

Exhibit A



US008495167B2

(12) **United States Patent**
Valjakka

(10) **Patent No.:** **US 8,495,167 B2**
(45) **Date of Patent:** **Jul. 23, 2013**

(54) **DATA COMMUNICATIONS NETWORKS, SYSTEMS, METHODS AND APPARATUS**

(76) Inventor: **Lauri Valjakka**, Espoo (FI)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 252 days.

(21) Appl. No.: **10/208,685**

(22) Filed: **Jul. 30, 2002**

(65) **Prior Publication Data**

US 2003/0093491 A1 May 15, 2003

(30) **Foreign Application Priority Data**

Aug. 2, 2001 (EP) 01660145

(51) **Int. Cl.**
G06F 15/167 (2006.01)

(52) **U.S. Cl.**
USPC **709/214**; 709/212; 709/213; 709/216

(58) **Field of Classification Search**
USPC 709/216, 218, 213, 214, 215, 217, 709/224
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,905,952 A	5/1999	Joensuu et al.	
6,038,296 A *	3/2000	Brunson et al.	379/100.11
6,157,965 A *	12/2000	Mohammed et al.	710/8
6,226,673 B1 *	5/2001	Yoshimoto	709/223
6,249,810 B1 *	6/2001	Kiraly	709/217
6,587,756 B2 *	7/2003	Moriguchi et al.	701/1
6,873,627 B1 *	3/2005	Miller et al.	370/466
6,879,982 B2 *	4/2005	Shirasaka	1/1
6,912,514 B2 *	6/2005	Matsushima et al.	705/52
6,950,431 B1 *	9/2005	Nozaki et al.	370/390

6,970,939 B2 *	11/2005	Sim	709/236
7,139,827 B1 *	11/2006	Iwayama et al.	709/227
7,222,186 B2 *	5/2007	Kobayashi	709/235
7,228,416 B2 *	6/2007	Nishizawa et al.	713/168
7,373,103 B2 *	5/2008	Sato et al.	455/7
2001/0011301 A1 *	8/2001	Sato et al.	709/219
2002/0010785 A1 *	1/2002	Katsukawa et al.	709/229

(Continued)

FOREIGN PATENT DOCUMENTS

EP	0709994	5/1996
EP	0726663	8/1996

(Continued)

OTHER PUBLICATIONS

“System and Method for communication” by Kiraly et al. International Publication # WO 00/65776.*

(Continued)

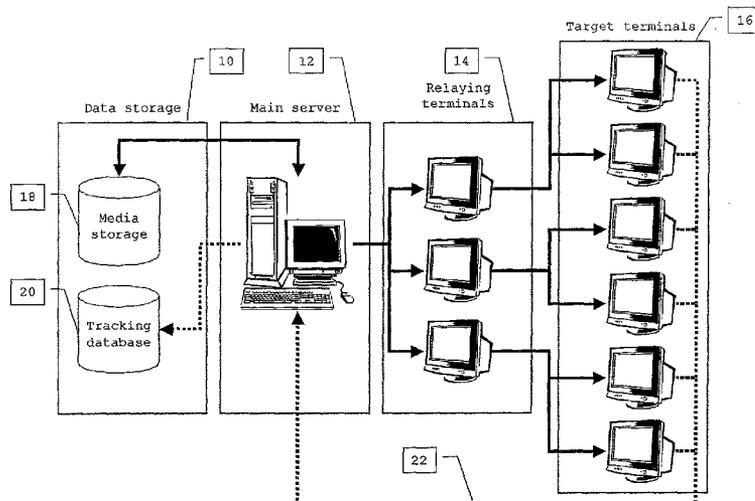
Primary Examiner — Dhairya A Patel

(74) *Attorney, Agent, or Firm* — Seppo Laine Oy; Joshua P. Wert

(57) **ABSTRACT**

A data communications network comprises a plurality of terminals and a main server adapted to manage selective retrieval of data from a first server by at least one target terminal. Some or all of the terminals are adapted to act as relay servers for serving data retrieved from the first server to at least one target terminal. The network includes a network information database and the main server selects at least one target terminal to act as a relay server for serving data to other target terminals on the basis of terminal performance information stored in the network information database. Terminals acting as relay servers also select further downstream target terminals to act as further relay servers on the basis of the relative performances of the further target terminals. The load on the main server is thus distributed among all of the relay servers, providing improved network performance.

21 Claims, 4 Drawing Sheets



US 8,495,167 B2

Page 2

U.S. PATENT DOCUMENTS

2002/0143977 A1* 10/2002 Togashi 709/231
2003/0009539 A1* 1/2003 Hattori 709/219
2004/0192275 A1* 9/2004 Kim 455/418
2006/0114350 A1* 6/2006 Shimada et al. 348/423.1

FOREIGN PATENT DOCUMENTS

EP 0863646 9/1998
HU 222 337 B1 2/2000

WO WO 00/65776 * 11/2000
WO WO 0065776 11/2000
WO WO 01/22688 A1 3/2001

OTHER PUBLICATIONS

“System and Method for communication” by Kiraly et al. Internal
Publication # WO 00/65776.*

* cited by examiner

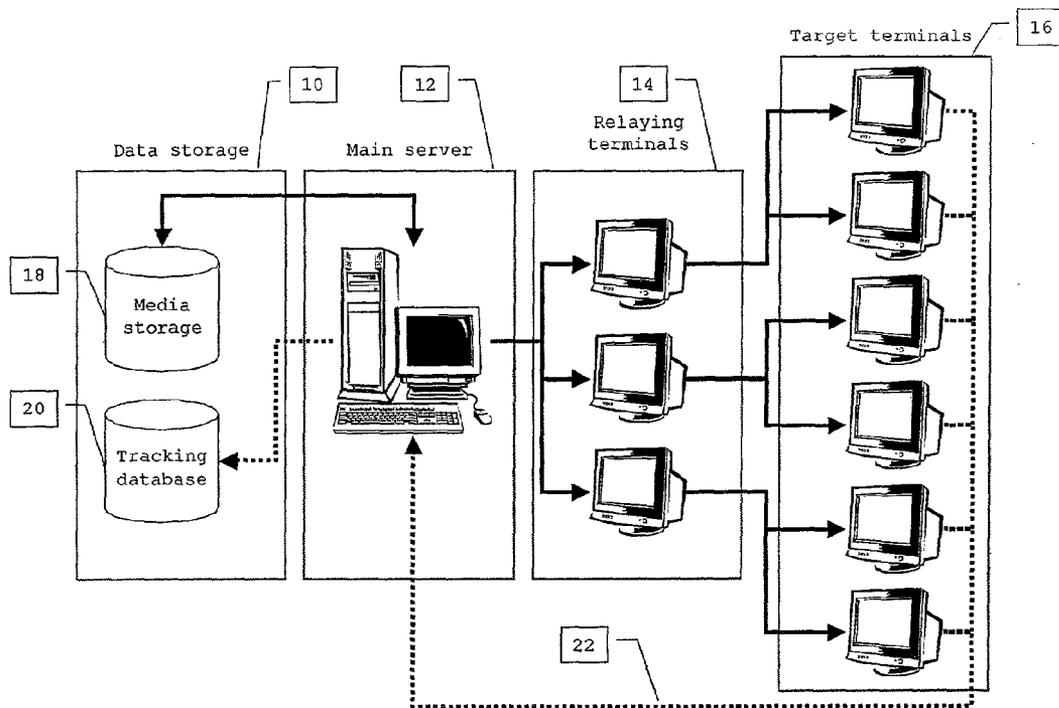


Fig. 1

Fig. 2A

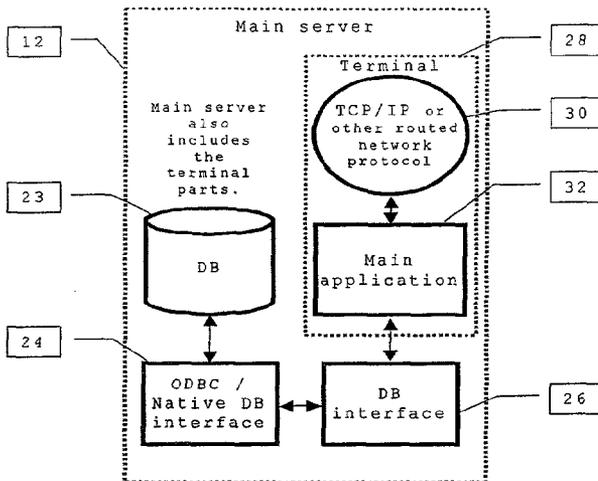


Fig. 2B

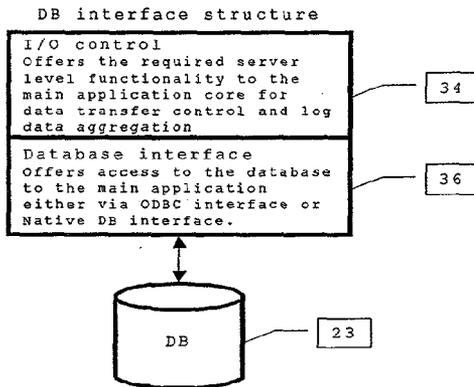
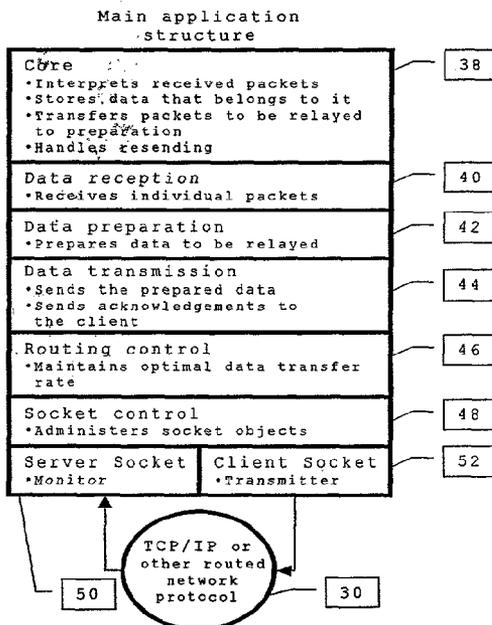


Fig. 2C



Routing

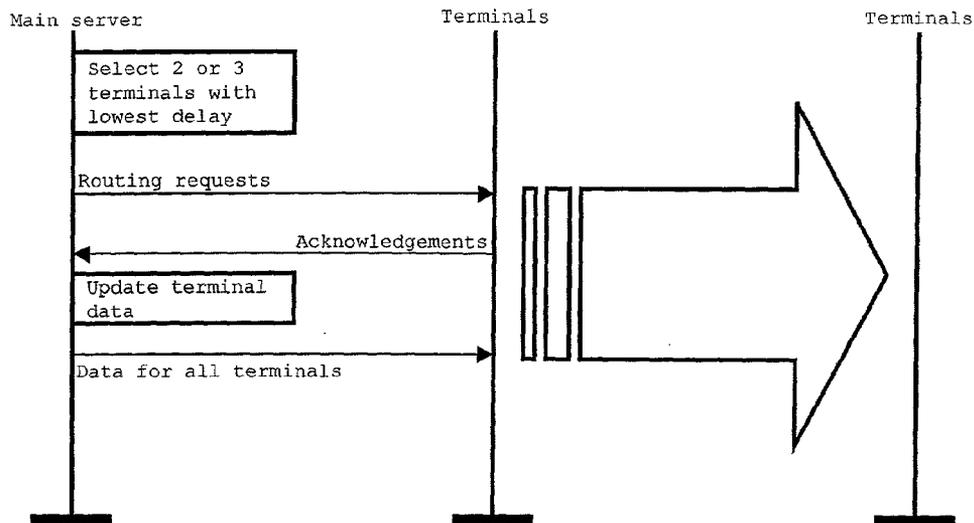


Fig. 3A

Data aggregate transfer process

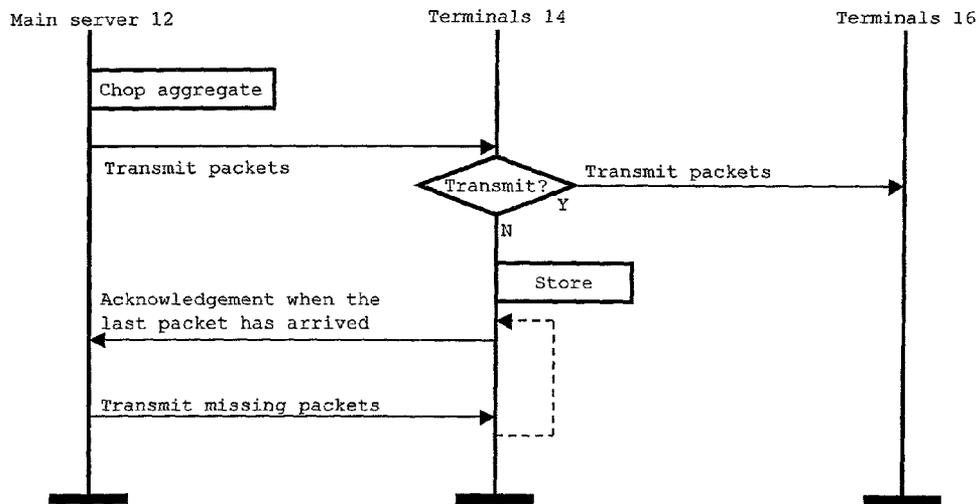


Fig. 3B

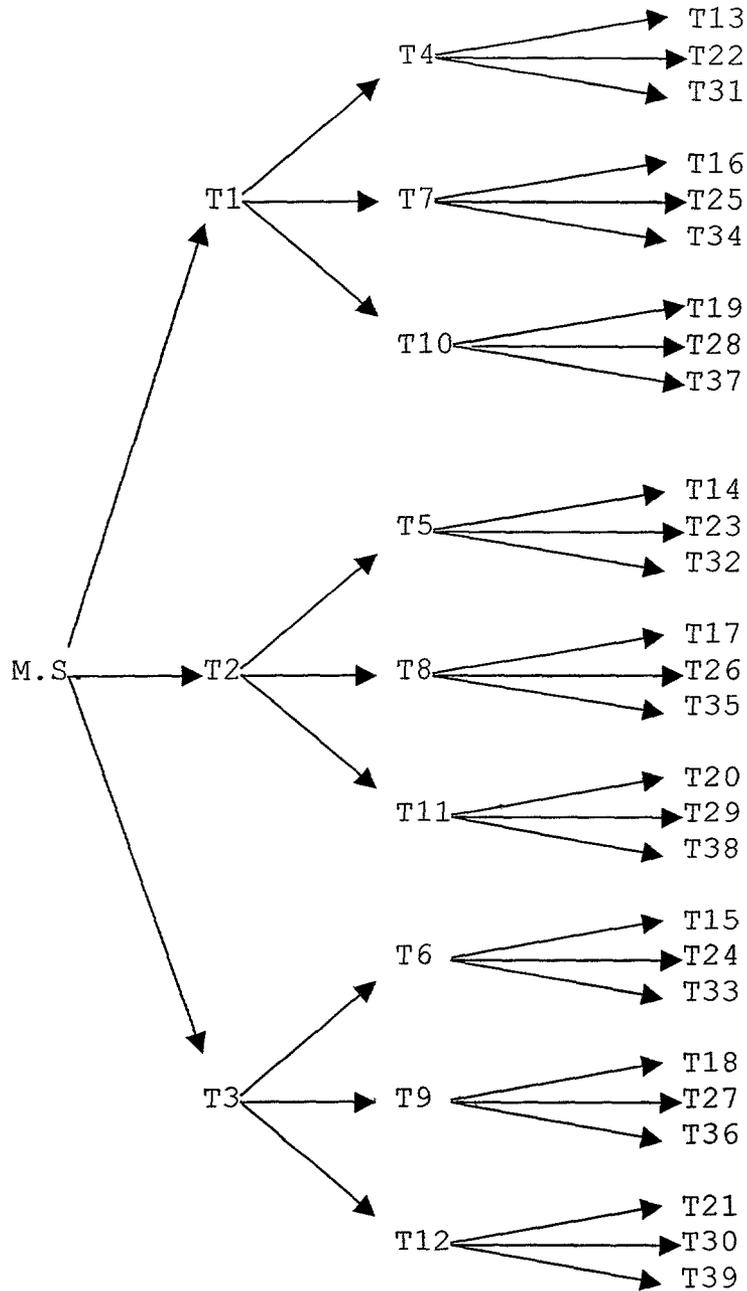


Fig. 4

US 8,495,167 B2

1

**DATA COMMUNICATIONS NETWORKS,
SYSTEMS, METHODS AND APPARATUS**

FIELD OF THE INVENTION

The present invention relates to improvements in data communications networks and to systems, methods and apparatus employed in such networks.

BACKGROUND TO THE INVENTION

In conventional client/server data networks, such as TCP/IP or other routed networks, a main server serves all terminals via a single server socket. This results in extreme spikes in the network load, especially when data is required to be transferred to a large number of clients simultaneously, causing delays in data transmission.

The present invention seeks to provide improved network systems, methods and apparatus whereby network performance is enhanced.

SUMMARY OF THE INVENTION

The invention provides improved data communications networks, methods of operating data communications networks, network servers, network terminals and computer programs as defined in the claims appended hereto.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

FIG. 1 is a diagram illustrating the operational model of a data communications network embodying the present invention;

FIGS. 2A, 2B and 2C are diagrams illustrating the operational structure of a main server and terminals employed in the network of FIG. 1;

FIGS. 3A and 3B are transaction diagrams illustrating routing and data transfer processes employed in the network of FIG. 1; and

FIG. 4 is a diagram illustrating one example of a scheme for distributing data from a main server to a number of target terminals in accordance with the invention.

DETAILED DESCRIPTION OF THE PREFERRED
EMBODIMENTS

Referring now to the drawings, FIG. 1 illustrates an operational model of a simplified exemplary embodiment of a data communications network in accordance with the invention. The network includes a data storage system 10, which in this embodiment includes media storage system 18 for data (i.e. "media" or "content") that is to be selectively distributed over the network, and a tracking database 20 that is used for managing the operation of the network as shall be described in more detail below. For convenience, data that is to be distributed from the media storage system 18 will be referred to herein as "content", which will be understood to include any type of data of interest to end users, including but not limited to text, graphics, video, audio, executable code etc. Content will generally comprise a data file of some type.

For the purposes of the present invention, "content" means files or parts of files or equivalents thereof that are stored on a server, downloaded from the server by a client and stored by the client for subsequent use, as distinct from digital broad-

2

cast media in which a data stream is transmitted by a broadcast server and is temporarily buffered by clients and, in some cases, by intervening relay units.

The network further includes a main server 12 that communicates with the media storage system 18 and tracking database 20, and controls the distribution of content from the media storage system 18. The network also includes a plurality of terminals 14 and 16, to which content is to be distributed. In accordance with the invention, when the same content is to be distributed to a number of terminals, at least some of the terminals 14 also act as "relay servers" in distributing the content to the remaining terminals 16 (i.e. some or all of the terminals may also be capable of acting as relay servers).

All transactions between the media storage system 18 and the terminals 14, 16 are controlled by the main server 12. In particular, all data downloads to the terminals from the media storage system 18 are managed by the main server 12. Generally, content is retrieved from the storage system by the main server and forwarded on to the terminals 14, 16 by the main server. In some cases, however, the main server does not itself retrieve and forward content, but manages the retrieval and forwarding of content by other servers.

The term "target terminal" used here means a terminal which is the intended recipient of content (a data file) from the media storage 18. Each terminal 14, 16 can be the target for a data file. In this embodiment, each of the first set of terminals 14 is also adapted to operate as a relay server by forwarding data to one or more of the second set of terminals 16 as described further below. The terminals 16 may also act as relay servers for relaying data to additional terminals (not shown) downstream thereof. It will be understood that not all of the terminals included in the network need operate as relay servers and the network may include terminal devices that are not suited for operation as relay servers.

The tracking database 20 keeps records of transactions between the main server 12 and the various terminals 14, 16. In particular, the tracking database monitors the performance (communication speed and/or other parameters such as reliability) of all terminals that also act as relay servers in the network. This information is available to the main server. In particular, the tracking database 12 is able to provide the main server with lists of terminal addresses ranked by their relative performances.

In operation of the network, when a content data file is to be distributed to particular target terminals, the main server 12 initiates a data transport operation by sending a transport request to the first set of terminals 14, which are selected as being the best terminals from the list of target terminals on the basis of the current performance data. The transport request includes:

Details of the file to be transported. These will generally include, for example, the file type and size, time stamps for activation and deactivation of the content, encryption and compression details, etc.

The addresses of relay servers and terminals that are to be involved in the distribution of the file.

The transport request sent from the main server 12 to the first set of terminals 14 instructs these terminals to retrieve the data from the main server 12 (or from another server address included in the transport request). The list of the remaining target terminal addresses is divided between the first terminals 14, so that each of the first terminals 14 acts as a relay server for distributing the data to a subset of the remaining target terminals.

In response to the transport request from the main server 12, each of the first terminals 14 begins to download the file from the main server 12. When one of the first terminals 14

US 8,495,167 B2

3

has received a predetermined number of bytes of the file, that terminal **14** sends a modified version of the original transport request to its subset of the target terminals **16**. The modified transport request identifies the relevant first terminal **14** as the server address from which its subset of the target terminals **16** should retrieve the data. Depending on the number of target terminals, the list of target terminals may be sub-divided a number of times. That is, each of the second set of terminals **16** may receive a list of further target terminals for which it is to act as a relay server. At each stage, it is preferred that the “best” terminals from the list of remaining targets are selected to act as relay servers for the remainder.

When each terminal **14** or **16** has downloaded the whole file, it sends a notification message direct to the main server **12**, as indicated by **22** in FIG. **1**.

The main server **12** is adapted to serve data requests from the first set of terminals **14**. If the terminal in the second set of terminals **16** cannot reach the terminal in the first set of terminals **14** it will send the data request to the main server **12**.

Generally, the main server and each downstream terminal acting as a relay server will only serve a small number (e.g. 2 to 5) of downstream terminals. If the number of target terminals is less than or equal to this number, the target terminals may all retrieve the data direct from the main server, or the main server may request the best of the target terminals to act as the relay server for the other(s).

It will be understood that the network may include many more terminals than are illustrated in FIG. **1**, arranged in a tree structure wherein each terminal is either a node (functioning as both a relay server and a target terminal) or a leaf (functioning only as a target terminal); i.e. there may be multiple node terminals in the downstream data transmission path between the main server and each target terminal. Preferably, there is also an upstream communication path **22** from each terminal **14**, **16** direct to the main server **12**. The upstream path **22** is used by target terminals to acknowledge receipt of data. These acknowledgements are sent directly from the target terminals to the main server **12** as illustrated. The upstream path **22** between the terminals **14** and the main server **12** has been omitted from FIG. **1** for clarity of illustration.

