

MEASURING RESPONSE TRENDS IN A DIGITAL TELEVISION NETWORK

PRIORITY CLAIM AND RELATED PATENT APPLICATION

[0001] This document claims the benefit of priority from United States Patent Application Serial Number 13/952,471, entitled “Measuring Response Trends In A Digital Television Network,” filed July 26, 2013. The entire disclosure of the above-referenced application is incorporated by reference as part of the specification of this PCT application.

[0002] This document relates to digital content networks, including monitoring a content network based on message trends.

BACKGROUND

[0003] With the advent of connectivity technologies such as the World Wide Web, multimedia content is ubiquitously present and accessible from practically any location where a user can be. To meet the increasing user demands, content service providers have deployed networks that have steadily increased in size and complexity. Being able to provide content that a user wants at the time and location the user wants it, and be able to bill the user correctly can provide competitive advantage to content service providers.

[0004] Better techniques for monitoring operation of a content network for any irregularities are desirable.

SUMMARY

[0005] Techniques and systems for providing a trend server outside a content provider network to communicate with the content provider network to build a trend record based on responses received from the content provider network for monitoring certain activity trend or tendency in the content provider network and provide an alert when the content provider network behaves abnormally.

[0006] In one disclosed aspect, a method of monitoring status of a content provider network is provided. The method may be implemented by a trending server and includes sending a first request message to the content provider network, where the first request message has an associated type value, receiving a first response message from the content provider network, determining whether or not the first response message is normal by comparing with a trend record for the type value associated with the first request message,

selectively updating the trend record based on the first response message and generating an alert when it is determined that the first response message is abnormal.

[0007] In another aspect, an apparatus for monitoring operation of a content provider's network is disclosed. The apparatus includes a status response reception module that receives status responses from the content provider's network, a response sorting module that assigns, for every received status response, a corresponding request type based on a status request for which the status response was received, a trending record-keeping module that maintains trend records based on a predetermined number of prior status responses received for each request type, a trending estimation module that estimates a trend for each request type based on the status response and a user setting for that request type, a trend checker module that checks, for each received status response, whether or not the received status response is within the trend for the corresponding request type, and a trend alert module that issues an alert when the received status response is not within the trend for the corresponding request type.

[0008] In yet another aspect, a computer program product comprising a computer-readable storage medium having code stored thereupon is disclosed. The code, when executed by a processor, causes the processor to implement a method of monitoring operation of a content provider's network. The method includes receiving a first message from the content provider's network, determining whether the first message is normal by comparing with a trend of previously received messages having a same message type as the first message, selectively updating, based on the first message, the trend for the message type of the first message, and generating an alert when the first message is determined to be not normal.

[0009] In a further aspect, a system for monitoring operation of a web service provider is disclosed. The system includes one or more web servers under management of the web service provider, the one or more web servers located inside the web service provider's network and a plurality of monitoring servers positioned external to the web service provider's network and geographically separated from each other. Each monitoring server is configured to transmit sync messages to the one or more web servers, maintain a record of expected responses to the sync messages, receive a status response messages from the one or more web servers, and issue an alert, when the status response message is not an expected response.

[0010] These, and other, aspects are further described below with reference to the drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

- [0011] FIG. 1 depicts an example of a content network configuration.
- [0012] FIG. 2 depicts an example of a content network configuration including a trending server.
- [0013] FIG. 3 is a block diagram representation of an example of a content network.
- [0014] FIG. 4 is a block diagram representation of an example of a web-services network.
- [0015] FIG. 5 is a flowchart representation of an example of a process implemented by a trending server.
- [0016] FIG. 6 is a flowchart representation of an example of a process of monitoring a content network.
- [0017] FIG. 7 is a block diagram representation of an example of a content network monitoring apparatus.
- [0018] FIG. 8 is a flowchart representation of an example of a process of monitoring a content network.
- [0019] FIG. 9 depicts examples of signals exchanged in a cloud based trending monitoring system.
- [0020] FIG. 10 depicts an example of a web page monitored by a trending server.
- [0021] FIGS. 11, 12, 13, 14 and 15 depict examples of contents of messages communicated in a trending server based monitoring system.
- [0022] Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0023] With ever-increasing availability of devices that a user can use to receive and view content, the complexity of content services provider systems is steadily going up. Today, users can access content over several networks, e.g., wired Internet (cable modem or high

speed phone), wireless networks (e.g., 3G/4G cellular networks, 802.11 local area networks, and so on), fiber networks (e.g., AT&T U-verse or Verizon FiOS), digital cable delivery services, and so on.

