



(43) International Publication Date
12 September 2014 (12.09.2014)

- (51) International Patent Classification:
H04L 29/06 (2006.01)
- (21) International Application Number:
PCT/US2014/017945
- (22) International Filing Date:
24 February 2014 (24.02.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
13/791,789 8 March 2013 (08.03.2013) US
- (71) Applicant: UNICORN MEDIA, INC. [US/US]; 24 West 5th Street, Tempe, Arizona 85281 (US).
- (72) Inventor: MCGOWAN, Albert, John; 4201 E. Camelback Road, Apt. 101, Phoenix, Arizona 85018 (US).
- (74) Agents: MCMILLAN, Scott et al.; Eighth Floor, Two Embarcadero Center, San Francisco, CA 94111 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:
— with international search report (Art. 21(3))

(54) Title: DYNAMIC CHUNKING FOR DELIVERY INSTANCES

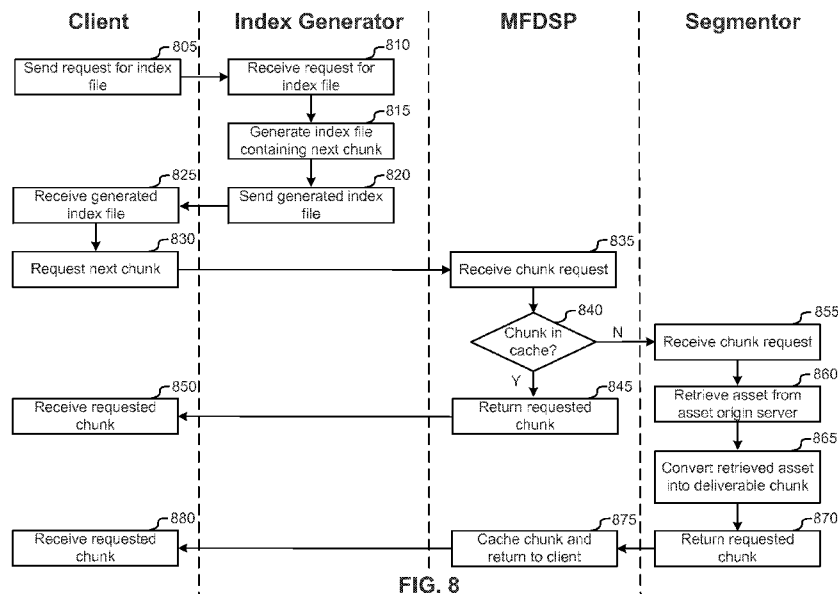


FIG. 8

(57) Abstract: Techniques are provided for dynamically creating index files for streaming media based on a determined chunking strategy. The chunking strategy can be determined using historical data of any of a variety of factors, such as Quality of Service (QoS) information. By using historical data in this manner, index files can be generated using chunking strategies that can improve these factors over time.

WO 2014/137639 A1

DYNAMIC CHUNKING FOR DELIVERY INSTANCES

BACKGROUND OF THE INVENTION

This disclosure relates in general to cloud-based computer processing and, but not by way of limitation, to providing media for use in media streaming.

5 The delivery of media over networks such as the Internet can be accomplished in many ways, including progressive downloading or streaming. Streaming a media file typically involves downloading “chunks,” or small segments, of the media file. Information including where chunks may be accessed can be stored in an index file (also known as a “manifest” file). This index file can be delivered to a client, such as a media player
10 application, for use in streaming. Additional information may also be provided, which can alter the appearance of the client.

BRIEF SUMMARY OF THE INVENTION

Systems and methods disclosed provide a technical solution improving quality of service of video streaming by using historical data of any of a variety of factors, such as
15 Quality of Service (QoS) information, to determine a chunking strategy. Index files can then be generated using chunking strategies, to help improve these factors over time.

An example method of providing media with a data network, according to the disclosure, includes determining historical data related to streaming media with the data network, and receiving a request having a universal source locator (URL), where the URL
20 includes information indicative of a requested media file. The method also includes determining one or more factors related to the request. The one or more factors related to the request include the historical data. The method further includes determining, with a processor, a chunking strategy based on the one or more factors related to the request, generating an index file having information for streaming the requested media file via the
25 data network, where generating the index file is based, at least in part, on the chunking strategy, and providing the index file.

An example server configured to provide media via data network, according to the disclosure, includes a network interface for communicating with the data network, a memory, and a processor communicatively coupled with the memory and the network interface. The
30 processor is configured to determine historical data related to streaming media with the data network, and receive a request having a universal source locator (URL), where the URL includes information indicative of a requested media file. The processor is further configured

to determine one or more factors related to the request, where the one or more factors related to the request include the historical data. The processor is also configured to determine a chunking strategy based on the one or more factors related to the request, generate an index file having information for streaming the requested media file via the data network, where
5 generating the index file is based, at least in part, on the chunking strategy, and provide the index file.

An example non-transitory computer-readable medium having instructions imbedded thereon for providing media with a data network, according to the disclosure, can have instructions, when executed by one or more computers, cause the one or more computers to
10 determine historical data related to streaming media with the data network, and receive a request having a universal source locator (URL), where the URL includes information indicative of a requested media file. The instructions can further cause the one or more computers to determine one or more factors related to the request, where the one or more factors related to the request include the historical data, and determine a chunking strategy
15 based on the one or more factors related to the request. The instructions can also cause the one or more computers to generate an index file having information for streaming the requested media file via the data network, where generating the index file is based, at least in part, on the chunking strategy, and provide the index file.

BRIEF DESCRIPTION OF THE DRAWINGS

20 The present disclosure is described in conjunction with the appended figures:

FIG. 1 illustrates a block diagram of a media servicing system.

FIG. 2A illustrates a block diagram of an embodiment of a kernel application center connected with application centers.

25 FIG. 2B illustrates a block diagram of an alternative embodiment of a kernel application center.

FIG. 3 illustrates a block diagram of an embodiment of an application center.

FIG. 4 illustrates a block diagram of processes and objects utilized by a cloud-hosted integrated multi-node pipelining system for media ingestion.

30 FIG. 5A illustrates a simplified block diagram of an embodiment of a system configured to provide dynamic indexing and chunking for media streaming.

FIG. 5B illustrates a simplified block diagram of another embodiment of a system configured to provide dynamic indexing and chunking for media streaming.

FIG. 5C illustrates a simplified block diagram of another embodiment of a system configured to provide dynamic indexing and chunking for media streaming, utilizing a
5 redirector.

FIG. 6 illustrates a simplified flowchart of an embodiment of a method for implementing a dynamic index for media streaming.

FIG. 7 illustrates a simplified flowchart of an embodiment of a method for dynamically chunking a media file for streaming.

10 FIG. 8 illustrates a simplified swim lane flowchart describing the interaction of components in a system configured to provide dynamic indexing and chunking for media streaming, according to one embodiment.

FIG. 9 illustrates a simplified flow chart providing a basic method for adapting chunking techniques for each delivery instance.

15 FIG. 10 illustrates a simplified flow chart illustrating a basic method for adapting chunking techniques for each delivery instance based on historical data, according to one embodiment.

In the appended figures, similar components and/or features may have the same reference label. Further, various components of the same type may be distinguished by
20 following the reference label by a dash and a second label that distinguishes among the similar components. If only the first reference label is used in the specification, the description is applicable to any one of the similar components having the same first reference label irrespective of the second reference label.

DETAILED DESCRIPTION OF THE INVENTION

25 The ensuing description provides preferred exemplary embodiment(s) only, and is not intended to limit the scope, applicability or configuration of the disclosure. Rather, the ensuing description of the preferred exemplary embodiment(s) will provide those skilled in the art with an enabling description for implementing a preferred exemplary embodiment. It is understood that various changes may be made in the function and arrangement of elements
30 without departing from the spirit and scope as set forth in the appended claims.

The increased availability of media content over data communications networks such as the Internet has mirrored the increased bandwidth for these networks. Because media has recently taken a more prominent role in data communications, the distribution of media and the data associated with such distribution has become increasingly important, particularly to
5 media media providers. Media streaming has become a widely-used method of media distribution, but the preprocessing associated with streaming can be burdensome. Certain protocols, including forms of Hypertext Transfer Protocol (HTTP) streaming, require chunking and storing media assets, and generating a corresponding index files. These requirements can deprive a media provider of the ability to dynamically insert additional
10 media such as advertisements into a media stream, and can consume a large amount of storage space to store chunks of media for a media asset, including chunks for any alternative sub-streams (e.g., streams with alternative bitrates, captions, alternative languages, etc.). Certain systems and methods can be utilized, however, to introduce the desired functionality back into the system.

15 A traditional approach to preprocessing media for streaming involves chunking and storing media assets, then creating corresponding index files to indicate where chunks may be located to download for streaming. Streaming protocols often provide for frequently updating an index file for instances where the corresponding media is frequently updated, such as during live streaming. Thus, an index file does not need to contain all chunks for a
20 requested media asset. In addition, because media files are frequently stored in a format that requires little additional processing to chunk, the chunks can be created in real time, during the streaming of a media file. The systems and methods disclosed herein take advantage of these features to enable dynamic index file creation and dynamic media file chunking.

For instance, rather than preprocess media assets for streaming by chunking and
25 indexing all files with relevant sub-streams prior to streaming the media, a server can dynamically create and update an index file during streaming. The dynamically-created index file can contain information regarding a next chunk of media in the various available sub-streams. The next chunk of media may not be cached at a location specified in the index file, in which case a chunk may be dynamically created by pulling all or part of the media file
30 of interest from a media file origin, chunking it, and making it available for download. The chunk also may be cached, thereby eliminating the need to create the chunk again if it is requested at some later time.

Because the chunks are created during streaming, a media provider and/or media distributor can have more information and control during the streaming process. Rather than

generate single index file for a given media asset, an instance of the index file generator may be created at the beginning of the media streaming to provide individualized media content to a particular end user and unique information regarding the streaming session to a media provider. The file index generator can vary the length of each chunk by, for example,
5 indicating starting and ending points in the index file. Thus, the file index generator may determine a uniform chunk length for a media asset, varying the length of the chunks for different media assets, or the file index generator may adjust the length of the chunks within a single media asset. The index file generator can further insert additional media, such as an advertisement, at any time during the streaming by specifying the location of the additional
10 media in the index file. The determination to insert advertisements can be based on any information, including data collected during the streaming session.

As the index file generator receives requests for and generates index files, it can further gather data regarding the streaming session for reporting to a media provider. Media providers often rely on beaconing data collected from media player applications to determine
15 when an end user stops, plays, pauses, skips, etc., the streaming media content. Such information can be vital in determining the value of the media.

Because not all media player applications provide this beaconing data, the data gathered by the index file generator can serve as a substitute for or complement to the beaconing data. For example, if a request is made for a chunk that does not immediately
20 follow a previously-requested chunk, a skip was made. If the amount of time elapsed between a previous request and a subsequent request exceeds the time for playback of the previously-requested chunk, a pause was made. If a request is not received within a certain time since a prior request, it can be determined that a stop was made.

As illustrated above, the state of a client may be determined from a variety of factors.
25 This can include when the request for the index file is received, when the index file is provided, a length of time to play back the segment of media for streaming, and/or the starting and/or ending point of the segment of media for streaming. The determined state of a client may also be based on whether the request for the index file has been received within a certain amount of time since receipt of a previous request for an index file, whether the
30 segment of media for streaming includes media other than the media file, and more. The state of a client and/or the data from which it was determined, may be used to create reporting data to serve as a substitute or complement to beaconing data from a client media player application. Because the index file generator can determine the length of the chunks, it therefore can determine the frequency of subsequent index file requests and the resolution of

the reporting data based on the requests. The index file generator may log the reporting data and/or transmit the reporting data over a network during streaming.

The determined state of a client may be used by the index file generator and/or other services for various purposes. For example, it may be used in behavioral advertisement targeting and enforcement of session advertisement behavior, adjusting advertisement content and playback based on the behavior of a user as determined by the stated of a client. The state of a client further may be used to support resume features on a per client basis, allowing a user to continue playback of a media asset from a point at which the user had previously stopped playback. The state of a client also may be used to support individual encryption keys in an encryption scheme and allow the index file generator to return secure URLs (e.g., time expiring or Internet Protocol (IP) allowed) for chunks to support functions such as payment services.

Additionally or alternatively, the tasks of generating the index file and providing a location a requested chunk can be split up, thereby enabling the system to determine which chunks are actually requested. For example, a system may be configured to dynamically create an index file having links to one or more redirectors on the system. These redirectors can be configured to issue the location of the chunk, which can be created dynamically. The redirectors can further determine which chunk is actually requested, thereby enabling, among other things, calculation of Quality of Service (QoS) metrics, an increase the accuracy of reporting data, a decrease the frequency of index file generation if efficient to do so, and the ability to more easily handle keys of an encryption scheme.

While the above embodiments may be implemented in a variety of different systems, some particular embodiments may be implemented as part of a media service system. **FIG. 1** is a block diagram illustrating a media servicing system 100, according to some embodiments of the present invention. The system may deliver media content to the end user device 140 through a network such as the Internet 120. The end user device 140 can be one of any number of devices configured to receive media over the Internet 120, such as a mobile phone, tablet computer, personal computer, portable media device, etc. A media asset provided by a media provider 130 can be processed and indexed by cloud-hosted integrated multi-node pipelining system (CHIMPS) 110, and further stored on media file delivery service provider (MFDSP) 150. Additionally or alternatively, the CHIMPS 110 may also be adapted to store the media asset.

The media servicing system further enables a media provider 130 or other entity to gather information regarding user behavior during media playback. For example, a media provider 130 can be provided with data indicating that end users tend to stop watching a video at a certain point in playback, or that users tended to follow links associated with certain advertisements displayed during playback. With this data, a media provider 130 can adjust factors such as media content, advertisement placement and content, etc., to increase revenue associated with the media content and provide the end user device 140 with a more desirable playback experience.