It should be understood that the operational model illustrated in FIG. **1** may be implemented using an existing, conventional network infrastructure (such as the Internet or equivalent) and does not require a new physical network. Servers and terminals may be connected to the network backbone by synchronous fixed connections such as ISDN, HSDL, T1 or T3 and the network may include dial-up connections, wireless connections etc. That is, FIG. **1** illustrates logical connections between the server and terminals, rather than physical connections. Further, the logical connections between the main server and terminals vary dynamically in use of the network, as shall be described further below.

The invention is particularly suited for use where all terminals are capable also of acting as relay servers as described and can be assumed to be permanently on-line. However, it will be understood that the invention may be adapted to accommodate terminals that do not also act as relay servers (such terminals would always be “leaves”, at the end of lists of target terminals).

The target terminal requests each packet to be transferred separately. The packet to be transferred includes the information about the type of the data to be transferred, size, compression, and the checksums required for the validation of the transferred data packet.

FIG. **2A** of the drawings illustrates the operational structure of the main server **12**, including a network database **23**

4

for storing network information including the addresses etc. of network terminals (this database may implement all or part of the functionality of the tracking database **20** of FIG. **1**; these database functions can be performed by one or more database systems on one or more computers/servers), database interface modules **24**, **26**, and a terminal module **28**. The terminal module **28** included in the main server **12** is also used in each of the network terminals/relay servers **14** and **16**, and includes a routed network protocol module **30** (preferably a TCP/IP module, but other routed protocols may be used) and a main application module **32**. As shown in FIG. **2B**, the database interface modules **24**, **26** provide I/O (input/output) control functions **34**, providing the required server level functionality to the core of the main application **32** for data transfer control and for logging data aggregation, and database interface functions providing access to the network database **23**. For example, this may be either via an ODBC (open database connectivity) interface or a database interface native to the network database system **23**.

As shown in FIG. **2C**, the main application **32**, as employed in both the main server and those terminals that also act as relay servers, comprises the following functional modules:

A core module **38** interprets received packets and stores data.

A data reception module **40** receives individual packets.

A data preparation module **42** prepares data to be relayed.

A data transmission module **44** sends data prepared by the preparation module **42** and sends acknowledgements to relevant clients.

A routing control module **46** maintains optimal data transfer rates.

A socket control module **48** administers socket objects.

A server socket **50** monitors data received via the TCP/IP module (or other routed network protocol module) **30**.

Client sockets **52** transmit data via the module **30**. The number of client sockets varies dynamically depending on the number of server connections required at any particular time.

In a conventional system, a server has a server-oriented connection for clients, comprising a server socket which is used to connect to the client’s server socket. In the present invention, the main application used in the main server **12** and in each terminal that also operates as a relay server contains a standard server socket **50** for receiving data from its clients. In addition to this, the main application also has client sockets **52** for downstream communications to the downstream terminals. The actual data to be transmitted to the target terminals is sent via these client sockets and acknowledgements are received from terminals via the server socket. When the required data has been sent by the server, the client socket created for the purpose of sending the data can be destroyed, so as not to consume network resources unnecessarily. By this method, received acknowledgements will not cause any interruptions in the outgoing data flow. Each terminal/server has two “hard-coded” sockets, one client socket **52** for serving other terminals/servers and one server socket **50** for main-server connection use only. Additional sockets can be created and used dynamically as required. Each socket has an independent processor thread controlling it so that sockets can be managed and controlled without interrupts and delays.

The opening and operation of sockets is handled dynamically using a C++ class-application which generates a new socket when it needs a new instance of this class. In this manner sockets can be managed dynamically and their number varied as necessary. Each thread owns and controls its own sockets. When a socket is no longer needed the controlling thread destroys the socket and then destroys itself.

US 8,495,167 B2

5

The operation of the network will now be described. FIG. 3A illustrates the routing process.

The main server selects a first set of a few (two or three) terminals, and sends the transport request (which includes the addresses of the relevant target terminals) to each of these. Each of this first set of terminals acknowledges its connection in the dynamic route by sending a message direct to the main server. The speed of this acknowledgement can be used to update the terminal data used for monitoring terminal performance. This first set of terminals is selected as being the “best” (“fastest”) terminals for use in transferring data to the particular target terminal, based on performance data previously acquired in operation of the network and stored in the network database.

When the data transfer is underway, data is transferred to the known terminals already registered as part of the network. If a new terminal is registered to the main server during the transfer it will be included in the next data transfer.

As previously described, the main server selects the terminals with the shortest response times. This information is obtained in the following manner: the primary recipient of routing acknowledgements from particular terminals is the “server role” application that sent the transport request to those terminals. When the transfer chain is completed, information is naturally relayed automatically to the main server. The performance of different terminals (network addresses) is measured simply by measuring the response time between different terminals and by selecting the terminals with shortest response times.

It is not necessary for the terminals to know the entire network address space of the network, since the target terminal addresses are included in the transport requests.

As part of the transport request, the main server sends the addresses of other target terminals to the first set of terminals/relay servers. Each terminal selects its own downstream terminals/relay servers and sends the rest of the target network addresses to these terminals/relay servers as part of the modified transport request. That is, each one of the first set of terminals selects a further two or three “best” terminals/relay servers from the addresses forwarded to it by the main server and passes the modified transport request on to these terminals, including the details of the other remaining target terminals. Because of this dynamic routing, the main server need not know explicitly which terminals deliver data and which terminals receive it. It is sufficient that it is ensured that each terminal in the route is accessible. If the delivery fails for one terminal for some reason, this is registered in the database and failed deliveries are repeated during the next transfer.

Once the route to a particular target has been established, the packets of the data file are passed along the defined route via the selected relay servers on the basis of the target terminal address in the handle/header of each packet.

Automatic routing evenly divides the load over a larger network region, reducing the time window required for any particular data transfer operation.

The data transfer process is illustrated in FIG. 3B.

As shown in FIG. 3B, in response to the transport request each target terminal requests the data from the main server or the upstream terminal acting as the server (as specified in the transport request) as packets, reassembles the file and, if necessary, relays the packets to downstream target terminals. When the target terminal has received the last packet of the file, it sends an acknowledgement to the main server.

It is preferred that all data is transferred in encrypted and compressed binary format. In this manner data security is improved as compared with transferring plain text and data transfer requires less time. Binary format data requires less

6

“intelligence” from the relevant application as there is no need to interpret the received data. It can be restructured directly to form a suitable data structure. All received data is primarily restructured to the base type (identified in the packet header), after which the information included in the base type indicates the oriented data type. This mechanism also provides for data verification: the size of each data type is predetermined and the amount of received data must correspond to the size of the data type.

Since the data delivered is binary only, and the size of the packets is quite small and the number of the packets may be quite large, there is no risk that the purpose of the delivered data may be determined in the event that some of the packets are accessed by unauthorised parties on delivery. It is very difficult to deduce the content of binary data without knowing its structure. Accordingly, this improves data security when using a public network.

In order to provide a better understanding of the invention, examples of data transfers will be described with reference to a preferred embodiment of a network in accordance with the invention. As previously described, the main server includes or has access to a network database that lists all of the currently active/registered terminals/relay servers in the network, ranked in order of their performance (speed). Assume that data to be transferred from the main server to one or more target terminals comprises a single data file.

As previously described, the transport request includes the address(es) of the/each target terminal and other information about the data to be transferred, including the number of packets etc.

As a first example, assume that the data is to be transferred to a single target terminal. The main server sends the transportation request direct to the target terminal. The target terminal acknowledges the request and then requests the main server to send each packet in turn. Each of these packets is compressed and encrypted individually. The target terminal acknowledges each packet. If a particular packet fails, it is only necessary to re-transmit that packet, rather than to begin the entire download from the beginning. In some circumstances, data transfers to a single target terminal using the invention might not be significantly faster than conventional download methods. However, the compression applied to the packets and the fact that failed packets do not require the download to be re-started mean that single target downloads are generally quicker and more reliable than conventional methods, particularly for very large files.

As a second example, referring to FIG. 4, assume that the data is to be transferred to thirty nine target terminals T1-T39, ranked in order of performance. Assume that the main server, M.S., and each terminal acting as a server will communicate directly downstream only with a predetermined number N of downstream terminals, and that $N=3$. The main server sends a first transport request to terminal T1, a second transport request to terminal T2, and a third transport request to terminal T3, each including one third of the complete list of target addresses. Since the terminal addresses are ranked in order of performance, in order to distribute the load evenly across the network the request sent to T1 comprises every $1+N$ th address (T1, 4, 7, 10, 13, 16, 19, 22, 25, 28, 31, 34, 37), the request sent to T2 comprises every $2+N$ th address (T2, 5, 8, 11, 14, 17, 20, 23, 26, 29, 32, 35, 38), and the request sent to T3 comprises every $3+N$ th address (T3, 6, 9, 12, 15, 18, 21, 24, 27, 30, 33, 36, 39). It can be seen how this approach may be applied for any value of N and any number of terminals.

Referring to T1 and its associated downstream addresses, upon receipt of the request from the main server, T1 acknowledges the request and can immediately begin downloading

US 8,495,167 B2

7

packets from the main server. T1 also relays the modified request to the next set of N fastest terminals (T4, T7 and T10) of the list of target addresses sent to T1. The request relayed to each of T4, T7 and T10 includes 1/Nth ($\frac{1}{3}$) of the remaining addresses originally sent to T1, distributed in a similar manner to that in which the complete target list was originally distributed among T1, T2 and T3 (i.e. T4 receives the addresses for T13, T22 and T31; T7 receives the addresses for T16, T25 and T34; and T10 receives the addresses for T19, T28 and T37). Each of the terminals T4, T7 and T10 acknowledges the request to the main server, begins downloading packets from T1 (this process can begin before the download from the main server to T1 is complete), and forwards the further modified request to the remaining terminals in its own address list, each of which acknowledges the request to the main server and begins downloading packets from its respective relay terminal. In this example, these are the final “leaf” terminals, but it can be seen how this process could be extended to any number of terminals through any number of relay stages. It can also be seen how the same scheme applies to the target address lists for T2 and T3.

It will be understood that the precise distribution scheme could be varied from that illustrated in FIG. 4. The important point is that relatively faster terminals are used at the beginning of the routes and the relatively slowest terminals are at the ends of the routes.

If a transfer to a particular terminal fails, that terminal is moved down the target list, so that the next fastest terminal in the relevant subset of the distribution list is “promoted” in the tree structure. For example, in FIG. 4, if the connection from T2 to T8 fails, T8 would be swapped with T17. If the new connection also failed then other options would be tried. If all available options fail then this is reported back to the main server.

It will also be understood that the distribution scheme in accordance with the invention could be implemented using different network architectures. The network database need not be on the same server/computer as the distribution management system (that generates the transport requests), but must be accessible to it. The data to be transferred need not be resident on or accessible to the same server/computer as the distribution management system. The transport request sent to the first set of terminals (T1, T2, T3 in FIG. 4) could include a further address of another server (a “distribution server”) from which the data is to be obtained. The distribution server may have substantially the same functionality as previously described for the main server and the relay servers.

The terminal downloading the data acknowledges the packets to the server from which it is downloading. When the download is complete it sends the acknowledgement to the main server.

The invention thus provides data communications systems, methods and apparatus with improved performance, in which some or all terminals also operate as relay servers, as necessary, dynamic routing and distributed data transfer ensures optimal or near-optimal data transfer rates to every terminal in the entire terminal network and dynamic routing ensures data delivery even if part of the network fails.

Improvements and modifications may be incorporated without departing from the scope of the invention as defined in the claims appended hereto.

The invention claimed is:

1. A data communication network comprising: a plurality of terminals; and
a main server adapted to manage selective retrieval of data from a first server by at least one target terminal selected

8

from said plurality of terminals, said main server being distinct from said first server; and
a network information database containing terminal performance information, wherein

at least two of said terminals are adapted to act as relay servers for serving data retrieved from said first server to at least one target terminal; and wherein

the main server is adapted to send transport requests direct to at least one first target terminal on the basis of said terminal performance information, and wherein the main server is further adapted to monitor response times of terminals in the network and in which terminals are selected to act as relay servers for a particular data transfers on the basis of their relative response times, and the first target terminal is adapted to act as relay server; and wherein each such transport request includes details of data to be retrieved, the address of the first server from which the data is to be requested by the first target terminal, the addresses of at least one second target terminal to which the data from the first server to be relayed by the first target terminal and an indication of a relative performance of a further target terminal based on the terminal performance information stored in the network information database; and

wherein terminals adapted to act as relay servers are adapted to modify transport requests received from the main server or from other relay servers and to transmit the modified transport request to selected target terminals from a set of target terminals identified in the transport request, wherein the modified transport request further includes addresses of further target terminals for which the recipient of the modified transport request is to act as relay server; and

wherein data to be retrieved by said target terminals are divided into a series of packets for transmission to said target terminals and each of said terminals is adapted to communicate directly with said main server to acknowledge receipt of the last packet of a series routed thereto.

2. The network as claimed in claim 1, wherein the modified transport request identifies the terminal transmitting the modified transport request as the server from which the recipients of the modified transport request should request the data.

3. The network as claimed in claim 1, wherein terminals acting as relay servers are adapted to select further downstream target terminals to act as further relay servers on the basis of their relative performances of the further target terminals indicated in said transport request.

4. The network as claimed in claim 1, wherein the first server is a terminal adapted to act as relay server.

5. The network as claimed in claim 1, wherein each of said terminals is adapted to communicate directly with said main server in an upstream direction.

6. The network as claimed in claim 1, wherein data is routed to said terminals as routed network protocol traffic such as TCP/IP traffic.

7. The network as claimed in claim 1, wherein said main server and each of said terminals includes a server socket for direct upstream communications between said terminals and said main server and at least one dynamically controlled and managed client socket for downstream data transfers between the main server and said terminals or between terminals acting as relay servers and other downstream terminals.

8. The network as claimed in claim 1, wherein data is transmitted in binary format.

9. A method of operating a data communication network, the data communication network comprising: a plurality of

US 8,495,167 B2

9

terminals, a network information database and a main server adapted to manage selective retrieval of data from a first server by at least one target terminal selected from said plurality of terminals; comprising

operating at least two of said terminals as relay servers for serving data retrieved from said first server to at least one target terminal, wherein said main server is distinct from said first server, and further comprising:

sending transport requests from the main server to at least one first target terminal based on terminal performance information stored in the network information database; and operating the first target terminal to act as relay server;

operating the main server to monitor the response times of terminals in the network and selecting terminals to act as relay servers for particular data transfer on the basis of their relative response times;

wherein each such transport request includes details of data to be retrieved, the address of the first server from which the data is to be requested by the first target terminal, addresses of at least one second target terminal to which the data retrieved from the first server is to be relayed by the first target terminal and an indication of a relative performance of a further target terminal based on the terminal performance information stored in the network information database;

operating terminals adapted to act as relay servers are adapted to modify transport requests received from the main server or from other relay servers and to transmit the modified transport request to selected target terminals from a set of target terminals identified in the transport request, wherein the modified transport request further includes addresses of further target terminals for which the recipient of the modified transport request is to act as relay server; and

wherein dividing data to be retrieved by said target terminals into a series of packets for transmission to said target terminals and wherein each of said terminals communicates directly with said main server to acknowledge receipt of the last packet of a series routed thereto.

10. The method as claimed in claim 9, including the modified transport request identifying the terminal transmitting the modified transport request as the server from which the recipients of the modified transport request should request the data.

11. The method as claimed in claim 9, including operating terminals acting as relay servers to select further downstream target terminals to act as further relay servers on the basis of the relative performances of the further target terminals indicated in said transport request.

12. The method as claimed in claim 9, wherein the first server is a terminal adapted to act as relay server.

13. The method as claimed in claim 9, wherein each of said terminals communicates directly with said main server in an upstream direction.

14. The method as claimed in claim 9, including routing data to said terminals as routed network protocol traffic such as TCP/IP traffic.

15. The method as claimed in claim 9, including transmitting said data in binary format.

16. A network server adapted to operate as a main server in a data communication network, the data communication network including:

a plurality of terminals, a network information database and a first server which from which data be retrieved by at least one target terminal from among said plurality of terminals, at least two of said terminals being adapted to

10

act as relay servers for serving data retrieved from said first server to at least one further target terminal based on terminal performance information stored in the network information database, said network server being distinct from said first server;

said network server being adapted to manage selective retrieval of data from said first server by at least one target terminal selected from said plurality of terminals; and wherein said network server being further adapted to monitor response times of terminals in the network and in which terminals are selected to act as relay servers for a particular data transfers on the basis of their relative response times,

said network server being further adapted to send transport requests direct to at least one first target terminal that is adapted to act as a relay server, each such transport request includes details of data to be retrieved, the address of the first server from which the data is to be requested by the first target terminal, the addresses of at least one second target terminal to which the data retrieved from the first server is to be relayed by the first target terminal and an indication of a relative performance of a further target terminal based on the terminal performance information stored in the network information database;

wherein terminals adapted to act as relay servers are adapted to modify transport requests received from said network server or from other relay servers and to transmit the modified transport request to selected target terminals from a set of target terminals identified in the transport request, wherein the modified transport request further includes addresses of further target terminals for which the recipient of the modified transport request is to act as relay server; and

wherein data to be retrieved by said target terminals are divided into a series of packets for transmission to said target terminals and each of said terminals are adapted to communicate directly with said main server to acknowledge receipt of the last packet of a series routed thereto.

17. A network terminal to operate as a relay server in a data communication network, the data communication network including:

a plurality of terminals, a network information database, a first server from which data may be retrieved by at least one target terminal from among said plurality of terminals; and

a main server adapted to manage selective retrieval of data from the first server by at least one target terminal selected from said plurality of terminals based on terminal performance data stored in the network information database, and wherein the main server is further adapted to monitor response times of terminals in the network and in which terminals are selected to act as relay servers for a particular data transfers on the basis of their relative response times;

said network terminal being adapted to act as relay server for serving data retrieved from said first server to at least one target terminal by receiving and responding to transport requests sent to said network terminal, each such transport request including details of data to be retrieved, the address of the first server from which the data is to be requested by the network terminal, the addresses of at least one second target terminal to which the data retrieved from the first server is to be relayed by the network terminal and an indication of a relative

11

performance of a further target terminal based on the terminal performance stored in the network information database;

wherein said network terminal adapted to act as relay server are further adapted to modify transport requests received from the main server or from other relay servers and to transmit the modified transport request to selected target terminals from a set of target terminals identified in the transport request, wherein the modified transport request further includes addresses of further target terminals for which the recipient of the modified transport request is to act as relay server; and

wherein data to be retrieved by said target terminals are divided into a series of packets for transmission to said target terminals and each of said terminals are adapted to communicate directly with said main server to acknowledge receipt of the last packet of a series routed thereto.

18. The network terminal as claimed in 17, wherein the modified transport request identifies the terminal transmitting the modified transport request as the server from which the recipients of the modified transport request should request the data.

19. A computer program product for enabling a network server to operate as a main server in a data communication network, the data communication network including:

a plurality of terminals, a network information database and a first server which from which data be retrieved by at least one target terminal from among said plurality of terminals, at least two of said terminals being adapted to act as relay servers for serving data retrieved from said first server to at least one further target terminal based on terminal performance information stored in the network information database, said main server being distinct from said first server, said computer program product comprising:

a non-transitory computer usable medium having computer readable program code means embodied in said non-transitory medium, said computer readable program code means including:

computer readable program code for causing said network server to manage selective retrieval of data from said first server by at least one target terminal selected from said plurality of terminals; and wherein said network server to monitor response times of terminals in the network and in which terminals are selected to act as relay servers for a particular data transfers on the basis of their relative response times,

computer readable program code for causing said network server to send transport requests direct to at least one first target terminal that is adapted to act as a relay server, each such transport request including details of data to be retrieved, the address of the first server from which the data is to be requested by the first target terminal, the addresses of at least one second target terminal to which the data retrieved from the first server is to be relayed by the first target terminal and an indication of a relative performance of a further target terminal based on the terminal performance information stored in the network information database;

a computer readable program code means for causing said network terminal to modify transport requests received from said network server or from other relay servers and to transmit the modified transport request to selected target terminals from a set of target terminals identified in the transport request, wherein the modified transport

12

request further includes addresses of further target terminals for which the recipient of the modified transport request is to act as relay server; and

wherein data to be retrieved by said target terminals are divided into a series of packets for transmission to said target terminals and each of said terminals are adapted to communicate directly with said main server to acknowledge receipt of the last packet of a series routed thereto.

20. A computer program product for enabling a network terminal to operate as a relay server in a data communication network, the data communication network including:

a plurality of terminals, a network information database, a first server from which data may be retrieved by at least one target terminal from among said plurality of terminals; and

a main server adapted to manage selective retrieval of data from the first server by at least one target terminal selected from said plurality of terminals based on terminal performance data stored in the network information database, and wherein the main server to monitor response times of terminals in the network and in which terminals are selected to act as relay servers for a particular data transfers on the basis of their relative response times; said computer program product comprising:

a non-transitory computer usable medium having computer readable program code means embodied in said non-transitory medium, said computer readable program code means including:

computer readable program code for causing said network terminal to act as relay server for serving data retrieved from said first server to at least one target terminal by receiving and responding to transport requests sent to said network terminal, each such transport request including details of data to be retrieved, the address of the first server from which the data is to be requested by the network terminal, the addresses of at least one second target terminal to which the data retrieved from the first server is to be relayed by the network terminal and an indication of a relative performance of a further target terminal based on the terminal performance stored in the network information database;

said computer readable program code for causing said network terminal to modify transport requests received from the main server or from other relay servers and to transmit the modified transport request to selected target terminals from a set of target terminals identified in the transport request, wherein the modified transport request further includes addresses of further target terminals for which the recipient of the modified transport request is to act as relay server; and

wherein data to be retrieved by said target terminals are divided into a series of packets for transmission to said target terminals and each of said terminals are adapted to communicate directly with said main server to acknowledge receipt of the last packet of a series routed thereto.

21. The computer program product as claimed in claim 20, said computer readable program code means further comprising computer readable program code whereby the modified transport request identifies the terminal transmitting the modified transport request as the server from which the recipients of the modified transport request should request the data.

Exhibit B



US010726102B2

(12) **United States Patent**
Valjakka et al.