[0024] To stay competitive, content providers today often carry thousands of movie, audio and television program titles in their inventory and make these available to users (subscribers or non-subscribers) either free or using some fee structure such as per-usage or subscription-based access. Furthermore, user requests may arrive at the content provider's servers practically from anywhere in the world. Content providers often restrict access to content based on a user's location or other user information (e.g., age) or charge a different amount based on where a user is located. For example, a newly released movie may be available for viewing in Australia but not in the United States.

[0025] Presently, there are no solutions available for effectively monitoring a content network with wide geographic reach and having a large number of content titles available for secure transmission to users. Many existing solutions often are reactive - e.g., a user brings a problem to the content provider's attention or disputes a billing charge, and then the content provider diagnoses the problem. Various currently deployed monitoring solutions that include a monitoring server in a content provider's network often have limited monitoring functions (e.g., response time of a server or throughput to a server, etc.) and tend to fail to detect any problems that can be seen "from outside" the content network, e.g., from a user's vantage point.

[0026] The techniques disclosed in the present document provide a trend server outside the content provider network to communicate with the content provider network to build a trend record based on responses received from the content provider network for monitoring certain activity trend or tendency in the content provider network and provide an alert when the content provider network behaves abnormally. The disclosed trend monitor server can be implemented to address one or more of the above-discussed technical limitations and others. In some implementations of the disclosed technology, a trending server may be deployed on the outside of a content provider's network. The trending server may be configured to monitor user interaction with the content provider and check messages going from the content provider to the user for any irregularities. To help with the monitoring of irregularities in a response, the trending server may automatically learn a response trend based on previously received responses. The trending server may either automatically set

acceptable values of normal range of parameters in the response or may be programmed by an operator with upper or lower limits to expect. In some implementations, user devices may be configured to communicate with the content network through the trending server. In some disclosed embodiments, new software functionalities and protocols to assist detection of unexpected trends in responses from web services for the purpose of handling unusual behavior are provided.

[0027] These, and other, features are disclosed in greater detail below.

[0028] FIG. 1 depicts a content network architecture 100 in which a service provider's network 110 and a publically available network such as the Internet 112 are separated by a boundary 104 (e.g., internet address sub-network, firewall, etc.). On the service provider network side, one or more application servers 198 might be inter-connected and configured to perform tasks such as content provisioning, content delivery, billing, user account access, and so on. User devices 104, located on the public network side, are able to communicate with the application server 108 using the above-described wired or wireless networks 106.

[0029] During operation, user devices 104 will communicate with the service provider's network 110 to perform various tasks such as requesting content or web pages, interacting with application servers 108 to download or upload information, and so on. During normal operation, certain information in a response from the service provider's network to a request from the user may be expected. For example, when the user accesses a top-level domain name, the home page is expected to be sending back to the user. Similarly, in a content provider's network, when a user access a particular movie title, a web page that includes a graphics image, movie metadata (e.g., cast, story line, release year, run time, etc.) and purchase information. Often, the type of information may be predictable, but the actual value of the information may change. For example, for a movie rental web page, it is predictable that a money amount will be displayed for renting the movie, but the actual value may not be known until the web page is served out by an application server 108.

[0030] Determining what is a normal response from a web service-oriented system has multiple complexity factors:

[0031] 1) local dependencies, e.g. work flow between services/databases, hardware, bus/connectivity and so on

[0032] 2) external factors, e.g. data input feeds, number of concurrent connections, Internet bandwidth, domain name server (DNS), etc.

[0033] 3) accidental or malicious modification of any of the above leading to unexpected behavior of the system.

[0034] Detecting when a normal response is regained also introduces complexity in applications. Sometimes, a no response may also be a normal response. For example, a particular web page, e.g., rental page for a new movie title, may not be available to all user devices 104, e.g., based on geographical location of the user. Therefore, a “title not found” message may be the normal response to a user outside the allowed geographic area, while the movie rental web page would be a normal response to another user in the allowed geographic area.