End user device 140 can request a media asset to stream with a client program executed by the end user device 140. The client program can be, for example, a media player, browser, or other application adapted to request and/or play media assets. In response to a request for a media asset, the CHIMPS 110 can utilize any number of application centers 112 and/or kernel application center(s) 111 to provide the client program with a data object concerning the requested media asset. The data object can include information about the media asset, including where the media asset can be located, such as within the MFDSP 150 or within the CHIMPS 150 itself. Location information may be provided by Universal Resource Indicator (URI), a Universal Resource Locator (URL) or other indicator. During playback of the media asset, the CHIMPS 150 can collect data regarding the playback through beaconing provided by a client program executed by the end user device 140 and/or indexing service from within the CHIMPS and/or MFDSP. The CHIMPS 150 can subsequently provide the data and/or any analytics information derived from the data to the media provider 130.

FIG. 2A is a block diagram illustrating an embodiment of a kernel application 111-1 center connected with application centers from within the CHIMPS 110-1. The kernel application center 111-1 and application centers 112 can be geographically distant and can be connected via the Internet 120, wide area network (WAN), and/or other data communication network. Because application centers can be geographically separated, DNS services (not shown) can be used to allow an end user device 140 to connect to the nearest available application center 112. The kernel application center 111-1 can connect with application centers 112 within the CHIMPS 110-1 through an internal interface 270, thereby enabling the application centers 112 access to the various components within the kernel application center 111-1.

Components within the kernel application center 111-1 can communicate through network 260 such as a local area network (LAN) and can include one or more origin servers

240 and a storage array 230 with which data objects and/or media assets may be stored and distributed. The storage array 230 may also be utilized by services running on processing server(s) 220 and/or transcoding server(s) 250 that may require temporary or long-term storage. Kernel server 210 can utilize processing server(s) 220, transcoding server(s) 250 to
5 provide various functional capabilities to the CHIMPS 110.

For example, as described in more detail below, the CHIMPS 110-1 can provide transcoding service for media assets provided by a media provider 130 for syndication. Such a service can allow a media provider 130 to upload a media asset to an application center 112, after which the application center 112 would notify the kernel server 210 that the media asset
10 has been uploaded. The kernel server can then notify services running on the processing server(s) 220 of the upload. These services can utilize transcoding server(s) to transcode the media asset, which can then be moved to a MFDSP and/or stored locally by storage array 230 and origin server(s) 240. Services running on the processing server(s) 220 can also update the associated data object stored by the storage array 230 and origin server(s) 240.

FIG. 2B is a block diagram illustrating an alternative embodiment of a kernel
15 application center 111-2. In addition to the components of the embodiment of FIG. 2A, this embodiment incorporates an application center 112 within the kernel application center 111-2. The application center 112 incorporated within kernel application center 111-2 may be located at or near the other components of the kernel application center 111-2, and can be
20 communicatively connected to the other components via network 260. The incorporated application center 112 can therefore have faster access to kernel application center functionality because it does not need to communicate over long distances. In consideration of this advantage, it will be understood that the CHIMPS 110 can include multiple kernel centers with one or more application centers incorporated therein. Additionally or
25 alternatively, components of the kernel application center may be incorporated into one or more application centers 112 in the CHIMPS 110 to provide quicker access to certain functionality.

FIG. 3 is a block diagram illustrating an embodiment of an application center 112. The application center 112 can include caching server(s) 330 and a storage array 310 for
30 storing and distributing data objects of media assets requested by end user devices through end user interface 360. Caching server(s) 330 and storage array 310 can also be used to collect, process, and/or store metrics information from beaconing data, media chunk requests, and/or other data sources, including data collected through end user interface 360. The application center can further include ingest server(s) 320 for ingesting uploaded media assets

from a media provider 130 through a media provider interface 370. The media assets may be stored on the storage array 310. As with the kernel application center 111, the components of the application center 112 can be communicatively linked through a network 340, such as a LAN. The application center can further include an internal interface 350, providing a communication link from the application center to the rest of the CHIMPS. It is through internal interface 350, for example, that media assets stored on storage array 310 can be made available to a kernel application center 111 for services such as transcoding.

FIG. 4 is a block diagram 400 of processes and objects utilized by the CHIMPS 110 for media ingestion, according to some embodiments. Although FIG. 4 further indicates the physical systems in which my execute or store these processes and objects, it will be understood that the processes and objects disclosed may be executed or stored on more than one system, including systems not disclosed in FIG. 4. In other words, the processes and objects shown in FIG. 4 allow for a variety of implementations through one or more of hardware, software, firmware, microcode, etc.

Media can be ingested into the CHIMPS 110 when a media provider 130 uploads a media asset to ingestion server(s) 410 in an application center 112 by utilizing a client 405. The client 405 can be a stand-alone application or browser based, for example, and can communicate with ingest server(s) 410 through an application programming interface (API) configured for the ingestion of media assets.

Ingest server(s) 410 can communicate with devices in the kernel application center 111 executing programs such as kernel server 425 and file replication service 430. The kernel server 425 can be configured organize the workflow among services such as transcoding 440 file system manager 435, and other services 445 (e.g., analytics, dynamic API, etc.) Upon a particular event, for example, the kernel server can be configured to notify the relevant services of the event, causing the services to process tasks associated with the event.

The file replication service 430, under direction of the kernel server 425, can coordinate the movement of the media assets between services. For example, retrieving the uploaded media asset from the ingest server(s) 410 and storing it on the file archive 450, or retrieving transcoded media assets from transcoding server(s) 460 and storing them in the media asset origin.

The data object updater 420 keeps the data object origin 415 up to date in response to any changes in the system. When, for example, a file is uploaded, transcoded, and stored in

media asset origin 455, the location and other metadata concerning the transcoded media assets need to be created or updated in the data object origin 415 to ensure an end user device that accesses the object in the data object origin 415 has the correct information regarding the related media asset. Because the data object updater 420 receives updates from the kernel server 425 (which is notified when a transcoded media asset is stored in the media asset origin 455, the system ensures the data objects in the data object origin are constantly up to date.

The upload of a media asset to the ingest server(s) 410, as described above, can provide an example of how the kernel server 425 may coordinate workflow. For instance, in response to the upload, the ingest server(s) 410 can notify the kernel server 425 that a media asset has been uploaded. The kernel server 425 informs the file replication service 430 of the uploaded media asset, and the file replication service 430 moves the uploaded media asset into the file archive 450 and notifies the kernel server 425 of the move. In response, the kernel server 425 notifies the file replication service 430, the file system manager 435, and the transcoding master 440 of the move. The file replication service 430 then will know it can delete the uploaded media asset from the ingest server(s) 410, the file system manager 435 will update the file system accordingly, and the transcoding master 440 will notify transcoding service(s) 460 of different transcoding tasks to be performed. The transcoding service(s) 460 can then retrieve the uploaded media asset from the file archive 450 to create transcoded media assets. The transcoding service(s) 460 notify the kernel server 425 once transcoding is complete, and the kernel server 425 relays this information to the file replication service 430. The file replication service 425 then takes the transcoded media assets from the transcoding services 460 and moves them to the media asset origin 455. Once the file replication service 430 notifies the kernel server 425 of the move, the kernel server 425, in turn, notifies the file replication service 430 and the data object updater 420. The data object updater 420 which updates the data object origin 415 accordingly, and the file replication service 430 deletes the transcoded media assets from the transcoding services 460.

The modular nature of the system enables all tasks associated with an event to be completed quickly. As illustrated in the example above, workflow relating to a particular event, such as a media asset upload, can be spread among the various services simultaneously. Moreover, because the system's modularity enables it to be scaled to accommodate differing hardware capacities, and because the system can be configured to dynamically allocate hardware to different services according to the needs of the system, the speed of completing tasks relating to a particular event can further be increased. For

example, a server of the CHIMPS 110 can be configured to dynamically switch its purpose based on external conditions such as load and overall system performance, providing functions such as transcode, upload, metrics collection, application web service, and more, on an as-needed basis.

5 Embodiments of such systems may include other systems that manage various requests from end users. For example, a system for dynamic index file generation and media file chunking. Referring to **FIG. 5A**, shows an embodiment of such a system 500-1. Media may be streamed to end user device 140 through a client 510. As mentioned above, the client 510 can be stand-alone media player, a plug-in, a browser, or other application, which can be
10 executed on a personal computer or other electronic device.

 An index file generator 530, as discussed previously, can be a program instantiated for media streaming to a particular client 510. The index file generator 530 can be executed on a server or other computing device within an application center 112 of the CHIMPS 110. Index files generated by the index file generator can include a wide variety of information
15 such as starting, ending, and or run times for media chunks and locations for media chunks. This information can be embedded in a single string of data, such as a URI or a URL. If media includes various sub-streams (e.g., streams with alternative bitrates, captions, alternative languages, etc.) the index file can include data for chunks corresponding to each of the alternative sub-streams, as well as information regarding the bitrate and/or other unique
20 information for each stream. Alternatively or in addition, index files indicating alternative sub-streams may be separate from index files indicating one or more media chunks for streaming.

 It should be understood that the index file can further comprise a wide variety of formats, which can depend on the particular protocol. HTTP streaming may, for example,
25 require index files to comprise one or more of M3U, M3U8, XML, and XML-based formats. Of course, other formats can be used in accordance with relevant streaming protocols.

Table 1 illustrates a simplified example of a generated index file in M3U9 format, indicating chunk of media for streaming. The index file in this example provides a URI for a chunk of media. The URI indicates the chunk is to be generated by dynamic segmentor 550,
30 the chunk being 10 seconds long, starting at 9 seconds into the media file and ending 19 seconds into the media file.

Table 1: Example Index File Contents

<pre>#EXTM3U #EXT-X-MEDIA-SEQUENCE:1 #EXT-X-TARGETDURATION:10 #EXTINF:10, http://video.example.com/seg/9/19/seg1.ts</pre>

Referring again to FIG. 5A, the index file generator 530 can also include an indicator within an index file to indicate whether a chunk of media is to be dynamically created. If, for example, it is determined that a requested media asset has not been chunked and that the asset will be chunked dynamically, the index file generator can include the indicator in data
5 corresponding to a chunk of media to be created. The indicator, which can be as simple as including the term “/seg/” in a URL, will indicate that a requested chunk of media needs to be generated.

The chunks of media can be generated during media streaming by a dynamic
10 segmentor 550, which can be incorporated into an HTTP service 540. The HTTP service 540, as well as the media asset origin 560 can be located within a kernel application center 111 of the CHIMPS 110 on, for example, a media asset origin server. The system 500-1 can be configured such that the kernel application center 111 provides dynamically-created chunks of media to a MFDSP 150 for delivery to client 510. The MFDSP 150 can store the
15 chunks locally in, for example, a media asset cache 520, thereby forgoing the need to dynamically create a chunk again if the same chunk is requested in the future.

In sum, the system for dynamic index file generation and media asset chunking 500-1 can, after receiving a request for an index file from a client 510, dynamically generate an index file with an index file generator 530. The index file can, among other things, indicate
20 where a next chunk of media may be located. A client can then request the chunk from the location indicated by the index file, which can comprise a media asset cache 520 in a MFDSP 150. If the chunk is not found in the media asset cache 520, the cache miss can redirect the request to a segmentor 550 of an HTTP service 540, which can dynamically generate the requested chunk of media by accessing the corresponding media asset in the media asset
25 origin 560. The requested media chunk can then be provided to the MFDSP 150 for storage in the media asset cache 520 and delivery to the client 510. If the same chunk is requested at a later point in time, the MFDSP 150 can deliver the chunk from the media asset cache 520, thereby forgoing the need to redirect the request to the segmentor 550 to regenerate the chunk.

FIG. 5B illustrates an alternative embodiment 500-2 of a system for dynamic index file generation and media file chunking. Rather than utilize a MFDSP, this embodiment 500-2 includes a media caching server within an application center 112 of the CHIMPS 110. The media caching server can receive chunk requests from and provide the corresponding chunks to a client. It will be understood that such a media caching server(s) or similar device(s) can be located anywhere within the CHIMPS and/or in a system(s) communicatively linked to the CHIMPS.

FIG. 5C illustrates an embodiment 500-3 of a system for index file generation used in conjunction with a redirector 590. As discussed above, an index file generator 530 may create an index file having one or more URIs or other location information directing the client 510 to one or more redirectors 590. Redirector 590, can then provide the URI of the requested chunk with a redirect, such as a redirect under HTTP status code 302. The URI can be located on a MFDSP 520 or other location (such as a media caching server 570) and/or dynamically created by the dynamic segmentor 550. It will be understood that there can be any number of redirectors 590, which can be at any location, including locations of the CHIMPS 110 such as the application center 112 (as shown in FIG. 5C) or the kernel application center 111. It also will be understood that the URI or other location information provided by redirector 590 can be generated by or provided to the redirector 590 in any number of ways. The URI can be generated, for example, based on the request received by redirector 590 from client 510. Finally, it will be understood that the CHIMPS 110 can be configured to dynamically implement any combination of embodiments 500-1, 500-2, and 500-3, further choosing whether to utilize one or more redirectors based on factors such as, for example, a detected type of client 510.

Embodiments utilizing one or more redirectors can have several advantages. For example, and not by way of limitation, if a certain client were implemented in such a way that it “reads ahead” to request chunks, it could result in incorrect reporting data. Thus, it would be advantageous to determine which chunk is actually requested by the client. Additionally or alternatively, where chunks are available in various sub-streams with different bitrates, determining the actual requested chunk can be useful in calculating Quality of Service (QoS) metrics. Furthermore, it there may be scenarios in which it is more efficient to create larger index files having many chunks comprising large segments of media, reducing the number of index files required to stream a media asset, and thereby reducing the processing requirements to create the index files. If encryption is used having, for example, a rotating key or a per client key encryption scheme in which a valid key might change during

playback of a media asset, it also may be advantageous to incorporate redirector(s) for handling legacy keys for some period of time.

FIG. 6 illustrates a simplified flowchart of an embodiment of a method 600 for implementing a dynamic index for media streaming. The method 600, which can be executed
5 by the index file generator 530, begins at block 610, where a request for an index file is received from a client 510. According to media streaming protocols contemplated by this embodiment, if data regarding a chunk of media is not provided in an initial index file, the client will continue to request or refresh an index file until it reaches an indicator in the index file that signals the end of a stream. Thus, the method can be assured to receive more than
10 one request for an index file from a client provided that an initial index file does not include an indicator signaling the end of the stream.

