 G06F17/60

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)
EPO-Internal, WPI Data, PAJ, INSPEC, COMPENDEX

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 1 085 444 A (MICROSOFT CORP) 21 March 2001 (2001-03-21) abstract figure 5 paragraph '0012! paragraph '0015! - paragraph '0016! paragraph '0043! paragraph '0045! --- -/--	1-23

[X] Further documents are listed in the continuation of box C.

[X] Patent family members are listed in annex.

° Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
26 May 2004	04/06/2004

Name and mailing address of the ISA	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL – 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Poppe, F

Form PCT/ISA/210 (second sheet) (January 2004)

INTERNATIONAL SEARCH REPORT

International Application No
PCT/GB2004/000805

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EDWARDS W K: "The design and implementation of the Montage multimedia mail system" COMMUNICATIONS FOR DISTRIBUTED APPLICATIONS AND SYSTEMS. CHAPEL HILL, APR. 18 - 19, 1991, PROCEEDINGS OF THE CONFERENCE ON COMMUNICATIONS SOFTWARE. (TRICOMM), NEW YORK, IEEE, US, vol. CONF. 4, 18 April 1991 (1991-04-18), pages 47-57, XP010039767 ISBN: 0-87942-649-7 abstract page 51, right-hand column, line 48 -page 52, left-hand column, line 2 ------	1,2,7,8, 10,11, 16,17, 19,22,23

Form PCT/ISA/210 (continuation of second sheet) (January 2004)

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No
PCT/GB2004/000805

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 1085444 A	21-03-2001	US 6704772 B1	09-03-2004
		EP 1085444 A2	21-03-2001

Form PCT/ISA/210 (patent family annex) (January 2004)

PATENT COOPERATION TREATY

REC'D 0 3 JUN 2004
WIPO PCT

From the
INTERNATIONAL SEARCHING AUTHORITY

To:

see form PCT/ISA/220

PCT

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY
(PCT Rule 43*bis*.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/GB2004/000805	27.02.2004	28.02.2003

International Patent Classification (IPC) or both national classification and IPC
H04L12/58, G06F17/60

Applicant
BLUESPACE GROUP LTD

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☒ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☐ Box No. VII Certain defects in the international application
- ☐ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA"). However, this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notifed the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of three months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

3. For further details, see notes to Form PCT/ISA/220.

Name and mailing address of the ISA:

European Patent Office - P.B. 5818 Patentlaan 2
NL-2280 HV Rijswijk - Pays Bas
Tel. +31 70 340 - 2040 Tx: 31 651 epo nl
Fax: +31 70 340 - 3016

Authorized Officer

Poppe, F

Telephone No. +31 70 340-4378

Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of the international application in the language in which it was filed, unless otherwise indicated under this item.

   ☐ This opinion has been established on the basis of a translation from the original language into the following language , which is the language of a translation furnished for the purposes of international search (under Rules 12.3 and 23.1(b)).

2. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application and necessary to the claimed invention, this opinion has been established on the basis of:

   a. type of material:

   ☐ a sequence listing

   ☐ table(s) related to the sequence listing

   b. format of material:

   ☐ in written format

   ☐ in computer readable form

   c. time of filing/furnishing:

   ☐ contained in the international application as filed.

   ☐ filed together with the international application in computer readable form.

   ☐ furnished subsequently to this Authority for the purposes of search.

3. ☐ In addition, in the case that more than one version or copy of a sequence listing and/or table relating thereto has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

4. Additional comments:

Box No. II Priority

1. ☒ The following document has not been furnished:

   ☒ copy of the earlier application whose priority has been claimed (Rule 43*bis*.1 and 66.7(a)).

   ☐ translation of the earlier application whose priority has been claimed (Rule 43*bis*.1 and 66.7(b)).

   Consequently it has not been possible to consider the validity of the priority claim. This opinion has nevertheless been established on the assumption that the relevant date is the claimed priority date.

2. ☐ This opinion has been established as if no priority had been claimed due to the fact that the priority claim has been found invalid (Rules 43*bis*.1 and 64.1). Thus for the purposes of this opinion, the international filing date indicated above is considered to be the relevant date.

3. Additional observations, if necessary:

Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	
	No:	Claims	1-23
Inventive step (IS)	Yes:	Claims	
	No:	Claims	1-23
Industrial applicability (IA)	Yes:	Claims	1-23
	No:	Claims	

2. Citations and explanations

see separate sheet

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

The examination is being carried out on the **following application documents**:

Description, pages:

1-38 as originally filed

Claims, No.:

1-23 as originally filed

Drawings, No.:

1-5 as originally filed

Reference is made to the following documents:
D1: EP-A-1085444
D2: W. K. Edwards, "The Design and Implementation of the MONTAGE Multimedia Mail System", in Proc. of TRICOMM 1991, pp. 47-57

1. The present application does not meet the criteria of Article 33(1) PCT, because the subject-matter of **claims 1, 10 and 19** is not new in the sense of **Article 33(2) PCT**.

1.1. The document D1 discloses (the references in parentheses applying to this document):

> a messaging system for providing messaging to end-users, the system comprising a data store module (113) arranged to store messages sent among the end-users, wherein each message includes one or more submessages (initial message 10, reply 412; par. 45 states that these can be presented to a recipient in one display, that is, as one message) and wherein the data store module is arranged to store the

messages and submessages in a relational manner (par. 15; fig. 5: storage of replies in association with the initial message).

The subject-matter of claim 1 is therefore not new.

1.2. Furthermore it is noted that the subject-matter of claim 1 is also not new over the disclosure of D2 (p. 51, right-hand col., l. 48 - p. 52, left-hand col., l. 2).

1.3. The same reasoning applies, mutatis mutandis, to the subject-matter of the corresponding independent computer program product claim 10 and computer-implemented method claim 19, which therefore are also considered not new.

2. Dependent claims 2-9, 11-18 and 20-23 do not contain any features which, in combination with the features of any claim to which they refer, meet the requirements of the PCT in respect of novelty, the reasons being as follows.

2.1. Referring to claims 2 and 11, a contents module adapted to store submessages of the messages sent among the end-users, wherein a message sent by a sender to a recipient includes one or more references to submessages in the content module is known from D1 (par. 15; fig. 5).

2.2. Referring to claims 3 and 12, certain ones of the submessages being created by different end-users at different times is known from D1 (reply 412 to initial message 10 is clearly created at a different time than initial message 10, and by a different end-user).

2.3. Referring to claims 4 and 13, storing only a single version of each message and/or submessage is known from D1 (implicit in the hierarchical tree model and the use of message IDs for referencing shown in fig. 5; see also par. 45).

2.4. Referring to claims 5, 6, 14, 15, 20 and 21, storing an attribute of the messages and/or submessages in the data store module, where this attribute is a length of time that a message and/or a submessage is valid is known from D1 (implicit in par. 43, which specifies the application of a rule based on this attribute).

2.5. Referring to claims 7, 8, 16, 17, 22 and 23, holding data describing relationships

among the messages and submessages, this data describing submessages within a message is known from D1 (fig. 5). Moreover this is also known from D2 (p. 51, right-hand col., l. 48 - p. 52, left-hand col., l. 2: the table of contents can be seen as data describing submessages (chapters) within a message).

2.6. Referring to claims 9 and 18, having a client interface module for interfacing with client applications utilized by the end-users to access the messaging system is known from D1 (par. 45: displaying the messages implies the presence of a client interface module).

INTERNATIONAL SEARCH REPORT

Inter ıal Application No
PCT/GB2005/001868

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/58 H04L29/06

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ, INSPEC, COMPENDEX

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 03/073711 A (RPOST INTERNATIONAL LIMITED) 4 September 2003 (2003-09-04) abstract page 12, line 3 - page 15, line 20 page 33, line 23 - page 35, line 4 page 44, line 2 - page 45, line 7 page 42, line 11 - page 43, line 28 figure 1 figure 2A figure 6 figure 7A figure 7B ----- -/--	1-21

[X] Further documents are listed in the continuation of box C.

[X] Patent family members are listed in annex.

° Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
8 August 2005	19/08/2005

Name and mailing address of the ISA	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Poppe, F

Form PCT/ISA/210 (second sheet) (January 2004)

INTERNATIONAL SEARCH REPORT

Inter nal Application No
PCT/GB2005/001868

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 2004/012415 A (SECURITY AND STANDARDS LIMITED; POPE, NICHOLAS, HENRY; ROSS, JOHN, GOR) 5 February 2004 (2004-02-05) abstract page 6, line 10 - page 10, line 14 figure 1 figure 2a figure 2b	1,9,17
A	EP 1 085 444 A (MICROSOFT CORP) 21 March 2001 (2001-03-21) abstract figure 5 paragraph '0012! paragraph '0015! - paragraph '0016! paragraph '0041! paragraph '0043! paragraph '0045!	1-21
A	WO 03/036492 A (IDETIC, INC; STIERS, TODD) 1 May 2003 (2003-05-01) abstract page 2, line 20 - page 3, line 10 page 4, line 10 - page 6, line 22 figure 1	1-21
A	EDWARDS W K: "The design and implementation of the Montage multimedia mail system" COMMUNICATIONS FOR DISTRIBUTED APPLICATIONS AND SYSTEMS. CHAPEL HILL, APR. 18 - 19, 1991, PROCEEDINGS OF THE CONFERENCE ON COMMUNICATIONS SOFTWARE. (TRICOMM), NEW YORK, IEEE, US, vol. CONF. 4, 18 April 1991 (1991-04-18), pages 47-57, XP010039767 ISBN: 0-87942-649-7 abstract page 51, right-hand column, line 48 - page 52, left-hand column, line 2	1-21

Form PCT/ISA/210 (continuation of second sheet) (January 2004)

INTERNATIONAL SEARCH REPORT

Information on patent family members

Inter ial Application No
PCT/GB2005/001868

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 03073711 A	04-09-2003	US 2002144154 A1	03-10-2002
		AU 2003224616 A1	09-09-2003
		CA 2476176 A1	04-09-2003
		EP 1476995 A2	17-11-2004
		JP 2005518763 T	23-06-2005
		WO 03073711 A2	04-09-2003
WO 2004012415 A	05-02-2004	AU 2003248959 A1	16-02-2004
		WO 2004012415 A1	05-02-2004
		GB 2391438 A	04-02-2004
		EP 1387551 A1	04-02-2004
EP 1085444 A	21-03-2001	US 6704772 B1	09-03-2004
		EP 1085444 A2	21-03-2001
		US 2004148359 A1	29-07-2004
WO 03036492 A	01-05-2003	WO 03036492 A1	01-05-2003
		US 2003055907 A1	20-03-2003

Form PCT/ISA/210 (patent family annex) (January 2004)

PATENT COOPERATION TREATY

PRED

REC'D 16 AUG 2005
WIPO PCT

PCT 17/11

From the
INTERNATIONAL SEARCHING AUTHORITY

To:

see form PCT/ISA/220

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY
(PCT Rule 43*bis*.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/GB2005/001868	12.05.2005	

International Patent Classification (IPC) or both national classification and IPC
H04L12/58, H04L29/06

Applicant
BLUESPACE GROUP LTD

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☒ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☐ Box No. VII Certain defects in the international application
- ☐ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA"). However, this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notifed the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of three months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

3. For further details, see notes to Form PCT/ISA/220.

Name and mailing address of the ISA:

European Patent Office - P.B. 5818 Patentlaan 2
NL-2280 HV Rijswijk - Pays Bas
Tel. +31 70 340 - 2040 Tx: 31 651 epo nl
Fax: +31 70 340 - 3016

Authorized Officer

Poppe, F

Telephone No. +31 70 340-4378

Form (PCT/ISA/237) (Cover Sheet) (January 2004)

Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of the international application in the language in which it was filed, unless otherwise indicated under this item.

   ☐ This opinion has been established on the basis of a translation from the original language into the following language , which is the language of a translation furnished for the purposes of international search (under Rules 12.3 and 23.1(b)).

2. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application and necessary to the claimed invention, this opinion has been established on the basis of:

   a. type of material:

   ☐ a sequence listing

   ☐ table(s) related to the sequence listing

   b. format of material:

   ☐ in written format

   ☐ in computer readable form

   c. time of filing/furnishing:

   ☐ contained in the international application as filed.

   ☐ filed together with the international application in computer readable form.

   ☐ furnished subsequently to this Authority for the purposes of search.

