

Google LLC / Microsoft Corporation

v.

Sandpiper CDN, LLC

IPR2025-00969

IPR2026-00095

U.S. Patent No. 8,478,903



Patent Owner's Demonstratives

Sandpiper EX2021
Google/Microsoft v. Sandpiper
IPR2025-00969 1

Roadmap

- ▶ Overview of the '903 patent
- ▶ Overview of Kenner
- ▶ Kenner's "video ID" does not teach the claimed "alias name"
- ▶ Kenner does not direct resource requests to a repeater server based on the "alias name"
- ▶ Kenner does not teach a "table listing a plurality of alias names corresponding to content providers authorized to have resources delivered to client machines"
- ▶ Kenner does not teach "the at least one repeater server is addressable" using the "alias name" (claim 2) (LEAP)
- ▶ Kenner in view of Rekimoto does not teach "the first alias name is a domain name" (claim 34) (LEAP)

'903 Patent Overview

The '903 Patent: “Shared Content Delivery Infrastructure”

SUMMARY OF THE INVENTION

This invention provides a way for servers in a computer network to off-load their processing of requests for selected resources by determining a different server (a “repeater”) to process those requests. The selection of the repeater can be made dynamically, based on information about possible repeaters.

EX1001, '903 patent, 4:60-67.

It is necessary that the reflection step create a single URL containing the URL of the original resource, as well as the identity of the selected repeater. A special form of URL is created to provide this information. This is done by creating a new URL as follows:


D1. Given a repeater name, scheme, origin server name and path, create a new URL. If the scheme is “http”, the preferred embodiment uses the following format:

`http://<repeater>/<server>/<path>`

* * *

This URL is interpreted when received by the repeater.

EX1001, '903 patent, 8:29-49.



POR, 8; EX1001, '903 patent, FIG. 1.

The '903 Patent's Alias Names

Alias names have the structure of a host name

Placing the origin server name in the reflected URL is generally a good strategy, but it may be considered undesirable for aesthetic or (in the case, e.g., of cookies) certain technical reasons.

It is possible to avoid the need for placing both the repeater name and the server name in the URL. Instead, a “family” of names may be created for a given origin server, each name identifying one of the repeaters used by that server.

For instance, if `www.example.com` is the origin server, names for three repeaters might be created:

`wr1.example.com`

`wr2.example.com`

`wr3.example.com`

The name “`wr1.example.com`” would be an alias for repeater 1, which might also be known by other names such as “`wr1.anotherExample.com`” and “`wr1.example.edu`”.

EX1001, '903 patent, 9:40-55.

If the repeater can determine by which name it was addressed, it can use this information (along with a table that associates repeater alias names with origin server names) to determine which origin server is being addressed. For instance, if repeater 1 is addressed as `wr1.example.com`, then the origin server is “`www.example.com`”; if it is addressed as “`wr1.anotherExample.com`”, then the origin server is “`www.anotherExample.com`”.

EX1001, '903 patent, 9:56-63.

The '903 Patent's Alias Names

Alias names have the structure of a host name

The repeater can use two mechanisms to determine by which alias it is addressed:

1. Each alias can be associated with a different IP address. Unfortunately, this solution does not scale well, as IP addresses are currently scarce, and the number of IP addresses required grows as the product of origin servers and repeaters.
2. The repeater can attempt to determine the alias name used by inspecting the “host:” tag in the HTTP header of the request. Unfortunately, some old browsers still in use do not attach the “host:” tag to a request.

Reflectors would need to identify such browsers (the browser identity is a part of each request) and avoid this form of reflection.

EX1001, '903 patent, 9:64-10:11.

For example, the following is a URL identifying a file “F” in the path “A/B/C” on a computer at “www.uspto.gov”:

<http://www.uspto.gov/A/B/C/F>

In order to access the file “F” (the resource) specified by the above URL, a program (e.g., a browser) running on a user’s computer (i.e., a client computer) would have to first locate the computer (i.e., a server computer) specified by the host name. I.e., the program would have to locate the server “www.uspto.gov”. To do this, it would access a Domain Name Server (DNS), providing the DNS with the host name (“www.uspto.gov”). The DNS acts as a kind of centralized directory for resolving addresses from names. If the DNS determines that there is a (remote server) computer corresponding to the name “www.uspto.gov”, it will provide the program with an actual computer network address for that server computer. On the Internet this is called an Internet Protocol (or IP) address and it has the form “123.345.456.678”. The program on the user’s (client) computer would then use the actual address to access the remote (server) computer.

EX1001, '903 patent, 6:42-61.

The '903 Patent's Alias Names

Petitioner's expert Dr. Mowry confirmed that alias names have the structure of a host name

I mean `wr1.anotherexample.com` would be another potential alias.

Q Right. Or `wr1.example.edu` would be another alias name for repeater 1; is that right?

A Yes, that's right.

Q Okay. So are each of these alias names URLs?

A Well, the text is consistent with a host name. A URL usually begins with a prefix like `HTTP://`. So it may or may not be a URL, but these are consistent with host names.

Q Okay. And would you consider them to be Internet addresses?

A Well, the term Internet address more commonly would refer to an IP address. That's something else. So a host name and an Internet address are -- are not necessarily the same thing.

A host name is a text -- is a string of text with dots in it, typically, like the ones we just were discussing, `wr1.anotherexample.com`, that is consistent with the form of a host name, but an IP address is a set of numbers.

EX2017, Mowry Depo. Tr., 20:13-21:12.