At block 615, the method 600 additionally provides for receiving input from an advertising service. According to some embodiments, this input could be the availability of an advertisement, and can be provided by a service inside or outside the CHIMPS. In other
15 embodiments, the input could come from a service that factors in any of a variety of factors to indicate that a specific advertisement or type of advertisement should be shown. Or that any advertisement should be shown.

At block 620, a determination is made whether to include an advertisement in the next chunk. According to some embodiments, this determination can be made with or without
20 input from an advertisement service. It should be known that this determination can include the factors used by an advertisement service to provide the input of block 615. Whether the determination includes input from an advertisement service of block 615 or not, the determination can still include factors such as information about an end user collected before or during streaming of the media. This can include behavior of the end user during streaming
25 of the media (as determined, for example, by machine-based logic through beaconing data and/or requested chunks of media provided by a client 510). Factors can also include information regarding the media asset used for streaming (such as type of content or preferred points within the media for an advertisement), preference(s) and/or selection(s) of an end user, when a previous advertisement was shown, time of day, and more. It can further
30 include information regarding the source of a media asset, such as who created and/or provided the asset for viewing by an end user. It will be understood that other embodiments contemplate include secondary media, other than advertisements into the media stream in this manner. Moreover, the secondary media and/or advertisement can be of any length and also

may be chunked. Thus, it may be determined that the next chunk includes all, or a select portion, of an advertisement of any specific length.

An index file is created based on the request as well as the determination of whether media, such as an advertisement, should be streamed, indicated by block 625. As discussed above, the index file can assume a variety of formats and include any amount of information regarding a next chunk of media for streaming. For example, HTTP streaming can utilize index files having the URLs of available chunks of media. Information can be embedded in these URLs to indicate a location to download the corresponding chunk of media, starting point and/or ending point of a chunk of media, an indicator to indicate whether the chunk is to be dynamically created by a segmentor 550, a location of an advertisement to be streamed, and more. This information is included in the index file and sent to the client at block 630.

At block 635, reporting data can be created based on information included in the index file. As previously discussed, information included in an index file and/or index file request can indicate the behavior of an end user device 140 during streaming, such as a pause, stop, skip, play, etc. of the media. According to some embodiments, this information can be extracted from requests for an index file and/or providing the requested index file. The information can be gathered in addition to or as a substitute for beaconing data provided by a client 510. Moreover, if beaconing data is provided, the creation of reporting data may be omitted altogether.

Reporting data can include any amount of information regarding end user behavior, as indicated through index file requests and/or provided index files. This can include a particular action and when it was performed. Additionally or alternatively, the data may be kept in a more rudimentary form, depending on the application or embodiment, indicating the data included in a index file request and/or an index file. This reporting data may be stored in a log file for reporting after streaming and/or transmitted during streaming to a relevant service that collects such metrics.

As indicated by block 640, the reporting data may be sent to a metrics collector for analytics. A metrics collector, according to certain embodiments, may be an application executed by a server from within the application center 112 in which the index file generator 530 is executed, or it may be executed elsewhere, such as in a kernel application center 111 or in a system outside the CHIMPS 110. Depending on the form of the reporting data, the metrics collector can further process and store the information.

FIG. 7 illustrates a simplified flowchart of an embodiment of a method for dynamically chunking a media file for streaming 700. This method can be employed by a variety of systems and/or programs. For example, it may be executed by a dynamic segmentor 550 of an HTTP service 540 running on a server located in a kernel application center 111 of the CHIMPS 110.

The method 700 can begin at block 710, when a request for a chunk of media is received from a MFDSP 150. As discussed above, this request may be made in response to a cache miss at the MFDSP 150 and/or because an indicator was included in the request for the chunk of media that the chunk was to be created dynamically. As discussed herein, if the MFDSP 150 has the requested chunk cached from a prior request, the MFDSP 150 can provide the requested chunk and preclude the need to send the request to a dynamic segmentor 550 to generate the chunk. It should be understood that the request may come from sources other than a MFDSP 150 according to alternative embodiments. One such source includes the media caching server 570 of embodiment 500-2, as shown in FIG. 5B.

The starting and ending points of a requested chunk of media are then determined at block 715. This information can be included directly in the request or derived from the request, a previous request, and/or other sources. At block 720, the information, as well as information identifying the requested chunk of media, can be used to retrieve all or part of the relevant media asset from a media asset origin 560. The retrieved portion will include at least the relevant media from the starting point to the ending point of the requested chunk of media.

At block 725, the requested media chunk is generated by converting the relevant portion of the media asset into a deliverable chunk. The media asset, as stored in the media asset origin, may not be chunked; it may be stored in its entirety as a media file (or group of alternative files corresponding to alternative sub-streams). Generating the chunk therefore can require determining the starting and ending points from the retrieved portion of the media asset and converting the resulting segment of media into a deliverable chunk.

Although the generation of the deliverable chunk may involve transcoding, it may not. The media asset can be stored in a format where transcoding may not be needed, thereby reducing the processing requirements for creating chunks of media during streaming. For example, media assets may be stored such as H.264 or MPEG-4 video format and/or AAC, HE-AAC, or MP3 audio format. According to some streaming protocols, such as some forms of HTTP streaming, chunks of media in these formats would not need transcoding before

being wrapped in an MPEG-2 transport stream container format. Instead, such a conversion essentially would require the addition of metadata to create the streaming format from the format of the stored media asset. In other words, generating a deliverable chunk of media may only require identifying the stored media asset, extracting the relevant segment of the media from the media asset, and adding certain metadata in accordance with a container format. This process requires little processing power and can be easily performed on the fly during streaming. Once the deliverable chunk of media is generated, it is sent to the MF DSP 150 or other requesting entity, at block 730.

FIG. 8 illustrates a simplified swim lane flowchart describing the interaction of components in a system configured to provide dynamic indexing and chunking for media streaming, according to one embodiment. In this embodiment, a client can send a request for an index file 805, the request received by an index file generator 810. A particular request may be made to initiate the streaming of a media asset, or it may be during streaming. Depending on the streaming protocol, the request may be made while a client plays a chunk of media previously downloaded during streaming.

The index file generator generates an index file to indicate the next chunk of media 815. As described above, this chunk may include an advertisement, and the index file can include any amount of information about a chunk of media, including information regarding alternative sub-streams for streaming. The dynamic index file generator can include information regarding existing chunks of media, and, when used in conjunction with a dynamic segmentor may also include information regarding chunks of media that may need to be created. As detailed above, if a chunk of media is to be generated dynamically, the index file generator may indicate this by including an indicator in the generated index file, such as in a URL for one or more chunks described within the index file. Once the index file is generated, the index file generator sends the index file 820, which is received by the client 825.

Alternative embodiments may provide for the generation of index files containing more than a next chunk of media. For example, an index file generator may generate an index file containing information regarding several chunks of media, in which case the chunks of media can be dynamically generated by a dynamic segmentor when requested by the client. The determination of whether to include information regarding more than a next chunk of media can include factors such as whether the index generator is generating reporting data, the desired frequency of such reporting data, and more.

Using information contained in the index file, the client can then request the next chunk of media 830, and this request can be received by a MFDSP 835. The MFDSP then checks to see if the chunk is already stored in the cache 840. If so, the MFDSP can provide the requested chunk to the client, blocks 845 and 850. The requested chunk may be found in
5 a MFDSP's cache if the chunk was created and stored in the MFDSP during preprocessing or if the chunk was dynamically created and stored in the MFDSP from an earlier request.

If the chunk is not found on the MFDSP, the chunk can be requested of the dynamic segmentor, which receives the request 855 and retrieves the corresponding media asset from an asset origin server 860. As discussed above, the entirety of the relevant media asset does
10 not need to be retrieved as long as at least the portion containing the relevant segment for the requested chunk is retrieved. It will be understood that alternative embodiments can provide for the media asset being stored in a variety of locations accessible, directly or indirectly, to the dynamic segmentor.

The dynamic segmentor can then generate the requested chunk by converting the
15 retrieved media into a deliverable chunk. That is, the dynamic segmentor converts the retrieved media into an acceptable format for streaming, which can vary depending on the streaming protocol utilized. The dynamic segmentor can then return the requested chunk 870 to the MFDSP, which can cache the chunk and return it to the client 875. Once the chunk is received by the client 880, the client can play the chunk to an end user.

20 The techniques for creating index files and/or creating corresponding chunks of media can be adapted for each request. This allows the CHIMPS 110 (or similarly-enabled system) to provide optimal chunking for each delivery instance, which can differ even among similar devices, depending on any of a variety of factors. For example, a new version of a particular device type (e.g., a particular smart phone, set-top box, tablet, etc.) may have the same
25 operating system and/or browser as an older version of the device type, but with a larger buffer than the older version. As such, optimized chunking for delivery instances involving the old and new versions of the device type may differ in that delivery to the new version can include using chunk sizes that would cause a buffer overflow in the older version. By dynamically determining and implementing a chunking strategy for each delivery instance in
30 this manner, the CHIMPS 110 (or similarly-enabled system) can help provide the optimal user experience under any of a variety of circumstances.

FIG. 9 is a simplified flow chart illustrating a basic method 900 for adapting chunking techniques for each delivery instance. At block 910, a request for a media file is

received. In other embodiments, the request may be for a live stream or other media asset not stored as a media file. At block 920, one or more chunking considerations are determined. Because the streaming of a media file can involve multiple requests (and, correspondingly, multiple index files), the method 900 may be performed multiple times during the playback
5 of a media file.

Chunking considerations are factors that can impact any aspect of the chunking of a media file and/or advertisements included in the delivery of the media file. Such considerations can include, but are not limited to, bandwidth, codec, processor type, buffer size, memory and/or processor utilization, battery, network type, geographic location,
10 whether a device is moving, a location in the playback of the media file, and the like.

The bandwidth of a device, for example, can impact the size of the chunks delivered to the device. For a relatively low bandwidth, the chunk size may be reduced so that chunks are downloaded more quickly, which may reduce the chance that a media player would need to interrupt playback to while waiting for the next chunk to be delivered. Conversely, chunk
15 size can be increased for higher bandwidths. Similarly, the chunks may be adapted to accommodate a particular network type (e.g., WiFi, mobile wireless network, land line (wired), etc.), which can determine, or be indicative of, an available bandwidth.

Different aspects of the hardware of a device can also be considered when chunking a media file. In addition to the buffer size, as indicated above, a processor type can also impact
20 how a file is chunked—such as the size of the chunk and/or the codec used. Additionally or alternatively, QoS metrics can indicate memory and/or processor utilization, including during playback of the requested media file, which can impact the chunking of the media file.

Chunking considerations further can include a variety of factors related to mobile devices. For example, whether a device is situated in a particular geographic location, at a
25 particular event (where large numbers of mobile devices may be located), and/or the device is moving can impact the available bandwidth to the device. The battery level of a mobile device may also impact the bandwidth in certain circumstances. The CHIMPS 110 (or other chunking system) can anticipate changes in bandwidth by determining such chunking considerations and adjusting the chunking accordingly. Moreover, QoS, bandwidth, and/or
30 other information from multiple devices can be aggregated to help the CHIMPS 110 further adapt media chunking for related delivery instances. Thus, the chunking considerations for streaming to one device may include information regarding one or more other devices. For example the CHIMPS 110 may anticipate a sudden reduced bandwidth for a particular

delivery instance related to a cell phone where it is determined that other cell phones nearby on the same carrier have had sudden drops in bandwidth.

Chunking considerations further can include a type of codec used and/or a location in the playback of the media file. Certain codecs, for example, may require or prohibit the use of certain-sized chunks. Also, for example, if the playback of the media file has just started, smaller chunks may be used to help ensure that playback begins quickly, without the need to wait for larger chunks to be delivered. Also, chunk sizes can be altered to accommodate ad insertion or similar interleaving of media content at particular point(s) in the playback of the media file.

Chunking considerations may be determined from information included in the request, a database, and/or other sources, which may disclose data such as a URL, client ID, globally-unique identifier (GUID) or other identifier, a network type, a device type and/or capability, an operating system executed by the end user device, an application, application identifier, application state, or other application information, a location associated with the device, information associated with a user of the end user device 140, information regarding the requested media file (e.g., genre, length, rating(s), ownership, artist, etc.). Additionally or alternatively, the request may simply include information that enables one or more of these items to be determined. Additionally or alternatively, a repository may be stored, maintained or derived, and queried for authorization, authentication, validation, or selection; for example, a repository of application identifiers may be maintained and queried to determine whether an application is authorized to request the content and if so, to select further aspects of or for processing the content request. Additionally or alternatively, such stored or derived repository data may be used in conjunction with other data, either internally or externally identified, such as a secret key, shared key, public key, stored certificate, other stored data, or other data for authorization, authentication, validation, or selection, including data stored on another digital service, on another server, on the client device, in a device associated with the client device, in the operating system, in the application, in another application, in a network, or in another location from which it may be retrieved.

As indicated previously, the determination of chunking considerations can include utilizing information other than the information provided in the request. This may involve accessing information stored in one or more a databases or other data structures internal or external to the CHIMPS 110. It may also involve communicating with other entities and/or systems, such as a content owner or media provider 130. Additionally or alternatively,

chunking considerations can be gathered using data independent of information provided in the request, such as the time at which the request was received.

At block 930, a chunking strategy based on chunking considerations is determined. And at block 940, the corresponding index file and chunks are provided accordingly. As indicated above, the chunking strategy (i.e., the way in which the chunks are created and delivered) can be impacted by the chunking considerations, including the size of the chunks, the length of the playback of the chunks, and the choice of codec used. Furthermore, multiple considerations can be taken into account to determine the chunking strategy.

This chunking strategy can be implemented on the fly (i.e., immediately before or during playback of the media file) utilizing the dynamic indexing and chunking techniques discussed above. For example, an index file generator 530 can incorporate the chunking strategy into a dynamically-generated index file, and the dynamic segmentor 550 can produce the requested chunks accordingly. Because the contextual data and chunking strategies can vary for each delivery instance, the content of the index files corresponding to requests for the same media file can be different (e.g., indicate different chunk sizes, codecs to use, etc.), based on differing chunking strategies.