3. ☐ In addition, in the case that more than one version or copy of a sequence listing and/or table relating thereto has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

4. Additional comments:

Box No. II Priority

1. ☒ The validity of the priority claim has not been considered because the International Searching Authority does not have in its possession a copy of the earlier application whose priority has been claimed or, where required, a translation of that earlier application. This opinion has nevertheless been established on the assumption that the relevant date (Rules 43*bis*.1 and 64.1) is the claimed priority date.

2. ☐ This opinion has been established as if no priority had been claimed due to the fact that the priority claim has been found invalid (Rules 43*bis*.1 and 64.1). Thus for the purposes of this opinion, the international filing date indicated above is considered to be the relevant date.

3. Additional observations, if necessary:

Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	7,8,15,16
	No:	Claims	1-6,9-14,17-21
Inventive step (IS)	Yes:	Claims	
	No:	Claims	1-21
Industrial applicability (IA)	Yes:	Claims	1-21
	No:	Claims	

2. Citations and explanations

see separate sheet

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

Reference is made to the following documents:

D1: WO 03/073711 A
D2: WO 2004/012415 A
D3: EP 1 085 444 A
D4: WO 03/036492 A

1. The present application does not meet the criteria of Article 33(1) PCT, because the subject-matter of claims 1, 9 and 17 is not new in the sense of Article 33(2) PCT.

1.1. Document D1 discloses (referring in general to p. 33, l. 23 - p. 35, l. 4 and fig. 2A):

a computerized messaging server ("Rost Server 14") in an electronic messaging system, comprising:

- a messaging module adapted to control a message database storing messages exchanged in the messaging system ("the system creates a copy of the original message to be stored for later processing", step 203), each message comprising one or more message components ("file attachments");
- an interface module adapted to receive a message component from a messaging client ("the sender's message is transferred to the RPost server", step 202); and
- an audit module adapted to interact with an audit information database storing audit information describing usage of the messaging system ("the system creates a record in a database", step 204), and further adapted to generate tamper-detection data for the message component received from the messaging client and store the generated tamper-detection data in the audit information database ("the system generates and stores a message digest for the message body", step 206; "the system generates and stores a message digest for each attachment included in the message", step 207).

The subject-matter of claim 1 is therefore not new.

1.2. The same reasoning applies, mutatis mutandis, to the subject-matter of the corresponding independent claim 9, which therefore is also considered not new.

1.3. Furthermore document D1 discloses (referring in general to p. 44, l. 2 - p. 45, l. 7 and fig. 7A):

> a computer-implemented method of authenticating messages in an electronic messaging system, comprising:
> - receiving a message component identified as a relational reference in a message container (step 700);
> - receiving tamper-detection data for the message component ("hash value recorded in the body of the receipt"); and
> - authenticating the message component using the tamper-detection data ("generate a hash of the attached file and compare it to the hash value recorded in the body of the receipt", step 708-709).

The subject-matter of claim 17 is therefore not new.

1.4. For the sake of completeness it is noted that the subject-matter of claims 1, 9 and 17 also does not involve an inventive step over the disclosure of D2 (p. 6, l. 23 - p. 7, l. 15; fig. 1).

2. Dependent claims 2-8, 10-16 and 18-21 do not contain any features which, in combination with the features of any claim to which they refer, meet the requirements of the PCT in respect of novelty and/or inventive step, see documents D1, D2, D3 and D4 and the corresponding passages cited in the search report.

C3

Load Distribution for Scalable Web Servers:
Summer Olympics 1996 - A Case Study

Germán Goldszmidt and Andy Stanford-Clark
IBM Watson Research Center
{gsg, andysc}@watson.ibm.com

Abstract

The 1996 Atlanta Olympic Games were the first to have an official Web site. This site had more than 60 Web server hosts, spread across 5 locations in 4 countries, and registered over 192 million hits during the 17 days of the games. This paper describes some of the challenges of building, operating and managing such a large distributed system. In particular, it outlines the use of novel TCP/IP technologies for sharing the load of HTTP requests across local and global servers. Using these technologies enabled us to build a scalable distributed Web server site that was able to comfortably support a huge demand. We also describe some of the management tools used for monitoring the system, logging HTTP requests, and providing secure access to system resources.

1 Introduction

The 1996 Atlanta Olympic Games Web site, www.atlanta.olympic.org, was created and operated by IBM. The name *WOMplex* [1] was used as the umbrella code-name to describe a combination of technologies that included content and application servers, scalable parallel processing hardware, a content authoring environment, load sharing software, and several user-level applications and supporting sub-systems. The WOMplex contains collections of reusable objects that can be dispersed over the Internet. HTML pages are dynamically assembled using these objects.

The primary WOMplex was located in Southbury, Connecticut, and hosted the majority of the Web server nodes. The rest of the nodes were located in academic centers at Cornell (New York), Karlsruhe (Germany) and Keio (Japan), and at the IBM Laboratories in Hursley (UK). Each of these sites hosted mirror replicas of most of the content of the primary site. Some of the content was only served from the primary site, including the rapidly changing "Sneak Peek" images which provided a constantly changing montage of captured images from 38 video feeds from the Games in Atlanta. The WOMplex also delivered live audio from a local Atlanta radio station, and low bit-rate video and audio clips.

Before the games, some people wondered whether the Internet would "survive" the Olympics [6]. Bob Metcalfe, for instance, asked "Can IBM's Womplex handle Internet traffic of Olympian proportions?" [5]. The WOMplex technology, which was pioneered to support large sporting events on the Internet, proved adequate to cope with the demands made by so many visitors to the site.

[1] WOM is an acronym for "Web Object Management".

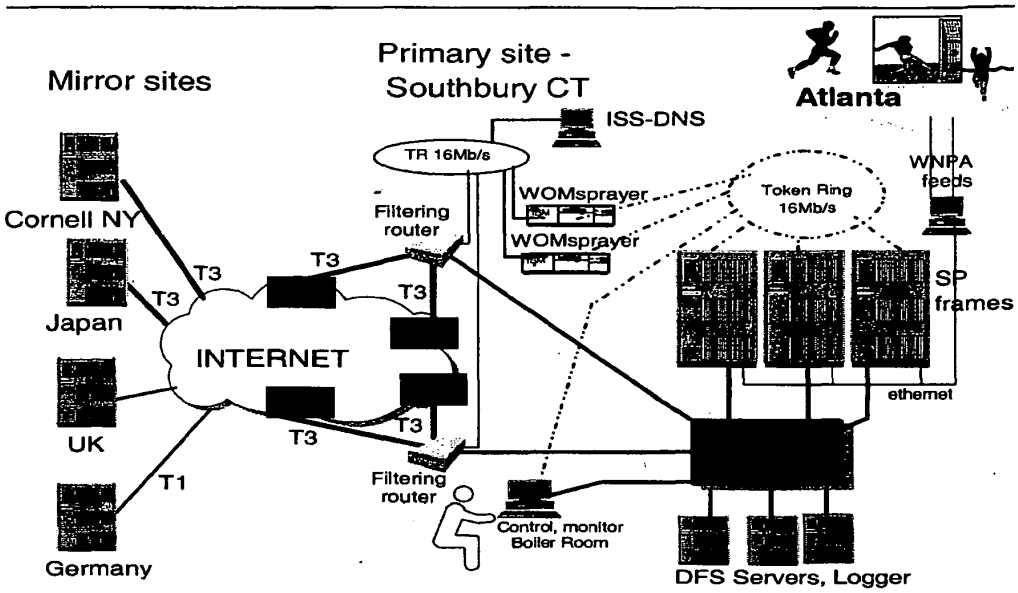


Figure 1: The Infrastructure of the WOMplex for the Atlanta 1996 Olympic Games.

This paper describes some of the challenges of building, operating and managing such a large distributed system. In particular, we describe the techniques and tools used to distribute the incoming Internet workload across a collection of servers, both locally at each location and globally across locations. At each site, incoming traffic directed towards a single IP address, is distributed across multiple back-end WOMserver machines. Between sites, a global load balancing scheme allows logical Web sites to be "moved" between physical server locations. The paper also outlines the WOMplex use of network and system management technologies for monitoring the state of the WOMplex, logging requests and securing access to its resources.

The rest of this paper is organized as follows. Section 2 describes the networking, hardware, and software infrastructure of the Atlanta 1996 WOMplex, Section 3 describes the WOM-DNS and WOMbat tools used to distribute requests across locations. Section 4 describes the WOMsprayer tool used to support distribution of the HTTP requests among the servers at each location. Section 5 outlines some of the approaches taken to deal with the problems of configuration, access security, logging requests, and monitoring in this complex distributed system, and Section 6 concludes the paper.

2 The WOMplex Infrastructure

The infrastructure of the WOMplex is shown in Figure 1. The Southbury site was connected to four major US Internet Network Access Points (NAPs) by DS3 circuits (T3 connections at 45Mb/s). These multiple circuits provided redundancy protection against failure of any of the links. Such failures occurred from time to time at various NAPs. When a link failed, the IP packets were either automatically routed over the

other links, or they could be manually routed round the trouble spots. The primary site was also connected to Atlanta by 56Kb/s leased lines which provided the feed of official games results by the *World News Press Agency (WNPA)* to the WOMplex databases.

Incoming HTTP requests from the Internet were forwarded via IP filtering routers to a WOMsprayer (Section 4) which then routed them to one of the servers. DNS requests were forwarded to a modified name server, WOM-DNS, (Section 3). The filtering IP routers, WOMsprayer and WOM-DNS were attached to an *"external"* 16Mb/s Token Ring.

The network bandwidth requirements for a Web site are highly asymmetrical - one incoming IP packet containing a single "GET" HTTP request can provoke the transfer of potentially many megabytes of outgoing data. We took advantage of this asymmetry by using two separate LANs for incoming and outgoing traffic. Incoming HTTP requests were forwarded by the WOMsprayer to the back-end servers via a second Token-Ring. Outbound traffic (HTTP responses) from the WOMplex servers were forwarded over a 100Mb/s ATM switch directly to the IP routers, and from there via the NAPs to the corresponding remote clients.

The infrastructure of the WOMplex was built around IBM RS/6000 Scalable Parallel systems (referred to hereafter as SP) at each of the mirror sites. Control access to the SP frames at the remote sites was arranged using secure Internet connections. Every processor in the WOMplex could be controlled - right down to power cycling - from Southbury. In the main site, all processors were connected to a Token Ring, an Ethernet, and an ATM network. Other sites used combinations of Ethernet, FDDI, Token Ring, and the SP Switch. Transarc's Distributed File System (DFS) was used to insure that all the servers at all the sites had access to the same data. A big advantage of DFS was the client-side caching: during the Olympics, most clients achieved a 98 % cache hit rate.

WOMplex separates logical services from physical servers, allowing a Web site to be hosted from any location in the WOMplex. IP aliasing is used within each location to separate logical servers from physical nodes. All externally visible IP addresses are "virtual", that is, they are aliases of several hosts. Each host also has its own private IP address, but these addresses are never visible to the outside world. In the event of a machine failure, the IP addresses that it supported can be readily transferred to another machine. This is done by adding IP aliases to the new machine, and flushing the ARP cache on the nearest IP router so that it is forced to re-establish the new IP to MAC address mappings. This operation can be automated using SNMP-based commands to the IP router.

3 Global load balancing

Two mechanisms were implemented to allow Web traffic to be effectively spread across the multiple sites comprising the WOMplex: WOM-DNS, and the WOMbat.

3.1 WOM-DNS

WOM-DNS is a load balancing DNS method that allows the DNS administrator to configure a table of Internet domain names to several IP address mappings. A WOMagent runs on each of the servers designated to serve any given name and collects metrics from each host. Incoming resolution requests for that name to WOM-DNS receive the IP address of the most appropriate server, according to the health, load, and the configured balancing policy for that name.

Two balancing policies can be configured: ROUND and FAIL. The ROUND policy does a weighted round-robin of the IP addresses in the list associated with that name, distributed according to two factors: the server activity profile, as determined from the inputs provided by the WOMagents, and a server weighting factor, which is used to make allowance for different server capacity at each site - e.g., different numbers of WOMserver nodes at each site. The FAIL policy attempts to use the first IP address in the list - the preferred site - until that site becomes unavailable for some reason. At that time it switches to a weighted round-robin of the remaining IP addresses in the list, the same as in the ROUND policy.