(10) **Patent No.:** **US 10,726,102 B2**
(45) **Date of Patent:** **Jul. 28, 2020**

(54) **METHOD OF AND SYSTEM FOR PROVIDING ACCESS TO ACCESS RESTRICTED CONTENT TO A USER**

(56) **References Cited**

U.S. PATENT DOCUMENTS

(71) Applicant: **SC Intelligent Holding Ltd.,**
Lappeenranta (FI)

(72) Inventors: **Lauri Valjakka,** Lappeenranta (FI);
Jukka-Pekka Jussila, Lappeenranta (FI);
Jari Tapio, Lappeenranta (FI)

(73) Assignee: **IPRA Technologies Oy Ltd.,**
Lappeenranta (FI)

8,184,811 B1	5/2012	Patten et al.	
2003/0172278 A1*	9/2003	Farnham	H04L 63/0435 713/176
2003/0210789 A1*	11/2003	Farnham	H04L 9/0844 380/270
2004/0139315 A1	7/2004	Tokutani et al.	
2005/0039034 A1	2/2005	Doyle et al.	
2005/0071280 A1	3/2005	Irwin et al.	
2005/0273629 A1	12/2005	Abrams et al.	
2009/0254751 A1*	10/2009	Fujiwara	H04L 9/12 713/171

(Continued)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

FOREIGN PATENT DOCUMENTS

CN	101739522 B	1/2013
EP	1189432 A2	3/2002

(Continued)

Primary Examiner — Yonas A Bayou
(74) *Attorney, Agent, or Firm* — Seppo Laine Oy

(21) Appl. No.: **14/591,952**

(22) Filed: **Jan. 8, 2015**

(65) **Prior Publication Data**
US 2015/0193601 A1 Jul. 9, 2015

(30) **Foreign Application Priority Data**
Jan. 8, 2014 (FI) 20145013

(51) **Int. Cl.**
G06F 21/10 (2013.01)

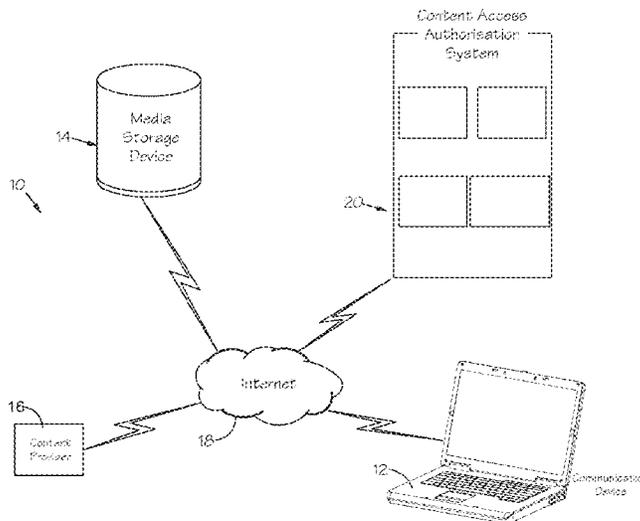
(52) **U.S. Cl.**
CPC **G06F 21/10** (2013.01)

(58) **Field of Classification Search**
CPC G06F 21/6245; G06F 21/6218; G06F 21/6209; G06F 21/62; H04L 63/102
See application file for complete search history.

(57) **ABSTRACT**

According to an example embodiment of the invention, there is provided a system for providing access to access restricted content to a user, the system including a communication arrangement operable to receive a content request message, the content request message including a content identifier, a processor configured to cause a first determination to be performed to yield a positive or a negative result, a validation module configured to, in response to the first determination yielding a positive result, obtain a first digital rights management key, the processor being further configured to cause a second determination to be performed to yield a positive or a negative result, and responsive to the first and second determinations yielding a positive result, the validation module is configured to cause access to the access restricted content to be provided to the user.

11 Claims, 8 Drawing Sheets



US 10,726,102 B2

Page 2

(56)

References Cited

U.S. PATENT DOCUMENTS

2010/0211779	A1*	8/2010	Sundaram	H04L 9/0847	713/168
2010/0250388	A1	9/2010	Lee		
2010/0268649	A1*	10/2010	Roos	G06F 21/10	705/50
2010/0275023	A1*	10/2010	Fujiwara	H04N 21/239	713/171
2011/0010298	A1	1/2011	Robert et al.		
2012/0144200	A1*	6/2012	Liu	H04L 9/0844	713/171
2012/0221853	A1	8/2012	Wingert et al.		
2012/0308008	A1*	12/2012	Kondareddy	H04L 63/0471	380/273
2013/0024701	A1	1/2013	Hwang et al.		
2013/0174272	A1	7/2013	Chevalier et al.		
2013/0268771	A1*	10/2013	Blankenbeckler	H04L 9/0866	713/189

2013/0325994	A1*	12/2013	Chai	H04L 67/1097	709/208
2014/0059156	A1	2/2014	Freeman, II et al.		
2014/0281576	A1*	9/2014	Suzuki	G06F 12/1408	713/189
2015/0101069	A1	4/2015	Stappenbeck et al.		
2015/0229473	A1*	8/2015	Klein	H04L 9/0841	713/171
2016/0374133	A1*	12/2016	Logue	H04W 8/005	

FOREIGN PATENT DOCUMENTS

EP	1326157	A2	7/2003
EP	2273409	A2	1/2011
JP	2003218851	A	7/2003
JP	2009507433	A	2/2009
JP	2011013714	A	1/2011
KR	20090000042	A	1/2009
WO	WO 2007028099	A2	3/2007