Historical data can also be utilized when determining a chunking strategy. Historical data can be based on any of a variety of previously-obtained information regarding, for example QoS, network conditions, client type, previous chunking strategies, content consumption analytics, and the like. For instance, QoS data regarding previous chunking strategies can be used to determine whether previous chunking strategies were successful. If so, the same or similar chunking strategy can be used again. Otherwise, a different chunking strategy can be used.

On particular type of information that can be useful when analyzing historical data is QoS information. QoS information -- from one or multiple devices -- can indicate memory and/or processor utilization, including during playback of a requested media file, which can be used to adjust the chunking of the media file. Additionally or alternatively, information from entities such as QoS providers can shed light on the type of client making the request for the media file, which can be taken into consideration when determining a chunking strategy. In some cases, content distribution network providers contract with QoS providers to deliver region-specific data indicating delivery performance and buffering over their tiered distribution architecture. This data can be used to tailor chunk length and boundaries within

the content to optimize on a per content delivery network, per region basis the ability to cache the content, deliver the content bytes, decrease time to first byte, and the like.

FIG. 10 is a simplified flow chart illustrating a basic method 1000 for adapting chunking techniques for each delivery instance based on historical data. As with other figures provided herein, FIG. 10 is provided as an example method for illustrative purposes. Alternative embodiments may implement similar methods by combining, separating, omitting, and/or adding to the blocks illustrated in FIG. 10. Additionally or alternatively, the functions illustrated in some or all of the blocks may be performed in a different order and/or simultaneously, depending on desired functionality. Moreover, as with other methods described herein, some or all of the method 1000 illustrated in FIG. 10 may be performed by one or more computers. Such computers can, for example, form part of the CHIMPS 110, Index File Generator 530, and/or other systems described herein above. A person of ordinary skill in the art will recognize many implementations and/or variations.

At block 1010, historical data related to streaming media with the data network is determined. The historical data can include, for example, data obtained from streaming one or more media files over time. Additionally or alternatively, the historical data can include information regarding network conditions, QoS, client type, previous chunking strategies, and the like, which may or may not be associated with the streaming of media files. Furthermore, as indicated previously, QoS information can be obtained from other entities, such as a QoS provider, and may include information regarding a type of client making the request.

Blocks 1020-1050 echo blocks 910-940 of FIG. 9, as described above. At block 1030, however, the historical data is included as one of the one or more factors related to the request which used, at block 1040, to provide at least part of the basis with which the chunking strategy is determined. In this manner, the method 1000 of FIG. 10 can allow the successes and/or failures of prior chunking strategies to improve chunking strategies going forward, ultimately improving QoS and/or other factors over time.

It should be noted that the methods, systems, and devices discussed above are intended merely to be examples. It must be stressed that various embodiments may omit, substitute, or add various procedures or components as appropriate. For instance, it should be appreciated that, in alternative embodiments, the methods may be performed in an order different from that described, and that various steps may be added, omitted, or combined. Also, features described with respect to certain embodiments may be combined in various other embodiments. Different aspects and elements of the embodiments may be combined in

a similar manner. Also, it should be emphasized that technology evolves and, thus, many of the elements are examples and should not be interpreted to limit the scope of the invention.

Specific details are given in the description to provide a thorough understanding of the embodiments. However, it will be understood by one of ordinary skill in the art that the
5 embodiments may be practiced without these specific details. For example, well-known circuits, processes, algorithms, structures, and techniques have been shown without unnecessary detail in order to avoid obscuring the embodiments. This description provides example embodiments only, and is not intended to limit the scope, applicability, or configuration of the invention. Rather, the preceding description of the embodiments will
10 provide those skilled in the art with an enabling description for implementing embodiments of the invention. Various changes may be made in the function and arrangement of elements without departing from the spirit and scope of the invention.

Also, it is noted that the embodiments may be described as a process which is depicted as a flow diagram or block diagram. Although each may describe the operations as
15 a sequential process, many of the operations can be performed in parallel or concurrently. In addition, the order of the operations may be rearranged. A process may have additional steps not included in the figure. Furthermore, embodiments of the methods may be implemented by hardware, software, firmware, middleware, microcode, hardware description languages, or any combination thereof. When implemented in software, firmware, middleware, or
20 microcode, the program code or code segments to perform the necessary tasks may be stored in a non-volatile computer-readable medium such as a storage medium. Processors may perform the necessary tasks.

Having described several embodiments, it will be recognized by those of skill in the art that various modifications, alternative constructions, and equivalents may be used without
25 departing from the spirit of the invention. For example, the above elements may merely be a component of a larger system, wherein other rules may take precedence over or otherwise modify the application of the invention. Also, a number of steps may be undertaken before, during, or after the above elements are considered. Accordingly, the above description should not be taken as limiting the scope of the invention.

WHAT IS CLAIMED IS:

1. A method of providing media with a data network, the method comprising:
 - determining historical data related to streaming media with the data network;
 - receiving a request having a universal source locator (URL), wherein the URL includes information indicative of a requested media file;
 - determining one or more factors related to the request, wherein the one or more factors related to the request include the historical data;
 - determining, with a processor, a chunking strategy based on the one or more factors related to the request;
 - generating an index file having information for streaming the requested media file via the data network, wherein generating the index file is based, at least in part, on the chunking strategy; and
 - providing the index file.
2. The method of providing media with the data network as recited in claim 1, wherein the historical data includes Quality of Service (QoS) information regarding a previously-streamed media file.
3. The method of providing media with the data network as recited in claim 2, further comprising obtaining the QoS information from a QoS provider.
4. The method of providing media with the data network as recited in claim 2, wherein the QoS information comprises information regarding a type of client making the request.
5. The method of providing media with the data network as recited in claim 2, wherein the QoS information is based on information obtained from streaming media to a plurality of devices.
6. The method of providing media with the data network as recited in claim 1, wherein the chunking strategy includes one or more of:
 - a size of a chunk for streaming the requested media file,
 - a length of playback of a chunk for streaming the requested media file, or
 - a codec to use for streaming the requested media file.

7. The method of providing media with the data network as recited in claim 1, wherein generating the index file occurs during playback of the requested media file.

8. A server configured to provide media via data network, the server comprising:

a network interface for communicating with the data network;

a memory; and

a processor communicatively coupled with the memory and the network interface, the processor further configured to:

determine historical data related to streaming media with the data network;

receive a request having a universal source locator (URL), wherein the URL includes information indicative of a requested media file;

determine one or more factors related to the request, wherein the one or more factors related to the request include the historical data;

determine a chunking strategy based on the one or more factors related to the request;

generate an index file having information for streaming the requested media file via the data network, wherein generating the index file is based, at least in part, on the chunking strategy; and

provide the index file.

9. The server configured to provide media via the data network as recited in claim 8, wherein the processor is configured to determine the historical data by obtaining Quality of Service (QoS) information regarding a previously-streamed media file.

10. The server configured to provide media via the data network as recited in claim 9, wherein the processor is further configured to obtain the QoS information from a QoS provider.

11. The server configured to provide media via the data network as recited in claim 9, wherein:

the QoS information comprises information regarding a type of client making the request; and

the processor is further configured to process the QoS information to determine the historical data.

12. The server configured to provide media via the data network as recited in claim 9, wherein the QoS information is based on information obtained from streaming media to a plurality of devices.

13. The server configured to provide media file via the data network as recited in claim 8, wherein the processor is configured to determine the chunking strategy by determining one or more of:

- a size of a chunk for streaming the requested media file,
- a length of playback of a chunk for streaming the requested media file, or
- a codec to use for streaming the requested media file.

14. The server configured to provide media file via the data network as recited in claim 8, wherein the processor is configured to generate the index file during playback of the requested media file.

15. A non-transitory computer-readable medium having instructions imbedded thereon for providing media with a data network, wherein the instructions, when executed by one or more computers, cause the one or more computers to:

- determine historical data related to streaming media with the data network;
- receive a request having a universal source locator (URL), wherein the URL includes information indicative of a requested media file;
- determine one or more factors related to the request, wherein the one or more factors related to the request include the historical data;
- determine a chunking strategy based on the one or more factors related to the request;
- generate an index file having information for streaming the requested media file via the data network, wherein generating the index file is based, at least in part, on the chunking strategy; and
- provide the index file.

16. The non-transitory computer-readable medium as recited in claim 15, further comprising instructions for causing the one or more computers to determine the historical data by obtaining Quality of Service (QoS) information regarding a previously-streamed media file.

17. The non-transitory computer-readable medium as recited in claim 16, further comprising instructions for causing the one or more computers to obtain the QoS information from a QoS provider.

18. The non-transitory computer-readable medium as recited in claim 16, wherein:

the QoS information comprises information regarding a type of client making the request; and

the instructions further include instructions for causing the one or more computers to process the QoS information to determine the historical data.

19. The non-transitory computer-readable medium as recited in claim 16, wherein the QoS information is based on information obtained from streaming media to a plurality of devices.

20. The non-transitory computer-readable medium as recited in claim 15, wherein the wherein instructions for causing the one or more computers to determine the chunking strategy include instructions for determining one or more of:

a size of a chunk for streaming the requested media file,

a length of playback of a chunk for streaming the requested media file, or

a codec to use for streaming the requested media file.