This system allows a single logical Web site (defined by a specific DNS name) to be distributed across multiple physical locations, each of which could be a cluster of servers controlled by a WOMsprayer. Additional servers can be dynamically switched-in, if they are required, to cope with excess load. Much of the flexibility of WOM-DNS (in contrast to standard DNS) is that the name server can be dynamically re-configured through both script commands and a graphical interface. WOM-DNS also enables automatic fail-over to other servers in the event of server down-time. This can either take the form of fail-over to other servers from a preferred, or primary, server; or the omission from the round-robin schedule of any of a set of servers which is unresponsive.

3.2 WOMbat

The WOMbat is a mechanism which uses ICMP ECHO packets (i.e. like the PING program) to echo-locate a client machine, and thereby determine which of the WOMplex sites it is nearest to, in network round-trip terms. With this knowledge, the WOMserver can then divert the client to its nearest server site, where they will hopefully receive faster service thanks to the closer network connectivity to that site.

When a client browser visits a WOMplex site for the first time, the client does not have a "Browser Cookie" for that site. The WOMserver generates a Cookie, (a unique key that identifies the client) and sends it to the browser for use whenever it visits that site. The Cookie is also stored in the *Global Profile Daemon (GPD)*, which is a database of profile information, keyed by browser Cookie.

At this time, the WOMserver also invokes, asynchronously in the background, the *WOMbat Driver*, before composing the HTML page requested by the client and returning it to them. The WOMbat Driver sends details of the client's IP address and its browser Cookie to a WOMbat daemon at each geographical site in the WOMplex. Each WOMbat daemon sends out a volley of ICMP echo request packets to the client IP address, and measures the round-trip time of any responses that return within a pre-specified timeout period.

Many client IP addresses are associated with proxy firewalls, which will typically return the ICMP echo response. Since such firewalls are usually located at the same site or in geographical proximity to the client, such responses are good approximations. Unfortunately, some sites may block the ICMP echo requests, e.g., via IP packet filtering routers. In this case, an approximation of the network distance could be obtained via other (slower) tools like *traceroute*.

After collecting the ICMP echo responses from the client's network, the WOMbat daemons send to the WOMbat Driver their average network response time. The WOMbat Driver waits for a specified amount of time for the WOMbat daemons to respond, and then determines which WOMserver site is closest to the client by comparing their responses. Server availability information can also be factored into this calculation, so that the closest site might be rejected in favour of a more distant but less busy site. The identity of the preferred site is stored in the user's profile in the GPD.

Next time that a client browser requests a page, the WOMserver interrogates the

GPD, using the browser Cookie as a key. If the preferred site is the local site, the requested page is built with embedded local links and returned to the client. But if the preferred site for that client is not the local site, the page is built with all embedded URL links pointing to the preferred site.

Assume, for example, that the preferred site for a French client was given by the URL `www.hursley.atlanta.olympic.org`, and a local link containing data that is mirrored to all sites would normally have been: `<IMG SRC="/data/i.gif">`. Then the link that is sent back to this client is re-written as:
`<IMG SRC="http://www.hursley.atlanta.olympic.org/data/i.gif">`. When the client browser follows any of these URL links, they are retrieved from the preferred site, and from that point on, all mirrored URL links will point back to that site. Therefore, the service for this client has effectively been "moved" to the preferred site.

4 Local Load Sharing

The purpose of load-sharing is to reduce the average response time by smoothing out transient peak overload on some of the nodes of a distributed system [4]. The WOMsprayer is a software TCP connection router, that serves as a load-sharing front-end to a collection of back-end servers. The WOMsprayer host offers a collection of *Virtual IP Addresses (VIPA)* to the outside world. For example, the primary site exported 2 different VIPAs, one for www.atlanta.olympic.org, (204.146.176.111) and another one for results.atlanta.olympic.org (204.146.176.113). These VIPAs are shared by a dynamically changing set of servers which alias some or all of these addresses but do not advertise them via ARP. The WOMsprayer allocates requests based on their port numbers. For instance, some servers of www.atlanta.olympic.org supported only port 80, others only port 8080, and some supported both. The allocation of servers and ports was dynamically configurable.

All requests arriving to these VIPAs were routed via the WOMSprayer. When a new TCP connection request is made to any VIPA, the WOMsprayer chooses one server among the set of active back-end servers for the particular `<VIPA, port>` combination. The WOMSprayer modifies the MAC address field of the incoming IP packet to that of the chosen server's interface card, and then forwards the packet via the corresponding LAN (e.g. the internal Token Ring). The IP packet itself is not changed.

The server interface will pick up the connection request IP packet. The server then does the required accept to create the TCP connection, returning a response IP packet directly to the remote client. The network path from the server back to the client (e.g. Web browser) is via the back-end server's default route, (e.g. the ATM network). Hence the returning server-to-client packets are not intercepted by the WOMsprayer. The WOMsprayer maintains a table of currently active connections, and any further incoming packets for that TCP session, such as packet acknowledgements, are forwarded to the corresponding back-end server. Notice that the WOMsprayer only sees the client-to-server half of each TCP connection. Hence it must handle several cases of connection termination by timestamps and garbage collection.

The WOMsprayer mantains several counters that summarize the state of the connections, e.g., the number of active connections for each server. These counters, together with a set of configuration variables, could be used to prepare an SNMP MIB for monitoring the WOMsprayer. We used a combination of scripted commands, which could also be invoked from a Web browser. One of the main uses of the configuration commands was to temporarily add and remove servers from servicing a giving IP address. This allowed to dynamically upgrade the software and content of the servers without disrupting the service and without dropping existing connections.

A probabilistic weighted round-robin algorithm, implemented as part of the kernel extension to the TCP/IP stack, chooses the back-end server for each incoming TCP connection. The weights associated with each server correspond to the proportion of new connection requests that will be allocated. These weights are dynamically assigned in real-time by a user-level process, the WOMsprayer *Manager*, based on current state information. The WOMsprayer Manager takes into account several sources of information about the state of the servers, including feedback host metrics, local counters for input metrics and service metrics. Input metrics provide a current estimate of the rate of incomming connections, e.g., number of new connections in the last t seconds. Host metrics represent the load of a server, and are based on system-wide estimates, like the number of active processes or the number of allocated mbufs. Service metrics are specific to each application, for instance, the time measured to retrieve an HTML page.

When the algorithm finds that there is an imbalance in the allocation of resources it changes the weights to correct the problem. This allocation method implements a dynamic feedback loop that automatically controls the proper spread of the service requests. During the Olympic games, it reacted effectively to the bursty nature of the requests traffic, discovered and routed around server failures, and was very efficient in balancing the load among all the active servers.

The two main strategies for improving performance using load distribution are *load-balancing* and *load-sharing* [1]. Load-balancing strives to equalize the workload among servers. Load-sharing attempts to smooth out transient peak overload periods on some nodes [3]. Load-balancing strategies typically require many more resources and are better suited for large batch jobs. The WOMplex HTTP requests were a mixed load, with many short requests for a page and some very long requests (e.g. radio). The resources necessary to achieve an approximate load-balance among the WOMplex servers would have been so high that they would have outweighted any potential benefits. Hence we designed the WOMsprayer Manager to provide effective and efficient load-sharing.

Notice that at any particular point in time, there maybe some imbalance in the utilization state of the servers. This is unavoidable as the processes that consume resources (e.g. dynamic page creation) work much faster than the WOMsprayer feedback. To forward feedback fast enough to observe all relevant events would have required too much cycles and bandwidth, far outweighting any potential benefits. Still, the algorithm was very sensitive to host overload and achieved good load balancing. The local WOMsprayer feedback provided by the connection counters gave a sufficient indication of load with low overhead. Whenever a server host became overloaded ("over-heated"), it was quickly quiesced, giving it time to "cool-down" (i.e. complete all in-progress requests). When the server host was again relatively "cool" (usually a few seconds later), it was automatically returned to the working set. Any server that becomes overloaded has a large impact on the quality of service of the site, as the number of enqueued requests grows quickly. Therefore, automatically removing, cooling and returning a server to service was of paramount importance.

Similarly, any server which experienced failure problems was dynamically and automatically excluded from the WOMsprayer schedule, and no new connections were routed to it until the problem was resolved. For instance, a server may develop some hardware or software problems that resulted in connection requests accumulating in its queue. The WOMsprayer would automatically detect that the connection count for that server reached a high threshold and react by lowering the weight of the server, even to 0. Hence, this method effectively prevents incoming traffic from hitting a non-responsive server and getting "Connection Refused" error messages.

Each server host runs a WOMagent daemon, which executes a host metric script,

which determines how busy, in some sense, the server is. The metric can be any script that returns an integer. For example, the metric "`ps -ef | wc -l`" reports the number of processes currently running on the system. More elaborate metrics can be contrived, to evaluate other host performance limits. One server host metric that was used for the Olympics counted the number of ESTABLISHED socket connections on port 80 of each server. Of course, compiled programs can be used to generate the metric values. The WOMagents periodically report to a central WOMmonitor process, which collates agent responses, starts and stops the remote WOMagents, and reacts appropriately if the WOMagents do not report-in for an extended period of time. The WOMmonitor transfers a load profile of server activity to the WOMsprayer Manager, which uses it to compute the weights for each server. Any server whose WOMagent is not responding is removed from the load profile, and hence is given a non-active weight by the WOMsprayer.

For reliability, the WOMplex at the primary site included two WOMsprayers, but only one of them was active at any given time. If the active WOMsprayer host failed, the filtering router ARP table would be updated so that new requests will arrive to the secondary WOMsprayer. A detailed presentation of the *NetDispatcher* technology that was used in the WOMsprayer is given in [2]. A theoretical qualititative and quantitative analysis of adaptive load sharing algorithms is given in [3].

5 Operational Management

Operational Management of the WOMplex involved configuration of all the sites, monitoring and control of all subsystems, logging and securing all operations and systems.

5.1 Remote Configuration and Control

Network Configuration. The global configuration of the Web site required keeping configuration files for the WOMsprayers, DNS servers, and the IP address aliases of more than 80 server nodes. This was achieved using a master database, where the configuration rules where defined (e.g. which site name on which IP address was going to be served by which subset of nodes at which site). A program then took this data and built the set of configuration files, customised to each machine as required. This customisation was defined by site-specific rules, which included site-specific DNS suffixes and IP addresses. The configuration files, once built, were than copied using an encrypted protocol across the internet to the target machines. When all files were in place, a series of remote shell operations (also encrypted) triggered the various daemons to re-read the new configuration files.

Control. Operators in the Web Operations Centre (WOC) in Southbury, CT, USA, were able to control all physical aspects of the SP2 frames located across the world, through the X-windows based GUI to PSSP, a system control package. Most usefully, this gave us the ability to reset or power-cycle nodes, without having to rely on anyone being physically present at the remote site. With an operation of this size and complexity, it was essential to have everything controlled by scripts, partly so that they could easily be re-used in other situations, and partly for auditability: if the operation failed, we could look at the script to see exactly what sequence of steps had been performed.

5.2 Monitoring

System Logs. All Unix system logs (systlogd) were channeled to a single machine, where a set of filters were applied to the reams of information to make sure that the ones which were important were displayed on the screen and were not lost in a morass of trivial output. We also attached a console beep to the most important messages, which effectively gave us an audible alert when something critical was happening, such as a detected intrusion onto one of the systems, or a port scan in progress, or critical failure of hardware or software.

Performance. We positioned HTTP monitors at strategic locations round the world to frequently request pages from our various servers, and to measure the response time. After comparison with other major sites on the Internet, an alarm would be raised (Email to a pager) if a certain threshold was exceeded on the response time. This mechanism brought problems to our attention even when from our view (the server end of the operation) everything was looking good.

5.3 Security

TCP/IP Security. Such a high visibility site was an obvious target for hackers, and hence security was a major concern. The IBM Global Security Lab (*the "Ethical Hackers"* team) performed thorough penetration testing of the site before the games, and provided valuable advice on security procedures. Filtering IP routers provided the main protection against invasion. Inter-working with the remote mirror sites involved setting filter rules at the remote sites which would allow only certain IP packets in and out from certain places. The primary site (Southbury) had to have extensive access to systems inside the DMZ (De-Militarized Zone) at the remote sites, but the rest of the Internet could only have HTTP port 80 access. Also, the trusted mirror sites had to have access to hosts inside the Southbury DMZ.