Claim 1 of the '903 Patent

1. A content delivery system operative in a computer network for delivering resources associated with a plurality of content providers to multiple client machines, the system comprising:

at least one shared repeater server constructed and adapted to replicate at least some of the resources associated with a first content provider of said plurality of content providers, and to replicate at least some of the resources associated with a second content provider of said plurality of content providers, wherein the second content provider is distinct from the first content provider;

at least one table listing a plurality of alias names corresponding to content providers authorized to have resources delivered to client machines via the at least one shared repeater server, wherein the plurality of content providers comprises the first content provider and the second content provider,

wherein at least a first resource associated with the first content provider is associated with a first alias name of the plurality of alias names, said first alias name being associated with said at least one shared repeater server; wherein, requests for the first resource from said client machines are directed to the at least one shared repeater server, based at least in part on said first alias name;

wherein at least a second resource associated with the second content provider is associated with a second alias name of the plurality of alias names, said second alias name being associated with said at least one shared repeater server;

wherein requests for the second resource made from said client machines are directed to the at least one shared repeater server, based at least in part on said second alias name; and

wherein the at least one shared repeater server is further constructed and adapted to analyze, using the table, an alias name received with a client request for a particular resource to determine a content provider associated with the particular resource.

Claim 28 of the '903 Patent

28. A method, in a content delivery system operative in a computer network for delivering content to client machines, the computer network comprising a plurality of origin servers, each of said origin servers having resources associated therewith, and the content delivery system comprising at least one shared repeater server operable to replicate resources associated with the plurality of origin servers, the method comprising:

associating the at least one repeater server with a first alias name, wherein requests for a first resource located on a first origin server are directed, based at least in part on said first alias name, to the at least one repeater server for delivery of the first resource from said at least one repeater server;

associating the at least one repeater server with a second alias name, wherein requests for a second resource located on a second origin server are directed, based at least in part on said second alias name, to the at least one repeater server for delivery of the second resource from said at least one repeater server, wherein the second origin server is distinct from the first origin server;

providing a table listing origin servers having content located thereon, wherein said content is authorized for delivery to client machines via the at least one shared repeater server, and wherein the origin servers comprise the first origin server and the second origin server; and wherein the at least one repeater server is further constructed and adapted to analyze, using the table, an alias name received with a client request for a particular resource to determine an origin server associated with the particular resource.

Claim 37 of the '903 Patent

37. A method, in a content delivery system operative in a computer network for delivering content to client machines and comprising at least one shared repeater server operable to replicate resources stored on a plurality of origin servers, the method comprising:

associating a plurality of alias names with the at least one shared repeater server, each of said plurality of alias names being associated with an origin server, wherein a first alias name of said plurality of alias names is associated with a first origin server, and at least a second alias name of said plurality of alias names is associated with a second origin server distinct from said first origin server,

providing at least one table associating alias names with origin servers having content located thereon, wherein said content is authorized for delivery to client machines via the at least one shared repeater server, wherein the origin servers comprise the first origin server and the second origin server;

wherein requests for a first resource originating on the first origin server are directed, based at least in part on said first alias name, to the at least one shared repeater server for delivery of the first resource from said at least one repeater server; and

wherein requests for a second resource originating on the second origin server are directed, based at least in part on said second alias name, to the at least one shared repeater server for delivery of the second resource from said at least one repeater server; and

wherein the at least one repeater server uses the at least one table and an alias name received with a client request for a resource to determine an origin server associated with the requested resource.

Claim 40 of the '903 Patent

40. A server operative in a computer network for delivering resources associated with a plurality of content providers to multiple client machines, wherein the plurality of content providers comprise at least a first content provider and a second content provider distinct from the first content provider, the server comprising:

- cache storage;
- a table listing content providers of said plurality of content providers having content that is authorized for delivery to client machines via the server,

wherein the server is associated with a first alias name corresponding to said first content provider, and

wherein the server is associated with a second alias name corresponding to said second content provider, said second alias name being distinct from the first alias name, and

wherein the server is operable to replicate at least some of the resources associated with the first content provider, and at least some of the resources associated with the second content provider, and

wherein, in response to a request for a resource, the server is operable to use at least the table to analyze an alias name associated with the request to determine a content provider associated with the resource.

Claim 43 of the '903 Patent


43. A server operative in a computer network for delivering resources to multiple client machines, the server comprising:
cache storage;
a table listing a plurality of origin servers having content that is authorized for delivery to client machines via the server,
wherein the server is associated with a first alias name corresponding to a first origin server of the plurality of origin servers, and
wherein the server is associated with a second alias name corresponding to a second origin server of the plurality of origin servers, said second origin server being distinct from the first origin server and said second alias name being distinct from the first alias name, and
wherein the server is operable to replicate at least some of the resources located on the first origin server, and at least some of the resources located on the second origin server, and
wherein, in response to a request for a resource, the server uses at least the table to analyze an alias name associated with the request to determine an origin server associated with the particular resource.

Roadmap


- ▶ Overview of the '903 patent
- ▶ Overview of Kenner
- ▶ Kenner's "video ID" does not teach the claimed "alias name"
- ▶ Kenner does not direct resource requests to a repeater server based on the "alias name"
- ▶ Kenner does not teach a "table listing a plurality of alias names corresponding to content providers authorized to have resources delivered to client machines"
- ▶ Kenner does not teach "the at least one repeater server is addressable" using the "alias name" (claim 2)
- ▶ Kenner in view of Rekimoto does not teach "the first alias name is a domain name" (claim 34)

Overview of Kenner

Kenner's Architecture



EX1005, Kenner, FIG. 4 (annotated).



EX1005, Kenner, FIG. 1 (annotated).