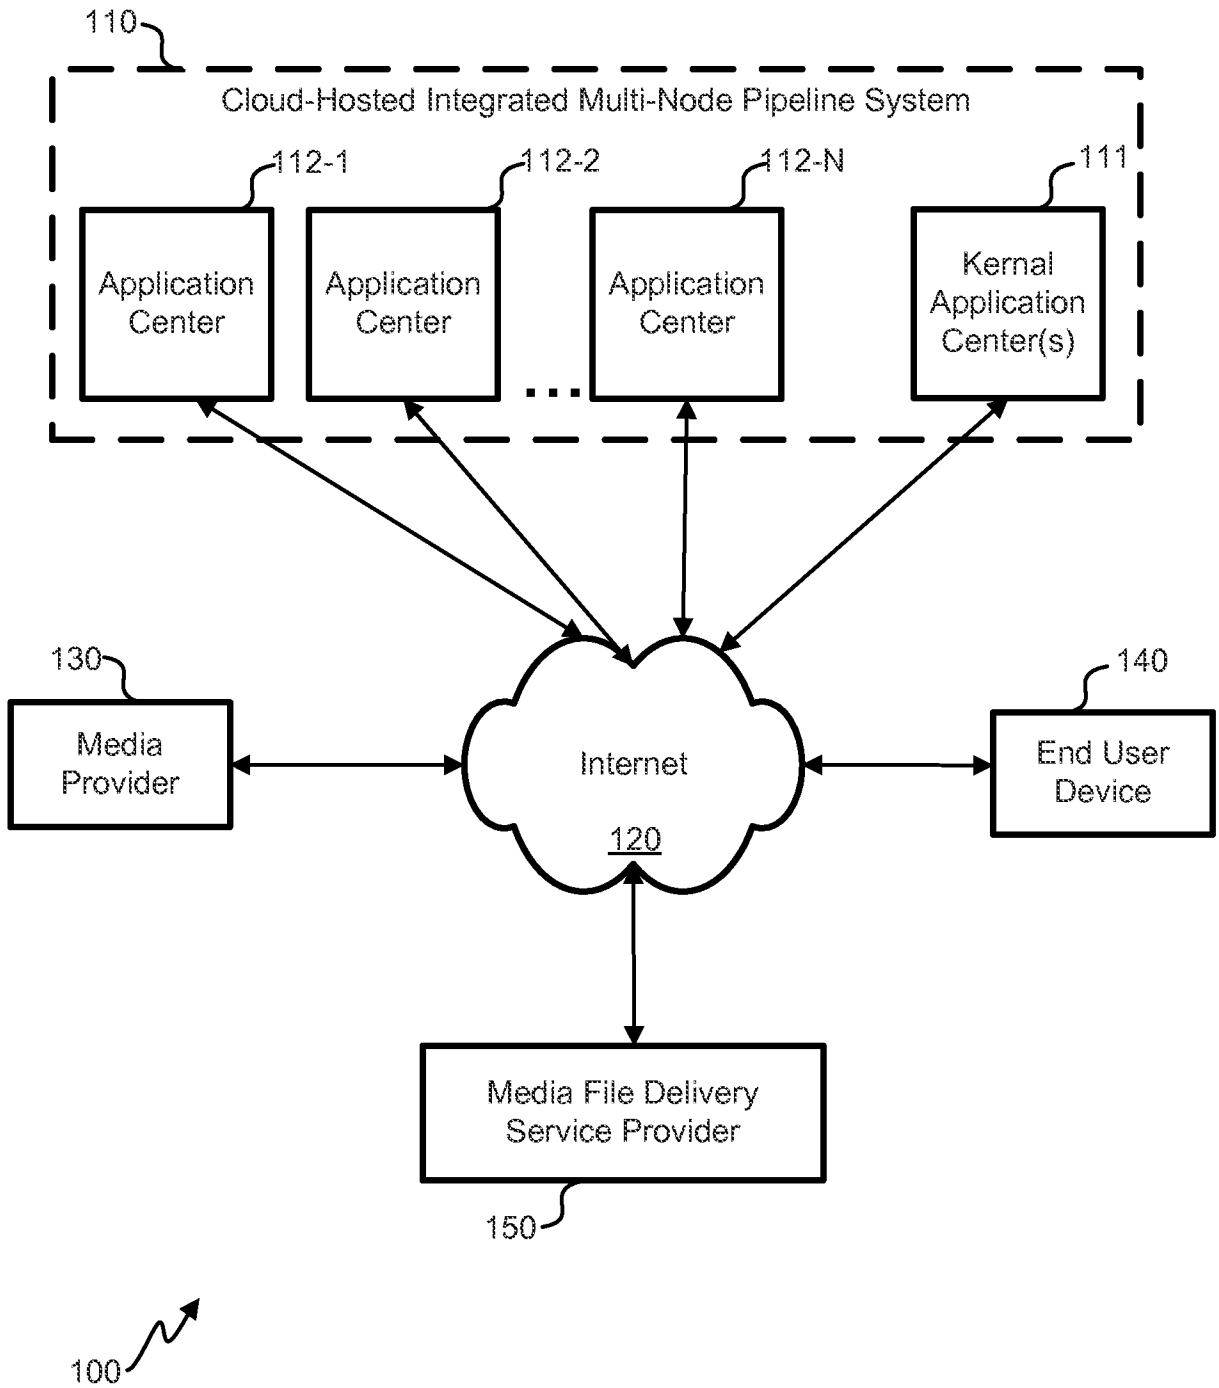


FIG. 1

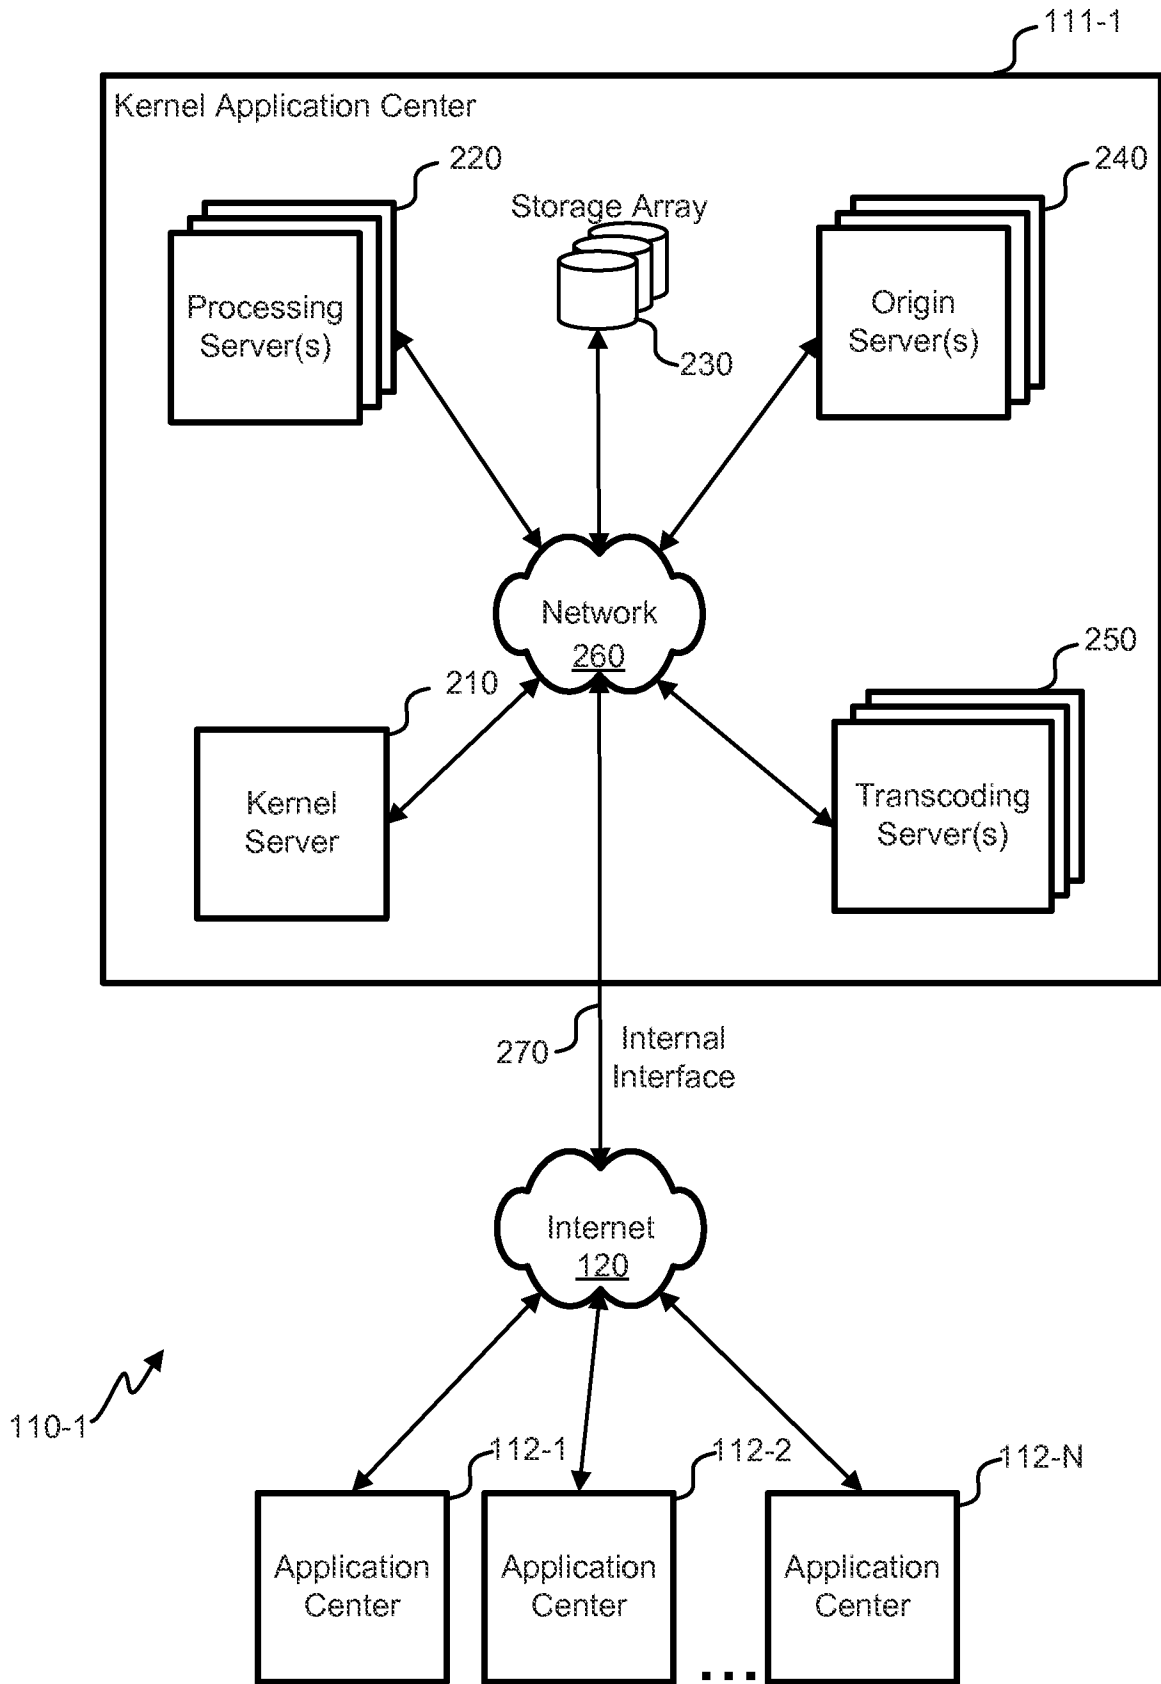


FIG. 2A

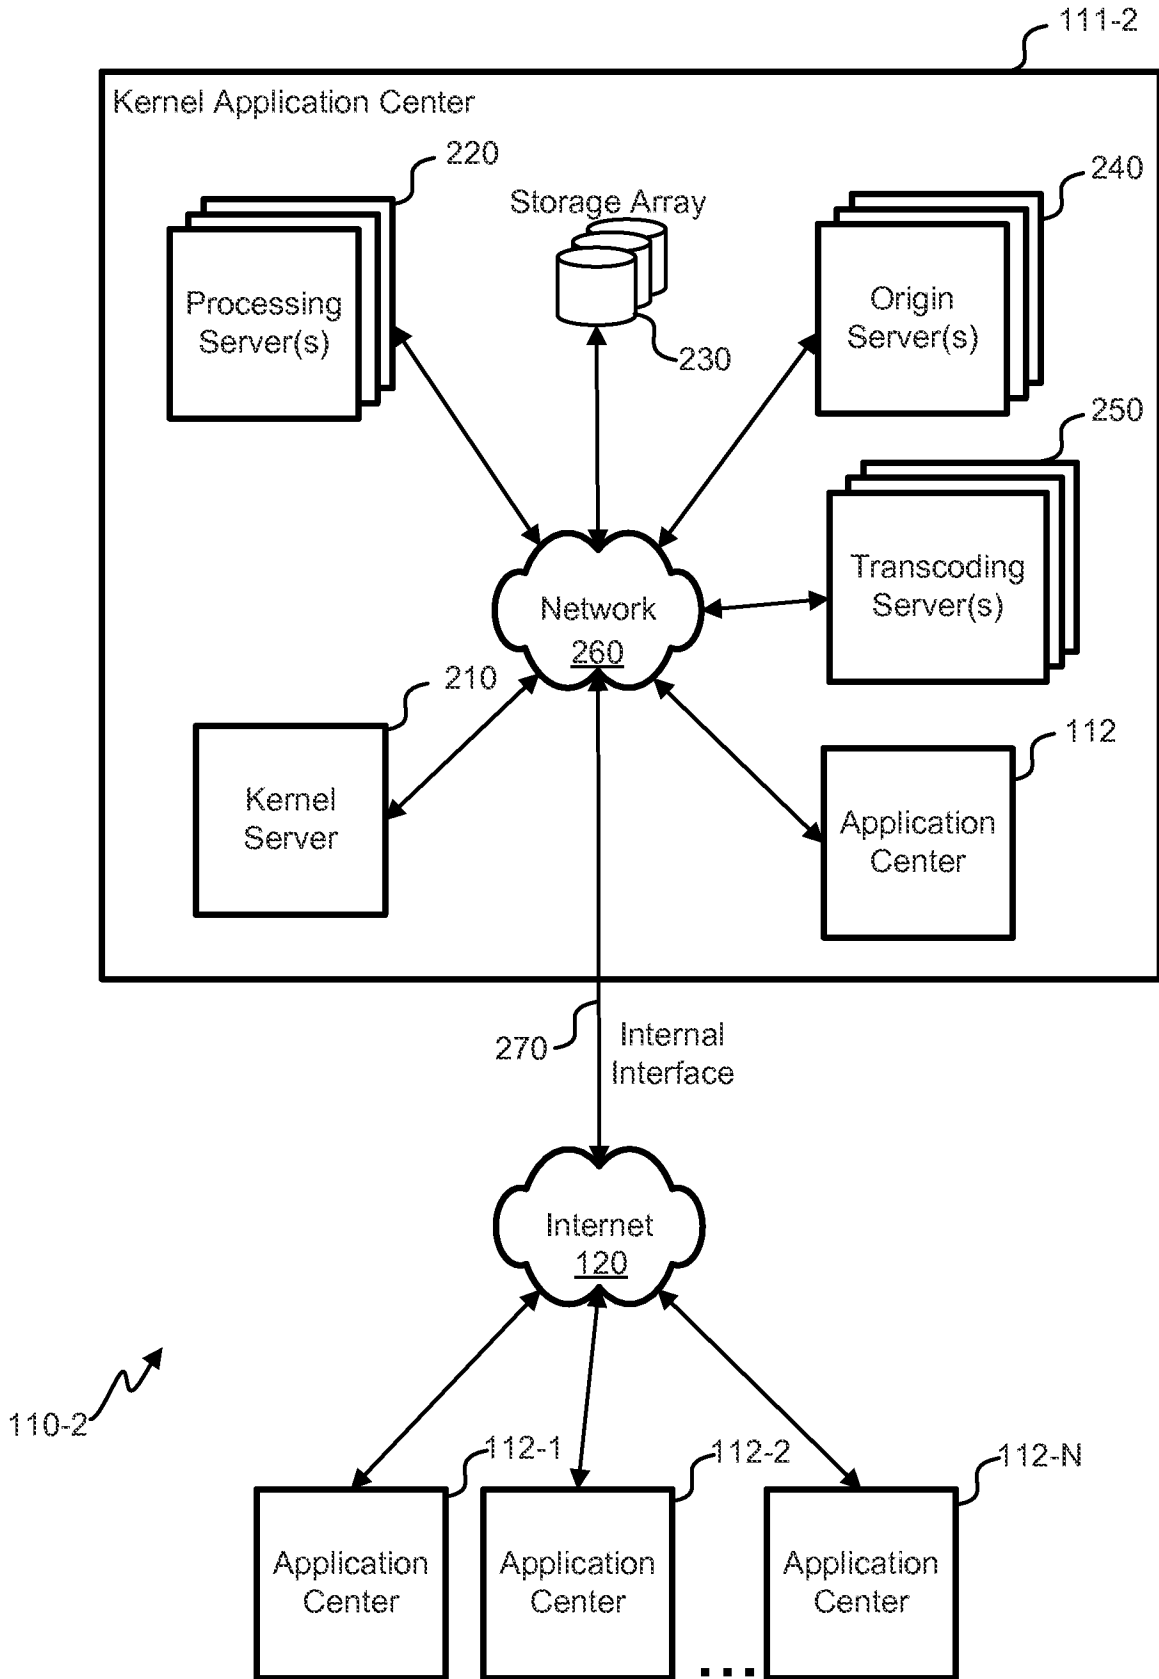