DNS Security. WOM-DNS as the primary name server provided another level of security. WOM-DNS was configured to have authority over one or more specified sub-domains. A name resolution request may arrive for either a sub-domain over which WOM-DNS does not have authority or for a name which is not being used for load-balancing. In both cases, WOM-DNS passes the request over to one of two back-end, ordinary, name servers. WOM-DNS is configured with a table of sub-net specifications for "trusted" domains. The name resolution request is matched against this table. If there is a match, then the client is deemed to be internal, and is passed to one back-end DNS server, otherwise it is deemed to be external and is sent to the other DNS server. Even with a distributed, multi-site operation, it is possible for the client machines that need to have access to all of the main site's internal server names and addresses to do so. At the same time, the rest of the Internet is only able to resolve the limited set of names and addresses which the DNS administrator chooses to make available. The WOM-DNS also blocks the DNS Zone Transfer operation, which is a popular method of gaining access to lists of internal machines to be explored as hacking targets.

5.4 Centralized HTTP Logging

An important aspect of WOMplex operation is the collection of HTTP access and error logs from the servers distributed across the sites, and bringing them to a central location for monitoring and analysis. The *Grim Reaper* tool is a tree-structured "plumbing" of

sockets and Unix pipes. This tool streamed log records from each server to a confluence point at each site, and then out across the Internet to a central collection point on a primary site machine called the *Logger*. When all the log data is streaming nicely, the task of the Grim Reaper is quite straightforward. The problems start occurring when one of three things happens: (1) the servers produce log data faster than the local Grim Reaper can take it from them; (2) network delays or failures cause hold-ups in the Grim Reaper plumbing which cause a back-log of data; and (3) the final receiver of the logs can not cope with the rate at which they are arriving from the Grim Reaper. Any of these scenarios, coupled with the simple mandate that no log records may be discarded, make the Grim Reaper's job difficult.

At any stage, log streams may be diverted into a file on local disk, to hold data which for some reason can not be sent onwards. Later, when connections are restored, or congestion eases, the files are spooled out back into the network. The Grim Reaper also supported off-line transfer of batched-up log data (for example late at night when there was less network traffic), and also on-the-fly compression of data for transmission across slower network connections. After arrival at the logger machine, the log streams were transferred across a high-speed link to a parallel mainframe where they were pre-processed by a series of filters and then the individual log records were added to a DB2 database. From here, SQL queries could be used to extract useful information from the logs, including hit counts and rates, online data mining, and real-time Click Tracking. The raw log data was also spooled to tape for later analysis.

A good indication of the site performance as a whole was given by a large screen with dozens of windows displaying the httpd log files from a subset of the servers. If these windows were all in motion, then everything was good - the WOMplex was serving pages across all our server locations. If any window stopped, or if the whole screen froze, then we knew there was trouble: most often a network problem which stopped people from reaching some part of the WOMplex.

5.5 The WOMplex Boiler Room

The Boiler Room system allowed statistics to be gathered from various parts of the WOMplex, its applications, and sub-systems. The system was structured as follows. Each application to be monitored was instrumented with a funneld (funnel daemon) client. The protocol used throughout this system was IRC (Internet Relay Chat), transmitting data in a textual format, with each monitored system using a different IRC channel number. The client code connected to a funneld instance, running on each server. The funneld instances connected to other funneld's in a tree structure. At the root of the tree of funnel daemons was an IRC server.

The Boiler Room monitors were Java applets embedded in HTML pages on the WOMplex Web site. These applets used IRC client classes, written in Java, which connected back to the IRC server machine. The monitor applications then tuned-in to the IRC channels that they required to construct their displays, and waited for information to arrive. A variety of dials, counters and indicators were implemented as Java classes which could be combined in various ways to display the statistics for any system. The Boiler Room system also had facilities for tailing log files, and receiving data through sockets and IPC message queues, so it was easy to hook up monitors for almost anything. A special chart recorder client was used to log IRC channel data for subsequent analysis.

Boiler Room monitors delivered vital statistics on the loading and health of the various servers in the system. The monitoring effort was distributed amongst a number of people, as there were too many moving bars and flashing lights for any one person to take in for long periods of time. Since the Olympics, a lot of our efforts have been ap-

plied to making the management of the WOMplex-type infrastructure more integrated, using Tivoli Managed Environment products. To permit cost-effective 24x7 operation, a greatly simplified monitoring system is essential, with all commonly required operations fully scripted and automated wherever possible.

6 Conclusions

Exponential growth in Internet traffic and the increasing complexity of CPU-intensive Web-based applications, pose increasing challenges to popular Web sites. As many main Internet links (e.g. across the Atlantic) are saturated for long periods of time, geographic distribution is necessary to serve an international audience with adequate response times. Anticipating such problems for the Olympic Games Web site, we devised a combination of technologies based on TCP connection routing and name resolution load balancing. The combined technologies enable almost unlimited scalability both at the local site level (LAN) and across multiple geographical locations (WAN). These technologies were crucial in the successful global deployment of the Atlanta 1996 Olympic WOMplex, which served a record number of requests world-wide.

Operating and managing this large and complex distributed site was a very challenging task. It was essential to have everything controlled by scripts, that could easily be re-used and provided an audit trail. Automated notification of critical events by filtering of system logs enabled a small group of operators at the main site to manage the entire WOMplex. The efficiency of the WOMplex management components and tools was crucial to the success of the global Web site. Since this event, many of these technologies have been upgraded and succesfully integrated and used at many other large-scale Web sites worldwide.

Acknowledgments. The success of this Web site was due to the combined efforts of many people. Unfortunately, we cannot list all of them in the space available here, so we only acknowledge those who worked most closely with us in the technologies described in this paper. These include Michael Factor, Ed Geraghty, Dave Grossman, Jane Harper, Guerney Hunt, Sean Martin, Chet Murthy, Scott Penberthy, Frank Schwichtenberg, Graham Wallis, and Ron Woan.

References

[1] D. L. Eager, E.D. Lazowska, and J. Zahorjan. Adaptive Load Sharing in Homogeneous Distributed Systems. *IEEE Transactions on Software Engineering*, 12(5):662–675, May 1986.

[2] Germán Goldszmidt and Guerney Hunt. NetDispatcher: A TCP Connection Router. Technical report, IBM Research, Hawthorne, New York, July 1997. RC 20853.

[3] O. Kremien and J. Kramer. Methodical Analysis of Adaptive Load Sharing Algorithms. *IEEE Transactions on Parallel and Distributed Processing*, 3(6), November 1992.

[4] O. Kremien, J. Kramer, and J.N. Magee. Scalable and Adaptive Load Sharing for Distributed Systems. *IEEE Parallel and Distributed Technology*, 1(3), 1993.

[5] Bob Metcalfe. Can IBM's Womplex handle Internet traffic of Olympian proportions?, July 1 1996. InfoWorld.

[6] Bob Metcalfe. Summer Olympics to use IBM's 'extranet' 390s, but beware the Ides of July, April 8 1996. InfoWorld.

PTO/SB/06 (07-06)
Approved for use through 1/31/2007. 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.

PATENT APPLICATION FEE DETERMINATION RECORD

Substitute for Form PTO-875

Application or Docket Number	Filing Date	☐ To be Mailed
11/623,653	01/16/2007	

APPLICATION AS FILED – PART I

	(Column 1)	(Column 2)	SMALL ENTITY ☒		OR	OTHER THAN SMALL ENTITY	
FOR	NUMBER FILED	NUMBER EXTRA	RATE ($)	FEE ($)		RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A			N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A			N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A			N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*	X $ =		OR	X $ =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*	X $ =			X $ =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $250 ($125 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).						
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))							
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL			TOTAL	

APPLICATION AS AMENDED – PART II

		(Column 1)		(Column 2)	(Column 3)	SMALL ENTITY		OR	OTHER THAN SMALL ENTITY	
AMENDMENT	01/15/2010	CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)		RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	* 22	Minus	** 22	= 0	X $26 =	0	OR	X $ =	
	Independent (37 CFR 1.16(h))	* 3	Minus	***3	= 0	X $110 =	0	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	0	OR	TOTAL ADD'L FEE	

		(Column 1)		(Column 2)	(Column 3)					
AMENDMENT		CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	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 number found in the appropriate box in column 1.

Legal Instrument Examiner:
/FREDERICK E. BRISCOE/

This collection of information is required by 37 CFR 1.16. 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 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.

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

758 7590 02/12/2010

FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041

EXAMINER	
HOCKER, JOHN P	
ART UNIT	PAPER NUMBER
2158	

DATE MAILED: 02/12/2010

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
11/623,653	01/16/2007	Justin Marston	23682-12381	5243

TITLE OF INVENTION: DETERMINING RELEVANCE OF ELECTRONIC CONTENT

APPLN. TYPE	SMALL ENTITY	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	YES	$755	$300	$0	$1055	05/12/2010

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 SMALL ENTITY status shown above.

If the SMALL ENTITY is shown as YES, verify your current SMALL ENTITY status:

A. If the status is the same, pay the TOTAL FEE(S) DUE shown above.

B. If the status above is to be removed, check box 5b on Part B - Fee(s) Transmittal and pay the PUBLICATION FEE (if required) and twice the amount of the ISSUE FEE shown above, or

If the SMALL ENTITY is shown as NO:

A. Pay TOTAL FEE(S) DUE shown above, or

B. If applicant claimed SMALL ENTITY status before, or is now claiming SMALL ENTITY status, check box 5a on Part B - Fee(s) Transmittal and pay the PUBLICATION FEE (if required) and 1/2 the ISSUE FEE shown above.

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 and an extra copy of the form should be submitted. 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: Utility patents issuing on applications filed on or after Dec. 12, 1980 may require payment of maintenance fees. It is patentee's responsibility to ensure timely payment of maintenance fees when due.

PTOL-85 (Rev. 08/07) Approved for use through 08/31/2010.

PART B - FEE(S) TRANSMITTAL

Complete and send this form, together with applicable fee(s), to: Mail **Mail Stop ISSUE FEE**
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450
or Fax (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)

758 7590 02/12/2010

FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041

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 facsimile transmitted to the USPTO (571) 273-2885, on the date indicated below.

(Depositor's name)

(Signature)

(Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
11/623,653	01/16/2007	Justin Marston	23682-12381	5243

TITLE OF INVENTION: DETERMINING RELEVANCE OF ELECTRONIC CONTENT

APPLN. TYPE	SMALL ENTITY	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	YES	$755	$300	$0	$1055	05/12/2010

EXAMINER	ART UNIT	CLASS-SUBCLASS
HOCKER, JOHN P	2158	707-002000

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/SB/122) attached.

❑ "Fee Address" indication (or "Fee Address" Indication form 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, 1_____

(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. 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 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.

(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. The following fee(s) are submitted:

❑ Issue Fee

❑ Publication Fee (No small entity discount permitted)

❑ Advance Order - # of Copies _____

4b. Payment of Fee(s): **(Please first reapply any previously paid issue fee shown above)**

❑ A check is enclosed.

❑ Payment by credit card. Form PTO-2038 is attached.

❑ The Director is hereby authorized to charge the required fee(s), any deficiency, or credit any overpayment, to Deposit Account Number _____ (enclose an extra copy of this form).

5. **Change in Entity Status** (from status indicated above)

❑ a. Applicant claims SMALL ENTITY status. See 37 CFR 1.27.

❑ b. Applicant is no longer claiming SMALL ENTITY status. See 37 CFR 1.27(g)(2).

NOTE: The Issue Fee and Publication Fee (if required) will not be accepted from anyone other than the applicant; a registered attorney or agent; or the assignee or other party in interest as shown by the records of the United States Patent and Trademark Office.

Authorized Signature _____ Date _____

Typed or printed name _____ Registration No. _____

This collection of information 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 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, 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.

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.
11/623,653	01/16/2007	Justin Marston	23682-12381	5243

758 7590 02/12/2010

FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041

EXAMINER	
HOCKER, JOHN P	
ART UNIT	PAPER NUMBER
2158	

DATE MAILED: 02/12/2010

Determination of Patent Term Adjustment under 35 U.S.C. 154 (b)
(application filed on or after May 29, 2000)

The Patent Term Adjustment to date is 184 day(s). If the issue fee is paid on the date that is three months after the mailing date of this notice and the patent issues on the Tuesday before the date that is 28 weeks (six and a half months) after the mailing date of this notice, the Patent Term Adjustment will be 184 day(s).

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 Customer Service Center of the Office of Patent Publication at 1-(888)-786-0101 or (571)-272-4200.

PTOL-85 (Rev. 08/07) Approved for use through 08/31/2010.

Notice of Allowability	Application No.	Applicant(s)	
	11/623,653	MARSTON ET AL.	
	Examiner	Art Unit	
	JOHN P. HOCKER	2158	

-- 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 *Interview and RCE(filed: 01/15/2010)*.