Kenner's Primary Index Manager ("PIM")

(57) Abstract

A video clip storage and retrieval system whereby video clips, stored locally and/or at a more remote location, can be requested and retrieved by a user at the user's multimedia terminal. When the user requests a desired video clip, the request is processed by a primary index manager ("PIM") via a Local Search and Retrieval Unit ("SRU"). Before the message is communicated to the PIM, the local SRU checks its own storage to see whether the requested video clips are available locally. If some of the video clips are local, the local SRU still forwards the request to the PIM so that the PIM may determine specific video clip usage. The PIM determines the extended SRU where the audio-visual data is stored and passes this information to a Data Sequencing Interface ("DSI"). The DSI collects the video clips and downloads the clips to the user's terminal. The user may then view, copy, or print the video clip as desired. In a preferred embodiment, a distributed digital video clip delivery system, according to the invention, provides video clips stored at local and/or remote locations, which can be requested from the Internet and retrieved at the user's multimedia terminal. When the user requests a desired video clip shown on a Web page, the request is diverted to a primary index manager ("PIM"). The PIM attempts to locate the closest server containing the requested clip, from which the download is completed. The system further includes means for uploading and distributing clips to geographically diverse servers, dynamic load balancing, subscription management mechanisms, and protection means to discourage unauthorized duplication of video clips.

EX1005, Kenner, Abstract.

Kenner's Video ID

The PIM 64 also maintains information on the audio-visual clips stored on its extended SRUs 66 in a clip database (corresponding to the "audio-visual data index" previously discussed). In one embodiment, the clip database contains the following information for each clip:

Item Name	Format	Description
Counter	numeric	Primary index for the records. Each record represents one video clip.
Video ID	text	The globally unique name of the video clip, as specified above.
Extended SRUs	IP array	The IP addresses of all the extended SRUs 66 which contain the file.
Copyright	boolean	A flag to indicate that the file is copyrighted and must be protected.

EX1005, Kenner, 34:3-13.

Kenner's Video ID

In one embodiment, the video ID is constructed from the following portions: a text name of the file as defined by the content provider; the content provider's account number as provided by the organization running the subscription service; a category coordinate, possibly a representation of a hierarchical portion of a category tree, a geographic coordinate used to determine where the file is relevant (e.g. a region, state, or city); a time stamp, and a time period over which the file is relevant. The foregoing characteristics are known at the time a file is made available to the present system, and will be discussed in detail below.

EX1005, Kenner, 36:32-37:6.

Kenner's Video ID

In a preferred embodiment, the local SRU 51 passes the video ID to the PIM 64 in the following form: a "virtual URL" is constructed in the form "http://" plus the Internet address of the PIM 64, plus the user's subscriber ID number, plus the video ID. If the desired clip was located on the local SRU 51, then the virtual URL request will contain a further attribute specifying that a particular version of the file has already been located by the local SRU 51. The Netscape Navigator procedure NPN_GetURLNotify is available to accomplish sending this virtual URL to the PIM 64; the virtual URL is constructed to be in a format that will be accepted by Netscape Navigator and other network browsers. Upon receipt by the PIM 64, the virtual URL is decomposed into the video ID and subscriber ID components, which are then used to access the PIM's internal databases.

EX1005, Kenner, 37:7-17.

Roadmap

- ▶ Overview of the '903 patent
- ▶ Overview of Kenner
- ▶ Kenner's "video ID" does not teach the claimed "alias name"
- ▶ Kenner does not direct resource requests to a repeater server based on the "alias name"
- ▶ Kenner does not teach a "table listing a plurality of alias names corresponding to content providers authorized to have resources delivered to client machines"
- ▶ Kenner does not teach "the at least one repeater server is addressable" using the "alias name" (claim 2)
- ▶ Kenner in view of Rekimoto does not teach "the first alias name is a domain name" (claim 34)

Kenner's "video ID" does not teach the claimed "alias name"

The Petition's Mapping

The Petition alleges that Kenner's video ID teaches the claimed *first and second alias name* and that PIM 64 teaches the claimed *at least one shared repeater server*

Kenner's first clip (first resource) is associated with a first alias name of the plurality of alias names because each clip is associated with a corresponding video ID. Mowry, ¶84. In uploading a clip, the source IM registers the clip along with a video ID (alias name). Kenner, 44:6-14. And the video ID "is a unique reference identifier for [the] clip." Id., 22:1-2, 43:27-29 ("[E]ach clip ... ha[s] a unique video ID").

Pet., 20.

As explained, *Kenner* discloses a network of IMs. §VI.A.2.a. When a clip is requested, the request is directed to a primary IM (*shared repeater server*) associated with the user's terminal. *Kenner*, 10:24-31, 16:2-9; §VI.A.1.a; Mowry, ¶69.

Pet., 14.

Kenner's "video ID" does not teach the claimed "alias name"

Kenner's video ID is not associated with a repeater server

Video ID

text

The globally unique name of the video clip, as specified above.

EX1005, Kenner, 34:8-13.

In one embodiment, the video ID is constructed from the following portions: a text name of the file as defined by the content provider; the content provider's account number as provided by the organization running the subscription service; a category coordinate, possibly a representation of a hierarchical portion of a category tree, a geographic coordinate used to determine where the file is relevant (e.g. a region, state, or city); a time stamp, and a time period over which the file is relevant. The foregoing characteristics are known at the time a file is made available to the present system, and will be discussed in detail below.

EX1005, Kenner, 36:32-37:6.

Kenner's "video ID" does not teach the claimed "alias name"

**Kenner's video ID is not associated with a repeater server because
Kenner's virtual URL identifies the PIM**

In a preferred embodiment, the local SRU 51 passes the video ID to the PIM 64 in the following form: a "virtual URL" is constructed in the form "http://" plus the Internet address of the PIM 64, plus the user's subscriber ID number, plus the video ID. If the desired clip was located on the local SRU 51, then the virtual URL request will contain a further attribute specifying that a particular version of the file has already been located by the local SRU 51. The Netscape Navigator procedure NPN_GetURLNotify is available to accomplish sending this virtual URL to the PIM 64; the virtual URL is constructed to be in a format that will be accepted by Netscape Navigator and other network browsers. Upon receipt by the PIM 64, the virtual URL is decomposed into the video ID and subscriber ID components, which are then used to access the PIM's internal databases.