* cited by examiner

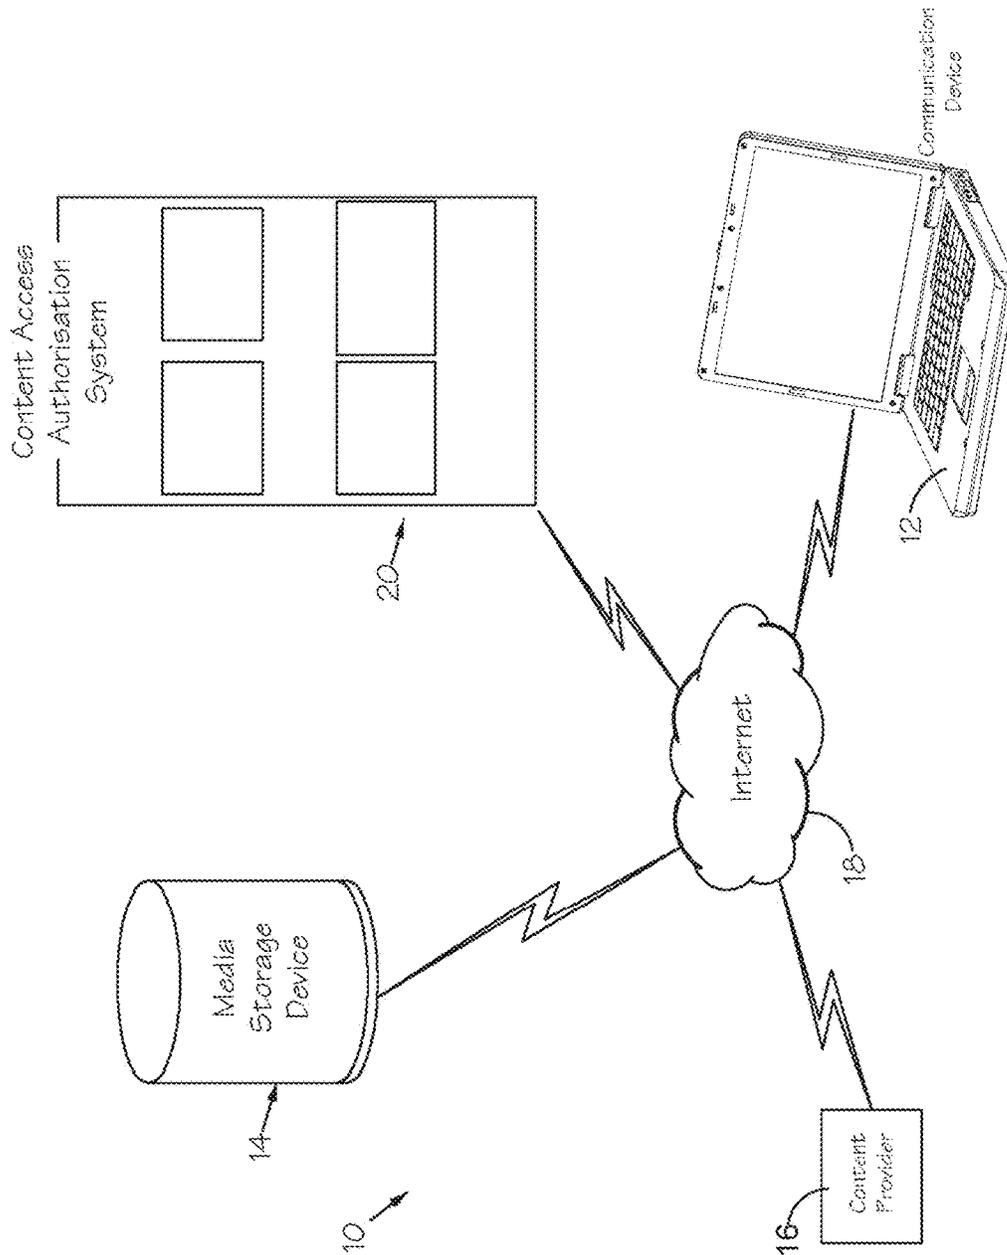


FIG 1

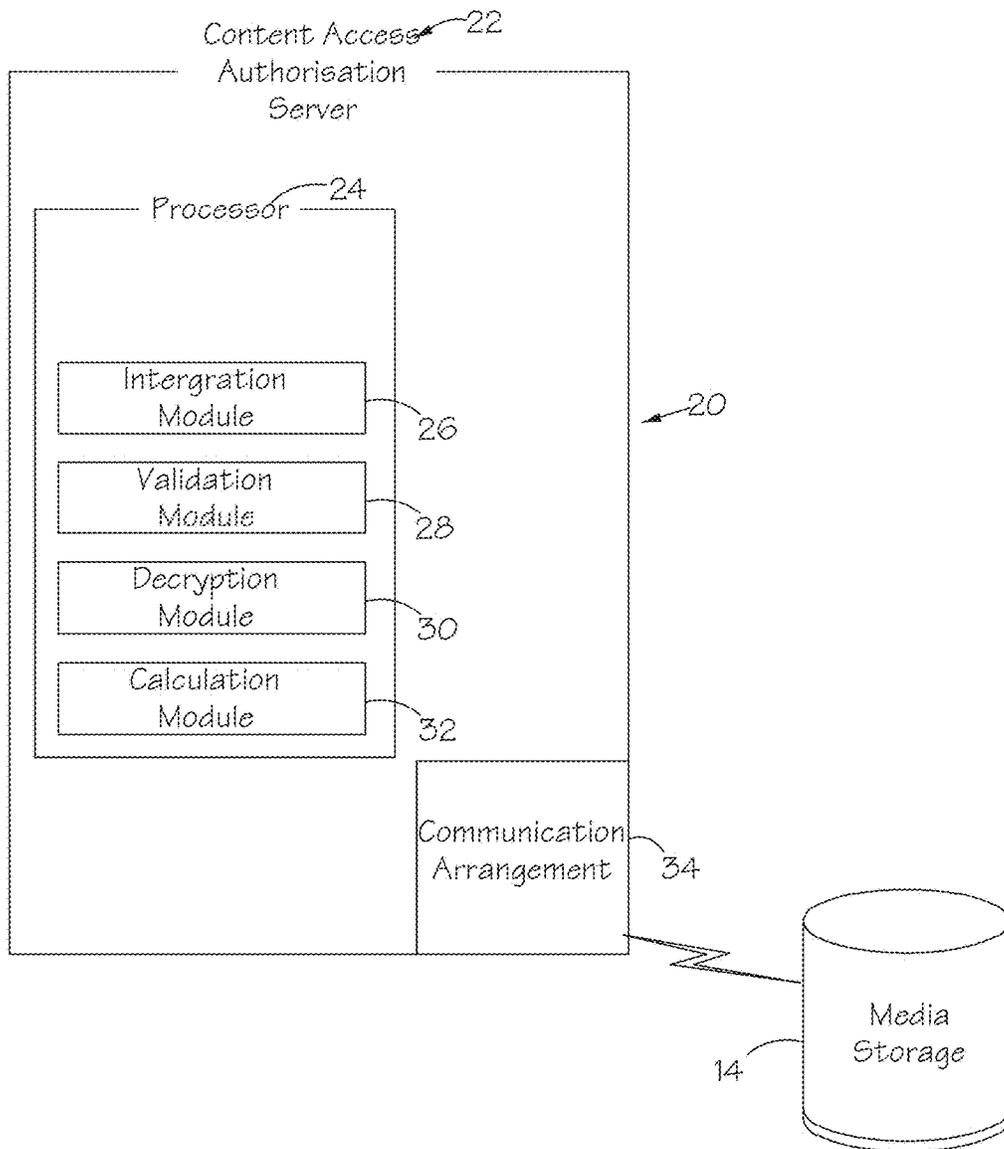


FIG 2

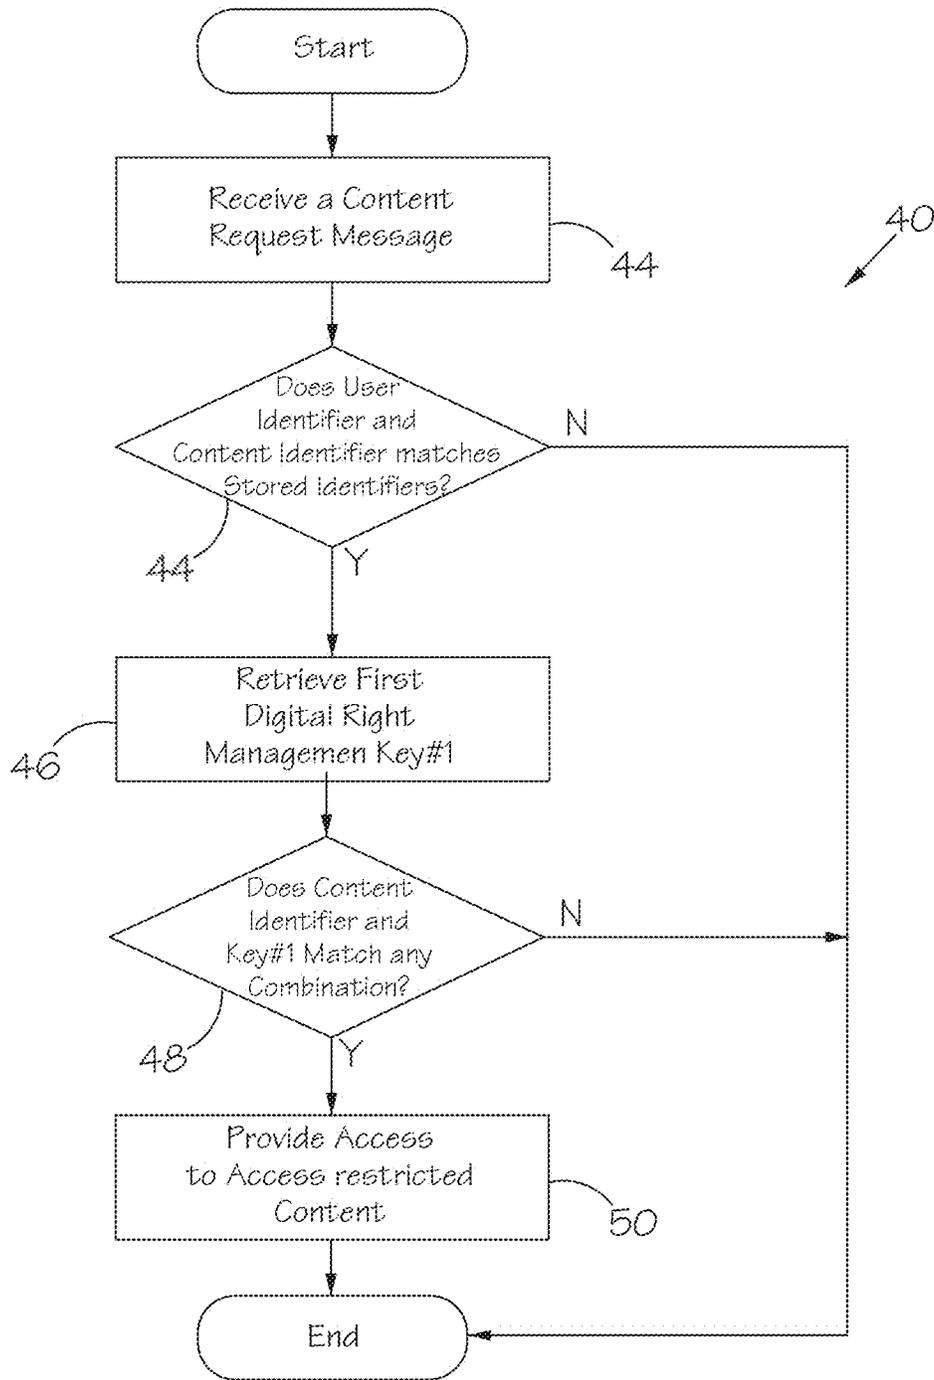


FIG 3

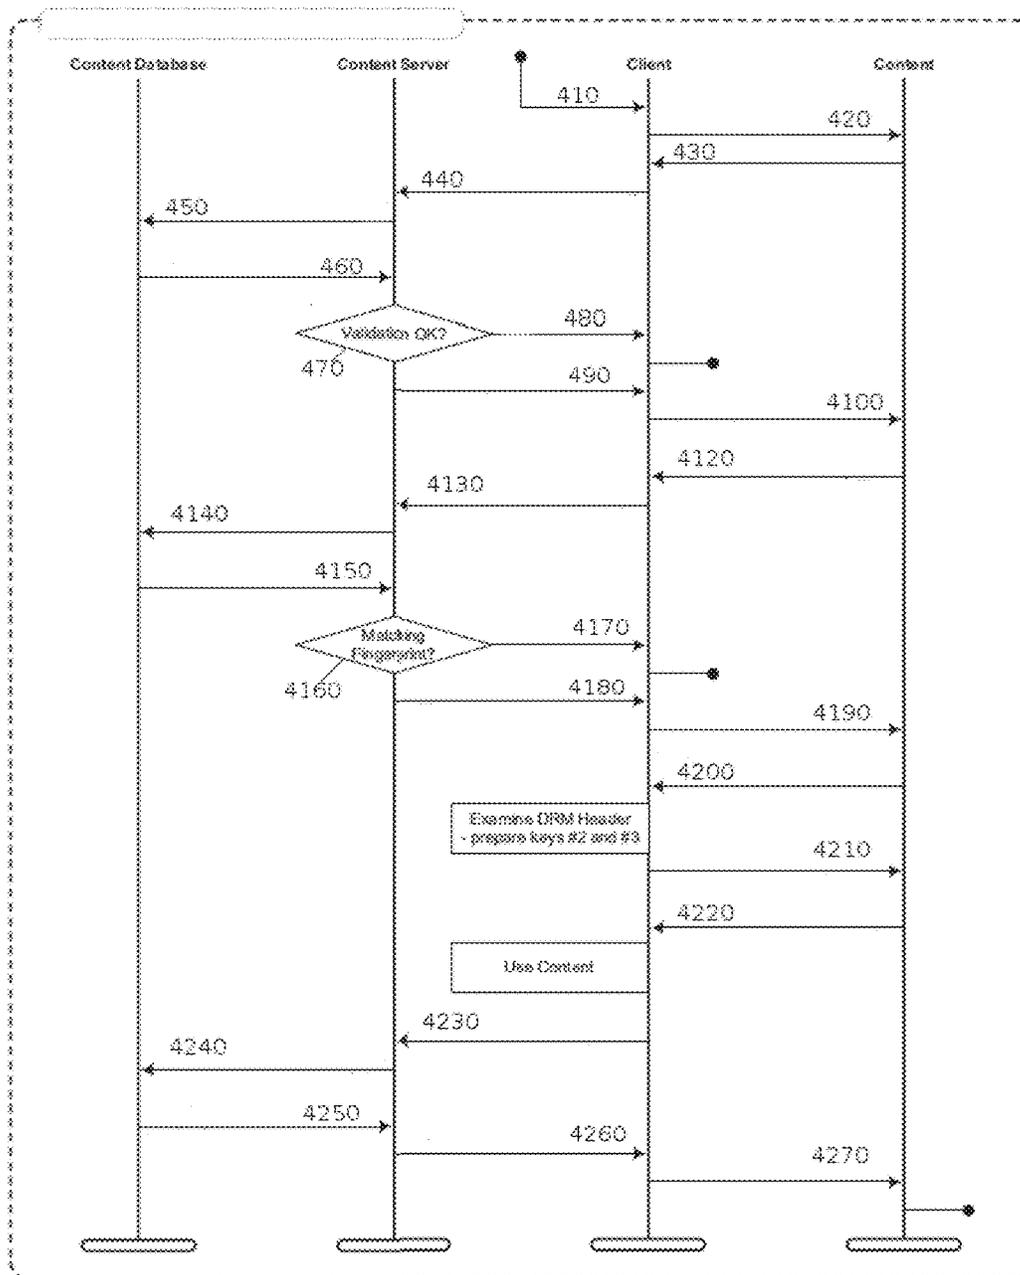


FIGURE 4

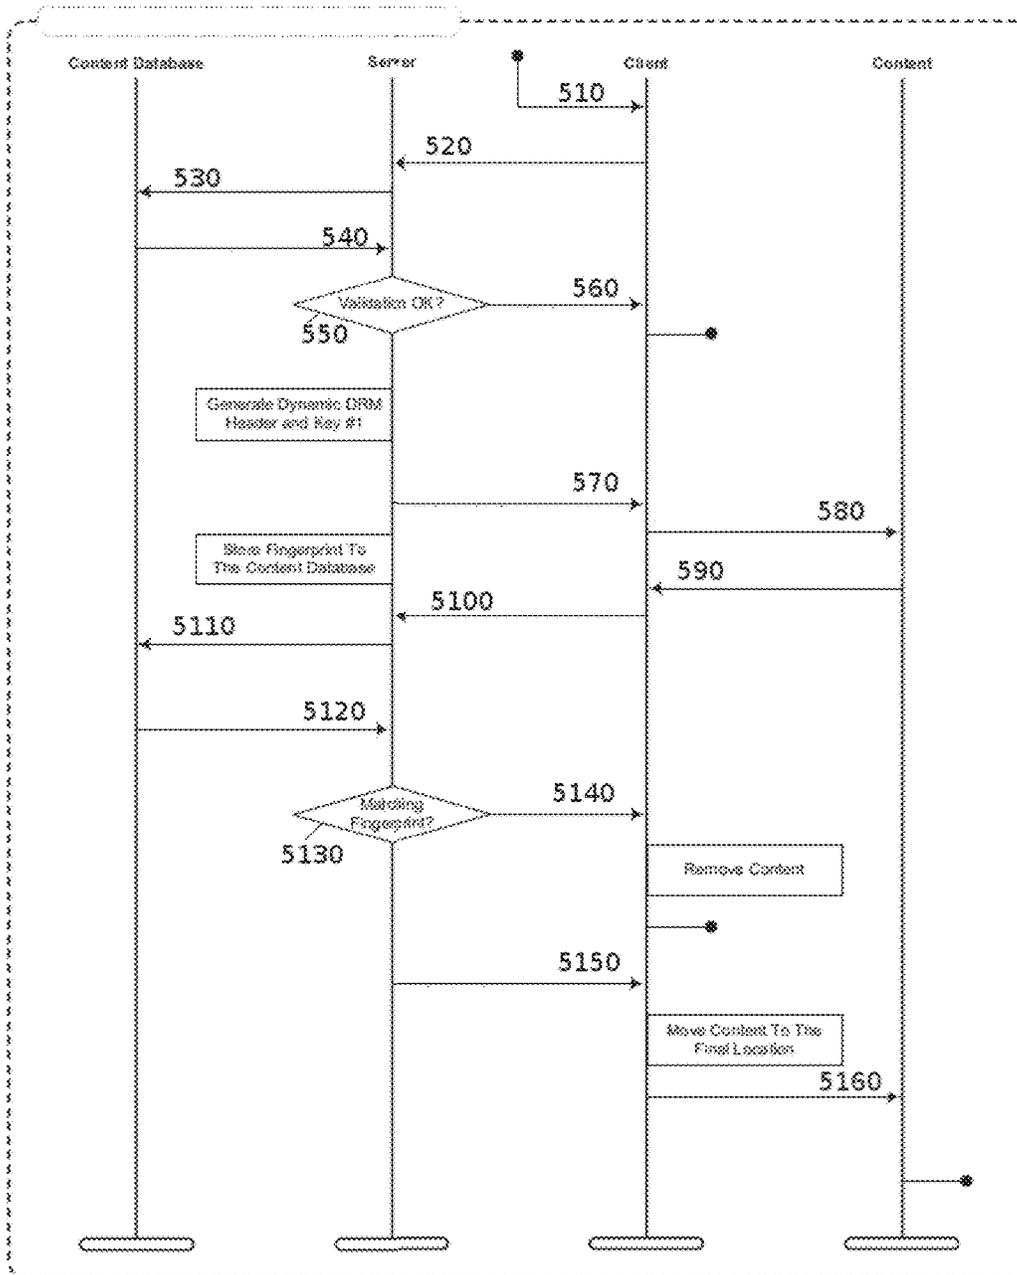


FIGURE 5

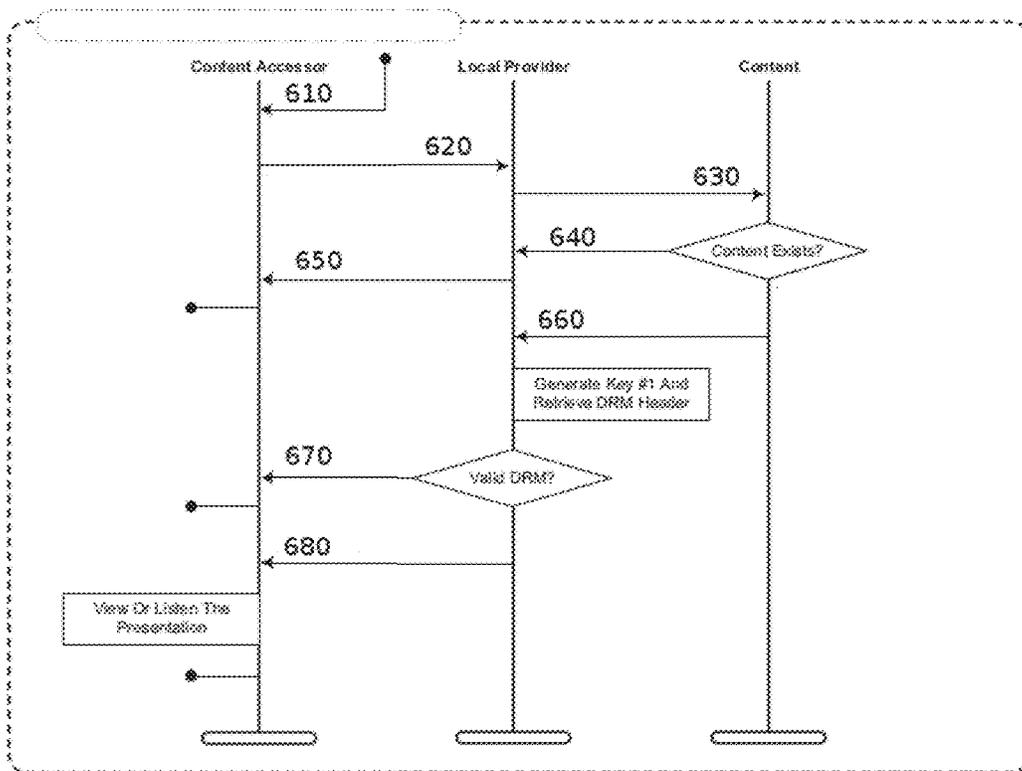
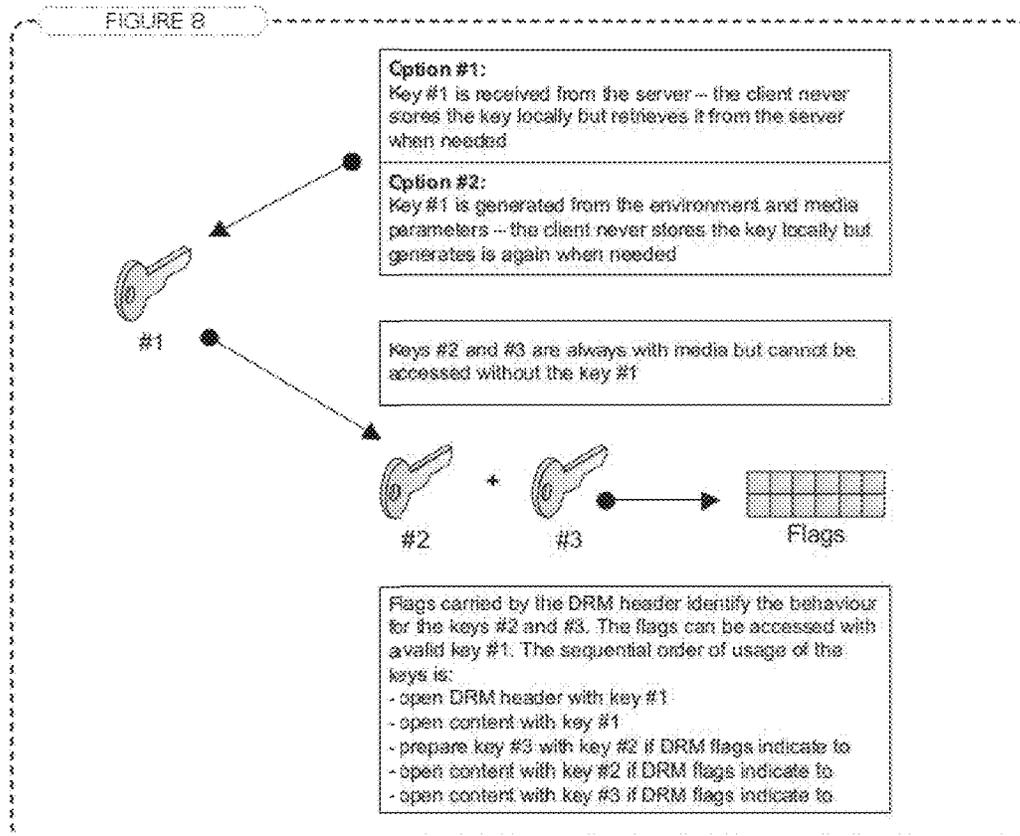
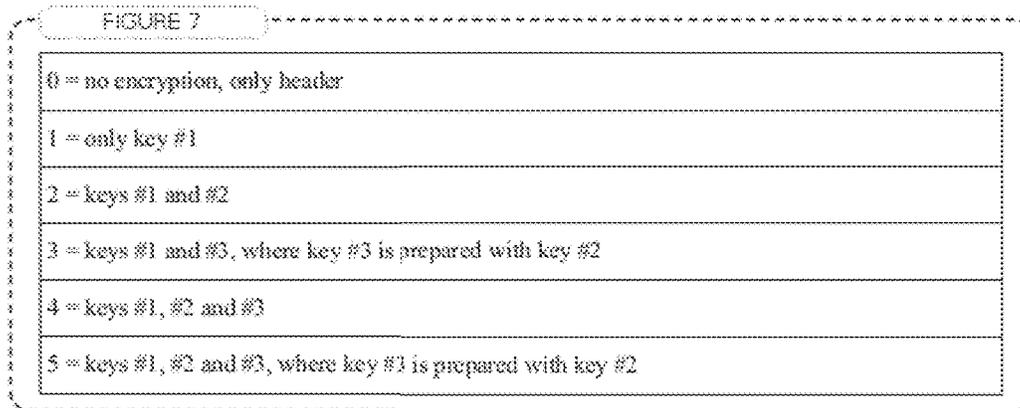
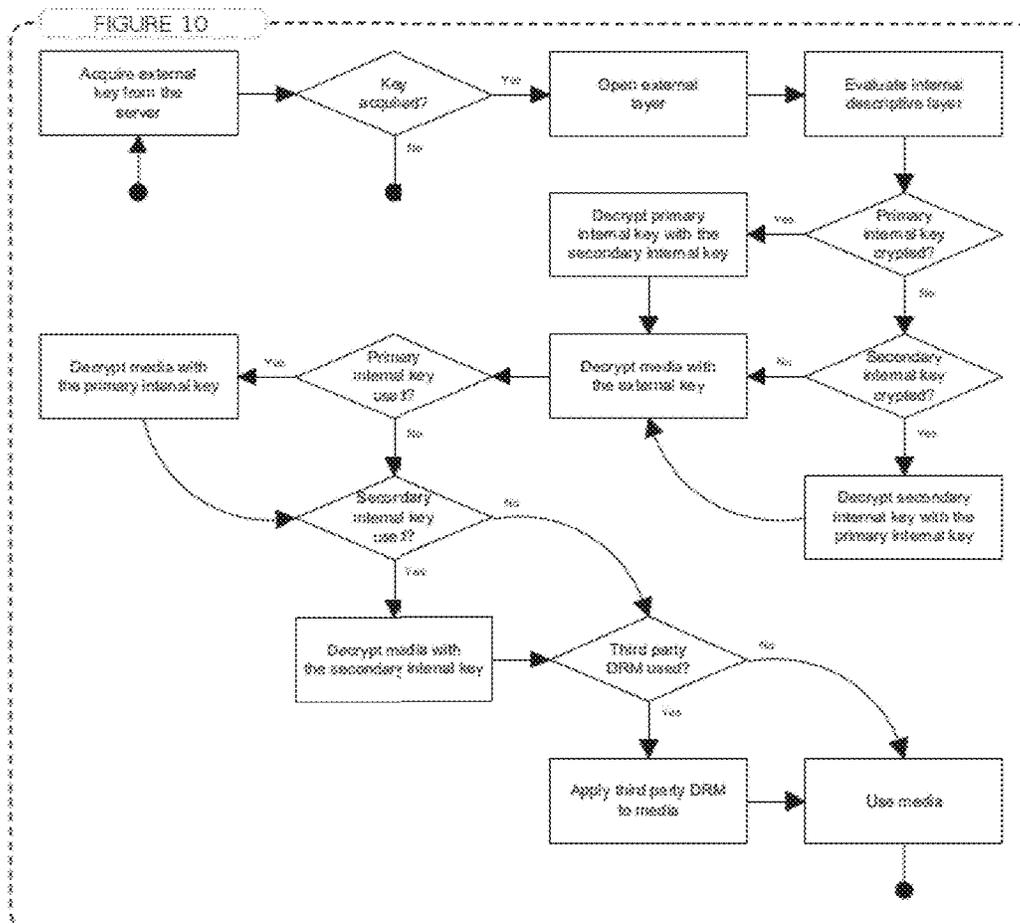
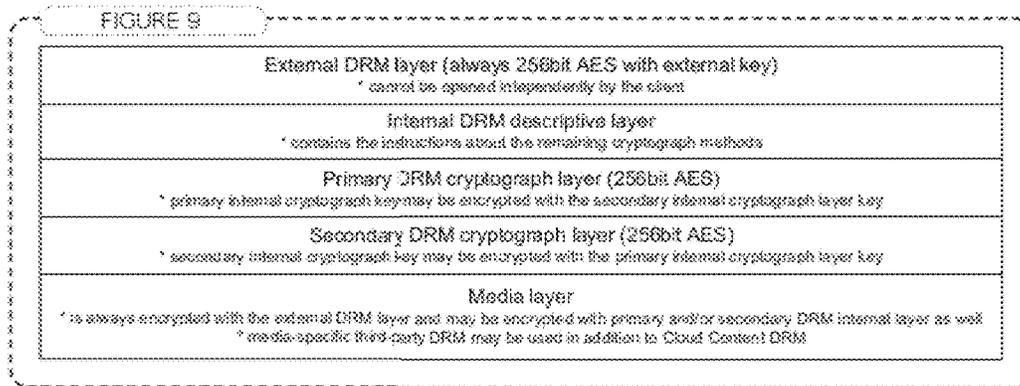


FIGURE 6





US 10,726,102 B2

1

**METHOD OF AND SYSTEM FOR
PROVIDING ACCESS TO ACCESS
RESTRICTED CONTENT TO A USER**

FIELD OF THE INVENTION

This invention relates generally to a method of using digital rights management keys to provide access to access restricted content. In particular, this invention relates to an apparatus, method and/or system for providing access to access restricted content to a user and a method thereof.

BACKGROUND TO THE INVENTION

Many publishers, copyright holders, and individuals wish to control the use of digital content and devices after sale. There are numerous ways of controlling and protecting such digital content, for example, using digital rights management methods. However, such digital rights management methods are in general not effective.

The aim of the present invention is thus to provide an alternative method of and a system for providing access to access restricted content to a user.

SUMMARY OF THE INVENTION

According to a first aspect of the invention, there is provided a system for providing access to access restricted content to a user, the system including a communication arrangement operable to receive a content request message, the content request message including a content identifier, a processor configured to cause a first determination to be performed to yield a positive or a negative result, a validation module configured to, in response to the first determination yielding a positive result, obtain a first digital rights management key, the processor being further configured to cause a second determination to be performed to yield a positive or a negative result, and responsive to the first and second determinations yielding a positive result, the validation module is configured to cause access to the access restricted content to be provided to the user.

According to a second aspect of the present invention, there is provided a method, comprising receiving a content request message, the content request message including a content identifier of an access restricted content, causing a first determination to be performed to yield a positive or a negative result, obtaining, in response to the first determination yielding a positive result, a first digital rights management key, causing a second determination to be performed to yield a positive or a negative result, and responsive to the first and second determinations yielding a positive result, causing access to the access restricted content to be provided to the user.

In a first set of embodiments of the invention in accordance with the second aspect, the method comprises causing transmission of the access restricted content to the user, wherein the access restricted content comprises in encrypted form at least one of a second and a third digital rights management key, wherein the at least one of the second and a third digital rights management key is obtainable from the access restricted content by using the first digital rights management key.

In a second set of embodiments of the invention in accordance with the second aspect, the second determination is based at least in part on a fingerprint of the access restricted content.

2

According to a first set of variants of the second set of embodiments in accordance with the second aspect, the method comprises performing the second determination by comparing a first fingerprint received in the apparatus from the user to a second fingerprint received in the apparatus from a content database.

According to a third aspect of the present invention, there is provided an apparatus, comprising at least one processor and a memory comprising program instructions, the processor, memory and program instructions configured to cause the apparatus at least to obtain an access restricted content from at least one of a content database and a content providing server, obtain a first digital rights management key, derive, using the first digital rights management key, from the access restricted content information describing encryption properties of the access restricted content, and to derive, using the information describing encryption properties of the access restricted content, from the access restricted content at least one of a content payload and a second digital rights management key.

According to a fourth aspect of the present invention, there is provided a method, comprising obtaining an access restricted content from at least one of a content database and a content providing server, obtaining a first digital rights management key, deriving, using the first digital rights management key, from the access restricted content information describing encryption properties of the access restricted content, and deriving, using the information describing encryption properties of the access restricted content, from the access restricted content at least one of a content payload and a second digital rights management key.

According to a fifth aspect of the present invention, there is provided a non-transitory computer readable medium having stored thereon a set of computer readable instructions for causing an apparatus to perform actions, the computer readable instructions comprising code for receiving a content request message, the content request message including a content identifier of an access restricted content, code for causing a first determination to be performed to yield a positive or a negative result, code for obtaining, in response to the first determination yielding a positive result, a first digital rights management key, code for causing a second determination to be performed to yield a positive or a negative result, and code for causing, responsive to the first and second determinations yielding a positive result, access to the access restricted content to be provided to the user.

According to a sixth aspect of the present invention, there is provided an apparatus, comprising at least one processor and a memory comprising program instructions, the processor, memory and program instructions configured to cause the apparatus at least to receive a content request message, the content request message including a content identifier of an access restricted content, cause a first determination to be performed to yield a positive or a negative result, obtain, in response to the first determination yielding a positive result, a first digital rights management key, cause a second determination to be performed to yield a positive or a negative result, and responsive to the first and second determinations yielding a positive result, to cause access to the access restricted content to be provided to the user.

In some embodiments of the invention in accordance with the first, second, third, fourth, fifth and/or sixth aspects the first digital rights management key is unique to a specific session.

In response to a determination that the received content identifier and user identifier matches with any combination of the stored user identifiers and content identifiers, the

US 10,726,102 B2

3

validation module may generate a first digital rights management key and a header associated with the content identifier.

The communication arrangement may be operable to receive a content usage request message, the content usage request message including the user identifier, the first digital rights management key and the content identifier.

In an embodiment, the interrogation module may further be operable to interrogate the content database in order to determine whether or not the first digital rights management key and content identifier matches with any combination of first digital rights management key and the content identifier stored in the content database; and if the answer is affirmative, the validation module may provide content access to access restricted content to the user.

The header may be associated with the particular content and the first digital rights management key are used to obtain access to the access restricted content.

The validation module may further analyze the header associated with the first digital rights management key and may prepare a second digital rights management key.

The system may include a decryption module operable to use the first digital rights management key and second and third digital rights management keys in order to decode the content, thereby allowing the user use of content.

The validation module may further analyze the second digital rights management key in order to prepare a third digital rights management key.

In this embodiment, the interrogation module may be operable to determine whether or not the second digital rights management key is used to prepare the third digital rights management key; and if the answer is affirmative, the decryption module may use the first digital rights management key and the third digital rights management key to decode the content, thereby allowing the user use of content.

If the answer is negative, the decryption module may use the first digital rights management key, the second digital management key and the third digital rights management key to decode the content, thereby allowing the user use of content.

The content may be compressed. In some embodiments, the content is media content.

In this embodiment, the system may include an extraction module being operable to extract the compressed content, thereby allowing the user use of the content.

In addition, the system may include a calculation module being operable to calculate a time period indicative of time in which the user uses the content. The validation module may stop the use of the content by the user, in response to a determination that the calculated time period is equal to a pre-defined time period.

In an alternative embodiment, the system may include a content registration module being operable to register the status content usage against the user identifier and the content identifier on the content database. The status content usage may include the calculated time period associated with the user identifier and the content identifier.

In at least some embodiments, the user identifier is an identifier associated with a communication device of the user such as a MSISDN number of the communication device.

The invention further extends to a non-transitory computer readable medium having stored thereon a set of computer readable instructions for a causing a processor to provide access to access restricted content to a user comprising the computer implemented steps of;

4

receiving a content request message, the content request message including a unique identifier associated with the user and a content identifier;

interrogating a content database of content identifiers, digital management keys and user identifiers in order to determine whether or not the received content identifier and user identifier matches with any combination of the stored user identifiers and content identifiers;

in response to a determination that there is a match, retrieving a first digital rights management key and a header associated with the user identifier;

interrogating the content database in order to determine whether or not the content identifier and the first digital rights management key matches with any combination of content identifier and first digital rights management key stored in the content database and associated with the user identifier; and

if the answer is affirmative, providing the content to the user.

In response to a determination that the received content identifier and user identifier matches with any combination of the stored user identifiers and content identifiers, the computer readable instructions may include the computer implemented step of generating a first digital rights management key and a header associated with the content identifier.

The computer readable instructions may include the computer implemented steps of;

receiving a content access request message, the content request message including the user identifier, the first digital rights management key and the content identifier;

interrogating the content database in order to determine whether or not the first digital rights management key and content identifier matches with any combination of first digital rights management key and the content identifier stored in the content database; and if the answer is affirmative, the providing content access to access restricted content to the user.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention will be described, by way of example only, with reference to the accompanying drawings in which:

FIG. 1 shows an example of network topology including a system for providing access to access restricted content to a user;

FIG. 2 shows the system of FIG. 1 in more detail;

FIG. 3 shows a flowchart representing an example method of providing access to access restricted content to a user, according to another aspect of the present invention;

FIG. 4 shows an example flow graph of a first method in accordance with at least some embodiments of the invention;

FIG. 5 shows an example flow graph of a second method in accordance with at least some embodiments of the invention;

FIG. 6 shows an example flow graph of a third method in accordance with at least some embodiments of the invention;

FIG. 7 illustrates different DRM levels;

FIG. 8 illustrates DRM key handling options;

FIG. 9 illustrates DRM layers, and

FIG. 10 illustrates an example DRM process sequence.

DETAILED DESCRIPTION

The following description of the invention is provided as an enabling teaching of the invention. Those skilled in the

US 10,726,102 B2

5

relevant art will recognize that many changes can be made to the embodiment described, while still attaining the beneficial results of the present invention. It will also be apparent that some of the desired benefits of the present invention can be attained by selecting some of the features of the present invention without utilizing other features. Accordingly, those skilled in the art will recognize that many modifications and adaptations to the present invention are possible and can even be desirable in certain circumstances, and are a part of the present invention. Thus, the following description is provided as illustrative of the principles of the present invention and not a limitation thereof.

In FIG. 1 of the drawings, reference numeral 10 refers generally to an example of network topology including a system for providing access to access restricted content to a user. Referring also to FIG. 2, an example system for providing access to access restricted content to a user is indicated by reference numeral 20.

The topology 10 includes a communication device 12 belonging to or used by a user (not shown) who intends to obtain access to access restricted content. It will be appreciated that the communication device 12 is may comprise a personal computer located at the premises of the user, a smart phone or a Personal Digital Assistant (PDA), for example. However, the communication device 12 can be a mobile telephone of the user or any other device with suitable communication capability. Communication device 12 may comprise any suitable device with communication capability, examples including tablet devices, set-top boxes, video game consoles etc.

In simple terms, the user (not shown) attempts to obtain access to access restricted content through the use of the user's communication device 12. The access restricted content may be stored in a database which is indicated as media storage device with reference numeral 14. Prior to the user obtaining access to the access restricted content, the system 20 may be configured to ensure that the user or communication device 12 of the user has particular access, i.e., the communication device 12 is authorized to obtain access to such access restricted content.