2. ☒ The allowed claim(s) is/are *Claims 1, 7 – 11, 17 and 23 – 37 are allowed (Renumbered as claims 1 – 22).*.

3. ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
   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.

4. ☐ A SUBSTITUTE OATH OR DECLARATION must be submitted. Note the attached EXAMINER'S AMENDMENT or NOTICE OF INFORMAL PATENT APPLICATION (PTO-152) which gives reason(s) why the oath or declaration is deficient.

5. ☐ CORRECTED DRAWINGS (as "replacement sheets") must be submitted.
   (a) ☐ including changes required by the Notice of Draftsperson's Patent Drawing Review (PTO-948) attached
      1) ☐ hereto or 2) ☐ to Paper No./Mail Date _____.
   (b) ☐ 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. ☐ Notice of Draftperson's Patent Drawing Review (PTO-948)
3. ☒ Information Disclosure Statements (PTO/SB/08), Paper No./Mail Date 01/15/2010
4. ☐ Examiner's Comment Regarding Requirement for Deposit of Biological Material
5. ☐ Notice of Informal Patent Application
6. ☒ Interview Summary (PTO-413), Paper No./Mail Date *01/2010* .
7. ☐ Examiner's Amendment/Comment
8. ☒ Examiner's Statement of Reasons for Allowance
9. ☐ Other _____.

/JPH/ (Art Unit:2158)	/Mohammad Ali/ Supervisory Patent Examiner, Art Unit 2158

Interview Summary	Application No.	Applicant(s)	
	11/623,653	MARSTON ET AL.	
	Examiner	Art Unit	
	JOHN P. HOCKER	2158	

All participants (applicant, applicant's representative, PTO personnel):

(1) *JOHN P. HOCKER*. (3)_____.

(2) *Chris King*. (4)_____.

Date of Interview: *Januaroy 14th, 201*.

Type: a)☒ Telephonic b)☐ Video Conference
c)☐ Personal [copy given to: 1)☐ applicant 2)☐ applicant's representative]

Exhibit shown or demonstration conducted: d)☐ Yes e)☐ No.
If Yes, brief description: _____.

Claim(s) discussed: *Claims 1, 11 and 17 (and all their dependents)*.

Identification of prior art discussed: *Malik*.

Agreement with respect to the claims f)☒ was reached. g)☐ was not reached. h)☐ N/A.

Substance of Interview including description of the general nature of what was agreed to if an agreement was reached, or any other comments: *Examiner and Applicant successfully negotiated claim lanuage that embodied the applicant's claimed inventive concept, and at the same time was distinguishable over the prior art. A preliminary agreement was reached that the proposed amended claims were allowable (as amended. Examiner stated an updated search would be performed prior to allowance.)*.

(A fuller description, if necessary, and a copy of the amendments which the examiner agreed would render the claims allowable, if available, must be attached. Also, where no copy of the amendments that would render the claims allowable is available, a summary thereof must be attached.)

THE FORMAL WRITTEN REPLY TO THE LAST OFFICE ACTION MUST INCLUDE THE SUBSTANCE OF THE INTERVIEW. (See MPEP Section 713.04). If a reply to the last Office action has already been filed, APPLICANT IS GIVEN A NON-EXTENDABLE PERIOD OF THE LONGER OF ONE MONTH OR THIRTY DAYS FROM THIS INTERVIEW DATE, OR THE MAILING DATE OF THIS INTERVIEW SUMMARY FORM, WHICHEVER IS LATER, TO FILE A STATEMENT OF THE SUBSTANCE OF THE INTERVIEW. See Summary of Record of Interview requirements on reverse side or on attached sheet.

. /JPH/ (Art Unit: 2158)	/Mohammad Ali/ Supervisory Patent Examiner, Art Unit 2158

Summary of Record of Interview Requirements

Manual of Patent Examining Procedure (MPEP), Section 713.04, Substance of Interview Must be Made of Record
A complete written statement as to the substance of any face-to-face, video conference, or telephone interview with regard to an application must be made of record in the application whether or not an agreement with the examiner was reached at the interview.

Title 37 Code of Federal Regulations (CFR) § 1.133 Interviews
Paragraph (b)

In every instance where reconsideration is requested in view of an interview with an examiner, a complete written statement of the reasons presented at the interview as warranting favorable action must be filed by the applicant. An interview does not remove the necessity for reply to Office action as specified in §§ 1.111, 1.135. (35 U.S.C. 132)

37 CFR §1.2 Business to be transacted in writing.
All business with the Patent or Trademark Office should be transacted in writing. The personal attendance of applicants or their attorneys or agents at the Patent and Trademark Office is unnecessary. The action of the Patent and Trademark Office will be based exclusively on the written record in the Office. No attention will be paid to any alleged oral promise, stipulation, or understanding in relation to which there is disagreement or doubt.

The action of the Patent and Trademark Office cannot be based exclusively on the written record in the Office if that record is itself incomplete through the failure to record the substance of interviews.

It is the responsibility of the applicant or the attorney or agent to make the substance of an interview of record in the application file, unless the examiner indicates he or she will do so. It is the examiner's responsibility to see that such a record is made and to correct material inaccuracies which bear directly on the question of patentability.

Examiners must complete an Interview Summary Form for each interview held where a matter of substance has been discussed during the interview by checking the appropriate boxes and filling in the blanks. Discussions regarding only procedural matters, directed solely to restriction requirements for which interview recordation is otherwise provided for in Section 812.01 of the Manual of Patent Examining Procedure, or pointing out typographical errors or unreadable script in Office actions or the like, are excluded from the interview recordation procedures below. Where the substance of an interview is completely recorded in an Examiners Amendment, no separate Interview Summary Record is required.

The Interview Summary Form shall be given an appropriate Paper No., placed in the right hand portion of the file, and listed on the "Contents" section of the file wrapper. In a personal interview, a duplicate of the Form is given to the applicant (or attorney or agent) at the conclusion of the interview. In the case of a telephone or video-conference interview, the copy is mailed to the applicant's correspondence address either with or prior to the next official communication. If additional correspondence from the examiner is not likely before an allowance or if other circumstances dictate, the Form should be mailed promptly after the interview rather than with the next official communication.

The Form provides for recordation of the following information:

- Application Number (Series Code and Serial Number)
- Name of applicant
- Name of examiner
- Date of interview
- Type of interview (telephonic, video-conference, or personal)
- Name of participant(s) (applicant, attorney or agent, examiner, other PTO personnel, etc.)
- An indication whether or not an exhibit was shown or a demonstration conducted
- An identification of the specific prior art discussed
- An indication whether an agreement was reached and if so, a description of the general nature of the agreement (may be by attachment of a copy of amendments or claims agreed as being allowable). Note: Agreement as to allowability is tentative and does not restrict further action by the examiner to the contrary.
- The signature of the examiner who conducted the interview (if Form is not an attachment to a signed Office action)

It is desirable that the examiner orally remind the applicant of his or her obligation to record the substance of the interview of each case. It should be noted, however, that the Interview Summary Form will not normally be considered a complete and proper recordation of the interview unless it includes, or is supplemented by the applicant or the examiner to include, all of the applicable items required below concerning the substance of the interview.

A complete and proper recordation of the substance of any interview should include at least the following applicable items:

1) A brief description of the nature of any exhibit shown or any demonstration conducted,
2) an identification of the claims discussed,
3) an identification of the specific prior art discussed,
4) an identification of the principal proposed amendments of a substantive nature discussed, unless these are already described on the Interview Summary Form completed by the Examiner,
5) a brief identification of the general thrust of the principal arguments presented to the examiner,
   (The identification of arguments need not be lengthy or elaborate. A verbatim or highly detailed description of the arguments is not required. The identification of the arguments is sufficient if the general nature or thrust of the principal arguments made to the examiner can be understood in the context of the application file. Of course, the applicant may desire to emphasize and fully describe those arguments which he or she feels were or might be persuasive to the examiner.)
6) a general indication of any other pertinent matters discussed, and
7) if appropriate, the general results or outcome of the interview unless already described in the Interview Summary Form completed by the examiner.

Examiners are expected to carefully review the applicant's record of the substance of an interview. If the record is not complete and accurate, the examiner will give the applicant an extendable one month time period to correct the record.

Examiner to Check for Accuracy

If the claims are allowable for other reasons of record, the examiner should send a letter setting forth the examiner's version of the statement attributed to him or her. If the record is complete and accurate, the examiner should place the indication, "Interview Record OK" on the paper recording the substance of the interview along with the date and the examiner's initials.

DETAILED ACTION

1. Prior to allowance, an updated search was performed by Examiner.

2. Claims 2 – 6, 12 – 16 and 18 – 22 are cancelled.

3. Claims 1, 7 – 11, 17 and 23 – 37 are allowed (Renumbered as claims 1 – 22).

Information Disclosure Statement

The information disclosure statements (IDS) submitted on 01/15/2010 has been considered by examiner. The submission is in compliance with the provisions of 37 CFR 1.97. Accordingly, the information disclosure statement is being considered by the examiner.

Reasons for Allowance

Claims 1, 7 – 11, 17 and 23 – 37 are allowed (Renumbered as claims 1 – 22).

The following is an examiner's statement of reasons for allowance:

The closest prior art cited is: "Malik" (US Publication No. 2004/0078447)

Malik discloses a system and program (see fig 1,2A and 2B), comprising:

Although Malik teaches the general state of the art in "relevany" determinations, specifically that a: **(i) a monitoring module configured to track actions by a plurality of users associated with an application for managing electronic content, (ii) a relevance analysis module coupled to the monitoring module and configured to determine a relevance to a user of a received piece of electronic content, (iii) the relevance being determined dependent on: a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content and (iv) a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content** is known in the art

The prior art of record does not teach alone nor reasonably suggest in combination all of the limitations together with independent claims 1, 11 and 17 (and therefore dependent claims also). Malik does not teach alone nor reasonably suggest each and every limitation of claims 1, 11 and 17. For more complete explanation of the teachings of Liang in light of the claims please see entire prosecution history for US application no. 11/623653.

The closest prior art is Malik

Malik teaches a computer method, comprising::

(i) a monitoring module configured to track actions by a plurality of users associated with an application for managing electronic content (see para [0035] and [0045]),

(ii) a relevance analysis module coupled to the monitoring module and configured to determine a relevance to a user of a received piece of electronic content (user(s) profile, see abstract and para [0013], [0014], [0040]-[0046], [0072-0074] and [0079 – 0080]),

(iii) the relevance being determined dependent on: a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079] - 0080),

and (iv) a measure of a relevance of the piece of electronic content as determined dependent on one or more actions of at least one other of the plurality of users that is a recipient of the piece of electronic content (see para [0013], [0014], [0040]-[0046], [0072-0074] and [0079 - 0080]).

Examiner's Note: Although Malik teaches the general state of the art, Malik does not teach each and every limitation in claims 1, 11 and 17. Please see entire prosecution history for complete details on the prior art of record. Malik at least differ from claim 1 because claim 1 teaches:

a first relevance measurement module coupled to the monitoring module and configured to determine a first measure of [[a]] relevance of the piece of electronic content as determined dependent on one or more actions of the user on the piece of electronic content; ~~*, and*~~

a second relevance measurement module coupled to the monitoring module and configured to:

identify at least one other of the plurality of users that is a recipient of the piece of electronic content and that is cognate to the user, and

determine a second measure of [[a]] relevance of the piece of electronic content as determined dependent at least in part on one or more measures derived from the group consisting of: a local importance describing an importance of the sender to the user, a global importance describing an importance of the sender within an organization with which the plurality of users are associated, a content relevance determined responsive to content of the piece of electronic content, and

one or more actions on the piece of electronic content of the at least one other of the plurality of users that is a recipient of the piece of electronic content and that is cognate to the user; and a relevance analysis module configured to determine the relevance score of the piece of electronic content sent from the sender to the user, the relevance score being determined based at least in part on the first determined measure of relevance and on the second determined measure of relevance.

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 John P. Hocker whose telephone number is 5712705004. The examiner can normally be reached on Mon - Fri 10:30 - 3:00 EST (Alternating Fridays Off).