[0035] Using Nagra MediaLive Service Delivery Platform / Content Management System (SDP/CMS) as an example, this system has a complexity where incorrect data entry for a popular video asset might cascade quickly to faults being reported by subscribers calling the customer call center. If the root cause is not immediately evident and later failing to indicate when a fault is resolved leads to service level agreements not being met as well as costs whilst the fault remains, increased churn risk and ultimately unhappy customers.

[0036] In another operational scenario, cybersecurity of web-based information systems such as content networks can also be complex. A CMS for example might offer a URL to company press releases. Such a web resource normally contains changes only to article data during business hours. An untested or unmonitored configuration change might lead to resource being moved away, a hacked system might display an entirely different web resource due to tampering with internal data or external factors such as ISP DNS modification (URL redirect), overload of the service (DDoS - distributed denial of service). Early detection of such compromises is paramount and not always apparent to internal monitoring systems.

[0037] To host a monitoring web service adds to the complexity of the overall system as this requires either local connectivity to the system to avoid complexities such as cross origin (CORS) or adds more security holes (hack points) to support it. An internal system might not see outside of its own box or become unavailable outside.

[0038] FIG. 2 depicts an example of architecture of a content network 200 (in general, a web services network). Compared to the system 100 depicted in FIG. 1, the content network 200 includes a trending server 202. While not explicitly shown in FIG. 2, in some embodiments, multiple trending servers 202 may be deployed spread out over a wide region. Such a trending server 202 is located external to the service provider's network, and also on the same side as the user devices 104. An external location could mean, e.g., in a network not controlled by the service provider or based on IP address of the trending server 202 or on the outside of the firewall 102 that controls access to the service provider's network 110.

[0039] As service providers address the need for a seamless multiscreen user experience across devices and networks, many service providers are looking for a fast time-to-market and full control on their costs in a market where requirements are evolving at an increasing pace. Nagra tools and products are designed to meet those and other requirements. For example, Nagra MediaLive Service Platform offers software tools to manage, secure and operate, from a single platform, a wide range of media services across multiple delivery networks (broadcast, managed IP, over-the-top and mobile), devices (set-top boxes, connected TVs, PCs, tablets and phones) and content types. The MediaLive Service Platform unifies server-side functionalities across devices, networks or content types and throughout the entire content lifecycle, from ingestion to consumption and reporting. The MediaLive Service Platform can be accessed over the internet, from across the world, and enforces an operator's conditional access rules regarding which users are allowed to access which content.

[0040] The MediaLive Service Platform supports flexible service definition enabling service providers to easily and efficiently monetize their premium entertainment. It manages linear channels and on-demand content delivered across any network to any device. It supports a wide range of business models from subscription, pay-per-view, pay-per-time, "à-la-carte" to event-based purchase models, promotions and addressable advertising.

[0041] The MediaLive Service Platform maximizes the device reach of multiscreen services. It provides a sophisticated account and device management capability including personalization and authentication, a unified transaction and content access management including rights and billing event management, domain management, device messaging and many more functionalities. The MediaLive Service Platform makes available an extensive library of interactive applications ranging from extended EPGs, streamed/push/pull/ video-

on-demand, catch-up TV, start-over, interactive subscription, to multi-store services, cross-device purchases and bookmarking, social networks integration, recommendations and other interactive services.

[0042] FIG. 3 is a block diagram representation of functional blocks and messages exchanged according to a system configuration 300 suitable for implementing content or application servers based on MediaLive Service Platform or other platforms. A content provider's network 302, which can be implemented by, e.g., previously described Nagra MediaLive enabled network, is accessible by a plurality of user devices 204, 306 and 308. For example, user device 304 may be a digital-set top box such as an OpenTV5 set-top box. The user device 304 may also be embodied in equipment such as a cable-ready television or another consumer device having a return path to be able to communicate with the content provider network 302.

[0043] In a given deployment, some set-top boxes may be configured to directly communicate messages 310 with the application server 302 over the service provider's network. These messages 310 may include, e.g., content viewing or purchase requests and download or broadcast of the corresponding content. Some set-top boxes 304 may be configured to communicate with the application server through a trending server 202. This may be accomplished, e.g., by providing the IP address of the trending server 202 as the destination addresses for the corresponding messages sent out by the set-top box 304. The trending server 202 may in turn monitor requests from set top boxes 304, and receive, process and pass-through responses received from the application server 302 back to the requesting set-top box 304. The processing in the trending server 202 may include determining trends and raising alerts for any deviations from normal responses, as is described in the present document.