FIG. 2B

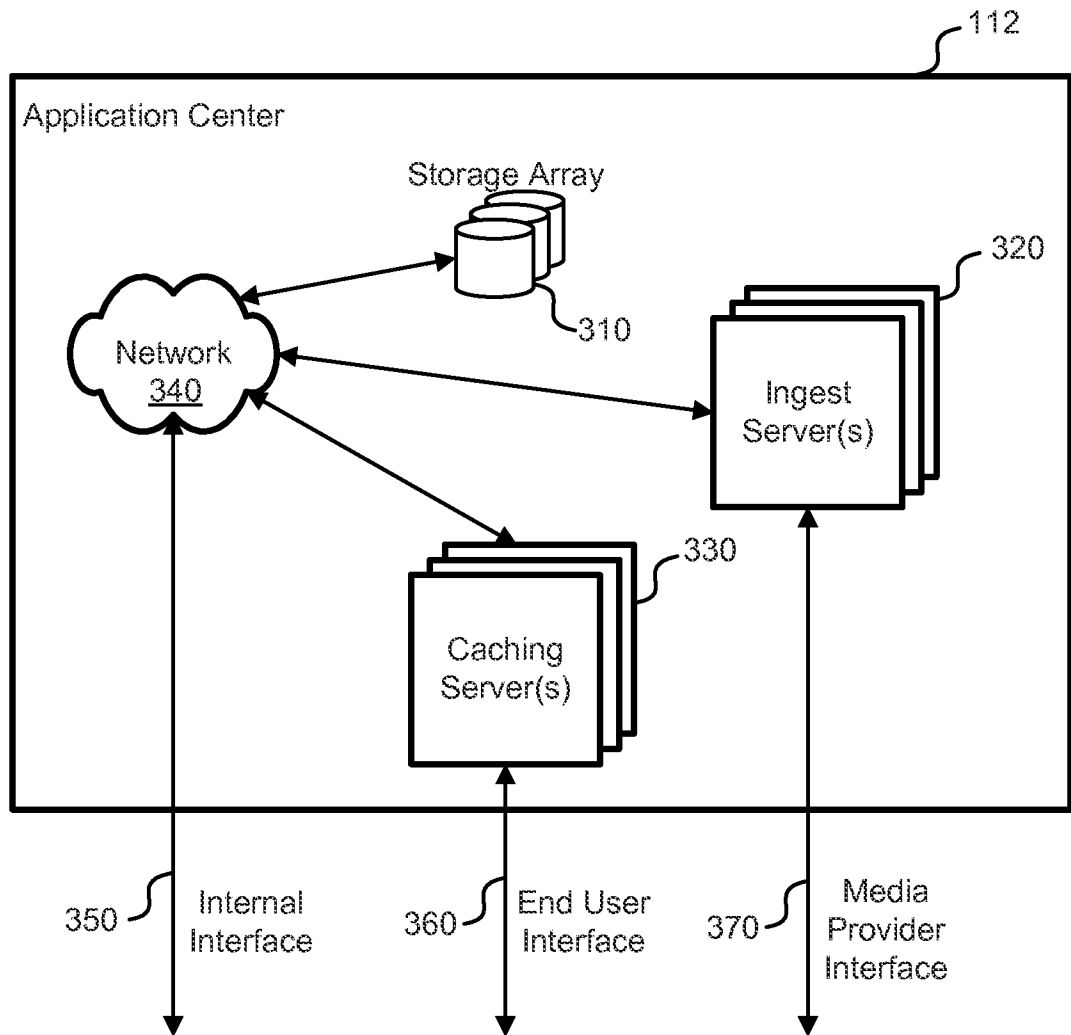
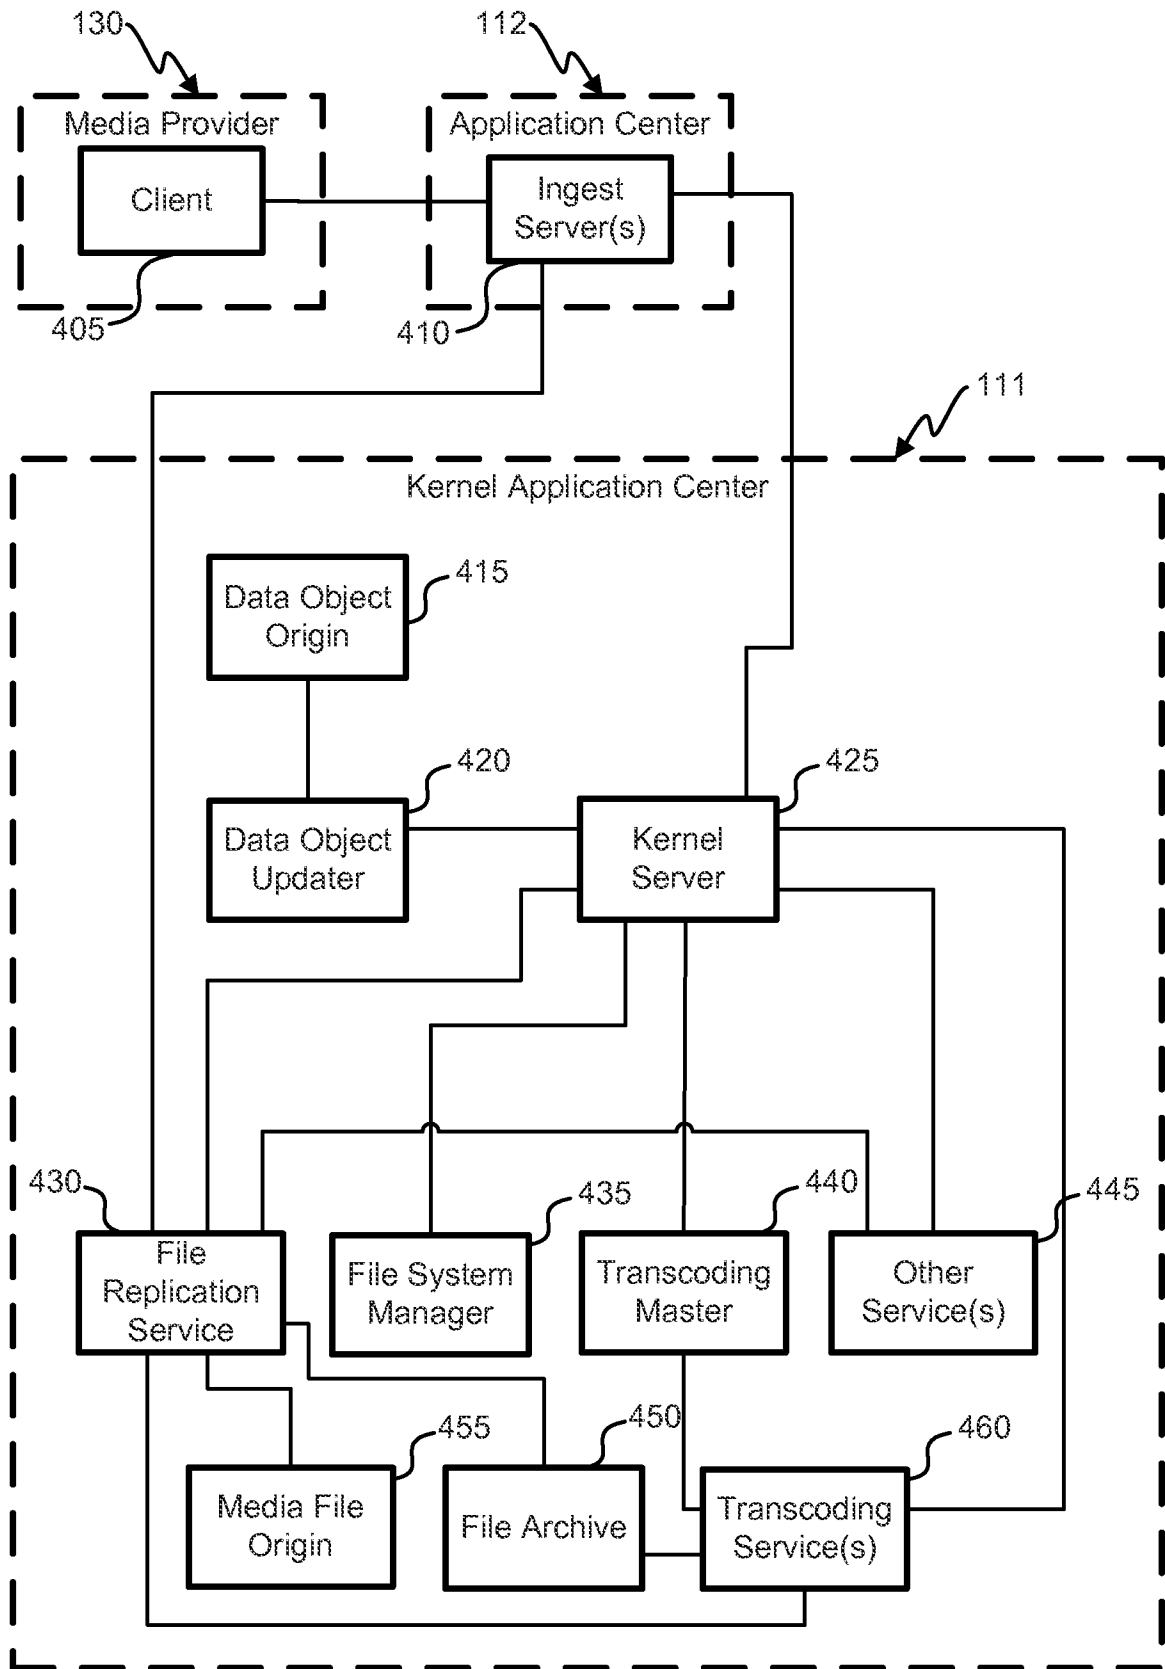
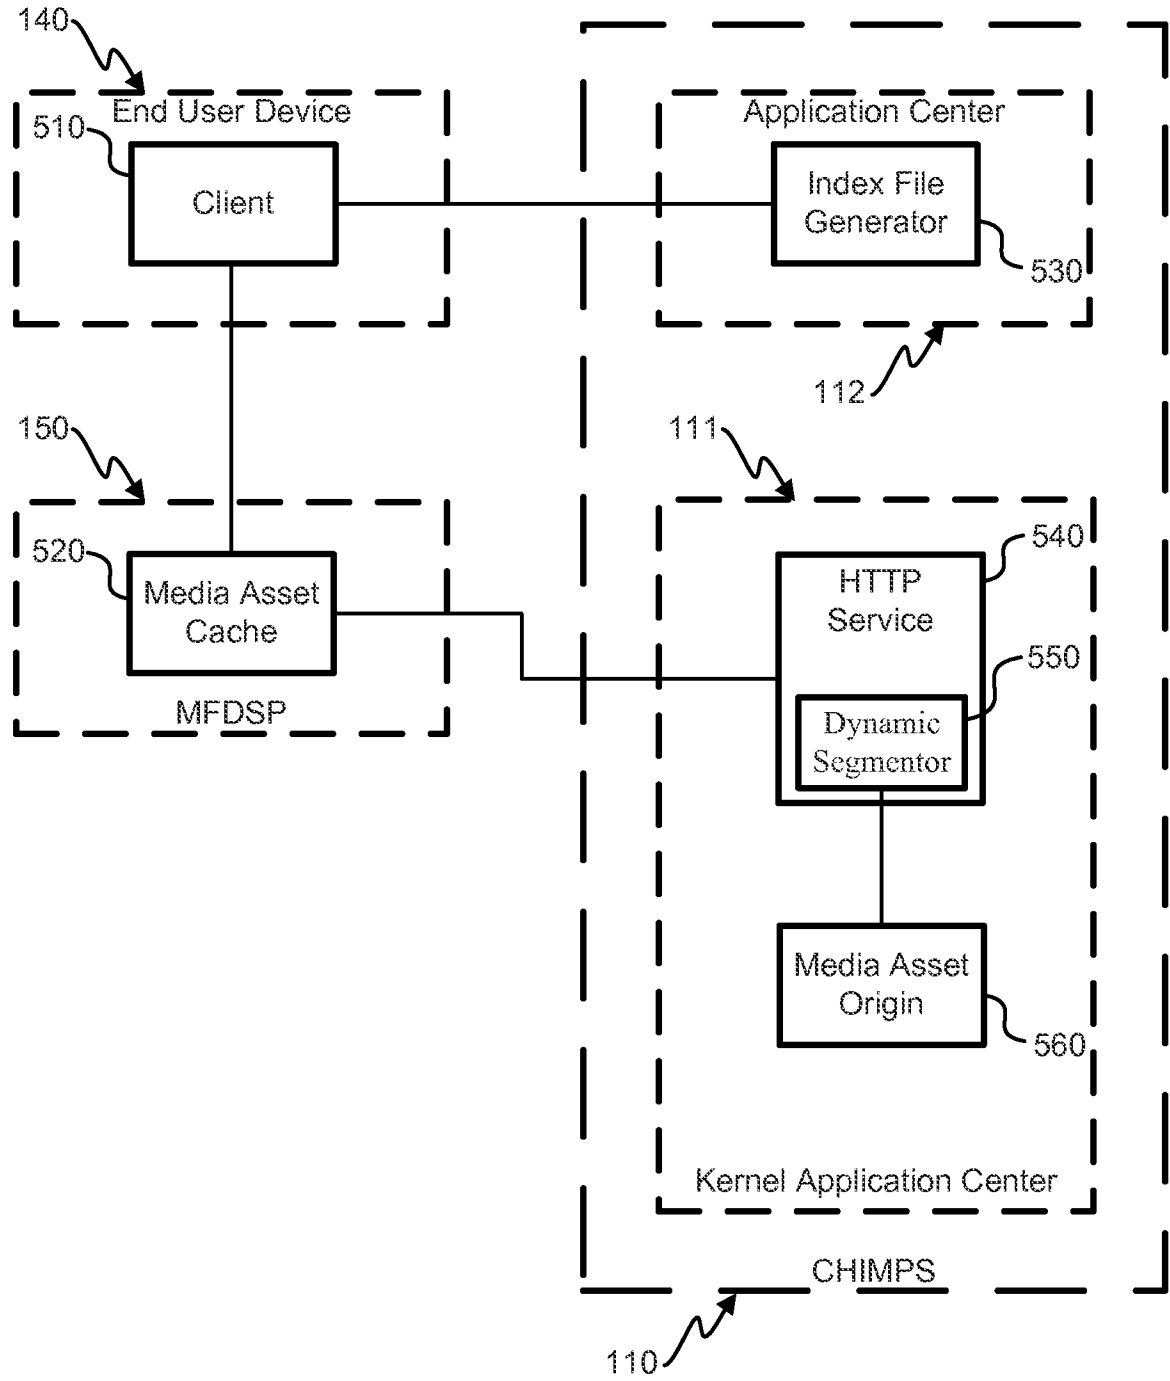