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Mohammad Ali can be reached on 571-272-4105. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Mohammad Ali/

Supervisory Patent Examiner, Art Unit 2158

/JPH/ Art Unit 2158

Notice of References Cited	Application/Control No. 11/623,653	Applicant(s)/Patent Under Reexamination MARSTON ET AL.	
	Examiner JOHN P. HOCKER	Art Unit 2158	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	Classification
*	A	US-2003/0074409 A1	04-2003	Bentley, Richard M.	709/206
*	B	US-2003/0126120 A1	07-2003	Faybishenko et al.	707/3
*	C	US-2004/0078447 A1	04-2004	Malik et al.	709/206
*	D	US-2007/0100806 A1	05-2007	Ramer et al.	707/003
*	E	US-2007/0100650 A1	05-2007	Ramer et al.	705/001
*	F	US-2007/0244867 A1	10-2007	Malandain et al.	707/3
*	G	US-2009/0216696 A1	08-2009	Downs et al.	706/20
	H	US-			
	I	US-			
	J	US-			
	K	US-			
	L	US-			
	M	US-			

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	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.

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	11623653	MARSTON ET AL.
	Examiner	Art Unit
	JOHN P HOCKER	2158

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	11/09/2008	05/07/2009	01/30/2010						
	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			=						

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	11623653	MARSTON ET AL.
	Examiner	**Art Unit**
	JOHN P HOCKER	2158

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	11/09/2008	05/07/2009	01/30/2010						
	37			=						

Substitute for form 1449A/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

Complete if Known	
Application No.	11/623,653
Filing Date	January 16, 2009
First Named Inventor	Justin Marston
Art Unit	2158
Examiner Name	John P. Hocker
Attorney Docket Number	23682-12381

Sheet 1 of 4

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Document No. Number – Kind Code[2] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document
	A1	US-2007/0038942 A1	02-15-2007	Chen et al.
	A2	US-2006/0168029 A1	07-27-2006	Fitzpatrick et al.
	A3	US-2005/0149498 A1	07-07-2005	Lawrence et al.
	A4	US-2005/0086340 A1	04-21-2005	Kang et al.
	A5	US-2005/0054334 A1	03-10-2005	Brown et al.
	A6	US-2004/0260761 A1	12-23-2004	Leaute et al.
	A7	US-2004/0177271 A1	09-09-2004	Arnold et al.
	A8	US-2004/0153515 A1	08-05-2004	Touboul et al.
	A9	US-2004/0064515 A1	04-01-2004	Hockey
	A10	US-2004/0054498 A1	03-18-2004	Shipp
	A11	US-2003/0140223 A1	07-24-2003	Desideri
	A12	US-2003/0055907 A1	03-20-2003	Stiers
	A13	US-2003/0009595 A1	01-09-2003	Collins
	A14	US-2002/0144154 A1	10-03-2002	Tomkow
	A15	US-2002/0007453 A1	01-17-2002	Nemovicher
	A16	US-7,412,437 B2	08-12-2008	Moody et al.
	A17	US-7,233,954	06-19-2007	Horvitz
	A18	US-7,103,794	09-05-2006	Malcolm
	A19	US-7,058,687	06-06-2006	Kucherawy
	A20	US-6,970,913	11-29-2005	Albert et al.
	A21	US-6,959,382	10-25-2005	Kinnis et al.
	A22	US-6,850,968	02-01-2005	Pfeffer et al.
	A23	US-6,789,107 B1	09-07-2004	Bates et al.
	A24	US-6,704,772	03-09-2004	Ahmed et al.
	A25	US-6,704,768	03-09-2004	Zombek
	A26	US-6,694,358 B1	02-17-2004	Swildens et al.
	A27	US-6,581,090 B1	06-17-2003	Lindbo et al.
	A28	US-6,578,076 B1	06-10-2003	Putzolu

Examiner Signature	/John Hocker/	Date Considered	01/30/2010

*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.

[1]Applicant's unique citation designation number (optional). [2]See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3]Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4]For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5]Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST. 16 if possible. [6]Applicant is to place a check mark here if English language Translation is attached.

23682/12381/DOCS/2183841.1

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /J.H./

Substitute for form 1449A/PTO				Complete if Known	
INFORMATION DISCLOSURE STATEMENT BY APPLICANT				**Application No.**	11/623,653
				Filing Date	January 16, 2009
				First Named Inventor	Justin Marston
				Art Unit	2158
				Examiner Name	John P. Hocker
Sheet	2	of	4	Attorney Docket Number	23682-12381

U.S. PATENT DOCUMENTS				
		Document No.		
Examiner Initials*	Cite No.[1]	Number – Kind Code[2] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document
	A29	US-6,553,413 B1	04-22-2003	Leighton et al.
	A30	US-6,502,125 B1	12-31-2002	Kenner et al.
	A31	US-6,484,143 B1	11-19-2002	Swildens et al.
	A32	US-6,480,893 B2	11-12-2002	Kriegsman
	A33	US-6,314,565 B1	11-06-2001	Kenner et al.
	A34	US-6,282,569 B1	08-28-2001	Wallis et al.
	A35	US-6,421,726 B1	07-16-2002	Kenner et al.
	A36	US-6,415,280 B1	07-02-2002	Farber et al.
	A37	US-6,405,252 B1	06-11-2002	Gupta et al.
	A38	US-6,370,571 B1	04-09-2002	Medin, Jr.
	A39	US-6,185,598 B1	02-06-2001	Farber et al.
	A40	US-6,185,551	02-06-2001	Birell et al.
	A41	US-6,182,136	01-30-2001	Ramanathan et al.
	A42	US-6,181,867 B1	01-30-2001	Kenner et al.
	A43	US-6,178,160 B1	01-23-2001	Bolton et al.
	A44	US-6,154,744	11-28-2000	Kenner et al.
	A45	US-6,134,582 A	10-17-2000	Kennedy
	A46	US-6,134,598 A	10-17-2000	Raman
	A47	US-6,122,632 A	09-19-2000	Botts et al.
	A48	US-6,112,239	08-29-2000	Kenner et al.
	A49	US-6,035,327	03-07-2000	Buckley et al.
	A50	US-6,029,175	02-22-2000	Chow et al.
	A51	US-6,009,462	12-28-1999	Birrell et al.
	A52	US-5,938,735	08-17-1999	Malik
	A53	US-5,938,732	08-17-1999	Lim et al.
	A54	US-6,003,030	12-14-1999	Kenner et al.

Examiner Signature	/John Hocker/	Date Considered	01/30/2010

*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.

[1]Applicant's unique citation designation number (optional). [2]See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3]Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4]For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5]Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST. 16 if possible. [6]Applicant is to place a check mark here if English language Translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /J.H./

Substitute for form 1449A/PTO				Complete if Known	
INFORMATION DISCLOSURE STATEMENT BY APPLICANT				**Application No.**	11/623,653
				Filing Date	January 16, 2009
				First Named Inventor	Justin Marston
				Art Unit	2158
				Examiner Name	John P. Hocker
Sheet	3	of	4	Attorney Docket Number	23682-12381

FOREIGN PATENT DOCUMENTS					
		Foreign Patent Document			
Examiner Initials*	Cite No.[1]	Country Code[3] – Number[4] Kind Code[5] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	T[6]
	B1	EP 0739115 A2	10-23-1996	AT&T IPM CORP	
	B2	WO 03/036492 A1	05-01-2001	Idetic, Inc.	
	B3	WO 03/073711 A2	09-04-2003	RPOST International Limited	
	B4	WO 04/012415 A1	02-05-2004	Security and Standards Limited	

OTHER REFERENCES – NON-PATENT LITERATURE DOCUMENTS			
Examiner Initials*	Cite No.[1]	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, page(s), volume-issue number(s), publisher, city and/or country where published	T[6]
	C1	BRUSSEE, ROGIER et al, "Content Distribution Networks", Telematica Institute, June 1, 2001, pp. V-IX and 1-88.	
	C2	DAVISON, BRIAN, "Brian Davison's Web-Caching Bibliography", Department of Computer Science, Rutgers, the State University, New Jersey, June 28, 2000, pp. 1-16.	
	C3	GOLDSZMIDT, GERMAN et al., "Load Distribution for Scalable Web Servers: Summer Olympics 1996 – A Case Study", In Proceedings of the 8th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management, Sydney, Australia, October 1997, pp. 1-10.	
	C4	GWERTZMAN, JAMES et al., "The Case for Geographical Push-Caching", *In* Proceedings of the Fifth Workshop on Hot Topics in Operating Systems, IEEE Computer Society, Orcas Island, WA , May 1995, pp. 51-55.	
	C5	GUPTA, G., et al., "Digital Forensics Analysis of E-Mails: A Trusted E-Mail Protocol," International Journal of Digital Evidence, Spring 2004, Vol. 2, Is. 4, 11 Pages.	
	C6	MAMBO, M., et al., "Proxy Signatures for Delegating Signing Operation," Conference on Computer and Communications Security, Proceedings of the 3rd ACM conference on Computer and Communications Security, 1996, Pages 48-57.	
	C7	STANFORD-CLARK, ANDY, "Atlanta Olympics WOMplex", Presented at Get Connected Technical Interchange '96, IBM, Hursley, U.K., October 1996, pp. IBM 000026-000037.	
	C8	VAN STEEN, MAARTEN et al, "Locating Objects in Wide-Area Systems", IEEE Communications Magazine Vo 36, No. 1, Jan. 1998, pp. 104-109.	
	C9	PCT International Preliminary Report on Patentability and Written Opinion, PCT/GB2005/001845, November 29, 2006, 6 pages.	
	C10	International Search Report and Written Opinion, PCT/GB2005/001845, August 24, 2005, 8 Pages.	
	C11	International Search Report and Written Opinion, PCT/GB2004/004609, February 22, 2005, 9 pages.	

Examiner Signature	/John Hocker/	Date Considered	01/30/2010

*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.

[1]Applicant's unique citation designation number (optional). [2]See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3]Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4]For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5]Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST. 16 if possible. [6]Applicant is to place a check mark here if English language Translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /J.H./

Substitute for form 1449A/PTO				Complete if Known	
INFORMATION DISCLOSURE STATEMENT BY APPLICANT				**Application No.**	11/623,653
				Filing Date	January 16, 2009
				First Named Inventor	Justin Marston
				Art Unit	2158
				Examiner Name	John P. Hocker
Sheet	4	of	4	Attorney Docket Number	23682-12381

OTHER REFERENCES – NON-PATENT LITERATURE DOCUMENTS			
Examiner Initials*	Cite No.[1]	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, page(s), volume-issue number(s), publisher, city and/or country where published	T[6]
	C12	International Search Report and Written Opinion, PCT/GB2004/000805, June 4, 2004, 9 pages.	
	C13	International Search Report and Written Opinion, PCT/GB2005/001868, August 19, 2005, 8 pages.	
	C14	FADEN, G., "Solaris Trusted Extensions," Architectural Overview, Sun Microsystems, April 2006, [online] [Retrieved on December 19, 2008] Retrieved from the internet <URL:http://www.opensolaris.org/os/community/security/projects/tx/TrustedExtensionsArch.pdf>.	
	C15	FADEN, G., "Multilevel Filesystems in Solaris Trusted Extensions," SACMAT'07, June 20-22, 2007, [online] [Retrieved on December 19, 2008] Retrieved from the internet URL:http://www.opensolaris.org/os/community/security/projects/tx/sacmat04s-faden-1.pdf>.	

Examiner Signature	/John Hocker/	Date Considered	01/30/2010

*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.

[1]Applicant's unique citation designation number (optional). [2]See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3]Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4]For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5]Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST. 16 if possible. [6]Applicant is to place a check mark here if English language Translation is attached.