EX1005, Kenner, 37:7-17.

Kenner's "video ID" does not teach the claimed "alias name"

The content coordinate data in a video ID identifies types of files and does not create an association with an IM

The video ID consists of a multidimensional set of content-characterization coordinates plus a unique file name. The content coordinates enable a match with the regional IMs that have subscribed to that type of file, as will be discussed later.

EX1005, Kenner, 36:28-31.

Each IM, including the PIM 64, maintains a database of all other IMs connected to the system. For each IM, the database includes the Internet address of the IM, and information on the types of files stored by the IM's SRUs. The file type information can be stored in the form of content coordinate data, as previously discussed. Accordingly, by consulting the database, the PIM 64 can determine which IMs are likely to have the desired clip, and query only those IMs.

EX1005, Kenner, 41:1-6.

expiration date, charge information, etc. For clips not on an extended SRU 66, another IM database maintains a reduced level of information about every IM, namely the Internet address of the IM and the content coordinates of all audio-video files that it maintains. For example, a library of "news:sports:baseball" clips might be maintained by several nearby and distant regional ISPs, including the one comprising the user's region 81. As previously indicated, the PIM 64 also has a user database

EX1005, Kenner, 38:3-8.

Kenner's "video ID" does not teach the claimed "alias name"

Dr. Mowry confirms there is no explicit association between a video ID and a particular IM

Q Sure. Do the content coordinates in Kenner identify a set of IMs?

A Well, the content coordinate data is when -- when -- say an IM does not have content, it uses the content coordinate data to identify IMs that are likely to have the clip, that might be one or more than one, or even possibly zero, other IMs.

EX2017, Mowry Depo. Tr., 126:5-12.

Kenner's "video ID" does not teach the claimed "alias name"

Kenner's video ID is not a host name

A video ID consists of a number of different pieces of information, so it's -- I don't think one would normally characterize that as being a host name, but it's something that's used -- it's associated with -- it's associated with a content provider -- with a content provider and with a repeater in this implementation.

EX2017, Mowry Depo. Tr., 129:16-22.

Kenner's "video ID" does not teach the claimed "alias name"

Kenner's video ID does not follow a host name structure

Video ID

text

The globally unique name of the video clip, as specified above.

EX1005, Kenner, 34:8-13.

In one embodiment, the video ID is constructed from the following portions: a text name of the file as defined by the content provider; the content provider's account number as provided by the organization running the subscription service; a category coordinate, possibly a representation of a hierarchical portion of a category tree, a geographic coordinate used to determine where the file is relevant (e.g. a region, state, or city); a time stamp, and a time period over which the file is relevant. The foregoing characteristics are known at the time a file is made available to the present system, and will be discussed in detail below.

EX1005, Kenner, 36:32-37:6.

Kenner's "video ID" does not teach the claimed "alias name"

Kenner's virtual URL confirms that Kenner's video ID is not an alias name – the virtual URL explicitly identifies the "Internet address of the PIM 64" (i.e., alleged *at least one shared repeater server*)

In a preferred embodiment, the local SRU 51 passes the video ID to the PIM 64 in the following form: a "virtual URL" is constructed in the form "http://" plus the Internet address of the PIM 64, plus the user's subscriber ID number, plus the video ID. If the desired clip was located on the local SRU 51, then the virtual URL request will contain a further attribute specifying that a particular version of the file has already been located by the local SRU 51. The Netscape Navigator procedure NPN_GetURLNotify is available to accomplish sending this virtual URL to the PIM 64; the virtual URL is constructed to be in a format that will be accepted by Netscape Navigator and other network browsers. Upon receipt by the PIM 64, the virtual URL is decomposed into the video ID and subscriber ID components, which are then used to access the PIM's internal databases.

EX1005, Kenner, 37:7-17.

Roadmap

- ▶ Overview of the '903 patent
- ▶ Overview of Kenner
- ▶ Kenner's "video ID" does not teach the claimed "alias name"
- ▶ Kenner does not direct resource requests to a repeater server based on the "alias name"
- ▶ Kenner does not teach a "table listing a plurality of alias names corresponding to content providers authorized to have resources delivered to client machines"
- ▶ Kenner does not teach "the at least one repeater server is addressable" using the "alias name" (claim 2)
- ▶ Kenner in view of Rekimoto does not teach "the first alias name is a domain name" (claim 34)

Kenner does not direct resource requests to “the at least one repeater server” based on the “alias name”

The '903 Patent's Claims

The claims recite directing a request from a client machine to the at least one shared repeater server based at least in part on the alias name

wherein, requests for the first resource from said client machines are directed to the at least one shared repeater server, based at least in part on said first alias name;

wherein requests for the second resource made from said client machines are directed to the at least one shared repeater server, based at least in part on said second alias name; and

EX1001, Claim 1 (excerpt).

The '903 Patent's Claims

The claims recite directing a request from a client machine to the at least one shared repeater server based at least in part on the alias name

associating the at least one repeater server with a first alias name, wherein requests for a first resource located on a first origin server are directed, based at least in part on said first alias name, to the at least one repeater server for delivery of the first resource from said at least one repeater server;

associating the at least one repeater server with a second alias name, wherein requests for a second resource located on a second origin server are directed, based at least in part on said second alias name, to the at least one repeater server for delivery of the second resource from said at least one repeater server, wherein the second origin server is distinct from the first origin server;

EX1001, Claim 28 (excerpt).

wherein requests for a first resource originating on the first origin server are directed, based at least in part on said first alias name, to the at least one shared repeater server for delivery of the first resource from said at least one repeater server; and

wherein requests for a second resource originating on the second origin server are directed, based at least in part on said second alias name, to the at least one shared repeater server for delivery of the second resource from said at least one repeater server; and

EX1001, Claim 37 (excerpt).