The illustrated topology includes a content provider indicated as reference numeral 16. The content provider may be a publisher or a copyright holder, or any suitable person who owns rights to the content and wishes to restrict access to such content. In other cases, the content provider may be the copyright owner, while a content distributor can be another party which has been authorized to manage access to the content on behalf of the content provider 16. The restricted content is digital content in a form of digital media. The digital media can be of any suitable form, for example, text, audio, video, graphics, animations or images.

The system 20 (further described in FIG. 2) is communicatively coupled to a communications network in the form of the Internet 18. Also communicatively coupled to the Internet is the communication device 12. The communication device 12 is thus connected to the system 18 via the Internet, for example. In use, the user may obtain access to the media content through Internet 18. The topology further illustrates a media storage device 14 which may form part of the system 20. Alternatively, the system 20 may be connected to the media storage device 14 via Internet 18 (as shown in FIG. 1).

Referring now to FIG. 2, the system 20, hereafter referred to as a content access authorization system 20, includes a content access authorization server 22 which, in turn, includes a processor 24 defining a plurality of modules 26, 28, 30 and 32 which correspond to functional tasks per-

6

formed by the processor 24. The processor 24 includes an interrogation module 26, a validation module 28, a calculation module 30 and a decryption module 32. Modules 26, 28, 30 and 32 may be comprised of software modules configured to cause processor 24 to perform corresponding functions, or the modules may comprise hardware and software elements. For example, a decryption module may comprise decryption software and hardware features of processor 24 designed to facilitate decryption of data in processor 24. Processor 24 may comprise an Intel Atom processor, for example. Server 22 may comprise memory, which is not illustrated, the memory being operable to store computer instructions processor 24 may execute to cause server 22 to perform various actions.

The content access authorization server 22 further includes a communication arrangement 34 operable to connect to the Internet 18. Communication arrangement 34 may comprise, for example, an Ethernet, fiber optic or wireless data interface. The authorization system 20 is operably connectable to a content database, that is, media storage device 14 for storing media content, content identifiers, digital management keys and/or user identifiers. The data storage device 14 can form part of the server 22 or be comprised as a standalone device external to server 22. In particular, there may be a plurality of media storage devices located within premises of various content providers and/or content distributors and communicatively coupled to the media storage device 14. The media storage device 14 stores therein a plurality of content identifiers and associated content, digital rights management keys and/or user identifiers. For example, each user may be able to have access to a particular content associated with a particular content identifier. Such user will be identified through the user identifier. In particular, the user identifier may comprise an identifier associated with the personal computer 12 of the user. In this case, the user may only be able to obtain access to the access restricted content, if the user attempts to access the access restricted content using the personal computer 12. In an embodiment where the communication device 12 is a mobile telephone, the identifier associated with the user may comprise an MSISDN number or session initiation protocol SIP identity associated with that particular mobile telephone.

Although described herein primarily with reference to an authorization server, communication device 12 may have similar structure, in particular communication device 12 may comprise a processor, memory and a communications arrangement which may each be similar to those described above.

Referring now also to FIG. 3 which shows a high-level method for providing access to access restricted content to a user, in accordance with at least one embodiment of the invention. The example method 40 is not necessarily dependent on the system 20 and/or the network topology 10, and vice versa.

In the illustrated embodiment, the communication arrangement 34 receives (at block 44) a content request message from the personal computer 12 of the user. The content request message includes a unique identifier associated with the user and a content identifier. Therefore, the user can use a user interface, e.g., a keyboard and input, e.g., serial number of the personal computer 12 and a predefined number of the content to which user wishes to obtain access. The communication arrangement can include a receiver module (not shown) operable to receive the content request message.

The interrogation module 26 interrogates (at block 44) the media storage device 14 whether or not the received content

US 10,726,102 B2

7

identifier and user identifier matches with any combination of the stored user identifiers and content identifiers. This is so as to check whether the user, for example, the communication device **12** of the user is registered to be able to have access to the media content. The media storage device **14** may have been populated during a period when the user pre-registered to have access to the access restricted content. For example, when the user purchased the media content, the user could have been requested to pre-register then. The pre-registration process can take any conventional format. For example, pre-registration may comprise that certain subscriber classes in a cellular network are granted access to certain classes of content, wherein granting access to a class of content comprises granting access to each content item comprised in the class of content.

The validation module **28** retrieves (at block **46**) a first digital rights management key (Key #1) and a header associated with the user identifier in response to a determination that the received content identifier and user identifier matches with at least one combination of the stored user identifiers and content identifiers. Key #1 is retrieved from the media storage device **14**, alternatively, Key #1 can be retrieved from a server which is located at the premises of the content distributor and/or content provider. In other embodiments, the Key #1 can be located at a server located at the user's premises. The interrogation module (at block **48**) further interrogates the media storage device **14** in order to determine whether or not the content identifier and the Key #1 matches with any combination of content identifier and Key #1 stored in the media storage device **14** and associated with the user identifier. If the answer is affirmative, the validation module (at block **50**) provides the access restricted media content to the user. Therefore, the Key #1 may be associated with the personal computer **12** of the user. In this instance, the Key #1 is used to provide the user access to access restricted media content. The Key #1 only provides access to the content i.e. the user is not able to use the content. Therefore, if the user attempts to obtain access to the access restricted content using a different personal computer, the user will not be able to obtain such access. This will prohibit users from providing the Key #1 to any other party in order for that party to access the access restricted content at another personal computer. The header is associated with the particular media content and the Key #1 and the header are used to obtain access to the access restricted content.

In an embodiment, in response to a determination that the received content identifier and user identifier matches with any combination of the stored user identifiers and content identifiers, the validation module **28** may generate (not shown) a first digital rights management key and a header associated with the content identifier. The Key #1 can be generated from the environment and media content parameters. Generating a key from environment and media content parameters may comprise, for example, using parameters relating to communication device **12**, a subscription of the user or aspects of the content in a key generation process. As a specific example, where communication device **12** comprises a cellular telephone, the key generation process may use as input information relating to a secret stored on a subscriber identity module, SIM, card. Therefore, each time when the user requires access to access restricted content, new Key #1 may be generated.

Once the user has access to the access restricted media content, the user may need to use the media content. In that instance, the communication arrangement **34** may receive (not shown) a content usage request message from the

8

personal computer **12** of the user. In simple terms, the user uses the keyboard to indicate that he/she requires usage of the restricted media content. The content usage request message includes the user identifier, the Key #1 and the content identifier. Therefore, the user will use the generated/retrieved Key #1 in order to be allowed to use the restricted media content. The user may wish to use the restricted media content by, for example, copying the media content, listening to the media content or editing (if allowed) the media content. The validation module **28** analyses the header associated with the Key #1 and prepares a second digital rights management key (Key #2). The preparation of the Key #2 may comprise, for example, performing a cryptographic operation on at least part of the access restricted content, wherein the cryptographic operation may employ Key #1. The cryptographic operation may comprise decrypting Key #2 by using Key #1. A decryption module (**30**) may be configured to use Key #1 and Key #2 to decode the media content, thereby allowing the user use of media content.

The validation module may further use Key #1 in order to prepare a third digital rights management key (Key #3). It will be appreciated that at least in some embodiments Key #2 and Key #3 cannot be prepared without Key #1. In these embodiments, the interrogation module **26** determines whether or not Key #1 is used to prepare the Key #2, and if the answer is affirmative, the decryption module **30** uses Key #1 and Key #2 to decode the media content, thereby allowing the user use of content. Alternatively, the decryption module may use Key #1, Key #2 and Key #3 to decode the media content, thereby allowing the user use of content. Key #3 may be obtained from the access restricted content using Key #1 in a similar way as described above in connection with obtaining Key #2.

In an example embodiment, the media content is compressed. In this embodiment, the system can include an extraction module (not shown). The extraction module is operable to extract the compressed media content, thereby allowing the user use of the content.

A calculation module **32** calculates (not shown) a time period indicative of time in which the user uses the content. The use of the media content may be available for a particular time period. For example, a user can be allowed to use the media for only one (1) hour. Therefore, the calculation module **32**, as the user uses the media content, can calculate the user's usage period. When the calculated time period reaches a pre-defined usage time, e.g., one (1) hour, the validation module **28** stops the use of the media content by the user.

In an alternative embodiment, the system **20** includes a content registration module (not shown). The content registration module registers the status content usage against the user identifier and the content identifier on the media storage device **14**. The status content usage includes the calculated time period associated with the user identifier and the content identifier. Therefore, it is possible to detect the rate of usage of the media content for each user, that is, the number of times the media content was accessed. The status content usage will also be able to provide an indication of the last time the media content was accessed by the user.

At least one of the first, second and third digital rights management keys may be, depending on the embodiment, arranged to be session-specific in the sense that it is generated dynamically for use in a single session. This is advantageous since if a session-specific key is compromised, it cannot be used to gain access to content in a subsequent session.

In general there is provided an apparatus, comprising at least one processor and a memory comprising program instructions. The apparatus may comprise a server, for example. The processor, memory and program instructions configured to cause the apparatus at to receive a content request message, the content request message including a content identifier of an access restricted content. The content request message may further comprise an identifier associated with the user. The apparatus may be caused to cause a first determination to be performed to yield a positive or a negative result and to obtain, in response to the first determination yielding a positive result, a first digital rights management key. The first determination may comprise a query, such as a query transmitted to a content database, the query comprising the content identifier and the identifier associated with the user. Alternatively to a query, the first determination may comprise a determination as to whether the access restricted content can be found in accordance with the content request.

The apparatus may be caused to cause a second determination to be performed to yield a positive or a negative result, and responsive to the first and second determinations yielding a positive result, to cause access to the access restricted content to be provided to the user. The second determination may be based at least in part on a fingerprint of the access restricted content. The second determination may comprise a comparison between a first fingerprint received in the apparatus from the user to a second fingerprint received in the apparatus from a content database.

The second determination may comprise a check as to whether the user has rights to access the access restricted content.

In response to the first determination yielding a positive result, for example when a content database returns a positive result to a query, the apparatus is in at least some embodiments configured to obtain a header associated with the content identifier and wherein the header associated with the content identifier and the first digital rights management key are usable to at least in part obtain access to the access restricted content. The header may be obtained by the apparatus, for example, by receiving it from the content database.

In general there is provided a second apparatus comprising at least one processor and a memory comprising program instructions, the processor, memory and program instructions configured to cause the apparatus at least to obtain an access restricted content from at least one of a content database and a content providing server. The second apparatus may be configured to obtain the access restricted content by receiving it over a cellular or Ethernet connection, for example. The second apparatus may be configured to store the access restricted content, at least in part, in a memory comprised in the second apparatus.

The second apparatus may be configured to obtain a first digital rights management key, to derive, using the first digital rights management key, from the access restricted content information describing encryption properties of the access restricted content and derive, using the information describing encryption properties of the access restricted content, from the access restricted content at least one of a content payload and a second digital rights management key. Where the second apparatus is caused to derive a second digital rights management key, it may be further configured to use the first and second digital rights management keys to obtain access to the content payload. The content payload may comprise, for example, a media file such as an audio or video recording.

FIG. 4 shows an example flow graph of a first method in accordance with at least some embodiments of the invention. On the vertical axes are, from left or right, a content database, a content server, a client and content. The illustrated method begins in phase 410 and proceeds to phase 420, where the client performs a content location, for example by searching with at least one keyword. As a response, in phase 430 the client receives a content identifier. In phase 440, the client transmits a message to the content server, the message comprising the content identifier received in phase 430 and an identifier of the client. The message of phase 440 may comprise a content request message. Responsive to receiving the message of phase 440, the content server may transmit, in phase 450, a query to the content database, the query comprising the content identifier and client identifier.

Responsive to the query, the content database may reply, in phase 460, to the content server with a message comprising a validation result wherein the validation result may comprise a first DRM key. The content server determines whether the validation was successful in phase 470. In case the validation was unsuccessful, for example, where the client does not have access to the content, processing advances from phase 470 to phase 480 and ends. On the other hand where the validation was successful and the message of phase 460 comprises a first DRM key, processing advances from phase 470 to phase 490 where the content server transmits the first DRM key to the client.

Responsive to receipt of the first DRM key in phase 490, the client may access the content using at least in part the first DRM key. This is illustrated as phase 4100. The client may obtain, in phase 4120, a fingerprint of the content wherein the obtaining may be based at least in part on the first DRM key. In phase 4130, the client may transmit the fingerprint to the content server, optionally with the content identifier and in phase 4140, the content server may query the content database for the content fingerprint. The query may comprise the content identifier. In phase 4150, the content database may responsively provide the fingerprint to the content server. In phase 4160, the server may compare the fingerprints received in phases 4150 and 4130. In case of mismatch, the processing advances to phase 4170 and ends. In case the fingerprints match, processing advances to phase 4180 where the client is provided with a positive validation result.

Responsive to the positive validation result of phase 4180, the client in phase 4190 proceeds to access the content to retrieve a DRM header, and optionally also to apply the first DRM key to the header, responsive to which the client gains access, phase 4200, to an open DRM header of the content. Using the header the client may be enabled to prepare second and third DRM keys, and, optionally, to apply at least one of the second and third DRM keys to retrieve payload of the content. This retrieval is illustrated as phases 4210 and 4220.

After using the payload of the content, the client may inform the content server of this, phase 4230, and the server may inform the content database of this, in phase 4240. In phase 4250, the content database may inform the content server of a registration of the content, the message of phase 4250 optionally comprising a result code. The content server may notify the client of this, phase 4260 and the client may modify the content accordingly, phase 4270.

When the first method is used, at least one of the following may apply:

US 10,726,102 B2

11

1) the client has an on-line connection with the server, and 2) the client has acquired the content before the validation starts.

FIG. 5 shows an example flow graph of a second method in accordance with at least some embodiments of the invention. The vertical axes are identical to those in FIG. 4. The process of FIG. 5 begins in phase 510 and proceeds to phase 520, where the client requests content from the server. The message of phase 520 may comprise a content identifier and client identifier. The server may, phase 530, query the content database with the content identifier and client identifier. The content database replies in phase 540 whether the client should be granted access to the content.

In phase 550 the server determines whether the validation was successful, in other words whether the client is to be granted access to the content. In case no, processing advances to phase 560 and the process ends. In case the answer is yes, the server generates a DRM header, for example a dynamic DRM header, and a first DRM key. In phase 570 the server encrypts the content and streams it to the client, as well as calculates a content fingerprint.

In phase 580 the client saves the received content, optionally in a temporary location and in phase 590 the client obtains from the content a content fingerprint, which is sent along with the content identifier to the server in phase 5100.

In phase 5110 the server queries the content database for a content fingerprint of the content, the query comprising the content identifier and the client identifier. The server receives the fingerprint from the database in phase 5120. In phase 5130 the server determines whether the fingerprints received in phases 5120 and 5100 are the same. In case they are not, the server informs the client of this in phase 5140, and responsively the client removes the content it stored in phase 580. In case the fingerprints match, phase 5150, the client is informed of this. In embodiments where the content was stored in a temporary location in phase 580, the client in phase 5160 moves the content to a final location. In embodiments where the content was stored in a non-temporary location in phase 580, the client may validate the content in phase 5160.

When the second method is used, at least one of the following may apply:

1) the client has an on-line connection to the server for receiving the streamed content, and 2) the client has been registered to the server before the streaming is started.

FIG. 6 shows an example flow graph of a third method in accordance with at least some embodiments of the invention. On the vertical axes are illustrated, from left to right, a content accessor, who may be a client, a local provider and the content being accessed. The process begins in phase 610 and proceeds to phase 620, where the content accessor requests content from a local provider with a content identifier comprised in the request message. In phase 630, the local provider locates the content using the content identifier. In case the content is determined to not exist, processing advances to phase 640 and 650, where the local provider informs the accessor of this and the process ends. On the other hand if the content is determined to exist, processing advances to phase 660. The local provider generates a first DRM key and obtains a DRM header. The DRM is verified, and in case the verification fails processing advances to phase 670 where the content accessor is informed of this and processing ends. On the other hand if DRM verification succeeds, processing advances to phase 680 and the content accessor is granted access to the content, which may comprise, for example, an audio, video, or audiovisual presentation.

12

When the third method is used, at least one of the following may apply:

1) the content has been delivered to the content accessor from the server or from an external supplier. 2) the client is able to use the content without leaving it in an unprotected state. 3) the content is provided in a format that cannot be stored for later access, and 4) the content DRM is accessed only when the content needs to be shown or played.

FIG. 7 illustrates different DRM levels. DRM levels may in this example be numbered from zero to five, with each level providing for different encryption and key management functionalities. Which level is used for a specific content item, may be determined from a DRM header of the content, for example. The keys described in the Figure may comprise DRM keys, for example.

FIG. 8 illustrates DRM key handling options.

FIG. 9 illustrates DRM layers in an example embodiment of the invention. The illustrated layers may be applied in sequence to the media layer, which may comprise a payload of the content item. To obtain access to the payload, a user may need to apply decryption operations in reverse order to the order in which encryption operations were performed. In different embodiments, at least one of the encryption layers illustrated in FIG. 9 may be omitted.

FIG. 10 illustrates an example DRM process sequence in an example embodiment of the invention. In the illustrated example, a set of DRM encryption layers such as the one illustrated in FIG. 9 is opened in phases to ultimately allow access to the media, that is the content payload.

It should be noted that in the above, a "server" may take different forms depending on the embodiment. In particular, alternatively to a fixed computer residing in a network, in some embodiments a server may comprise a peer device to the client device, such as for example where the client device is a tablet or smartphone device, the server may also be a tablet or smartphone device.

The invention claimed is:

1. An apparatus, comprising at least one processor and a memory comprising program instructions, the processor, memory and program instructions configured to cause the apparatus at least to:

- obtain an access restricted content from at least one of a content database and a content providing server;
- obtain a first digital rights management key from the content database, wherein the obtaining is based at least in part on a query, the query comprising the content identifier and an identifier associated with the user;
- using the first digital rights management key, obtain a fingerprint of the access restricted content wherein the obtaining is based at least in part on the first digital rights management key,
- cause the content providing server to validate the fingerprint, and, if the validation is successful, access the access restricted content and
- derive a second and third digital rights management key from the access restricted content using the digital rights management header of the access restricted content

wherein the second and third digital rights management keys are applied to retrieve the payload of the access restricted content and wherein at least one of the second or third digital rights management key is used to encrypt the other key of the second or third digital rights management key wherein the content is usable without being in an unprotected state.

2. The apparatus according to claim 1, wherein the apparatus is caused to obtain the first digital rights manage-

US 10,726,102 B2

13

ment key either by receiving it from a server or by generating it at least in part from at least one of environmental parameters and parameters of the access restricted content.

3. The apparatus of claim 1, wherein the first digital rights management key is unique to a specific session.

4. The apparatus of claim 1, wherein the access restricted content comprises a first layer that is decryptable using the first digital rights management key, the first layer comprising information on which key may be used to obtain fuller access to the access restricted content.

5. The apparatus of claim 1, wherein the first, second and third digital rights management keys are each encrypted with 256 bit AES encryption.

6. The apparatus of claim 1, wherein at least one of the content database and the content providing server comprise a content registration module, wherein said content registration module registers the usage of the content on the content database with respect to the user identifier and the content identifier.

7. The apparatus of claim 1, wherein said status content usage includes the calculated time period associated with the user identifier and the content identifier.

8. The apparatus of claim 1, wherein the apparatus retrieves the first digital rights management key from at least one of the content database and the content providing server when required.

9. The apparatus of claim 1, wherein the apparatus generates the first digital rights management key when required, the generation being based on environment and media parameters.

14

10. A method, comprising:

obtaining an access restricted content from at least one of a content database and a content providing server;

obtaining a first digital rights management key from a content database, wherein the obtaining is based at least in part on a query, the query comprising the content identifier and an identifier associated with the user;

deriving, using the first digital rights management key, from the access restricted content a fingerprint of the access restricted content wherein the obtaining is based at least in part on the first digital rights management key,

causing the content providing server to validate the fingerprint, and, if the validation is successful, accessing the access restricted content and

information describing encryption properties of the access restricted content, and

deriving, using the digital rights management header of the access restricted content, from the access restricted content a second and third digital rights management key,

wherein the second and third digital rights management keys are applied to retrieve the payload of the access restricted content and wherein at least one of the second or third digital rights management key is used to encrypt the other key of the second or third digital rights management key, wherein the content is usable without being in an unprotected state.

11. The method of claim 10, wherein the first digital rights management key is unique to a specific session.

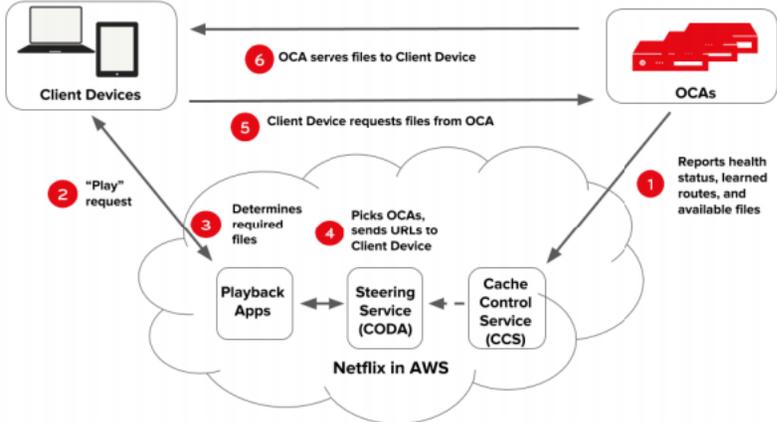
* * * * *

Exhibit C

CONFIDENTIAL ATTORNEY-CLIENT AND WORK PRODUCT PRIVILEGED

Claim Chart for U.S. Patent 8,495,167 - “Data communications networks, systems, methods and apparatus”

US Patent 8,495,167
 Filing Date: Jul. 30, 2002
 Priority Date: Jul 30, 2002

Claim Portion	'167 Patent	Netflix
Claim 1		
[1a]	A data communication network comprising: a plurality of terminals; and	<p>Netflix uses a system called Open Connect to deliver Netflix TV shows and movies to members world-wide.</p> <p>The building blocks of Open Connect are our suite of purpose-built server appliances, called Open Connect Appliances (OCAs). <i>See</i> Open Connect Overview, p. 2. These are deployed directly inside ISP networks. Netflix provides the server hardware.</p>  <p>The diagram illustrates the Open Connect system architecture. It shows Client Devices (laptops and a smartphone) on the left and OCA servers (red server racks) on the right. The central cloud represents 'Netflix in AWS' and contains three main components: Playback Apps, Steering Service (CODA), and Cache Control Service (CCS). The process is numbered 1 through 6:</p> <ul style="list-style-type: none"> 1: OCA servers report health status, learned routes, and available files to the Cache Control Service (CCS). 2: Client Devices send a "Play" request to the Playback Apps. 3: Playback Apps determine required files and send this information to the Steering Service (CODA). 4: The Steering Service (CODA) picks OCAs and sends URLs to the Client Device. 5: Client Devices request files from the OCA servers. 6: OCA servers serve files to the Client Device.