23682/12381/DOCS/2183841.1

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /J.H./

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	11623653	MARSTON ET AL.
	Examiner	Art Unit
	JOHN P HOCKER	2169

SEARCHED			
Class	Subclass	Date	Examiner
707	1-3	11/09/2008 & 01/30/2010	JPH
707	726, 749	11/09/2008 & 01/30/2010	JPH

SEARCH NOTES		
Search Notes	Date	Examiner
Inventive concept: " User Profiles for electronic content viewing (e.g. email)"	11/09/2008 & 01/30/2010	JPH
1: IDS Reviewed for Prior Art.	11/09/2008 & 01/30/2010	JPH
2: Examiner performed inventor name search.	11/09/2008 & 01/30/2010	JPH
3: Examiner searched 'Computer Technology' NPL (e.g. IEEE (Institute of Electrical and Electronics Engineers), Association for Computing Machinery, Wikipedia and/or Google Scholar)	11/09/2008 & 01/30/2010	JPH
4: Examiner searched International Technical & Patent Databases(US-PGPUB, USPAT, USOCR, FPRS, EPO, JPO, DERWENT and IBM_TDB).	11/09/2008 & 01/30/2010	JPH
During Patent Searches, Examiner searched class/subclasses (as listed above) combined with text of the inventive concept. Please see East Search History for text searched.	11/09/2008 & 01/30/2010	JPH

INTERFERENCE SEARCH			
Class	Subclass	Date	Examiner
707	1-3	11/09/2008 & 01/30/2010	JPH
707	726, 749	11/09/2008 & 01/30/2010	JPH

EAST Search History

EAST Search History (Prior Art)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	23790	(707/1-3).ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2010/01/30 22:49
L2	2246	second with (relevancy or relevance)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/30 22:58
L3	408	1 and 2	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/30 22:58
L4	227	3 and (score or scoring)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/30 22:58
L5	87	4 and (global)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/30 22:58
L6	82	5 and (content or mail)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/30 22:59
L7	151958	(email or (electronic adj3 mail))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2010/01/30 22:59

L8	937	based with (other or another) with (action or read or forward or reply or delete)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/30 22:59
L9	212	7 and 8	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/30 22:59
L10	0	9 and 6	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/30 23:00
L11	0	9 and 3	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/30 23:00
L12	24	7 and 5	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/30 23:00
L13	24	7 and 6	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/30 23:00
S1	0	("20040078447").PN.	USPAT; USOCR	OR	OFF	2008/11/09 11:58
S2	2	"20040078447"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/11/09 11:58
S3	19391	(707/1-3).ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/11/09 15:27

S4	3609	S3 and (email or (electronic adj3 mail))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/11/09 15:28
S5	1159	S4 and (profile)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/11/09 15:28
S6	1151	S5 and user	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/11/09 15:29
S7	757	S6 and (user near5 profile)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/11/09 15:29
S8	656	S7 and (search)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/11/09 15:29
S9	404	S8 and (track or tracking)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/11/09 15:30
S10	8	("20020122543" | "20020136279" | "20030191806" | "20040133645").PN.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/11/09 15:31
S11	56467	relevance	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2009/04/19 13:56

S12	659	(relevance with action)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2009/04/19 13:56
S13	240	S12 and (based with action)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2009/04/19 13:57
S14	240	S12 and S13	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2009/04/19 13:57
S15	8	"user action" with relevance with based	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2009/04/19 13:58
S16	2	"20080288533"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2009/05/31 13:40
S17	2	"20080288457"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2009/05/31 13:45
S18	86589	("other" or "group" or another) with (action)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/04 14:44
S19	98	S18 same (relevance)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/04 14:45

S20	196524	("other" or "group" or another) with (read or forward or send or action)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/04 14:47
S21	146	S20 same (relevance)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/04 14:47
S22	196528	("other" or "group" or another) with ("based on" or read or forward or send or action)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/04 14:53
S23	84191	S22 and (importance or important or relevant or relevance)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/04 14:53
S24	36443	based with (other or another)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/04 14:54
S25	3371	S23 and S24	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/04 14:54
S26	930	based with (other or another) with (action or read or forward or reply or delete)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/04 14:54
S27	578	S23 and S26	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/04 14:54

S28	219	S27 and (email or mail)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/04 14:54
S29	34	S26 same (importance or important or relevant or relevance)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2010/01/04 14:55

1/30/2010 11:02:43 PM
C:\Documents and Settings\jhocker\My Documents\EAST\Workspaces\Hocker Workspaces\11623653.wsp

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	11623653	MARSTON ET AL.
	Examiner	**Art Unit**
	JOHN P HOCKER	2158

ORIGINAL		INTERNATIONAL CLASSIFICATION														
CLASS	SUBCLASS	CLAIMED									NON-CLAIMED					
707	728	G	0	6	F	7 / 00 (2006.01.01)										
CROSS REFERENCE(S)																
CLASS	SUBCLASS (ONE SUBCLASS PER BLOCK)															

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original
1	1	7	17	18	33										
	2		18	19	34										
	3		19	20	35										
	4		20	21	36										
	5		21	22	37										
	6		22												
2	7	8	23												
3	8	9	24												
4	9	10	25												
5	10	11	26												
6	11	12	27												
	12	13	28												
	13	14	29												
	14	15	30												
	15	16	31												
	16	17	32												

/JOHN P HOCKER/ Examiner.Art Unit 2158 (Assistant Examiner)	(Date)	**Total Claims Allowed:** 22	
/MOHAMMAD ALI/ Supervisory Patent Examiner.Art Unit 2158 (Primary Examiner)	01/31/2010 (Date)	O.G. Print Claim(s) 1 (1)1	O.G. Print Figure 1

PART B - FEE(S) TRANSMITTAL

Complete and send this form, together with applicable fee(s), to: Mail **Mail Stop ISSUE FEE**
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450
or Fax (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)

758 7590 02/12/2010

FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041

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 facsimile transmitted to the USPTO (571) 273-2885, on the date indicated below.

(Depositor's name)

(Signature)

(Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
11/623,653	01/16/2007	Justin Marston	23682-12381	5243

TITLE OF INVENTION: DETERMINING RELEVANCE OF ELECTRONIC CONTENT

APPLN. TYPE	SMALL ENTITY	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	YES	$755	$300	$0	$1055	05/12/2010

EXAMINER	ART UNIT	CLASS-SUBCLASS
HOCKER, JOHN P	2158	707-002000

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/SB/122) attached.

☐ "Fee Address" indication (or "Fee Address" Indication form 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 Fenwick & West LLP

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 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.

(A) NAME OF ASSIGNEE (B) RESIDENCE: (CITY and STATE OR COUNTRY)

Bluespace Software Corporation Austin, Texas

Please check the appropriate assignee category or categories (will not be printed on the patent) : ☐ Individual ☒ Corporation or other private group entity ☐ Government

4a. The following fee(s) are submitted:

☒ Issue Fee

☒ Publication Fee (No small entity discount permitted)

☐ Advance Order - # of Copies _____

4b. Payment of Fee(s): **(Please first reapply any previously paid issue fee shown above)**

☐ A check is enclosed.

☐ Payment by credit card. Form PTO-2038 is attached.

☒ The Director is hereby authorized to charge the required fee(s), any deficiency, or credit any overpayment, to Deposit Account Number 192555 (enclose an extra copy of this form).

5. **Change in Entity Status** (from status indicated above)

☐ a. Applicant claims SMALL ENTITY status. See 37 CFR 1.27.

☐ b. Applicant is no longer claiming SMALL ENTITY status. See 37 CFR 1.27(g)(2).

NOTE: The Issue Fee and Publication Fee (if required) will not be accepted from anyone other than the applicant; a registered attorney or agent; or the assignee or other party in interest as shown by the records of the United States Patent and Trademark Office.

Authorized Signature /Christopher King/ Date March 18, 2010

Typed or printed name Christopher-Charles P. King Registration No. 60,985

This collection of information 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 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, 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.

Electronic Patent Application Fee Transmittal	
Application Number:	11623653
Filing Date:	16-Jan-2007
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Filer:	Christopher-Charles P. King/Linda Carlson
Attorney Docket Number:	23682-12381

Filed as Small Entity

Utility under 35 USC 111(a) Filing Fees

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Utility Appl issue fee	2501	1	755	755
Publ. Fee- early, voluntary, or normal	1504	1	300	300

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Extension-of-Time:				
Miscellaneous:				
		Total in USD ($)		**1055**

Electronic Acknowledgement Receipt	
EFS ID:	7240069
Application Number:	11623653
International Application Number:	
Confirmation Number:	5243
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Customer Number:	00758
Filer:	Christopher-Charles P. King
Filer Authorized By:	
Attorney Docket Number:	23682-12381
Receipt Date:	18-MAR-2010
Filing Date:	16-JAN-2007
Time Stamp:	17:31:18
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	Deposit Account
Payment was successfully received in RAM	$1055
RAM confirmation Number	3890
Deposit Account	192555
Authorized User	
The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows: Charge any Additional Fees required under 37 C.F.R. Section 1.16 (National application filing, search, and examination fees) Charge any Additional Fees required under 37 C.F.R. Section 1.17 (Patent application and reexamination processing fees)	

Charge any Additional Fees required under 37 C.F.R. Section 1.19 (Document supply fees)

Charge any Additional Fees required under 37 C.F.R. Section 1.20 (Post Issuance fees)

Charge any Additional Fees required under 37 C.F.R. Section 1.21 (Miscellaneous fees and charges)

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)	12381_IssueFee.pdf	91250 16237d95c801a7f6825e095db38eaa91987739d6	no	1
Warnings:					
Information:					
2	Fee Worksheet (PTO-875)	fee-info.pdf	31707 22774afef7ce5690b85f13787a5ff03863182ae9	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	122957		

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 NO.	ISSUE DATE	PATENT NO.	ATTORNEY DOCKET NO.	CONFIRMATION NO.
11/623,653	05/11/2010	7716217	23682-12381	5243

758 7590 04/21/2010

FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041

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 184 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 Data Management (ODM) at (571)-272-4200.

APPLICANT(s) (Please see PAIR WEB site http://pair.uspto.gov for additional applicants):

Justin Marston, Austin, TX;
Paul Marston, Lancashire, UNITED KINGDOM;

IR103 (Rev. 10/09)

IN THE UNITED STATES

PATENT AND TRADEMARK OFFICE

APPLICANT:	Justin Marston et al.
PATENT NO.:	7,716,217 B2
ISSUE DATE:	May 11, 2010
SERIAL NO:	11/623,653
FILING DATE:	January 16, 2007
TITLE:	Determining Relevance of Electronic Content
ATTY. DKT. NO.:	23682-12381

CERTIFICATE OF ELECTRONIC (EFS-WEB) TRANSMISSION

I hereby certify that this correspondence is being transmitted via the Office electronic filing system in accordance with 37 C.F.R. § 1.8(a)(i)(C) from the **Pacific Time Zone** of the United States on the local date shown below.

Dated: July 9, 2010 By: /Christopher King/
Christopher P. King, Reg. No. 60,985

COMMISSIONER FOR PATENTS
P.O. BOX 1450
ALEXANDRIA, VA 22313-1450
ATTENTION: CERTIFICATE OF CORRECTION BRANCH

REQUEST FOR CERTIFICATE OF CORRECTION

SIR:

The following error, as more fully described below, appears in this patent.

Column 47, line 47, after "content," delete --as --.

Attached hereto is Form PTO-1050.

☒ No fee is due for the correction, since the errors were made by the Patent and Trademark Office.

Respectfully submitted,

Dated: July 9, 2010

By: /Christopher King/

Christopher P. King, Reg. No. 60,985
Attorney for Applicant
Fenwick & West LLP
801 California Street
Mountain View, CA 94041
Tel.: (650) 335-7633
Fax: (650) 938-5200

UNITED STATES PATENT AND TRADEMARK OFFICE

CERTIFICATE OF CORRECTION

PATENT NO. : 7,716,217 B2

DATED : May 11, 2010

INVENTORS : Justin Marston et al.

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Column 47, line 47, after "content," delete --as--.

Electronic Acknowledgement Receipt	
EFS ID:	7985833
Application Number:	11623653
International Application Number:	
Confirmation Number:	5243
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Customer Number:	00758
Filer:	Christopher-Charles P. King
Filer Authorized By:	
Attorney Docket Number:	23682-12381
Receipt Date:	09-JUL-2010
Filing Date:	16-JAN-2007
Time Stamp:	16:28:43
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	Request for Certificate of Correction	12381_RequestCertificateofCorrection.pdf	76104 5432b0c78ad98ba1174bdebc9333df03cba7fd5f	no	3
Warnings:					
Information:					

Total Files Size (in bytes):	76104

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

CERTIFICATE OF CORRECTION

PATENT NO. : 7,716,217 B2
APPLICATION NO. : 11/623653
DATED : May 11, 2010
INVENTOR(S) : Justin Marston et al.

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Column 47, line 47, after "content," delete "as".

Signed and Sealed this

Twenty-fourth Day of August, 2010

David J. Kappos

David J. Kappos
Director of the United States Patent and Trademark Office

PTO/AIA/81A (02-15)
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

PATENT - POWER OF ATTORNEY OR REVOCATION OF POWER OF ATTORNEY WITH A NEW POWER OF ATTORNEY AND CHANGE OF CORRESPONDENCE ADDRESS

Patent Number	7716217
Issue Date	05/11/2010
First Named Inventor	Justin Marston
Title	Determining Relevance Of Electronic Content
Attorney Docket No.	

I hereby revoke all previous powers of attorney given in the above-identified patent.

☐ A Power of Attorney is submitted herewith.