Kenner does not direct resource requests to “the at least one repeater server” based on the “alias name”

**Kenner directs all content requests from terminal 50 to PIM 64 regardless of the video ID;
Kenner therefore does not direct requests to PIM 64 based on the video ID**

The type or format of the embedded file is analyzed by the browser 82. Typically, this is indicated by an extension on the filename of the embedded file and is known in the art. If the desired file is a video clip, the local SRU 51 belonging to the terminal 50 and the browser extension 84 are invoked to receive the data. First, the local SRU 51 intercepts a video ID, a unique identifier specifying the selected clip, which is stored within the EMBED field in the Web page. The local SRU 51 first determines if the desired clip is already stored locally. If not, the local SRU 51 passes the video ID to the PIM 64 associated with the user's terminal 50. The local SRU 51 then awaits authorization from the PIM 64 to proceed with a data transfer.

EX1005, Kenner, 36:19-27.

In a preferred embodiment, the local SRU 51 passes the video ID to the PIM 64 in the following form: a "virtual URL" is constructed in the form "http://" plus the Internet address of the PIM 64, plus the user's subscriber ID number, plus the video ID. If the desired clip was located on the local SRU 51, then the virtual URL request will contain a further attribute specifying that a particular version of the file has already been located by the local SRU 51. The Netscape Navigator procedure NPN_GetURLNotify is available to accomplish sending this virtual URL to the PIM 64; the virtual URL is constructed to be in a format that will be accepted by Netscape Navigator and other network browsers. Upon receipt by the PIM 64, the virtual URL is decomposed into the video ID and subscriber ID components, which are then used to access the PIM's internal databases.

EX1005, Kenner, 37:7-17.

Kenner does not direct resource requests to “the at least one repeater server” based on the “alias name”

**Kenner directs all content requests from terminal 50 to PIM 64 regardless of the video ID;
Kenner therefore does not direct requests to PIM 64 based on the video ID**

Q Okay. So regardless of whether or not the desired clip was found by local SRU 51, the video ID is always passed to PIM 64; correct?

A That's what's described in this paragraph, yes.

EX2017, Mowry Depo. Tr., 113:1-5.

Kenner does not direct resource requests to “the at least one repeater server” based on the “alias name”

A neighboring IM querying PIM 64 does not address the claimed directing of a request from a client machine to the PIM 64 based on an alias name

wherein, requests for the first resource from said client machines are directed to the at least one shared repeater server, based at least in part on said first alias name;

wherein requests for the second resource made from said client machines are directed to the at least one shared repeater server, based at least in part on said second alias name; and

EX1001, Claim 1 (excerpt).

It is possible to avoid the need for placing both the repeater name and the server name in the URL. Instead, a “family” of names may be created for a given origin server, each name identifying one of the repeaters used by that server.

For instance, if `www.example.com` is the origin server, names for three repeaters might be created:

`wr1.example.com`

`wr2.example.com`

`wr3.example.com`

The name “`wr1.example.com`” would be an alias for repeater 1, which might also be known by other names such as “`wr1.anotherExample.com`” and “`wr1.example.edu`”.

EX1001, '903 patent, 9:40-55.

Kenner does not direct resource requests to “the at least one repeater server” based on the “alias name”

A neighboring IM querying PIM 64 does not address the claimed directing of a request from a client machine to the PIM 64 based on an alias name

If the user's PIM 64 determines that the desired clip is not on any of the extended SRUs 66, and the PIM determines further that the user has subscription rights to access files from other regions, then the PIM 64 will query the closest IMs (e.g. IM 88) to determine if any of the remote SRUs 92 local to the other remote IMs have the desired clip. The query can be directed to those IMs which are likely to have the clip.

EX1005, Kenner, 40:29-33.

Each IM, including the PIM 64, maintains a database of all other IMs connected to the system. For each IM, the database includes the Internet address of the IM, and information on the types of files stored by the IM's SRUs. The file type information can be stored in the form of content coordinate data, as previously discussed. Accordingly, by consulting the database, the PIM 64 can determine which IMs are likely to have the desired clip, and query only those IMs.

EX1005, Kenner, 41:1-6.

Roadmap

- ▶ Overview of the '903 patent
- ▶ Overview of Kenner
- ▶ Kenner's "video ID" does not teach the claimed "alias name"
- ▶ Kenner does not direct resource requests to a repeater server based on the "alias name"
- ▶ Kenner does not teach a "table listing a plurality of alias names corresponding to content providers authorized to have resources delivered to client machines"
- ▶ Kenner does not teach "the at least one repeater server is addressable" using the "alias name" (claim 2)
- ▶ Kenner in view of Rekimoto does not teach "the first alias name is a domain name" (claim 34)

Kenner does not teach a “table listing a plurality of alias names corresponding to content providers authorized to have resources delivered to client machines”

The '903 Patent's Claims

The claims recite a table that correlates content providers who are authorized to have resources delivered to the client machines

...
at least one table listing a plurality of alias names corresponding to content providers authorized to have resources delivered to client machines via the at least one shared repeater server, wherein the plurality of content providers comprises the first content provider and the second content provider,

EX1001, Claim 1 (excerpt).

The '903 Patent's Claims

The specification emphasizes identifying content providers who are authorized to have resources delivered to the client machines

C2. The repeater uses an internal table to verify that the origin server belongs to a known “subscriber”. A subscriber is an entity (e.g., a company) that publishes resources (e.g., files) via one or more origin servers. When the entity subscribes, it is permitted to utilize the repeater network. The subscriber tables described below include the information that is used to link reflectors to subscribers.

EX1001, 10:24-31.