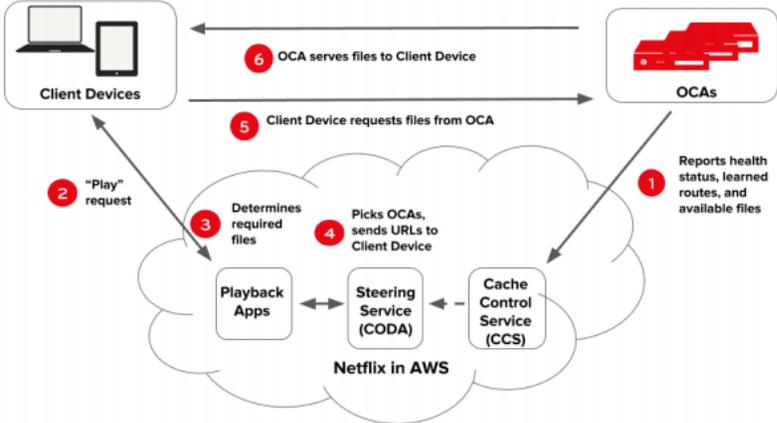
CONFIDENTIAL ATTORNEY-CLIENT AND WORK PRODUCT PRIVILEGED

		<p>Open Connect Appliances can be embedded in your ISP network. Embedded OCAs have the same capabilities as the OCAs that we use in our 60+ global data centers, and they are provided to qualifying ISP partners at no charge. Each embedded OCA deployment will offload a substantial amount of Netflix content traffic from peering or transport circuits. Multiple physical deployments can be distributed or clustered on a geographic or network basis to maximize local offload.</p> <p>Source: https://openconnect.netflix.com/en/#sample-architectures</p>
[1b]	<p>a main server adapted to manage selective retrieval of data from a first server by at least one target terminal selected from said plurality of terminals, said main server being distinct from said first server; and</p>	<p>Netflix runs the operation of Open Connect from an application in AWS (main server). When a Netflix user requests playback of a TV show or movie, the AWS application selectively retrieves data from the user device to send content through an internet service provider (ISP). <i>See</i> Open Connect Overview, p. 4-5.</p>

CONFIDENTIAL ATTORNEY-CLIENT AND WORK PRODUCT PRIVILEGED

		<ol style="list-style-type: none"> 1. OCA's periodically report health, routes they have learned, and content (file) availability to the cache control services in AWS. 2. A user on a client device requests playback of a title (TV show or movie) from the Netflix application in AWS. 3. The playback application services in AWS check user authorization and licensing, then determine which specific files are required to handle the playback request - taking individual client characteristics and current network conditions into account. <hr style="border: 2px solid black; margin: 10px 0;"/> <ol style="list-style-type: none"> 4. The steering service in AWS uses the information stored by the cache control service to pick OCAs that the requested files should be served from, generates URLs for these OCAs, and hands the URLs over to the playback application services. 5. The playback application services hand over URLs of the appropriate OCAs to the client device, and the OCA begins to serve the requested files.
[1c]	a network information database containing terminal performance information, wherein	<p>We constantly measure and analyze [OCA] performance and augment capacity as requirements evolve. <i>See</i> Open Connect Overview, p. 3.</p> <p>All OCA deployments are constantly monitored by the Open Connect Operations team to ensure reliability and efficiency. We troubleshoot and proactively fix most issues remotely with minimal input required from our ISP partners. <i>See</i> Open Connect Overview, p. 5.</p>

CONFIDENTIAL ATTORNEY-CLIENT AND WORK PRODUCT PRIVILEGED

		<p>Additionally, OCAs periodically report health. <i>Id.</i> at 4.</p> <p>Monitoring, Maintenance, and Updates</p> <p>All of our OCA deployments, whether in IXPs or embedded in ISP networks, are constantly monitored by the Open Connect Operations team to ensure reliability and efficiency. We troubleshoot and proactively fix most issues remotely with minimal input required from our ISP partners. If partners wish to monitor their own embedded OCAs' status and performance, we provide a Partner Portal where they can do so. If hardware performance degrades to the point where a server is no longer functioning in the range of our quality standards, we simply replace it - at no cost to our partners.</p>
<p>[1d]</p>	<p>at least two of said terminals are adapted to act as relay servers for serving data retrieved from said first server to at least one target terminal; and wherein</p>	<p>URLs of appropriate OCAs (relay servers) are given to the client device. The OCAs serve the requested files (Netflix content delivered through ISP) to the client device.</p>  <p>The diagram illustrates the Netflix in AWS architecture. It shows Client Devices (laptops and a smartphone) interacting with OCAs (Open Connect Appliances, represented by server racks). The process is numbered 1 through 6:</p> <ol style="list-style-type: none"> OCAs report health status, learned routes, and available files to the system. Client Devices send a "Play" request to Playback Apps. Playback Apps determine required files and send them to the Steering Service (CODA). Steering Service (CODA) picks OCAs and sends URLs to Client Devices. Client Devices request files from OCAs. OCAs serve files to Client Devices. <p>The central cloud is labeled "Netflix in AWS" and contains Playback Apps, Steering Service (CODA), and Cache Control Service (CCS).</p>

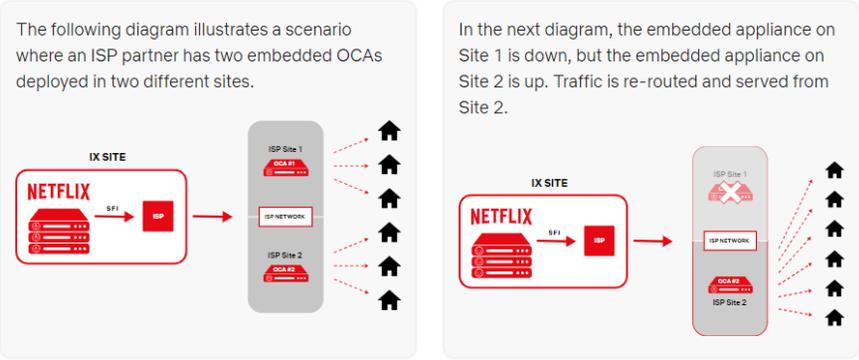
CONFIDENTIAL ATTORNEY-CLIENT AND WORK PRODUCT PRIVILEGED

[1e]	<p>the main server is adapted to send transport requests direct to at least one first target terminal on the basis of said terminal performance information, and wherein the main server is further adapted to monitor response times of terminals in the network and in which terminals are selected to act as relay servers for a particular data transfers on the basis of their relative response times, and the first target terminal is adapted to act as relay server; and</p>	<p>All OCA deployments are constantly monitored to ensure reliability and efficiency. Netflix makes use of non-peak bandwidth to download the vast majority of content updates to the OCAs in network during these configurable time windows. OCAs can also download updates from each other – minimizing significant usage of internet “backbone” capacity during the update cycle.</p> <p>Netflix also employs multiple failover scenarios where data transfer changes depending on response time of terminals. Further, Netflix can serve traffic directly to a user device.</p> <p>Source: https://openconnect.netflix.com/en/#sample-architectures</p>

CONFIDENTIAL ATTORNEY-CLIENT AND WORK PRODUCT PRIVILEGED

		Failover scenarios for embedded deployments
		<p>The following diagram illustrates a scenario where an ISP partner has two embedded OCAs deployed in two different sites.</p> <p>In the next diagram, the embedded appliance on Site 1 is down, but the embedded appliance on Site 2 is up. Traffic is re-routed and served from Site 2.</p> <p>In the final diagram, the embedded appliances on both Site 1 and Site 2 are down. Traffic is re-routed and served from Netflix IX appliances.</p>
[1f]	wherein each such transport request includes details of data to be retrieved, the address of the first server from which the data is to be requested by the first target terminal, the addresses of at least one second target terminal to which the data from the first server to be relayed by the first target terminal and an indication of a relative performance	Netflix sends URLs corresponding to the server from which data is requested by the user device along with other data such as specific files, client characteristics, and network conditions. In setups with multiple OCAs embedded in one ISP partner, data can be sent from one OCA to another in the event of an update or when traffic needs to be re-routed and served from one of the two OCAs because of performance of the network such as when an OCA goes down.

CONFIDENTIAL ATTORNEY-CLIENT AND WORK PRODUCT PRIVILEGED

	<p>of a further target terminal based on the terminal performance information stored in the network information database; and</p>	<p>Failover scenarios for embedded deployments</p> <p>The following diagram illustrates a scenario where an ISP partner has two embedded OCAs deployed in two different sites.</p>  <p>In the next diagram, the embedded appliance on Site 1 is down, but the embedded appliance on Site 2 is up. Traffic is re-routed and served from Site 2.</p>
<p>[1g]</p>	<p>wherein terminals adapted to act as relay servers are adapted to modify transport requests received from the main server or from other relay servers and to transmit the modified transport request to selected target terminals from a set of target terminals identified in the transport request, wherein the modified transport request further includes addresses of further target terminals for which the recipient of the modified transport request is to act as relay server; and</p>	<p>OCAs can also download updates from each other – minimizing significant usage of internet “backbone” capacity during the update cycle. The OCAs work in a network to distribute updates among each other and to include further OCAs to which updates and content can be sent. <i>See Open Connect Overview, p. 5.</i></p>

CONFIDENTIAL ATTORNEY-CLIENT AND WORK PRODUCT PRIVILEGED

[1h]	wherein data to be retrieved by said target terminals are divided into a series of packets for transmission to said target terminals and each of said terminals is adapted to communicate directly with said main server to acknowledge receipt of the last packet of a series routed thereto.	<h2>“What is TCP?</h2> <p>TCP is the de-facto transport protocol on the Internet today and one of the core protocols of the Internet Protocol (IP) suite. It guarantees in-order, error-checked, delivery of all content sent from one network device to another. TCP employs retransmissions to ensure that no portion of the content is lost. To that end, TCP breaks content into packets. Each packet has a sequence number that identifies its relative ordering. The sender transmits packets to the receiver and expects acknowledgements for in-order, correctly received, packets. If any packet is detected as lost, it is retransmitted. TCP is responsible for re-arranging the received packets and delivering the content to the application without errors or data gaps.</p> <p>...</p> <h2>Where does Netflix use TCP?</h2> <p>Netflix uses TCP for internet streaming to send packets of data for video. Additionally, Netflix specifically looks at the number of TCP connections to determine internet speeds in accordance with testing of OCAs.</p> <p><i>On information and belief, Netflix’s terminals also communicate with the Main server to acknowledge receipt of the data packets.</i></p>
------	--	---

CONFIDENTIAL ATTORNEY-CLIENT AND WORK PRODUCT PRIVILEGED

Claim Chart for U.S. Patent 10,726,102 - “Method of and System for Providing Access to Access Restricted Content to a User”

US Patent 10,726,102

Filing Date: Jan. 8, 2015

Priority Date: Jan 8, 2014

Claim Portion	‘102 Patent	Netflix
Claim 10		
[1a]	<p>A method, comprising:</p> <p>obtaining an access restricted content from at least one of a content database and a content providing server;</p>	<p>Netflix provides movie, television shows, and other media (content) available through the Netflix website on computers or any internet-connected device that offers the Netflix app. These titles are streamed, or downloaded, from Netflix servers hosting the titles (content database and/or content providing server).</p> <p>Before a Netflix user has logged in with a Netflix account, the user is unable to access any restricted content that is available for viewing.</p>

CONFIDENTIAL ATTORNEY-CLIENT AND WORK PRODUCT PRIVILEGED

		<div data-bbox="840 191 1785 695"> <p>Where can I watch? ✕</p> <p>Watch anywhere, anytime. Sign in with your Netflix account to watch instantly on the web at netflix.com from your personal computer or on any internet-connected device that offers the Netflix app, including smart TVs, smartphones, tablets, streaming media players and game consoles.</p> <p>You can also download your favorite shows with the iOS, Android, or Windows 10 app. Use downloads to watch while you're on the go and without an internet connection. Take Netflix with you anywhere.</p> </div> <div data-bbox="840 727 1785 1117"> <p>Source: https://www.netflix.com/</p> </div>
[1b]	obtaining a first digital rights management key from a content	On information and belief, in order for the content that is restricted to Netflix users to become available for streaming and/or download, Netflix must provide a digital rights management key for the purposes of licensing so that the protected

CONFIDENTIAL ATTORNEY-CLIENT AND WORK PRODUCT PRIVILEGED

	<p>database, wherein the obtaining is based at least in part on a query,</p>	<p>content is accessible. The digital rights management key must be based, at least in part, on a query of the Netflix user’s identity (email address), the content, and whether the content has been paid for (in order to grant a temporary license).</p>																																
<p>[1c]</p>	<p>the query comprising the content identifier and an identifier associated with the user;</p>	<p>For Netflix to provide a temporary license to a title that is within Netflix’s database, the Netflix user must (1) be logged in to the Netflix account (identifier associated with the user) and (2) have paid for access whether through a monthly payment or otherwise. Thus, Netflix must query an identifier of the Netflix user to check that the Netflix user is logged in and must also query at least an identifier of the content to determine whether the Netflix user has access to the content (either because of location based restrictions, multiple users on an account, or insufficient payment).</p> <p>Plans and Pricing</p> <p>Netflix offers a variety of plans to meet your needs. The plan you choose will determine the video quality and the number of screens you can watch Netflix on at the same time.</p> <p>With all of our plans, you can watch unlimited TV shows and movies, and play mobile games.</p> <table border="1" data-bbox="852 850 1648 1159"> <thead> <tr> <th></th> <th>Basic</th> <th>Standard</th> <th>Premium</th> </tr> </thead> <tbody> <tr> <td>Monthly cost* (United States Dollar)</td> <td>\$8.99</td> <td>\$13.99</td> <td>\$17.99</td> </tr> <tr> <td>Number of screens you can watch on at the same time</td> <td>1</td> <td>2</td> <td>4</td> </tr> <tr> <td>Number of phones or tablets you can have downloads on</td> <td>1</td> <td>2</td> <td>4</td> </tr> <tr> <td>Unlimited movies, TV shows and mobile games</td> <td>✓</td> <td>✓</td> <td>✓</td> </tr> <tr> <td>Watch on your laptop, TV, phone and tablet</td> <td>✓</td> <td>✓</td> <td>✓</td> </tr> <tr> <td>HD available</td> <td></td> <td>✓</td> <td>✓</td> </tr> <tr> <td>Ultra HD available</td> <td></td> <td></td> <td>✓</td> </tr> </tbody> </table>		Basic	Standard	Premium	Monthly cost* (United States Dollar)	\$8.99	\$13.99	\$17.99	Number of screens you can watch on at the same time	1	2	4	Number of phones or tablets you can have downloads on	1	2	4	Unlimited movies, TV shows and mobile games	✓	✓	✓	Watch on your laptop, TV, phone and tablet	✓	✓	✓	HD available		✓	✓	Ultra HD available			✓
	Basic	Standard	Premium																															
Monthly cost* (United States Dollar)	\$8.99	\$13.99	\$17.99																															
Number of screens you can watch on at the same time	1	2	4																															
Number of phones or tablets you can have downloads on	1	2	4																															
Unlimited movies, TV shows and mobile games	✓	✓	✓																															
Watch on your laptop, TV, phone and tablet	✓	✓	✓																															
HD available		✓	✓																															
Ultra HD available			✓																															
<p>[1d]</p>	<p>deriving, using the first digital rights management key, from the</p>	<p>To provide access past the restriction of the content, Netflix must provide unique access to the particular content for the Netflix user. The fingerprint of the</p>																																

CONFIDENTIAL ATTORNEY-CLIENT AND WORK PRODUCT PRIVILEGED

	access restricted content a fingerprint of the access restricted content wherein the obtaining is based at least in part on the first digital rights management key,	restricted content is Netflix's unique access key for the restricted content generated for the Netflix user that is based at least in part on the license to view the content (first digital rights management key).
[1e]	causing the content providing server to validate the fingerprint, an, if the validation is successful, accessing the access restricted content and information describing encryption properties of the access restricted content, and	On information and belief, the unique access is provided to Netflix's content providing server and if the validation is successful (for example, the Netflix user that is logged in to a compatible device has paid for the monthly subscription or does not have too many users on one account), then the restricted content is accessed with information describing the encryption properties so that the encrypted data can be sent to the Netflix user's compatible device.
[1f]	deriving, using the digital rights management header of the access restricted content, from the access restricted content a second and third digital rights management key,	On and information and belief, once the Netflix user's access to the content is validated, additional digital rights management keys are created by Netflix based on the header of the restricted content so that the restricted content can be sent to the Netflix user's device.
[1g]	wherein the second and third digital rights management keys are	The additional digital rights management keys are applied to begin the actual download of the restricted content (or streaming) from Netflix's content database

CONFIDENTIAL ATTORNEY-CLIENT AND WORK PRODUCT PRIVILEGED

	<p>applied to retrieve the payload of the access restricted content and</p>	<p>and/or server to the Netflix user’s device with the Netflix app or Netflix on the web.</p>																																
<p>[1h]</p>	<p>wherein at least one of the second or third digital rights management key is used to encrypt the other key of the second or third digital rights management key,</p>	<p>On information and belief, Netflix encrypts at least one of the digital rights management keys that enable download of an access restricted title because failure to do so would provide anyone with the ability to download access restricted content without having an account or paying simply by knowing the key.</p>																																
<p>[1i]</p>	<p>wherein the content is usable without being in an unprotected state.</p>	<p>Restricted content that has been accessed via Netflix is only available for streaming or download through Netflix websites or compatible devices. Additionally, the content is available only through the Netflix user’s account for a limited number of screens(devices). Thus, the content is usable without being in an unprotected state, the unprotected state being unlimited downloads, transfer, and viewing of the title.</p> <p>Plans and Pricing</p> <p>Netflix offers a variety of plans to meet your needs. The plan you choose will determine the video quality and the number of screens you can watch Netflix on at the same time.</p> <p>With all of our plans, you can watch unlimited TV shows and movies, and play mobile games.</p> <table border="1" data-bbox="852 1027 1650 1333"> <thead> <tr> <th></th> <th>Basic</th> <th>Standard</th> <th>Premium</th> </tr> </thead> <tbody> <tr> <td>Monthly cost* (United States Dollar)</td> <td>\$8.99</td> <td>\$13.99</td> <td>\$17.99</td> </tr> <tr> <td>Number of screens you can watch on at the same time</td> <td>1</td> <td>2</td> <td>4</td> </tr> <tr> <td>Number of phones or tablets you can have downloads on</td> <td>1</td> <td>2</td> <td>4</td> </tr> <tr> <td>Unlimited movies, TV shows and mobile games</td> <td>✓</td> <td>✓</td> <td>✓</td> </tr> <tr> <td>Watch on your laptop, TV, phone and tablet</td> <td>✓</td> <td>✓</td> <td>✓</td> </tr> <tr> <td>HD available</td> <td></td> <td>✓</td> <td>✓</td> </tr> <tr> <td>Ultra HD available</td> <td></td> <td></td> <td>✓</td> </tr> </tbody> </table>		Basic	Standard	Premium	Monthly cost* (United States Dollar)	\$8.99	\$13.99	\$17.99	Number of screens you can watch on at the same time	1	2	4	Number of phones or tablets you can have downloads on	1	2	4	Unlimited movies, TV shows and mobile games	✓	✓	✓	Watch on your laptop, TV, phone and tablet	✓	✓	✓	HD available		✓	✓	Ultra HD available			✓
	Basic	Standard	Premium																															
Monthly cost* (United States Dollar)	\$8.99	\$13.99	\$17.99																															
Number of screens you can watch on at the same time	1	2	4																															
Number of phones or tablets you can have downloads on	1	2	4																															
Unlimited movies, TV shows and mobile games	✓	✓	✓																															
Watch on your laptop, TV, phone and tablet	✓	✓	✓																															
HD available		✓	✓																															
Ultra HD available			✓																															

CONNEC

PEERING WITH OPEN CONNECT

OPEN CONNECT APPLIANCES

DEPLOYING EMBEDDED APPLIANCES

PARTNER PORTAL

SUPPORT



HOME

WELCOME TO OPEN CONNECT >

WHAT IS OPEN CONNECT >

HOW TO GET STARTED >

SAMPLE ARCHITECTURES >

WHAT'S HERE? >

Open Connect

HOME

WELCOME TO OPEN CONNECT >

WHAT IS OPEN CONNECT >

HOW TO GET STARTED >

SAMPLE ARCHITECTURES >

WHAT'S HERE? >

Welcome to Open Connect

The goal of the Netflix Open Connect program is to provide our millions of Netflix subscribers the highest-quality viewing experience possible. We achieve this goal by partnering with Internet Service Providers (ISPs) to deliver our content more efficiently. We partner with over a thousand ISPs to localize substantial amounts of traffic with Open Connect Appliance embedded deployments, and we have an open peering policy at our [interconnection locations](#). If you are an ISP with a substantial amount of Netflix traffic, review this information to learn more about the program.

For more information about Open Connect, see:

- [Overview of Open Connect \(PDF - English only\)](#)
- [Open Connect blog post](#)

Key links on this site:

CONNEC

PEERING
WITH
OPEN
CONNECTOPEN
CONNECT
APPLIANCESDEPLOYING
EMBEDDED
APPLIANCESPARTNER
PORTAL

SUPPORT



Fill out the [appliance request form](#) if you are interested in embedded appliance solutions.

What is Open Connect

The Netflix Open Connect program provides opportunities for ISP partners to improve their customers' Netflix user experience by localizing Netflix traffic and minimizing the delivery of traffic that is served over a transit provider.

There are two main components of the program, which are architected in partnership with ISPs to provide maximum benefit in each individual situation: embedded Open Connect Appliances and settlement-free interconnection (SFI).

Embedded Open Connect Appliances (OCAs)

Open Connect Appliances can be embedded in your ISP network. Embedded OCAs have the same capabilities as the OCAs that we use in our 60+ global data centers, and they are provided to qualifying ISP partners at no charge. Each embedded OCA deployment will offload a substantial amount of Netflix content traffic from peering or transport circuits. Multiple physical deployments can be distributed or clustered on a geographic or network basis to maximize local offload.

Netflix provides:

- Network architecture and technical turn-up expertise
- Ongoing monitoring and issue resolution

CONNEC

PEERING
WITH
OPEN
CONNECTOPEN
CONNECT
APPLIANCESDEPLOYING
EMBEDDED
APPLIANCESPARTNER
PORTAL

SUPPORT



The [OCA Deployment Guide](#) provides deployment details. If you are interested in OCA deployment at your ISP, please fill out our [OCA request form](#).