[0044] The system 300 may include, in some applications, user devices 306 that are configured to run web-based applications, e.g., hypertext markup language version 5 (HTML5) webapps. The user device 306 can communicate with the application server 302 over the connection 312 to the trending server 202, which further communicates with the application server 302 over connection 314 on behalf of the user device 306. In some embodiments the WebSockets protocol may be used. In one advantageous aspect, WebSockets allows a connection to remain open, thereby providing ability for the user device 306 to receive any changes to specific web pages monitored via the WebSockets interface. In

some embodiments, an HTML5 application running on the user device 306 could be alerted by a message from the trending server, when updates to web pages that are being monitored over the open WebSockets connection are considered to be irregular, or outside a range of normal deviation, by the trending server 202.

[0045] In some applications, the system 300 may further include user devices 308 equipped with a media player provided by the content service provider. For example, the devices 308 may be wireless or companion devices that users use to access content. In some implementations, the devices 308 may be tablets or mobile phones that run software platforms from Apple or Google (iOS or Android). These user devices 308 may also be configured to operate similar to the user devices 304 (which could be the primary devices for the companion devices 308).

[0046] FIG. 4 depicts a system configuration 400 in which a web-services network (e.g., a content management system 402) is being accessed by multiple user devices. A user may access a web service offered by the application server 402 using a user device 404 with a web browser or another suitable application or using another mechanism such as the previously described WebSockets based HTML5 mechanism. For multiple user devices 404 and 406 in a content network, some devices may be configured to communicate with server 402 through an intermediate trending server 202, while other devices 404 and 406 may be configured to communicate directly, i.e., without using the trending server 202 in the communication path. Furthermore, the devices 404 and 406 may be distributed over a wide geographic region and multiple networks (e.g., some devices may be wireless, some may be using cable modem access network, some others may be using a dial-up access network, etc.).

[0047] When a user makes a request to the server 402, the user may be provided with a web page, an extensible markup language (XML) document, etc. The trending server 202 may compare the returned web page or XML document for any differences from a trend previously calculated for documents of the same type or documents received from the same request type.

[0048] A message returned from the content provider's network may include various types of data. For example, audio, video, text and other multimedia data, metadata related to presentation of other data (e.g., formatting instructions), and so on.

[0049] Specific examples of trend calculations in the trending server are described below.

[0050] In some implementations, text may be compared against a dictionary of occurrences, frequency of occurrences etc.

[0051] Numbers in numerical field in a web page may be compared with a range of previously received numbers, a meaning associated with the numbers and any other input rules provided by a user.

[0052] Audio may be checked for file size, encoding parameters and other metadata associated with the audio stream.

[0053] Video may be checked for size, encoding parameters (resolution), any embedded messages or links, and so on.

[0054] Graphics may be checked for size, placement within a web page, pattern recognized for information in the graphics.

[0055] In some implementations, an operator may “seed” a trend by first providing a baseline web page and limits on how values in certain fields may change. The operator may also specify to ignore certain fields from trend checking.

[0056] In various applications, trending servers may be distributed in cloud based architecture, with individual trending servers being able to synchronize with each other. One advantage is that a service provider may be providing different content service to users located in different geographical region. For example, a service provider located in Australia may make content available for free to users in Australia, but may be charging viewing fees to users in Europe. Additionally, the content provider may not be making content available to users in Asia.

[0057] FIG. 5 is a flowchart representation of a process 500 of operation of a trending server 202. The trending server establishes a WebSocket connection with a user device (not shown in FIG. 5).

[0058] At 502, the process 500 parses an HTTP request received over the WebSocket from the user device.

[0059] At 504, the process 500 determines whether or not the HTTP request is familiar. For example, the process 500 uses certain fields of the request, or a type of request to

determine if a similar request was previously received, either from the same user device or from another user device.

[0060] When the process 500 determines that the request is familiar, at 504, the process access a trend record for the type of the request and appends the trend record with parameters and values from the request. In some implementations, the process also assigns a unique identifier value to the request for future tracking and report generation.