Adding Subscribers to the Repeater Network
When a new subscriber is added to the network, information about the subscriber is entered in a Subscriber Table at the master repeater and propagated to all repeaters in the network. This information includes the Committed Aggregate Information Rate (CAIR) for servers belonging to the subscriber, and a list of the repeaters that may be used by servers belonging to the subscriber.

EX1001, 17:50-57.

Kenner does not disclose a *table listing a plurality of alias names corresponding to authorized content providers*

Kenner checks end user subscription rights but does not describe a table that correlates content providers who are authorized to have resources delivered to the client machines

The PIM 64 checks the user's subscription rights in its user database, and if authorized and necessary, initiates a DSI process 58 to download the desired clip to the user's terminal 50. The PIM 64, having identified the clip corresponding to the video ID in its clip database, passes information to the DSI 58 regarding which extended SRUs 66 have the clip. The DSI oversees initiating the transfer process, ensuring that data is sent from the appropriate extended SRU 66 through the interface 54 to the user's terminal 50. The data is then downloaded from the appropriate extended SRU 66 to the user's terminal 50, at which time the local SRU 51 will intercept and transfer the data to the browser extension 84. The browser extension 84 can manipulate the data, store it in a storage area local to the terminal 50, or prepare it for viewing, as desired. If the PIM 64 determines that the user has not subscribed to the clip indicated by the selected video ID, no DSI process is invoked, and the local SRU 51 will be notified that either the clip must be downloaded directly from the Web page, as is traditionally done, or that the clip is entirely unavailable to that user.

EX1005, Kenner, 37:18-31.

Kenner does not disclose a *table listing a plurality of alias names corresponding to authorized content providers*

Kenner checks end user subscription rights but does not describe a table that correlates content providers who are authorized to have resources delivered to the client machines

If any of the foregoing database checks fail, the local SRU 51 will not receive authorization for the download, and the reason for the denial will be transmitted from the PIM 64 to the browser extension 84 and browser 82 user interface via the local SRU 51 to advise the user. If alternative versions of the desired clip are available which would be authorized given the user's subscription limitations, the user can be presented with the option to download the alternative versions.

EX1005, Kenner, 38:17-22.

For content that is available on a subscription basis, the invention limits access to those users who are designated as subscribers to the content. Users may directly subscribe to particular content, or ISPs may subscribe for their end users. In either case, the PIM 64 maintains in its user database a listing of all those authorized subscribers who are connected to the PIM 64, limiting downloading of clips through the invention to those who are authorized within that region.

EX1005, Kenner, 52:7-12.

Kenner does not disclose a *table listing a plurality of alias names corresponding to authorized content providers*

Kenner's clip database does not disclose a table listing alias names corresponding to content providers authorized to have resources delivered because Kenner's video ID does not correspond to any authorization

In one embodiment, the video ID is constructed from the following portions: a text name of the file as defined by the content provider; the content provider's account number as provided by the organization running the subscription service; a category coordinate, possibly a representation of a hierarchical portion of a category tree, a geographic coordinate used to determine where the file is relevant (e.g. a region, state, or city); a time stamp, and a time period over which the file is relevant. The foregoing characteristics are known at the time a file is made available to the present system, and will be discussed in detail below.

EX1005, Kenner, 36:32-37:6.

Dependent claims 2 and 34

Dependent claims 2 and 34

- ▶ Kenner does not teach “the at least one repeater server is addressable” using the “alias name” (claim 2)
- ▶ Kenner in view of Rekimoto does not teach the “alias name is a domain name” (claim 34)
 - ▶ Overview of Rekimoto
 - ▶ A POSA would not have been motivated to combine Rekimoto with Kenner
 - ▶ It would not have been obvious to structure Kenner’s video ID as a domain name

Kenner does not teach “wherein the at least one shared repeater server is addressable using the first alias name and the second alias name” as recited in dependent claim 2

Claim 2 of the '903 patent

2. The content delivery system as recited in claim 1, wherein the least one shared repeater server is addressable using the first alias name and the second alias name.

EX1001, '903 Patent, claim 2.

Claim 2 recites that the repeater be “addressable” using alias names

Alternate Approach

Placing the origin server name in the reflected URL is generally a good strategy, but it may be considered undesirable for aesthetic or (in the case, e.g., of cookies) certain technical reasons.

It is possible to avoid the need for placing both the repeater name and the server name in the URL. Instead, a “family” of names may be created for a given origin server, each name identifying one of the repeaters used by that server.

For instance, if `www.example.com` is the origin server, names for three repeaters might be created:

`wr1.example.com`

`wr2.example.com`

`wr3.example.com`

The name “`wr1.example.com`” would be an alias for repeater 1, which might also be known by other names such as “`wr1.anotherExample.com`” and “`wr1.example.edu`”.

EX1001, '903 patent, 9:39-55.

In order to access the file “F” (the resource) specified by the above URL, a program (e.g., a browser) running on a user’s computer (i.e., a client computer) would have to first locate the computer (i.e., a server computer) specified by the host name. I.e., the program would have to locate the server “`www.uspto.gov`”. To do this, it would access a Domain Name Server (DNS), providing the DNS with the host name (“`www.uspto.gov`”). The DNS acts as a kind of centralized directory for resolving addresses from names. If the DNS determines that there is a (remote server) computer corresponding to the name “`www.uspto.gov`”, it will provide the program with an actual computer network address for that server computer. On the Internet this is called an Internet Protocol (or IP) address and it has the form “`123.345.456.678`”. The program on the user’s (client) computer would then use the actual address to access the remote (server) computer.

EX1001, '903 Patent, 6:45-61.

Kenner's PIM is not "addressable" using a Video ID

2. The content delivery system as recited in claim 1, wherein the least one shared repeater server is addressable using the first alias name and the second alias name.

EX1001, '903 Patent, claim 2.