Settlement-free interconnection (SFI)

Connect via direct Private Network Interconnect (PNI) or IXP-based SFI peering to Netflix Open Connect Appliances in our data centers. Peering alone can be beneficial.

If you deploy embedded OCAs, you also set up SFI peering for additional resiliency and to enable nightly content fills and updates.

Netflix has the ability to interconnect at a number of global data center facilities and public Internet Exchange fabrics as listed on our [Peering Locations](#) page. We openly peer with any network at IXP locations where we are mutually present and we consider private interconnection as appropriate. If you are interested in interconnection, please review the information on the [Peering Locations](#) page.

ISPs who do not currently participate in public peering might want to consider that a single IX port can support multiple peering sessions, providing direct access to various content, cloud, and network providers. In addition to Netflix, many large organizations such as Akamai, Amazon, Facebook, and Google/YouTube widely participate in public peering and combine to deliver a substantial percentage of traffic to a typical ISP.

From a connectivity standpoint, IX ports can be reached locally in a data center or via transport. We recommend <http://peeringdb.com> as a detailed source of information that can help you find an IX that best meets your needs.

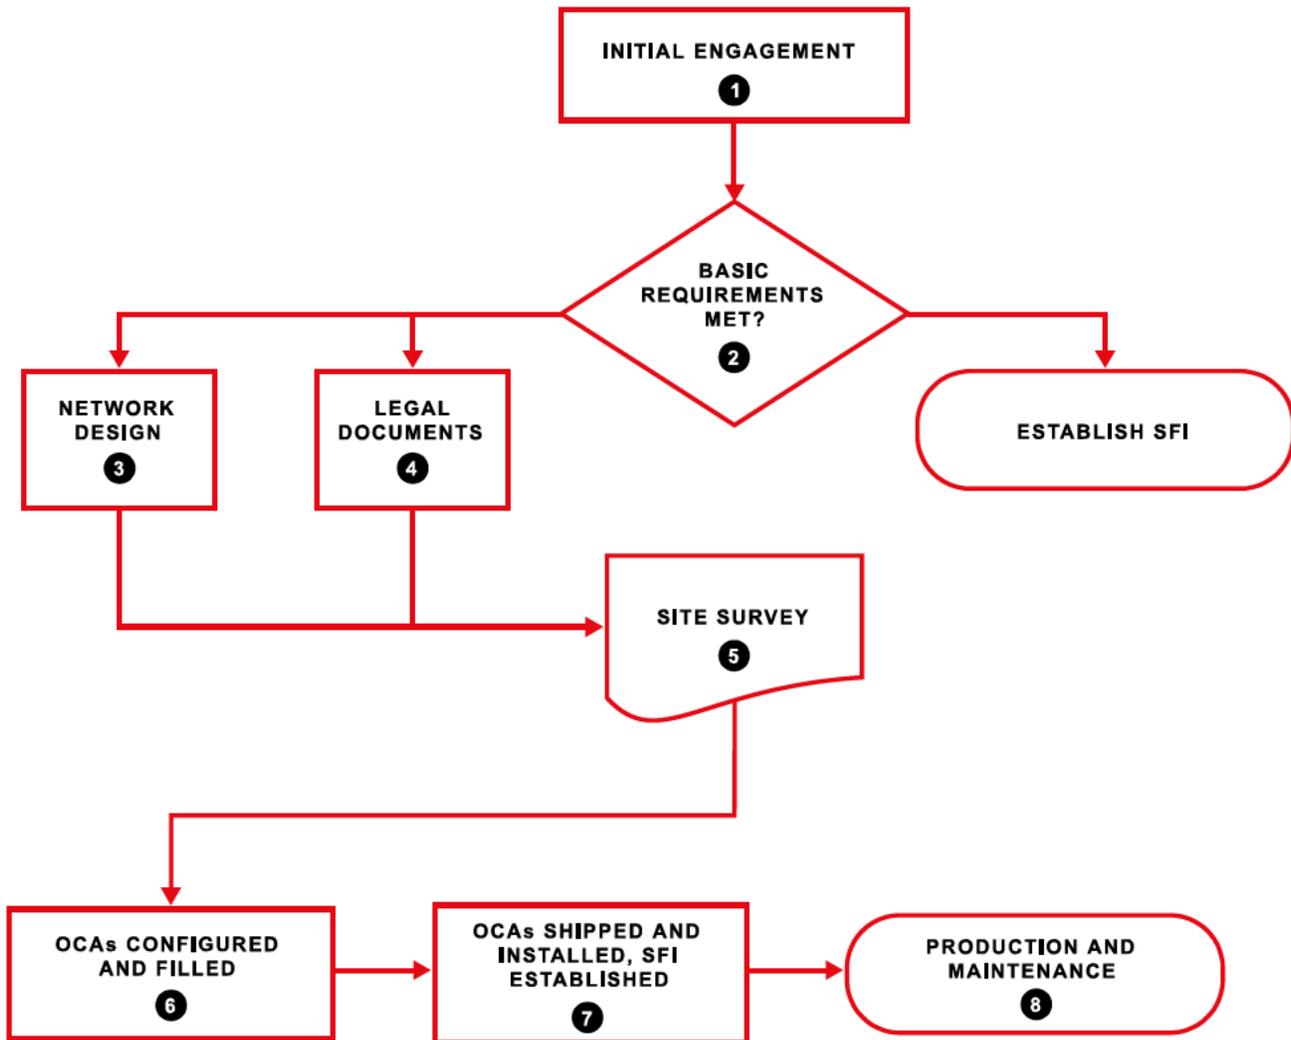


CONNEC

PEERING
WITH
OPEN
CONNECTOPEN
CONNECT
APPLIANCESDEPLOYING
EMBEDDED
APPLIANCESPARTNER
PORTAL

SUPPORT

A



1 To begin the engagement process, you complete one of the following actions:

- If you are engaging with Netflix for the first time, please submit an [online Open Connect Appliance request](#), and a Netflix Open Connect Partner Engagement Manager (PEM) will contact you.
- If you already have an established relationship with a Netflix PEM and you are interested in modifying an existing configuration, contact your PEM directly.

2 Netflix works with you to determine whether you meet the [basic requirements for deploying embedded OCAs](#). If embedded OCAs are not warranted or cannot be deployed for some reason, SFI peering alone can be established.

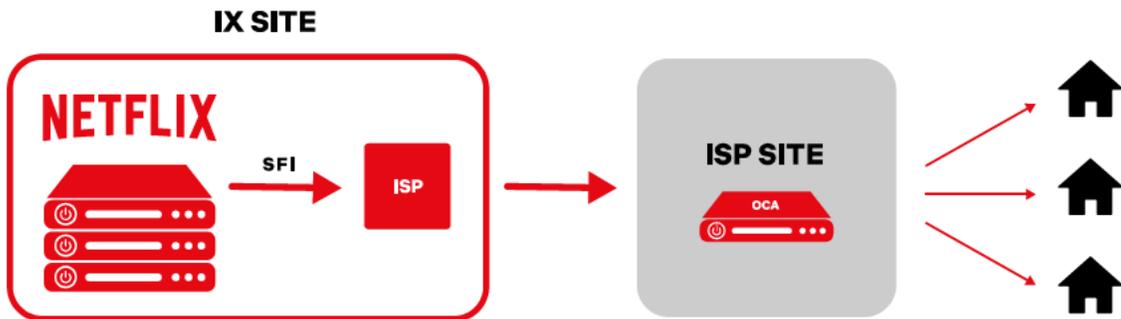
- 7 Netflix ships you the configured and filled appliance(s) for installation in your data center. SFI is also established to enable optimal resiliency and fill/updates. Shipping typically takes approximately **1 week**, depending on your geographical location.
- 8 After the appliance is enabled into production, Netflix performs ongoing **health monitoring and maintenance** to ensure smooth operation of the appliance.

Sample Architectures

The following examples illustrate how Open Connect deployments can be configured for resiliency and maximum benefit to the customer.

Basic Architectures

The following diagram shows an example of an OCA that is embedded in a partner network, in conjunction with SFI peering which is used to provide additional resiliency and to enable nightly content **fill and updates**.

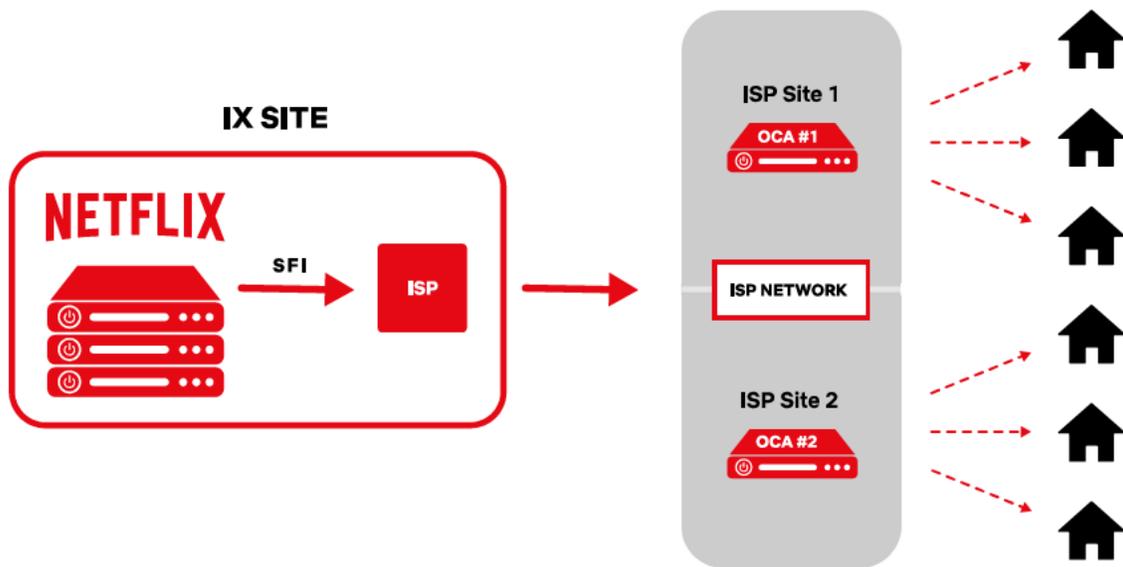


In contrast, the next diagram shows an example of SFI (peering) *without* the deployment of embedded OCAs in the partner network. In this scenario, traffic is delivered to end users via SFI from Netflix appliances that are located in local IXPs, to avoid both the cost and congestion that is associated with transit.



Failover scenarios for embedded deployments

The following diagram illustrates a scenario where an ISP partner has two embedded OCAs deployed in two different sites.

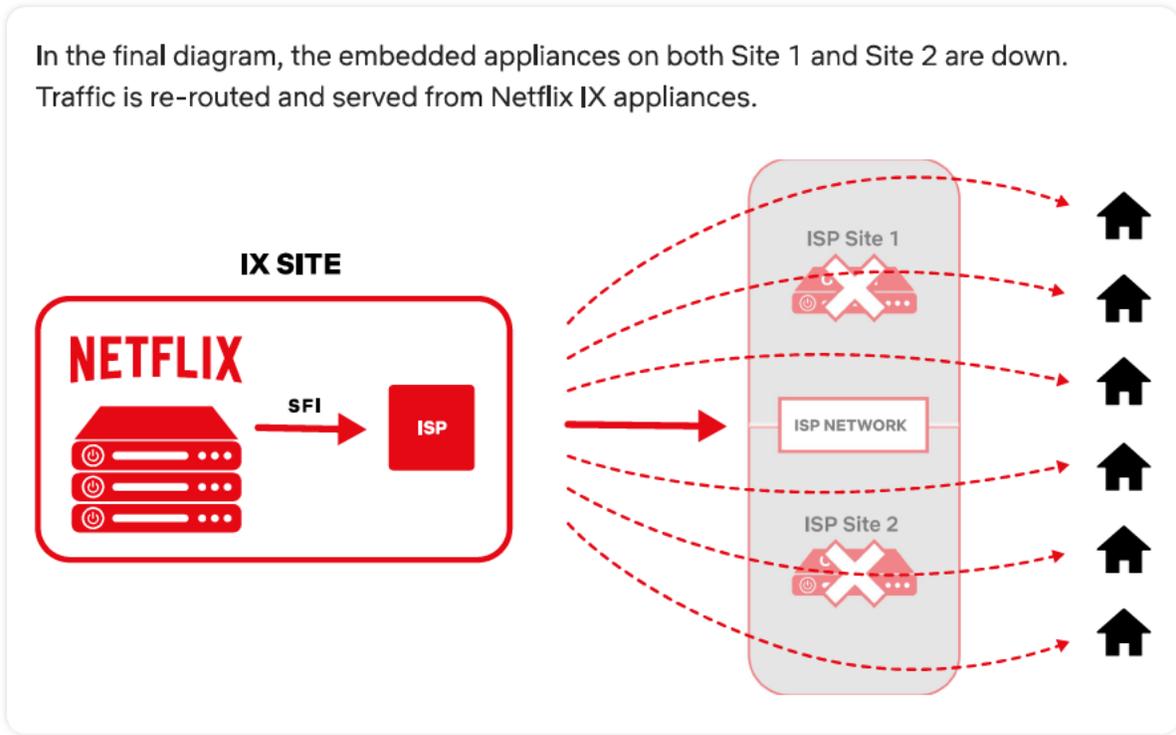


In the next diagram, the embedded appliance on Site 1 is down, but the embedded appliance on Site 2 is up. Traffic is re-routed and served from Site 2.





In the final diagram, the embedded appliances on both Site 1 and Site 2 are down. Traffic is re-routed and served from Netflix IX appliances.



CONNEX

PEERING
WITH
OPEN
CONNECT

OPEN
CONNECT
APPLIANCES

DEPLOYING
EMBEDDED
APPLIANCES

PARTNER
PORTAL

SUPPORT



Peering With Open Connect



Open Connect Appliances



Deploying Embedded Appliances



Partner Portal



Frequently Asked Questions



Jobs



CONNEC

PEERING
WITH
OPEN
CONNECT

OPEN
CONNECT
APPLIANCES

DEPLOYING
EMBEDDED
APPLIANCES

PARTNER
PORTAL

SUPPORT



Building fast.com



Netflix Technology Blog

Follow

Aug 9, 2016 · 7 min read

On [our company blog in May](#), we introduced [fast.com](#), our new internet speed test. The idea behind fast.com is to provide a quick and simple way for any internet user to test their current internet speed, whether they are a Netflix member or not. Since fast.com was released, millions of internet users around the world have run the test. We have seen a lot of interest in the site and questions about how it works. This blog will give a high-level overview of how we handled some of the [challenges inherent with measuring internet speeds](#) and the technology behind fast.com.

But first, some news — we are happy to announce a new FAST mobile app, available now for Android or Apple mobile devices. Get the free app from the [Apple App Store](#) or [Google Play](#).

Design goals

When designing the user experience for the fast.com application, we had several important goals in mind:

- Provide accurate, consistent results that reflect users' real-life internet use case
- Load and run as quickly as possible
- Provide simple results that are easy to understand
- Work on most devices from the browser without requiring installation of a separate application

We wanted to make sure that fast.com could be easily used and understood by the majority of internet users, without requiring them to have any prior knowledge of

computer networking, command line tools, and the like.

Technical goals

There are various ways to go about measuring internet speed and many variables that can impact any given measurement, some of which are not under our control. For example — configuration of the user's local or home network, device or router performance, other users on the network, TCP or network configuration on the device. However, we thought carefully about the variables that are under our control and how they would further our overall goal of a simple but meaningful test.

Variables that are under our control, and which can influence the results of the test, include things like:

- Server location
- Load on the server
- Number of TCP connections used
- Size and type of download content used
- Methodology used to aggregate measurements

One major advantage we have is our Open Connect CDN, a globally-distributed network of servers (Open Connect Appliances or OCAs) that store and serve Netflix content to our members — representing as much as 35% of last-mile internet peak traffic in some regions. Using our own production servers to test internet speed helps to ensure that the test is a good representation of the performance that can be achieved during a real-life user scenario.

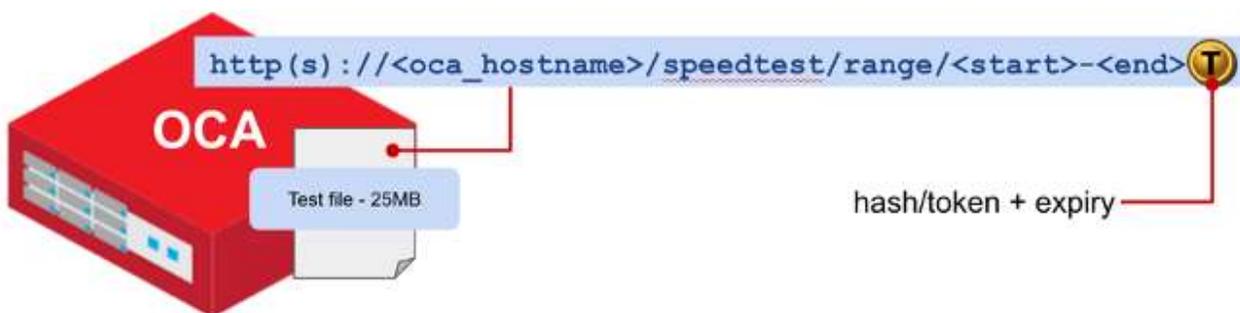
In pursuit of the design goal of simplicity, we deliberately chose to measure only download speed, measuring how fast data travels from server to consumer when they are performing activities such as viewing web pages or streaming video. Downloads represent the majority of activity for most internet consumers.

We also decided on the following high-level technical approaches:

- To open several connections for the test, varying the number depending on the network conditions
- To run the test on several of our wide network of Netflix production OCAs, but only on servers that have enough capacity to serve test traffic while simultaneously operating within acceptable parameters to deliver optimal video quality to members
- To measure long running sessions — eliminating connection setup and ramp up time and short term variability from the result
- To dynamically determine when to end the test so that the final results are quick, stable, and accurate
- To run the test using HTTPS, supporting IPv4 and IPv6

Architecture

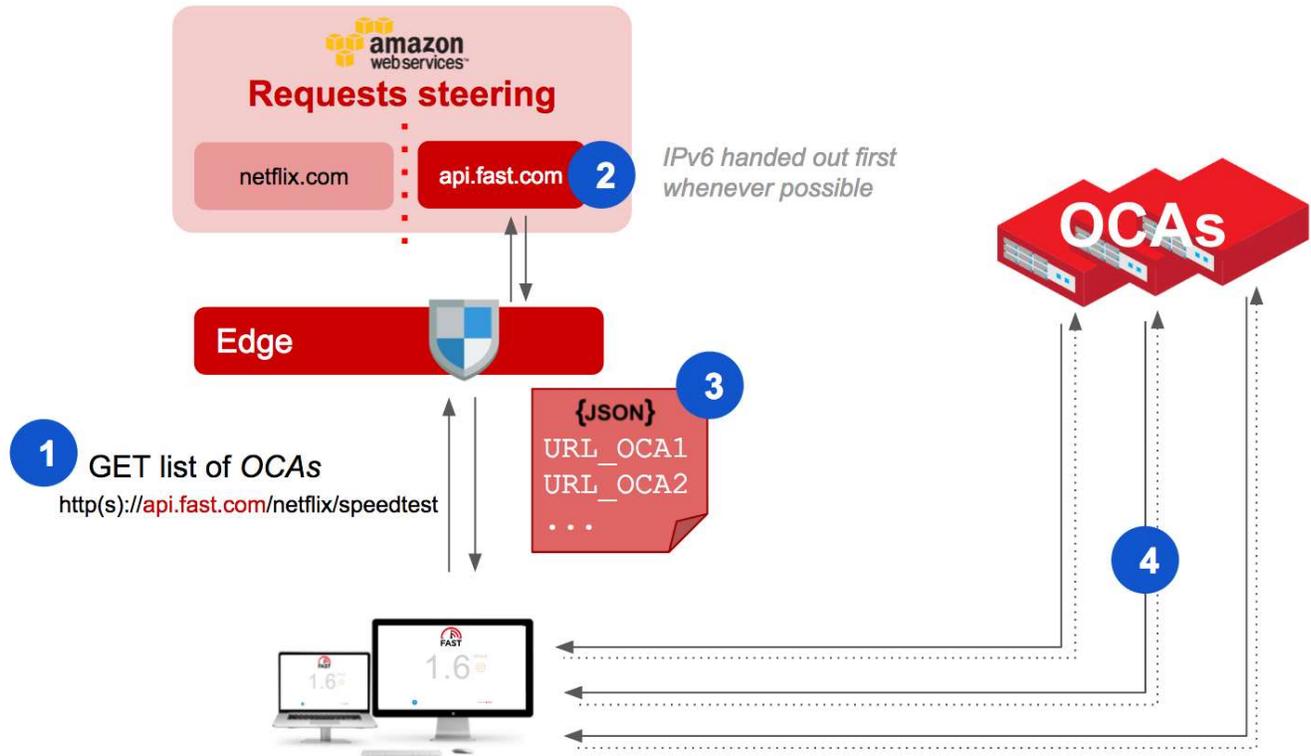
As mentioned above, fast.com downloads test files from our distributed network of Open Connect Appliances (OCAs). Each OCA server provides an endpoint with a 25MB video file. The endpoint supports a range parameter that allows requests for between a 1 byte to a 25MB chunk of content.



In order to steer a user to an OCA server, fast.com provides an endpoint that returns a list of several URLs for different OCAs that are best suited to run the test. To determine the list, the endpoint uses logic that is similar to the logic that is used to steer netflix.com video delivery. The OCAs that are returned are chosen based on:

- Network distance

- Traffic load for each OCA, which indicates overall server health
- Network structure — each OCA in the list belongs to a different cluster



As soon as the fast.com client receives the URLs, the test begins to run.

Estimating network speed

The test engine uses heuristics to:

- Strip off measurements that are collected during connection setup/ramp up
- Aggregate the rest of the collected measurements
- Decide how many parallel connections to use during the test
- Try to separate processing overhead from network time — because fast.com runs in the browser, it has limited visibility into timing of network events like DNS resolution time, processing of packets on the client side and latency to test server

- Make a decision about when the client has collected enough measurements to confidently present the final network speed estimate

We exclude initial connection ramp up, but we do take into account any performance drops during the test. Network performance drops might indicate a lossy network, congested link, or faulty router — therefore, excluding these drops from the test result would not correctly reflect issues experienced by users while they are consuming content from the internet.

Number of connections

Depending on network throughput, the fast.com client runs the test using a variable number of parallel connections. For low throughput networks, running more connections might result in each connection competing for very limited bandwidth, causing more timeouts and resulting in a longer and less accurate test.

When the bandwidth is high enough, however, running more parallel connections helps to saturate the network link faster and reduce test time. For very high throughput connections, especially in situations with higher latency, one connection and a 25MB file might not be enough to reach maximum speeds, so multiple connections are necessary.