OR

☒ I hereby appoint Practitioner(s) associated with the Customer Number identified in the box at right as my/our attorney(s) or agent(s) with respect to the patent identified above, and to transact all business in the United States Patent and Trademark Office connected therewith: 44088

OR

☐ I hereby appoint Practitioner(s) named below as my/our attorney(s) or agent(s) with respect to the patent identified above, and to transact all business in the United States Patent and Trademark Office connected therewith:

Practitioner(s) Name	Registration Number

Please recognize or change the correspondence address for the above-identified patent to:

☒ The address associated with the above-identified Customer Number.

OR

☐ The address associated with the Customer Number identified in the box at right:

OR

☐ Firm or Individual Name					
Address					
City		State		Zip	
Country					
Telephone		Email			

I am the:

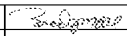
☐ Applicant.

OR

☒ Patent owner.

Statement under 37 CFR 3.73(c) (Form PTO/AIA/96) submitted herewith or filed on 10-28-2014.

SIGNATURE of Applicant or Patent Owner

Signature		Date	01/24/2022
Name	Brad Moore	Telephone	
Title and Company	CEO Sterling Computers Corporation		

NOTE: Signatures of all the applicants or patent owners of the entire interest or their representative(s) are required. If more than one signature is required, submit multiple forms, check the box below, and identify the total number of forms submitted in the blank below.

☐ A total of _____ forms are submitted.

This collection of information is required by 37 CFR 1.31, 1.32, and 1.33. The information is required to obtain or retain a benefit by the public, which is to update (and by the USPTO to process) the file of a patent or reexamination proceeding. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 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.

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.

Electronic Acknowledgement Receipt	
EFS ID:	44946585
Application Number:	11623653
International Application Number:	
Confirmation Number:	5243
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Customer Number:	758
Filer:	Brendan Bennett Dix/Jason Wormstadt
Filer Authorized By:	Brendan Bennett Dix
Attorney Docket Number:	23682-12381
Receipt Date:	08-FEB-2022
Filing Date:	16-JAN-2007
Time Stamp:	17:04:10
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	US7716217aia0081a.pdf	200449 7a210ad86aaa1b229cf48c89202ea57aebe3bfc0	no	2
Warnings:					

Information:	
Total Files Size (in bytes):	200449

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	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
11/623,653	01/16/2007	Justin Marston	23682-12381

CONFIRMATION NO. 5243

IMPROPER CPOA LETTER

758
FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041

OC000000131837671

Date Mailed: 02/16/2022

NOTICE REGARDING POWER OF ATTORNEY

This is in response to the power of attorney filed 02/08/2022. The power of attorney in this application is not accepted for the reason(s) listed below:

- The power of attorney is from an assignee and the statement required by 37 CFR 3.73(c) has not been received.

/dgela/

Office of Data Management, Application Assistance Unit (571) 272-4000, or (571) 272-4200, or 1-888-786-0101

44088
KAUFHOLD DIX PATENT LAW
P. O. BOX 89626
SIOUX FALLS, SD 57109
UNITED STATES

To: PTOC@Fenwick.com,,
From: PAIR_eOfficeAction@uspto.gov
Cc: PAIR_eOfficeAction@uspto.gov
Subject: Private PAIR Correspondence Notification for Customer Number 758

Feb 16, 2022 04:24:12 AM

Dear PAIR Customer:

FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041
UNITED STATES

The following USPTO patent application(s) associated with your Customer Number, 758 , 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.
11623653	N572	02/16/2022	23682-12381

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/AIA/81A (02-15)
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

PATENT - POWER OF ATTORNEY OR REVOCATION OF POWER OF ATTORNEY WITH A NEW POWER OF ATTORNEY AND CHANGE OF CORRESPONDENCE ADDRESS

Patent Number	7716217
Issue Date	05/11/2010
First Named Inventor	Justin Marston
Title	Determining Relevance Of Electronic Content
Attorney Docket No.	

I hereby revoke all previous powers of attorney given in the above-identified patent.

☐ A Power of Attorney is submitted herewith.

OR

☒ I hereby appoint Practitioner(s) associated with the Customer Number identified in the box at right as my/our attorney(s) or agent(s) with respect to the patent identified above, and to transact all business in the United States Patent and Trademark Office connected therewith: 44088

OR

☐ I hereby appoint Practitioner(s) named below as my/our attorney(s) or agent(s) with respect to the patent identified above, and to transact all business in the United States Patent and Trademark Office connected therewith:

Practitioner(s) Name	Registration Number

Please recognize or change the correspondence address for the above-identified patent to:

☒ The address associated with the above-identified Customer Number.

OR

☐ The address associated with the Customer Number identified in the box at right:

OR

☐ Firm or Individual Name					
Address					
City		State		Zip	
Country					
Telephone		Email			

I am the:

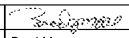
☐ Applicant.

OR

☒ Patent owner.

Statement under 37 CFR 3.73(c) (Form PTO/AIA/96) submitted herewith or filed on 10-28-2014.

SIGNATURE of Applicant or Patent Owner

Signature		Date	01/24/2022
Name	Brad Moore	Telephone	
Title and Company	CEO Sterling Computers Corporation		

NOTE: Signatures of all the applicants or patent owners of the entire interest or their representative(s) are required. If more than one signature is required, submit multiple forms, check the box below, and identify the total number of forms submitted in the blank below.

☐ A total of _____ forms are submitted.

This collection of information is required by 37 CFR 1.31, 1.32, and 1.33. The information is required to obtain or retain a benefit by the public, which is to update (and by the USPTO to process) the file of a patent or reexamination proceeding. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 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.

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.

PTO/SB/96 (11-18)
Approved for use through 11/30/2020. 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 displays a valid OMB control number.

STATEMENT UNDER 37 CFR 3.73(b)

Applicant/Patent Owner: Sterling Computer Corporation

Application No./Patent No.: 7716217 Filed/Issue Date: 05-11-2010

Titled: Determining Relevance of Electronic Content

Sterling Computer Corporation, a Corporation

(Name of Assignee) (Type of Assignee, e.g., corporation, partnership, university, government agency, etc.

states that it is:

1. ☒ the assignee of the entire right, title, and interest in;

2. ☐ an assignee of less than the entire right, title, and interest in (The extent (by percentage) of its ownership interest is _____ %); or

3. ☐ the assignee of an undivided interest in the entirety of (a complete assignment from one of the joint inventors was made)

the patent application/patent identified above, by virtue of either:

A. ☐ An assignment from the inventor(s) of the patent application/patent identified above. The assignment was recorded in the United States Patent and Trademark Office at Reel_____, Frame_____, or a copy* is attached.

OR

B. ☐ A chain of title from the inventor(s), of the patent application/patent identified above, to the current assignee as follows:

1. From: _____ To: _____

The document was recorded in the United States Patent and Trademark Office at

Reel_____, Frame_____, or a copy* is attached.

2. From: _____ To: _____

The document was recorded in the United States Patent and Trademark Office at

Reel_____, Frame_____, or a copy* is attached.

3. From: _____ To: _____

The document was recorded in the United States Patent and Trademark Office at

Reel_____, Frame_____, or a copy* is attached.

☐ Additional documents in the chain of title are listed on a supplemental sheet(s).

☐ *As required by 37 CFR 3.73(b)(1)(i), if a copy/copies is/are attached, the documentary evidence of the chain of title from the original owner to the assignee was, or concurrently is being, submitted for recordation pursuant to 37 CFR 3.11.

[NOTE: A separate copy (*i.e.*, a true copy of the original assignment document(s)) must be submitted to Assignment Division in accordance with 37 CFR Part 3, to record the assignment in the records of the USPTO. See MPEP 302.]

The undersigned (whose title is supplied below) is authorized to act on behalf of the assignee.

/brendanbdix/ 02/22/2022

Signature Date

Brendan B. Dix 37313

Printed or Typed Name Title or Registration Number

This collection of information is required by 37 CFR 3.73(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.

Electronic Acknowledgement Receipt	
EFS ID:	45053566
Application Number:	11623653
International Application Number:	
Confirmation Number:	5243
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Customer Number:	758
Filer:	Brendan Bennett Dix/Jason Wormstadt
Filer Authorized By:	Brendan Bennett Dix
Attorney Docket Number:	23682-12381
Receipt Date:	22-FEB-2022
Filing Date:	16-JAN-2007
Time Stamp:	15:00:37
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	US7716217aia0081a.pdf	200449 7a210ad86aaa1b229cf48c89202ea57aebe3bfc0	no	2
Warnings:					

Information:					
2	Assignee showing of ownership per 37 CFR 3.73	us7716217sb0096.pdf	175055 d53de78fe2bd05b5d2d883ebb17874a8407f70b5	no	2
Warnings:					
Information:					
			Total Files Size (in bytes):	375504	

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	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
11/623,653	01/16/2007	Justin Marston	23682-12381

CONFIRMATION NO. 5243

IMPROPER CPOA LETTER

758
FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041

OC000000132071345

Date Mailed: 03/01/2022

NOTICE REGARDING POWER OF ATTORNEY

This is in response to the power of attorney filed 02/22/2022. The power of attorney in this application is not accepted for the reason(s) listed below:

- The power of attorney is from an assignee and the statement required by 37 CFR 3.73(c) has not been received.

/hteffera/

Office of Data Management, Application Assistance Unit (571) 272-4000, or (571) 272-4200, or 1-888-786-0101

44088
KAUFHOLD DIX PATENT LAW
P. O. BOX 89626
SIOUX FALLS, SD 57109
UNITED STATES

To: PTOC@Fenwick.com,,
From: PAIR_eOfficeAction@uspto.gov
Cc: PAIR_eOfficeAction@uspto.gov
Subject: Private PAIR Correspondence Notification for Customer Number 758

Mar 01, 2022 04:45:59 AM

Dear PAIR Customer:

FENWICK & WEST LLP
SILICON VALLEY CENTER
801 CALIFORNIA STREET
MOUNTAIN VIEW, CA 94041
UNITED STATES

The following USPTO patent application(s) associated with your Customer Number, 758 , 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.
11623653	N572	03/01/2022	23682-12381

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/AIA/123 (08-12)
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.

CHANGE OF CORRESPONDENCE ADDRESS
Patent

Address to:
Mail Stop Post Issue
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Patent Number	7716217
Issue Date	05/11/2010
Application Number	11/623,653
Filing Date	01/16/2007
First Named Inventor	Justin Marston
Attorney Docket Number	

Please change the Correspondence Address for the above-identified patent to:

☑ The address associated with Customer Number: 44088

OR

☐ **Firm *or* Individual Name**

Address

City	**State**	**ZIP**

Country

Telephone	**Email**

This form cannot be used to change the data associated with a Customer Number. To change the data associated with an existing Customer Number use "Request for Customer Number Data Change" (PTO/SB/124).

This form will not affect any "fee address" provided for the above-identified patent. To change a "fee address" use the "Fee Address Indication Form" (PTO/SB/47).

I am the:

☐ Patentee.

☐ If the Patentee was not the applicant for patent (37 CFR 1.42), then a Statement under 37 CFR 3.73(c) (Form PTO/AIA/96 or equivalent) is enclosed or was filed on _____. See 37 CFR 3.71.

☐ Attorney or agent of record. Registration Number 37313.

☑ Patent practitioner acting in a representative capacity whose correspondence address is the correspondence address of record. Notice has been given to the patentee or owner. Registration Number 37313.

Signature	/brendanbdix/
Typed or Printed Name	Brendan B. Dix
Date 07/28/2022	Telephone 605-334-1571

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4(d) for signature requirements and certifications. Submit multiple forms if more than one signature is required, see below*.

☑ *Total of 1 forms are submitted.

This collection of information is required by 37 CFR 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: Mail Stop Post Issue, 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.

Electronic Acknowledgement Receipt

EFS ID:	46295551
Application Number:	11623653
International Application Number:	
Confirmation Number:	5243
Title of Invention:	DETERMINING RELEVANCE OF ELECTRONIC CONTENT
First Named Inventor/Applicant Name:	Justin Marston
Customer Number:	758
Filer:	Brendan Bennett Dix/Jason Wormstadt
Filer Authorized By:	Brendan Bennett Dix
Attorney Docket Number:	23682-12381
Receipt Date:	28-JUL-2022
Filing Date:	16-JAN-2007
Time Stamp:	15:29:02
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	Change of Address	us7716217aia0123.pdf	269067 569a167e1c2d2872c55c5a98ce6b3d9115e22d4d	no	2
Warnings:					

Information:	
Total Files Size (in bytes):	269067

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.