[0061] If the process 500 determines that the request is not familiar, at 506, a new trend record is established for the request by assigning a new trend record identification and extracting parameters of the request to uniquely identify a type of the request. In some embodiments, the values of the parameters received in the request are used as the initial trend values.

[0062] At 510, the process 500 passes on the original HTTP request to an HTTP server 520 such as a MediaLive SDP. In some embodiments, the original HTTP request is modified to the extent of routing information in headers of packets carrying the HTTP request.

[0063] At 512, the process 500 receives a response from the HTTP server 520.

[0064] At 514, based on the received response, the process 500 synchronizes the trend records for the response. For example, in some embodiments, the trend range may be expanded (or reduced) by including over an average of the range for a parameter value over past N number of responses of that type (N = a suitable integer, e.g., 10 to 50). In some embodiments, the received response may include text and any new words or messages in the received text may be added to the trend record for that type of message. In some embodiments, the received response may include graphics and the process 500 may update the trend record to include characteristics of the graphics (e.g., placement on a web page, resolution, color depth, size in bytes, etc.).

[0065] In some embodiments, the process 500 may further sync the received response with other trending servers 522, which may be spread out over a wide geography (e.g., other states, nations or continents).

[0066] In general, responses received may include text, numbers, and graphics and control commands. It may be unlikely that two responses received at two different times are

identically in every bit of the response. Based on the syncing in 514, the process 500 determines changes in the currently received response from the trend record. For changes that seem “normal” or “expected,” the process 500, at 518, may modify the representation of the trend record. For example, in one response message, a web page may include a “last updated” field that includes a date. In a subsequently received response message, the web page may include the “last updated” field by may include data and time. The process 500 may thus modify the expected representation for the web page to include a “time” field in addition to the “date” field.

[0067] At 524, the process 500 conveys the results of 518 to an HTML5 webapp.

[0068] FIG. 6 is a flowchart representation of a computer-implemented process 600 to monitor status of a content provider network.

[0069] At 602, the process 600 sends, from a trending server, a first request message to the content provider network, wherein the first request message has an associated type value. In some embodiments, the type value may be derived based on a URL included in the request. In some embodiments, the entire URL may correspond to the type.

[0070] At 604, the process 600 receives, at the trending server, a first response message from the content provider network. As previously discussed, the message may be formatted as a document, e.g., an XML document, or a web page (e.g., HTML script) and so on.

[0071] At 606, the process 600 determines, by the trending server, whether or not the first response message is normal by comparing with a trend record for the type value associated with the first request message. Various possible processing operations performed by a trending server in determining whether a response is normal or not (i.e., regular or irregular) have been previously disclosed.

[0072] At 608, the process 600 selectively updates the trend record based on the first response message. In some embodiments, the trend record may be updated only when it is determined that the first response message is normal. Acceptable values of certain fields may be appended or averaged to include the most recently received. For example, in some embodiments, previously received values for a certain integer field may be in the range between 1 and 5. The trending server may therefore consider the allowable values to be the set (1, 2, 3, 4, and 5). When a message is received that is determined to be normal, but

contains the value “6” in the field, and then the trending server may update the range of allowable values to (1, 2, 3, 4, 5, and 6).

[0073] At 610, the process 600 generates, when it is determined that the first response message is abnormal, an alert. For example, if the process 600 determines that a “movie rental cost” field in a web page, which is supposed to have a normal range between \$0.99 to \$7.99 now, has a value \$0 or \$100, then the process 600 raises an alert. The alert may be in the form of an audio/visual cue generated for the user device receiving the response. The alert may be a message that is displayed (or sounded) to an operator of the trending server. The alert may be a message that is logged in an alert log of the trending server.

[0074] FIG. 7 is a block diagram representation of an apparatus 700 for monitoring operation of a content provider’s network.

[0075] The module 702 a status request transmission module that transmits a plurality of status requests to the content provider’s network, each status request having an associated request type. In some embodiments, the status requests may be formatted as stateless queries such as web messages using the REST protocol (representational state transfer protocol). In some embodiments, the status requests may be generated autonomously by the trending server, i.e., without having received a previous request from a user device.