Size of downloads

For each connection, the fast.com client selects the size of the chunk of the 25MB file that it wants to download. In situations where the network layer supports periodical progress events, it makes sense to request the whole file and estimate network speed using download progress counters. In cases where the download progress event is not available, the client will gradually increase payload size during the test to perform multiple downloads and get a sufficient number of samples.

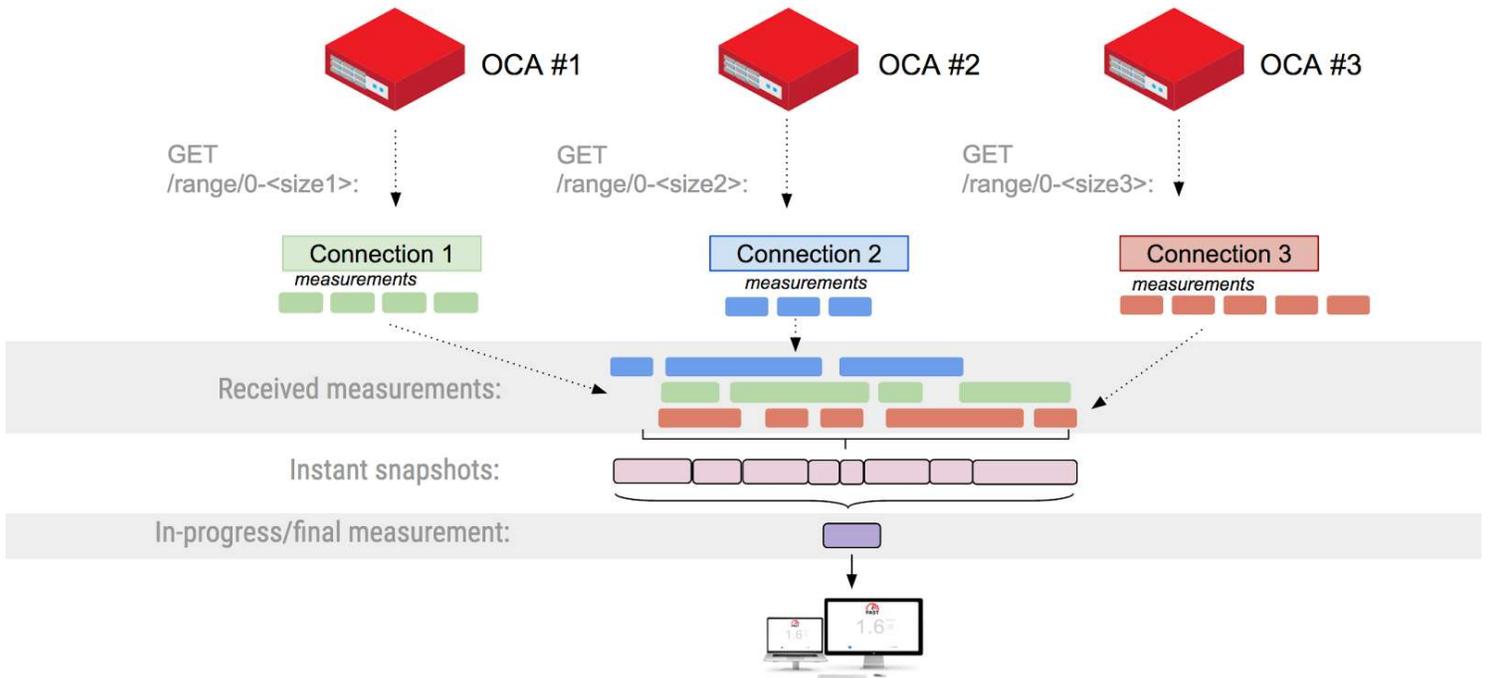
Computing the results

After the download measurements are collected, the client combines the downloaded content across all connections and keeps the snapshot speed.

The ‘instant’ network measurements are then passed to the results aggregation module. The aggregation module makes sure that:

- We exclude initial connection ramp up

- We take the rest and compute rolling average of the other measurements

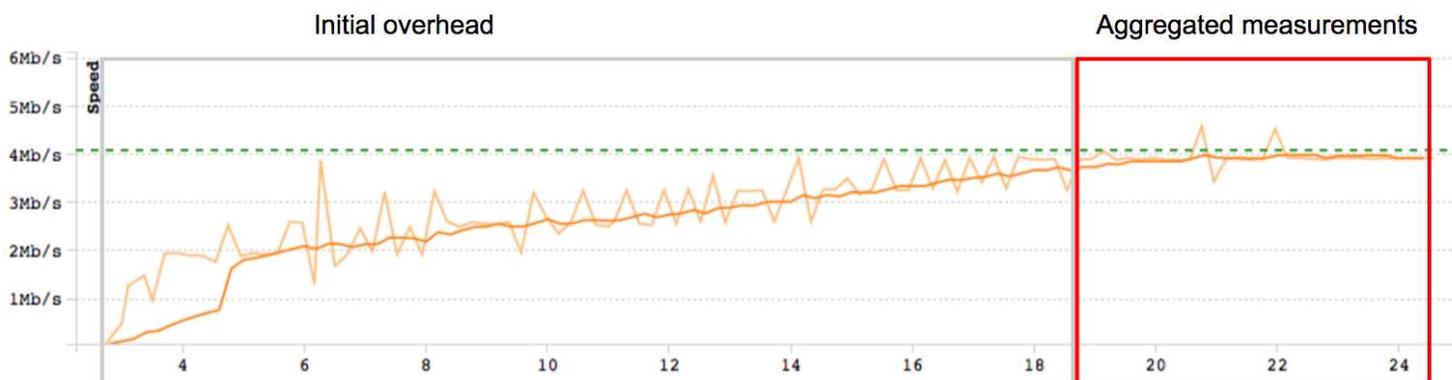


One of the primary challenges for the fast.com client is determining when the estimated speed measurements are ready to be presented as a final estimate. Due to the various environments and conditions that the fast.com test can be run under, the test duration needs to be dynamic.

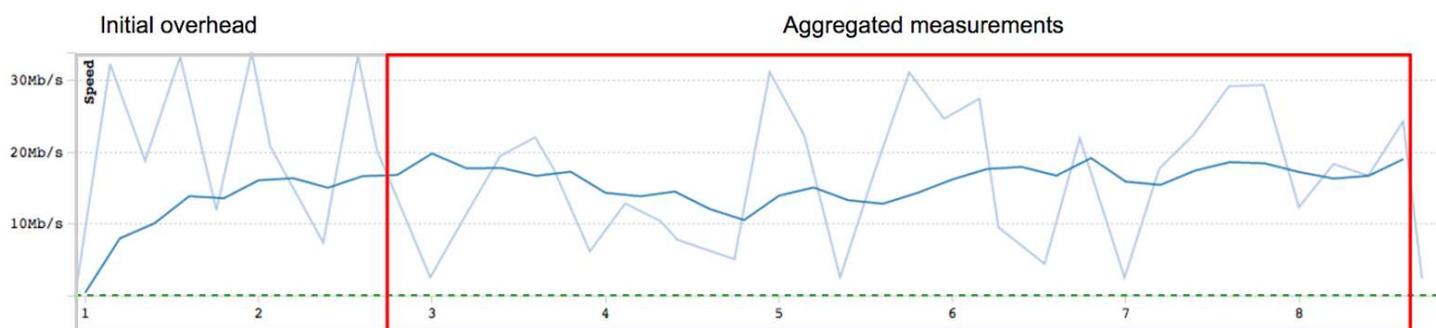
For stable low latency connections, we quickly see growth to full network speeds:



Higher latency connections take much longer to ramp up to full network speed:



Lossy or congested connections show significant variations in instant speed, but these instant variations get smoothed out over time. It is also harder to correctly identify the moment when connections have ramped up to full speed.



In all cases, after initial ramp up measurements are excluded, the ‘stop’ detection module monitors how the aggregated network speed is changing and makes a decision about whether the estimate is stable or if more time is needed for the test. After the results are stable, they are presented as a final estimate to the user.

Conclusion and Next Steps

We continue to monitor, test, and perfect fast.com, always with the goal of giving consumers the simplest and most accurate tool possible to measure their current internet performance. We plan to share updates and more details about this exciting tool in future posts.

— by Sergey Fedorov and Ellen Livengood

Originally published at techblog.netflix.com on August 9, 2016.

[Network Speed](#) [Internet Speed](#) [Ips](#)

[About](#) [Write](#) [Help](#) [Legal](#)

Get the Medium app





Open Connect Overview

What is Netflix Open Connect?

Open Connect is the name of the global network that is responsible for delivering Netflix TV shows and movies to our members world-wide. This type of network is typically referred to as a “Content Delivery Network” or “CDN” because its job is to deliver internet-based content (via HTTP/HTTPS) efficiently by bringing the content that people watch close to where they’re watching it. The Open Connect network shares some characteristics with other CDNs, but also has some important differences.

Netflix began the Open Connect initiative in 2011, as a response to the ever-increasing scale of Netflix streaming. We started the program for two reasons:

- 1) As Netflix grew to be a significant portion of overall traffic on consumer Internet Service Provider (ISP) networks, it became important to be able to work with those ISPs in a direct and collaborative way.
- 2) Creating a content delivery solution customized for Netflix allowed us to design a proactive, directed caching solution that is much more efficient than the standard demand-driven CDN solution, reducing the overall demand on upstream network capacity by several orders of magnitude.

Several years in, we are constantly learning - adjusting and evolving the program to ensure that our members continue to have a consistent high quality video experience from wherever they are in the world.

The overall mission of the program is to enable ISPs to provide a great Netflix experience for our mutual customers. We further this goal by localizing Netflix traffic as close as possible to our members, limiting the network and geographical distances that our video bits must travel during playback. This of course benefits Netflix members, but it also benefits ISPs and internet users in general. In short, we invest in efficiency

innovations and increasing the capacity of the internet to support playback requests for Netflix content - so that others don't have to.

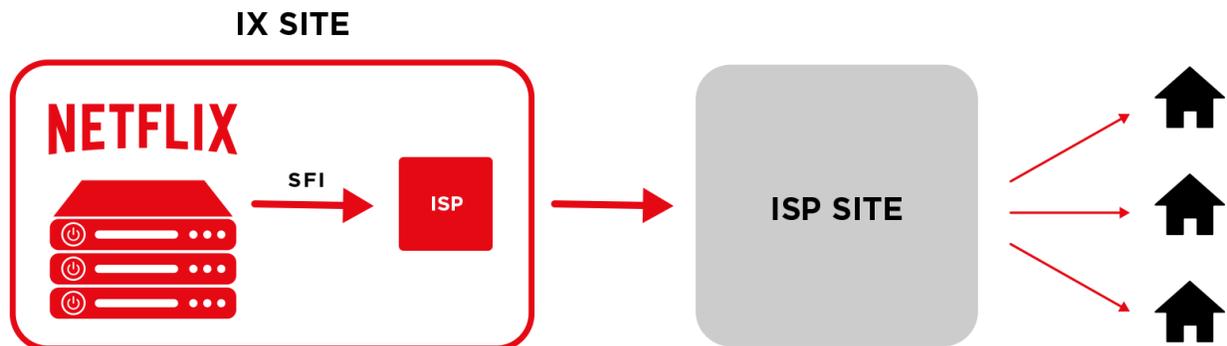
Open Connect Appliances

The building blocks of Open Connect are our suite of purpose-built server appliances, called Open Connect Appliances (OCAs). These appliances store encoded video/image files and serve these files via HTTP/HTTPS to client devices (for example: set top boxes, mobile devices, or smart TVs). OCAs have the sole responsibility of delivering playable bits to client devices as fast as possible.

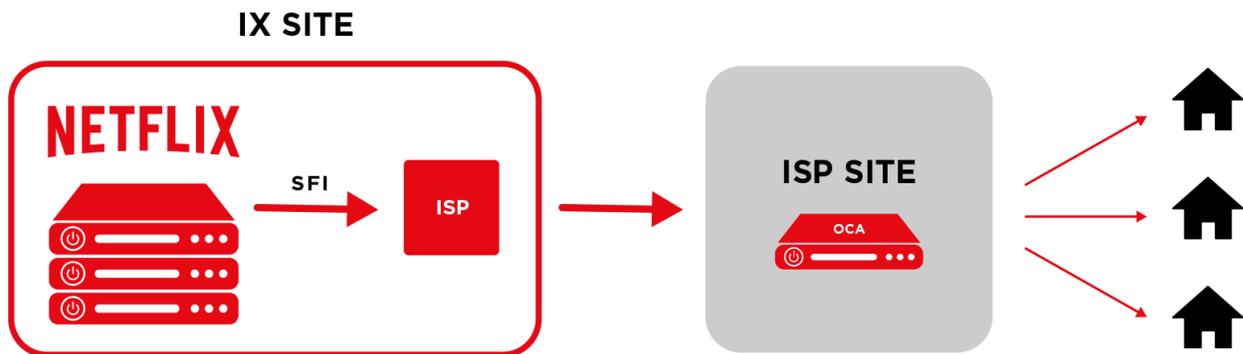
As with all facets of the Open Connect program, appliance design continues to expand and improve over time to keep up with our current and future needs. We are involved in the development of all layers of the software stack, and we make both the hardware design and the software available via open source for others to benefit from.

Our global network of thousands of OCAs are deployed in two ways:

- 1 We install OCAs within internet exchange points (referred to as IXs or IXPs) in significant Netflix markets throughout the world. These OCAs are interconnected with mutually present ISPs via settlement free public or private peering (SFI). Peering alone can be very beneficial to our ISP partners.



- We provide OCAs free of charge to qualifying ISPs. These OCAs, with the same capabilities as the OCAs that are in the IXPs, are deployed directly inside ISP networks. We provide the server hardware and the ISPs provide power, space, and connectivity. ISPs directly control which of their customers are routed to their embedded OCAs. ISP partners with embedded OCAs also use peering for resiliency and to enable efficient nightly fill and updates, described later.



Each individual OCA deployment site is custom tailored by the Open Connect team based on local network characteristics and other key capacity planning factors. After we deploy OCAs to a site, we constantly measure and analyze their performance and augment capacity as requirements evolve.

Interaction with Client Devices and Netflix AWS Services

OCAs do not store client data (for example - viewing history, DRM info, or member data). Essentially, OCA servers only do the following two things:

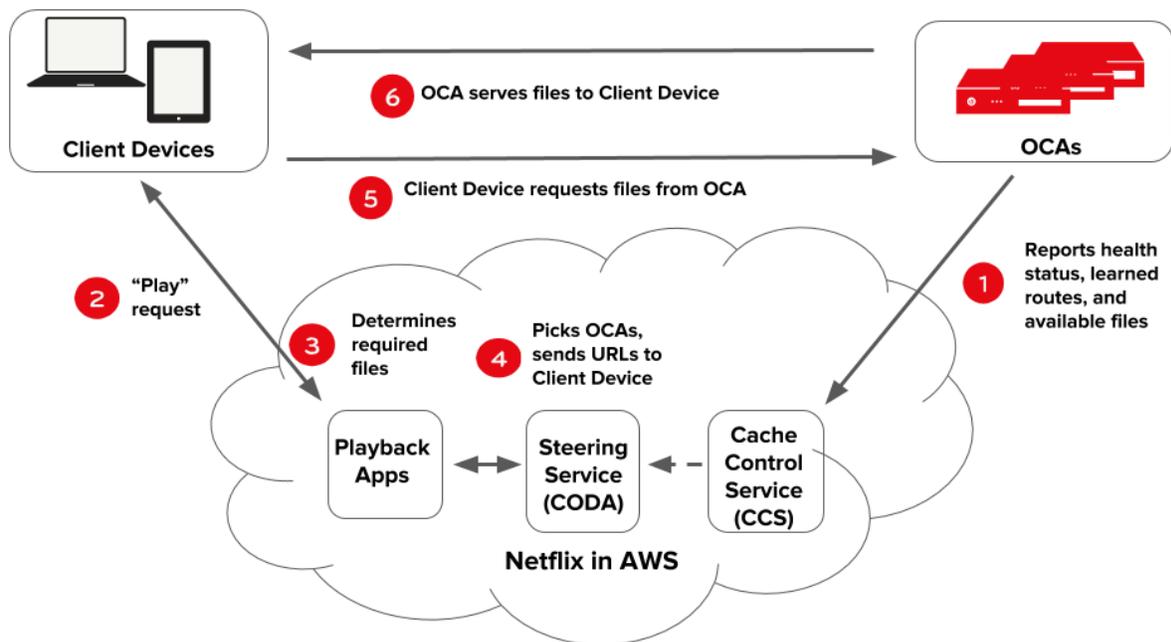
- Report their status to the Open Connect control plane services in Amazon Web Services (AWS). For example, they report health metrics, BGP routes they have learned from the BGP peer (router or switch) they have a configured BGP session with, and what files they have stored on disk.
- Serve content via HTTP/HTTPS when it is requested by a client device

The control plane services in AWS take the data that the OCAs report and use it to steer clients via URL to the most optimal OCAs given their file availability, health, and network

proximity to the client. The control plane services also control fill behavior (adding new files to OCAs nightly), compute optimal behavior for such things as file storage/hashing, and handle the storage and interpretation of relevant telemetry about the playback experience.

Open Connect also partners extensively with Netflix client teams to ensure that the content that is served by the OCAs is optimized dynamically by each client device based on its specific needs and the current network conditions.

The following diagram illustrates how the playback process works:



1. OCA's periodically report health, routes they have learned, and content (file) availability to the cache control services in AWS.
2. A user on a client device requests playback of a title (TV show or movie) from the Netflix application in AWS.
3. The playback application services in AWS check user authorization and licensing, then determine which specific files are required to handle the playback request - taking individual client characteristics and current network conditions into account

4. The steering service in AWS uses the information stored by the cache control service to pick OCAs that the requested files should be served from, generates URLs for these OCAs, and hands the URLs over to the playback application services.
5. The playback application services hand over URLs of the appropriate OCAs to the client device, and the OCA begins to serve the requested files.

Monitoring, Maintenance, and Updates

All of our OCA deployments, whether in IXPs or embedded in ISP networks, are constantly monitored by the Open Connect Operations team to ensure reliability and efficiency. We troubleshoot and proactively fix most issues remotely with minimal input required from our ISP partners. If partners wish to monitor their own embedded OCAs' status and performance, we provide a Partner Portal where they can do so. If hardware performance degrades to the point where a server is no longer functioning in the range of our quality standards, we simply replace it - at no cost to our partners.

We have the somewhat unique benefit of being able to deploy the majority of our content and software updates proactively during off-peak fill windows. Because we can predict with high accuracy what our members will watch and what time of day they will watch it, we make use of non-peak bandwidth to download the vast majority of content updates to the OCAs in our network during these configurable time windows. OCAs can also download updates from each other - minimizing significant usage of internet "backbone" capacity during the update cycle.

We use ever-evolving popularity algorithms and storage techniques in the control plane services to ensure that our content is distributed in ways that maximize offload efficiency and quality of experience, while minimizing churn in the form of updates to the content that is stored on the appliances.

More Information

For more information about Open Connect, see:

- [The Open Connect website](#)
- [Open Connect articles on the Netflix Tech Blog](#)
- [Open Connect articles on the Netflix Media Center](#)

Copyright 2016, 2019 by Netflix, Inc. All rights reserved. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or by any information storage or retrieval system, without express permission from Netflix, Inc.

Exhibit D



Gregory K. Peters
Chief Streaming and Partnerships Officer
Netflix
100 Winchester Circle,
Los Gatos, CA 95032

cc: David Hyman, General Counsel

September 29, 2014

Dear Mr. Peters:

SC Intelligent Holding Oy is the sole licensor of the U.S. Patent 8,495,167B2. According to our evaluation reports, Netflix, as the leading content delivery network, is a potential licensee of the aforementioned patent.

From public sources in the April 2014 time frame, it seemed apparent that Netflix was considering an R & D project on digital push distribution through P2P networks to ease the strain on data communication networks and reduce Netflix's reliance on third party technology.

In addition to our patent portfolio (U.S., Europe and China), we developed and field-tested a working digital distribution service, based on our patented solution, from 2002 – 2008 in Finland, with over 32,000 end-users connected to our P2P test service. The six year field test resulted in a modular system of approximately 75,000 well designed lines of code, most of which could be useful for Netflix to exploit today and for many years in the future.

Together with a functional, field-tested P2P distribution system, we could offer Netflix an opportunity to exploit these IP assets, decrease its dependency on CDN platform providers and strengthen its position as the leading content delivery network.

For your convenience, along with this letter, I have inserted the header pages of the patent, which gives you the ability to research this issue in full.



With regard to our CDN solution, we would be pleased to provide you with more detailed information under an NDA, in case you find this subject interesting.

Please let me know when would be a convenient time for us to begin a discussion over various options around the patent.

With Best Regards,

A handwritten signature in blue ink, appearing to read 'Lauri Valjakka'.

Lauri Valjakka
CEO
SC Intelligent Holding Oy
Lappeenranta
FINLAND

Tel: +358 50 467 0090
e-mail: lauri.valjakka@scih.fi

For delivery information, visit our website at www.usps.com®.

LOS GATOS CA 95032

Postage	\$	\$0.70
Certified Fee		\$3.30
Return Receipt Fee (Endorsement Required)		\$2.70
Restricted Delivery Fee (Endorsement Required)		\$0.00
Total Postage & Fees	\$	\$6.70

7014 2120 0002 8937 3504

15
Postmark Here
10/03/2014

Sent To

NETALIK ATTN: GREGORY PETERS
 Street & Apt. No.,
 or PO Box No. WINCHESTER CIRCLE
 City, State, ZIP+4 LOS GATOS, CA 95032

PS Form 3800, July 2014

See Reverse for Instructions

CARLSBAD MPO
 CARLSBAD, California
 920089998
 0567760008-0094
 10/03/2014 (800)275-8777 10:53:18 AM

Sales Receipt

Product Description	Sale Unit Qty	Price	Final Price
---------------------	---------------	-------	-------------

LOS GATOS CA 95032-1815 Zone-4			\$0.70
First-Class Mail Letter			
1.20 oz.			
Expected Delivery: Mon 10/06/14			
Return Rcpt (Green Card)			\$2.70
@@ Certified			\$3.30
USPS-Certified Mail #:			
70142120000289373504			

=====
 Issue Postage: \$6.70

LOS GATOS CA 95032 Zone-4			\$0.70
First-Class Mail Letter			
1.10 oz.			
Expected Delivery: Mon 10/06/14			
Issue Postage:			\$0.70

Total: \$7.40

Paid by:
 Cash \$8.00
 Change Due: -\$0.60

@@ For tracking or inquiries go to
 USPS.com or call 1-800-222-1811.

 BRIGHTEN SOMEONE'S MAILBOX. Greeting
 cards available for purchase at
 select Post Offices.

In a hurry? Self-service kiosks