Kenner's Video ID is not a host name

In one embodiment, the video ID is constructed from the following portions: a text name of the file as defined by the content provider; the content provider's account number as provided by the organization running the subscription service; a category coordinate, possibly a representation of a hierarchical portion of a category tree, a geographic coordinate used to determine where the file is relevant (e.g. a region, state, or city); a time stamp, and a time period over which the file is relevant. The foregoing characteristics are known at the time a file is made available to the present system, and will be discussed in detail below.

EX1005, Kenner, 36:32-37:6.

A video ID consists of a number of different pieces of information, so it's -- I don't think one would normally characterize that as being a host name, but it's something that's used -- it's associated with -- it's associated

EX2017, Mowry Depo. Tr., 129:16-20.

Kenner's PIM is not "addressable" using a Video ID

2. The content delivery system as recited in claim 1, wherein the least one shared repeater server is addressable using the first alias name and the second alias name.

EX1001, '903 Patent, claim 2.

Kenner's "virtual URL" includes "Internet address of an IM."

In a preferred embodiment, the local SRU 51 passes the video ID to the PIM 64 in the following form: a "virtual URL" is constructed in the form "http://" plus the Internet address of the PIM 64, plus the user's subscriber ID number, plus the video ID. If the desired clip was located on the local SRU 51, then the virtual URL request will contain a further attribute specifying that a particular version of the file has already been located by the local SRU 51. The Netscape Navigator procedure NPN_GetURLNotify is available to accomplish sending this virtual URL to the PIM 64; the virtual URL is constructed to be in a format that will be accepted by Netscape Navigator and other network browsers. Upon receipt by the PIM 64, the virtual URL is decomposed into the video ID and subscriber ID components, which are then used to access the PIM's internal databases.


EX1005, Kenner, 37:7-17.

Kenner's PIM is not “addressable” using a Video ID

2. The content delivery system as recited in claim 1, wherein the least one shared repeater server is addressable using the first alias name and the second alias name.

EX1001, '903 Patent, claim 2.

Kenner's video ID is passed to the PIM identified in the "virtual URL." The video ID does not impact the decision of which server receives the request.



EX1005, Kenner, FIG. 4 (annotated).

Dependent claims 2 and 34

- ▶ Kenner does not teach “the at least one repeater server is addressable” using the “alias name” (claim 2)
- ▶ Kenner in view of Rekimoto does not teach the “alias name is a domain name” (claim 34)
 - ▶ Overview of Rekimoto
 - ▶ A POSA would not have been motivated to combine Rekimoto with Kenner
 - ▶ It would not have been obvious to structure Kenner’s video ID as a domain name

**Kenner in view of Rekimoto does not teach the first
alias name being a “domain name” as recited in
dependent claim 34**

Claim 34 of the '903 patent

34. A method as recited in claim 28, wherein at least the first alias name is a domain name and the computer network further comprises at least one name server that provides domain name service (DNS) resolution, the method further comprising:

receiving, at the at least one name server, the first alias name as part of a client request; and

resolving, at the at least one name server, the received first alias name to identify the at least one shared repeater server to deliver the first resource in response to the client request.

EX1001, '903 Patent, claim 34.

A In this case, yes, the alias name is a domain name, and DNS returns an IP address, so it's taking the first alias name and converting that into an IP address.

EX2017, Mowry Depo. Tr., 61:17-20.

In order to access the file “F” (the resource) specified by the above URL, a program (e.g., a browser) running on a user’s computer (i.e., a client computer) would have to first locate the computer (i.e., a server computer) specified by the host name. I.e., the program would have to locate the server “www.uspto.gov”. To do this, it would access a Domain Name Server (DNS), providing the DNS with the host name (“www.uspto.gov”). The DNS acts as a kind of centralized directory for resolving addresses from names. If the DNS determines that there is a (remote server) computer corresponding to the name “www.uspto.gov”, it will provide the program with an actual computer network address for that server computer. On the Internet this is called an Internet Protocol (or IP) address and it has the form “123.345.456.678”. The program on the user’s (client) computer would then use the actual address to access the remote (server) computer.

EX1001, '903 patent, 6:45-61.

Kenner's video ID is not a domain name

Kenner's video ID is not a domain name

In one embodiment, the video ID is constructed from the following portions: a text name of the file as defined by the content provider; the content provider's account number as provided by the organization running the subscription service; a category coordinate, possibly a representation of a hierarchical portion of a category tree, a geographic coordinate used to determine where the file is relevant (e.g. a region, state, or city); a time stamp, and a time period over which the file is relevant. The foregoing characteristics are known at the time a file is made available to the present system, and will be discussed in detail below.

EX1005, Kenner, 36:32-37:6.

A video ID consists of a number of different pieces of information, so it's -- I don't think one would normally characterize that as being a host name, but it's something that's used -- it's associated with -- it's associated

A So the term domain name and the term host name are sometimes used somewhat loosely. A host name often refers to -- at least in the early days of the Internet, they refer to a specific machine typically. And a domain name would be some later part of a host name. For example, ESPN.com, that might be a domain name, but that's -- those terms are used sometimes a bit loosely, but in some cases a domain name might be a part of a host name. But, again, it would really depend on the context.

EX2017,
Mowry Depo. Tr.,
129:16-20.

EX2017,
Mowry Depo. Tr.,
22:12-22.

Rekimoto's DNS Server

34. A method as recited in claim 28, wherein at least the first alias name is a domain name and the computer network further comprises at least one name server that provides domain name service (DNS) resolution, the method further comprising:

- receiving, at the at least one name server, the first alias name as part of a client request; and
- resolving, at the at least one name server, the received first alias name to identify the at least one shared repeater server to deliver the first resource in response to the client request.

EX1001, '903 Patent, claim 34.