FIG. 3



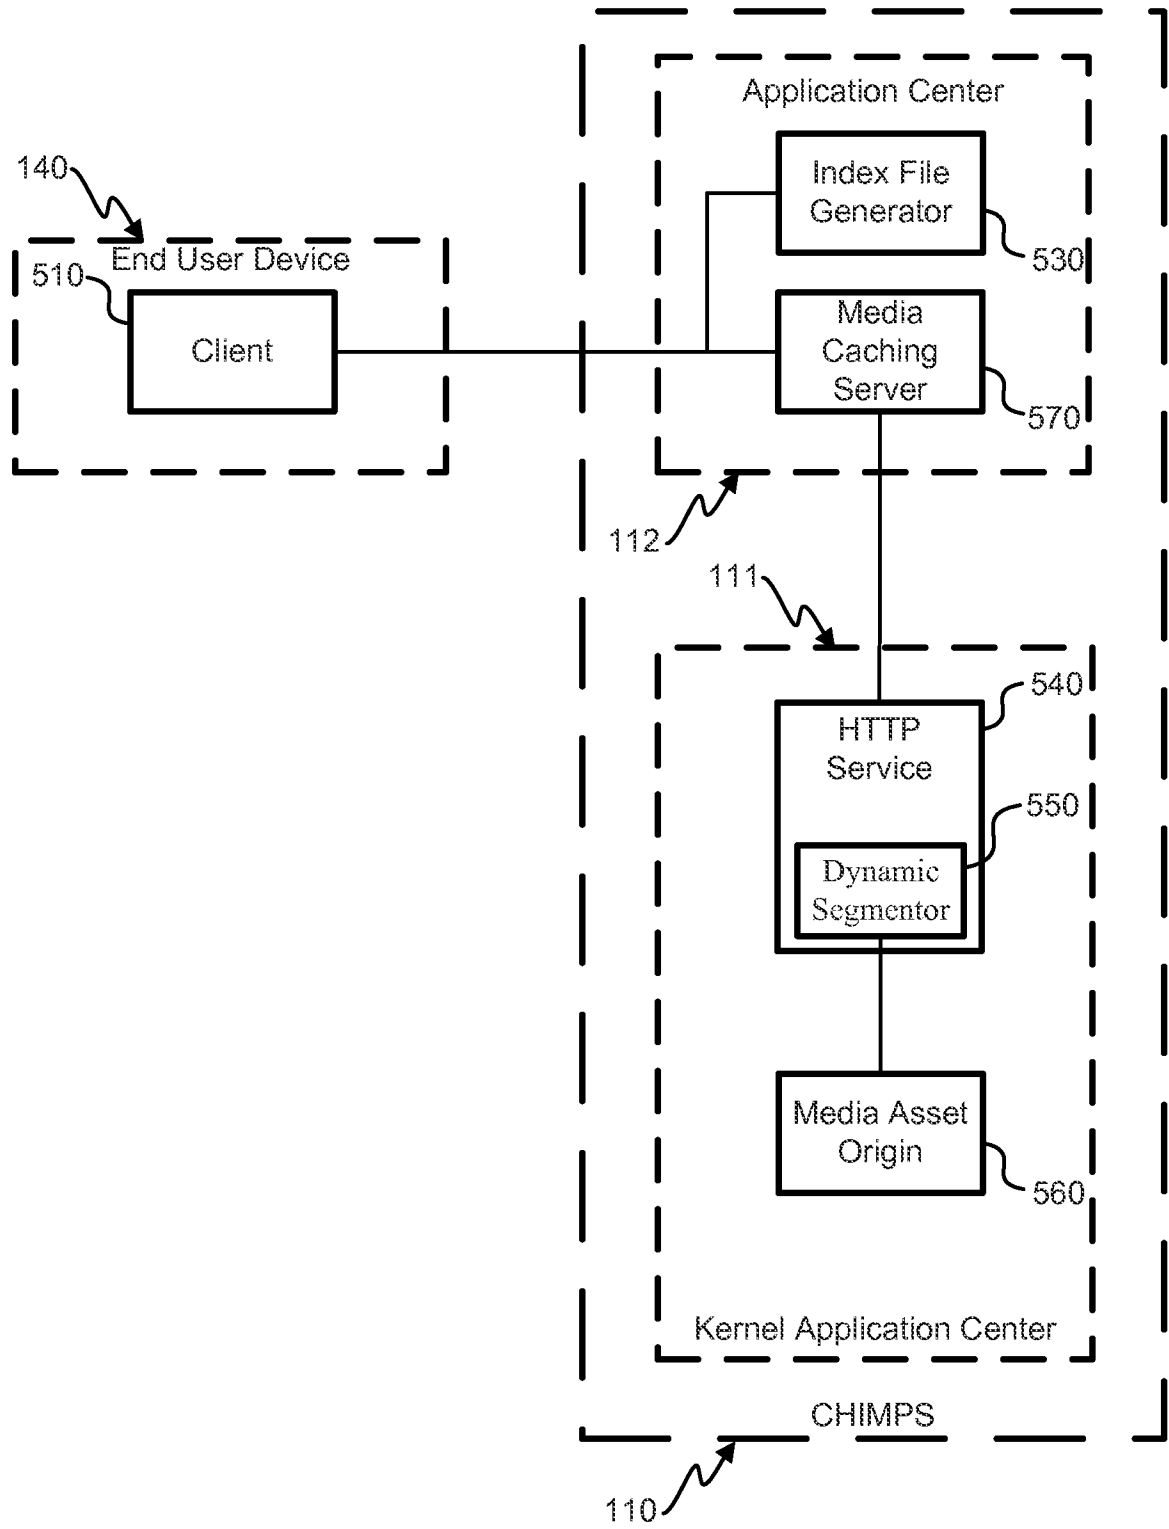
400 ↗

FIG. 4



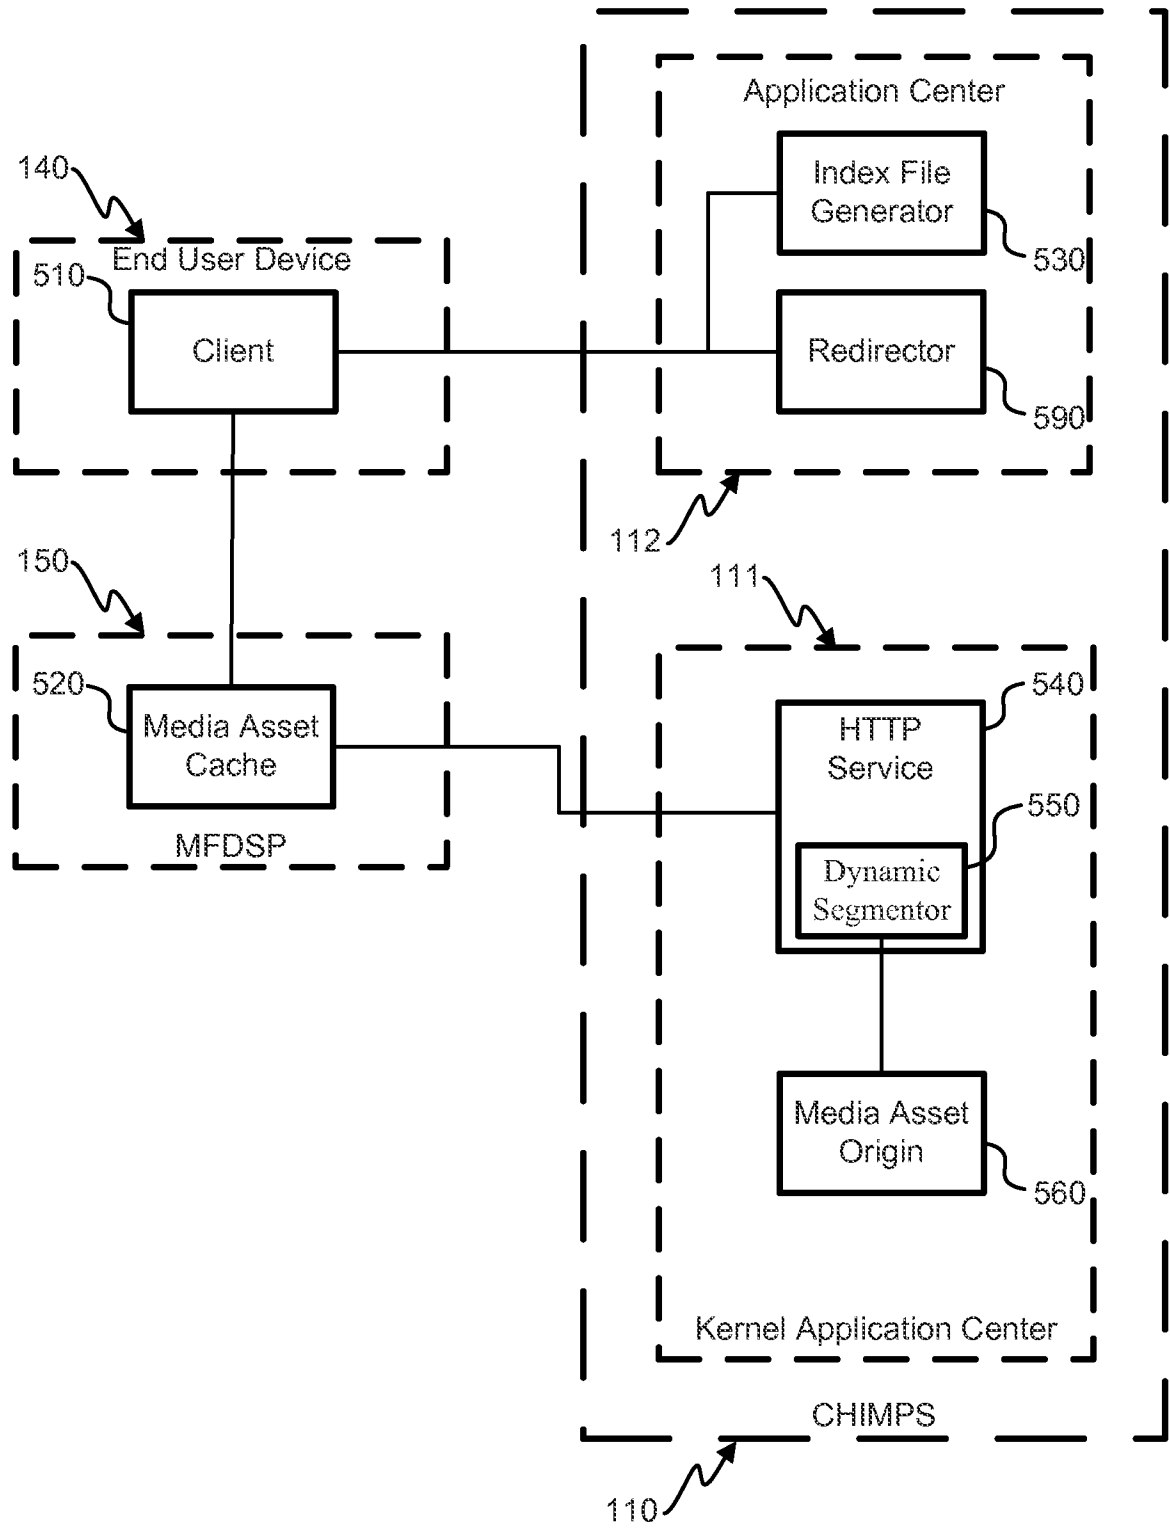
500-1

FIG. 5A



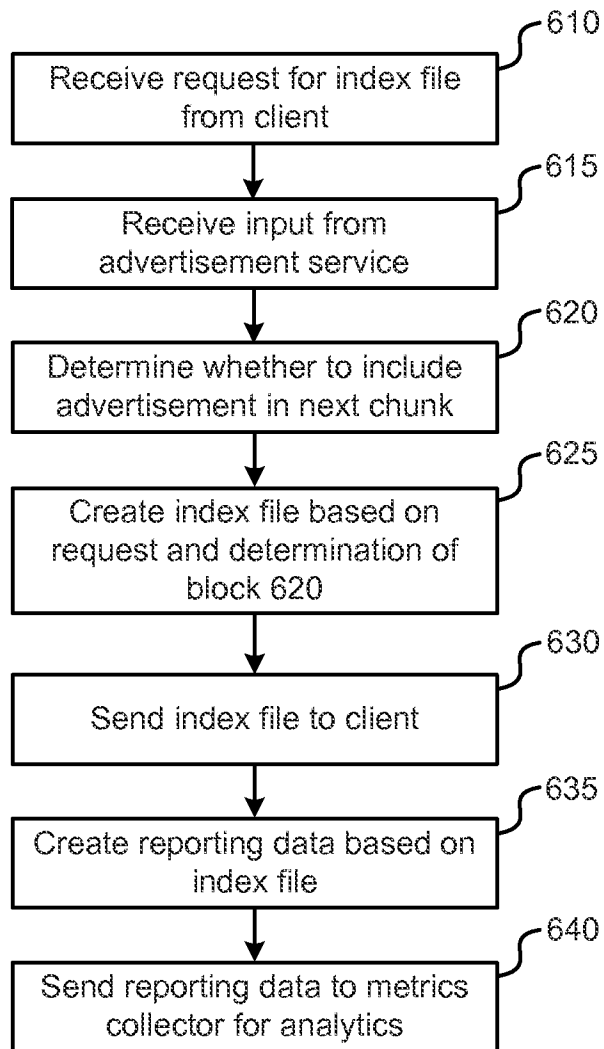
500-2 ↗

FIG. 5B



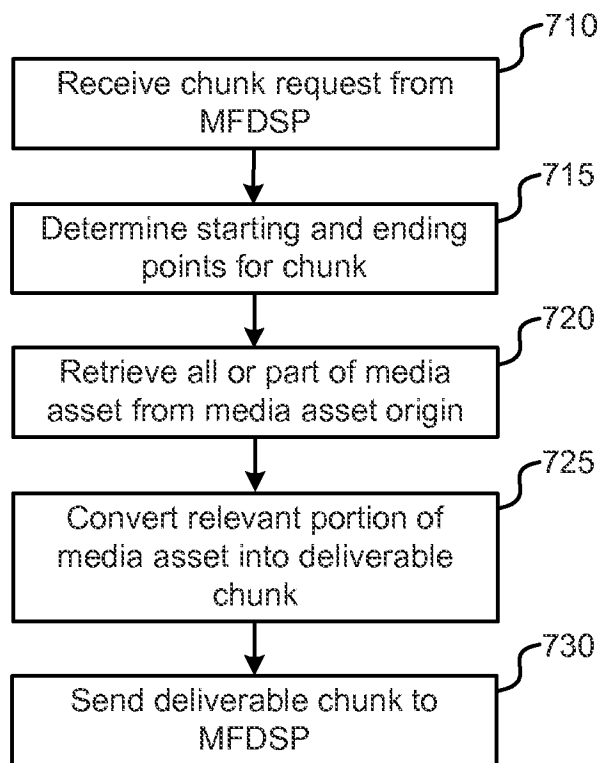
500-3 ↗

FIG. 5C



600 ↗

FIG. 6



700 ↗

FIG. 7

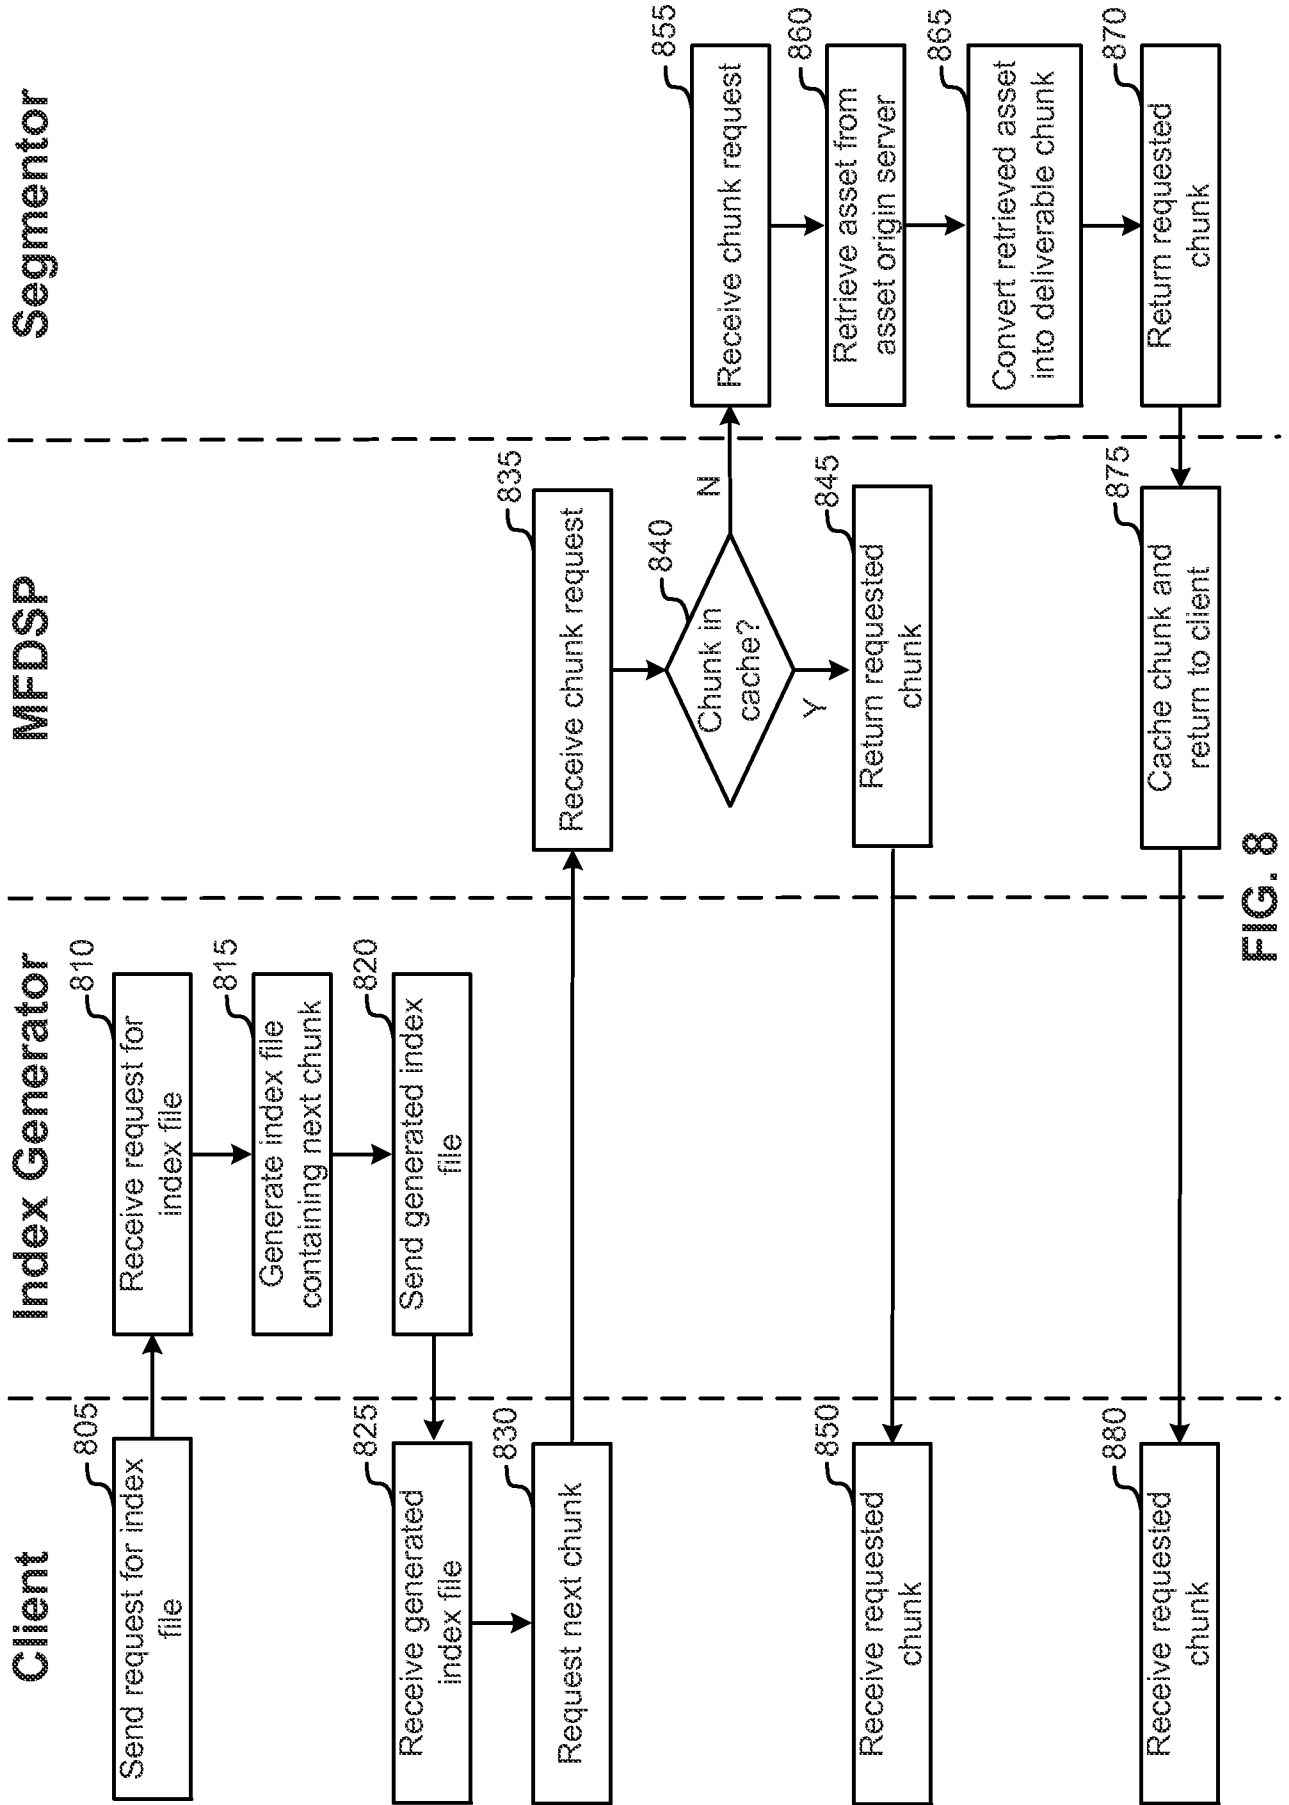
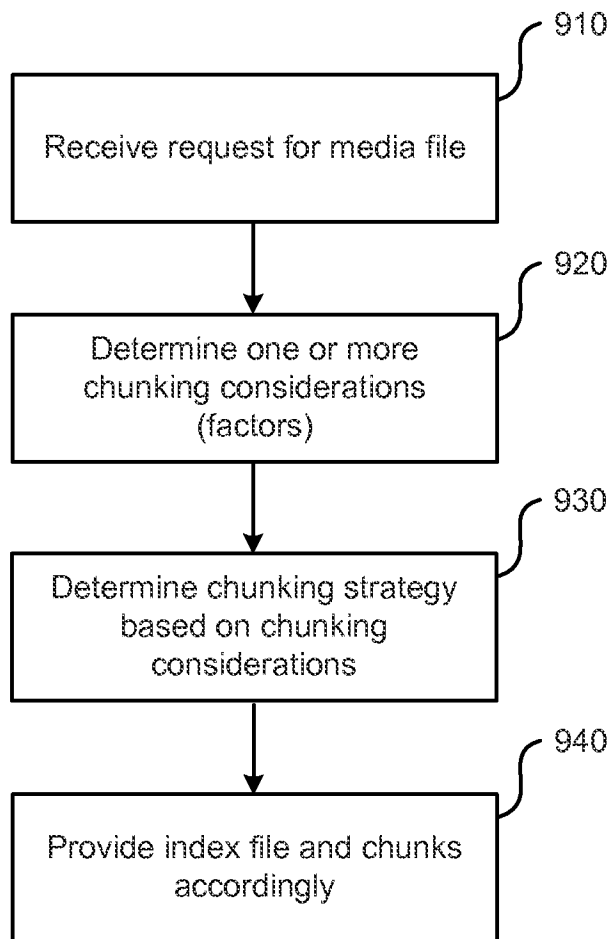
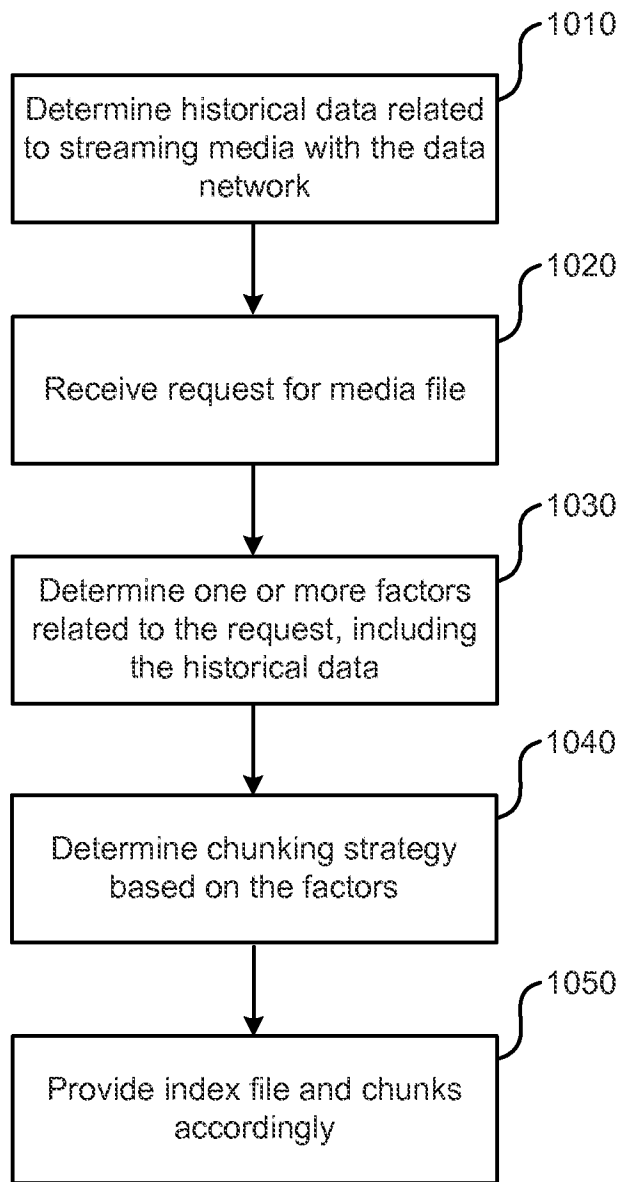


FIG. 8



900 ↗

FIG. 9



1000 ↗

FIG. 10

INTERNATIONAL SEARCH REPORT

International application No PCT/US2014/017945

A. CLASSIFICATION OF SUBJECT MATTER INV. H04L29/06 ADD.		
According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED		
Minimum documentation searched (classification system followed by classification symbols) H04L		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) EPO-Internal, WPI Data		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 2012/175145 A1 (ERICSSON TELEFON AB L M [SE]; PERKUHN HEIKO [DE]; KAMPMANN MARKUS [DE]) 27 December 2012 (2012-12-27) abstract page 6 - page 29; figures 1-13 -----	1-20
X	US 2011/082924 A1 (GOPALAKRISHNAN KUMAR [US]) 7 April 2011 (2011-04-07) abstract paragraph [0046] - paragraph [0112]; figures 1-12 paragraph [0167] - paragraph [0177] -----	1-20
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input checked="" type="checkbox"/> See patent family annex.		
* Special categories of cited documents :		
"A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier application or patent but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family	
Date of the actual completion of the international search	Date of mailing of the international search report	
7 May 2014	14/05/2014	
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Manthey, Axel	

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2014/017945

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 2012175145 A1	27-12-2012	CN 103748857 A	23-04-2014
		EP 2724515 A1	30-04-2014
		TW 201322699 A	01-06-2013
		WO 2012175145 A1	27-12-2012

US 2011082924 A1	07-04-2011	EP 2486491 A1	15-08-2012
		US 2011082924 A1	07-04-2011
		US 2011082946 A1	07-04-2011
		WO 2011044287 A1	14-04-2011