Meanwhile, an IP address is made up of 32 bits and normally expressed in a decimal notation delimited by dot in units of eight bits. For example, an IP is expressed in 43.0.35.117. This IP address provides an address which uniquely identifies a source or destination terminal connected to the Internet. Because an IP address expressed in four octets (32 bits) is difficult to remember, a domain name is used. The domain name system (DNS) is provided to control the relationship between the domain names assigned to the terminals all over the world and their IP addresses. The DNS answers a domain name for a corresponding IP address and vice versa. The DNS functions based on the cooperation of the domain name servers installed all over the world. A domain name is expressed in "hanaya@ipd. sony.co.jp" for example, which denotes a user name, a host name, an organization name, an organization attribute, and country name (in the case of US, the country name is omitted) in this order. If the country name of the first layer is "jp", that terminal is located in Japan. If there is no country name, that terminal is located in US.

EX1007, Rekimoto, 18:3-12.

A POSA would not have combined Kenner with Rekimoto

It would have been obvious to incorporate *Rekimoto's* DNS server and DNS resolution capabilities into *Kenner's* network to improve the response time for delivering resources. Mowry, ¶190. As *Kenner* explains, its network includes IMs distributed across different geographic regions. *Kenner*, 49:1-4, 54:12-16, Fig. 4. For example, *Kenner's* clip database includes the IP addresses for its own local SRUs (*id.*, 34:10), but queries neighboring IMs to identify remote SRUs storing the clip (40:29-33). In the combined *Kenner-Rekimoto* system, when the clip is not available on the local SRUs, the IM will direct the client request to the DNS server described in *Rekimoto* to identify the closest IMs (and associated remote SRUs) for handling the request. Mowry, ¶190. For example, the DNS server would map the URL containing the video ID to the remote IMs in the same manner described in the '903 patent. *Id.*; EX-1001, 21:56-64. The '903 patent admits DNS was known in the art. EX-1001, 6:45-61, 7:4-15. Mowry, ¶190.

Petition, 58.


Implementing Rekimoto's DNS capabilities in Kenner's system would not "reduce system latency" because Kenner's virtual URL already identifies the PIM. Such a combination would be inconsistent with how Kenner uses the video ID.

It would not have been obvious to structure Kenner's video ID as a domain name

A POSA would not have reformatted Kenner’s video ID as a domain name. Such a modification would render Kenner’s Video ID inoperable for its stated purpose of providing a “globally unique name of the video clip.”

In a preferred embodiment, the local SRU 51 passes the video ID to the PIM 64 in the following form: a "virtual URL" is constructed in the form "http://" plus the Internet address of the PIM 64, plus the user's subscriber ID number, plus the video ID. If the desired clip was located on the local SRU 51, then the virtual URL request will contain a further attribute specifying that a particular version of the file has already been located by the local SRU 51. The Netscape Navigator procedure NPN_GetURLNotify is available to accomplish sending this virtual URL to the PIM 64; the virtual URL is constructed to be in a format that will be accepted by Netscape Navigator and other network browsers. Upon receipt by the PIM 64, the virtual URL is decomposed into the video ID and subscriber ID components, which are then used to access the PIM's internal databases.

EX1005, Kenner, 37:7-17.



EX1005, Kenner, FIG. 4 (annotated).

Petitioner introduces a new theory in the Petitioner's Reply

Petition

In the combined system, a POSITA would have recognized that *Kenner's* video ID (alias name) may be used as a domain name for resolution through *Rekimoto's* DNS server. Mowry, ¶191. *Rekimoto* describes using a domain name in the format "'hanaya@ipd.sony.co.jp' for example, which denotes a user name, a host name, an organization name, an organization attribute, and country name." *Rekimoto*, 18:3-12. *Kenner's* video ID includes similar information, such as a "provider's account number ... [,] a category coordinate, ... [and] a geographic coordinate used to determine where the file is relevant." *Kenner*, 36:32-37:6. And *Kenner* describes formatting its video ID as a virtual URL including the internet address of the primary IM. *Id.*, 37:7-17. Accordingly, it would have been obvious to structure *Kenner's* video ID as a domain name to provide the DNS resolution advantages described above. Mowry, ¶191.

Petition, 58-59.

Petitioner's Reply

First, Patent Owner alleges that the *Kenner-Rekimoto* combination "introduces an internal conflict with how *Kenner's* system operates" because *Kenner's* "virtual URL already includes an 'Internet address.'" Resp., 53-54. This ignores the actual combination proposed in the Petition. Petitioner's proposed combination does not "generate a virtual URL including both the 'Internet address' and another IP address derived from applying DNS resolution to the video ID," as Patent Owner alleges. Resp., 54. Rather, the proposed combination reformats *Kenner's* virtual URL and video ID into *Rekimoto's* format to take advantage of DNS resolution. Pet., 55-60. There is no inconsistency as the video ID and virtual

Petitioner's Reply, 22.

APPENDIX

The '903 Patent's Claims

Despite relying on Claim 25, Petitioner has not shown that “host names” and “domains names” are interchangeable. Dr. Mowry also did not testify that these terms are interchangeable.

25. A content delivery system as recited in claim 1, wherein at least the first alias name is a domain name, the system further comprising:

at least one name server that provides domain name service (DNS) resolution on the first alias name to identify the at least one shared repeater server to deliver the first resource in response to a client request.

EX1001, Claim 25.

A So the term domain name and the term host name are sometimes used somewhat loosely. A host name often refers to -- at least in the early days of the Internet, they refer to a specific machine typically. And a domain name would be some later part of a host name. For example, ESPN.com, that might be a domain name, but that's -- those terms are used sometimes a bit loosely, but in some cases a domain name might be a part of a host name. But, again, it would really depend on the context.

EX2017, Mowry Depo. Tr., 22:12-22.