[0076] The module 704 is for receiving a status response from the content provider’s network. As previously discussed, a status response may use REST protocol, may be a web page or an XML document, and may include audio, video or other embedded information.

[0077] The module 706 is for assigning, for every received status response, a corresponding request type based on a status request for which the status response was received. Several techniques for assigning request types (e.g., based on URL fields) have been disclosed in the present document.

[0078] The module 708 is for maintaining records of a predetermined number of prior status responses received for each request type.

[0079] The module 710 is for estimating a trend for each request type, based on the predetermined number of prior status responses for that request type and a user setting for that request type. For example, in some implementations, a trend may be based on 10

previously received responses for the same message type, before receiving the most recent response message. In some implementations, a trend may be based on all previously received responses of a given type. In some implementations, a trend may be based on a number of previously received responses at a given time (e.g., on Mondays or in mornings, and so on).

[0080] The module 712 is for checking, for each received status response, whether or not the received status response is within the trend for the corresponding request type. As previously discussed, trend checking may be performed differently for different data types and may include range checking, comparison, and other processing.

[0081] The module 714 is for issuing an alert when the received status response is not within the trend for the corresponding request type. Techniques for issuing alerts have been discussed previously in this document.

[0082] FIG. 8 is a flowchart representation of a process 800 for monitoring operation of a content provider's network.

[0083] At 802, the process 800 receives a first message from the content provider's network. In some implementations, the first message may be a message that notifies the process 800 of a change on a web server via a WebSocket connection open between a trending server and the content provider's network.

[0084] At 804, the process 800 determines whether the first message is normal by comparing with a trend of previously received messages having a same message type as the first message.

[0085] At 806, the process 800 selectively updates, based on the first message, the trend for the message type of the first message. In some implementations, the trend updating is performed only when it is determined that the received message is normal. In some implementations, the trend updating may be performed regardless of whether the received message is normal or not.

[0086] At 808, the process 800 generates an alert when the first message is determined to be not normal. Various embodiments for generating the alert have been previously disclosed in this document.

[0087] In some implementation, multiple trending servers may be implemented as follows. A global trending server may act as a coordinator that coordinates activities of multiple trending servers in a trending server cloud geographically distributed over a region. The global trending server may communicate with the other trending servers and provide information about which web pages to monitor for trends. Periodically, each trending server requests and receives a web page being monitored. The trending server communicates the results of its query to the global trending server. For example, the trending server may either send outcome of its determination of whether or not the web page received was normal. Alternatively, the trending server may also pass on the received web page to the global trending server and the global trending server makes a determination of whether or not the received web page was normal.

[0088] With reference to FIG. 9 to FIG. 16, as a non-limiting example, a deployment in Sydney may monitor the status of a content network. A user device 902, an iPad with an HTML5 capable application, may be configured to monitor certain web pages, represented by their uniform resource identifier (URI).

[0089] FIG. 9 shows an example of a web page that may be monitored in a system that has multiple trending servers. In this example, the web page is formatted as an XML file that provides information about a movie available for rent and includes a number of fields and their corresponding description metadata.

[0090] A monitoring app running on the user device 902 (e.g., a tablet) may direct a local trending server 904, e.g., a server located in Sydney for a user in Sydney, to monitor a certain web page by providing a URI to the web page (912).

[0091] The local Trending Server 904 parses the HTTP Request and locates or creates a new Trending Record from the Trending database. The Sydney Trending Server 904 then triggers the global Trending Server cloud, comprising for example Trending Server 1 (906) and Trending Server 2 (908), to generate an HTTP Request to the URI data passed to it (914).

[0092] Each Trending Server in the cloud, on a specified schedule, requests status of the web page based on the URI (messages 916, 920 and 924). In some embodiments, the request times of the messages 916, 920, 924 may be synchronized to get multiple snapshots of the webpage from different locations at the same time. In some embodiments, each Trending Server may send its own HTTP request asynchronously from other Trending Servers.

[0093] An App server in the content network, e.g., Nagra MediaLive SDP, handles the HTTP Request and passes Response back to each Trending Server (responses 918, 922 and 926, respectively).

[0094] Each Trending Server 906, 908 in the global trending server cloud then compares the received response and generates a trending report 928, 930, which is then forwarded to the local Trending Server 904. The local Trending Server 904 generates its own report of whether the received response was normal or not.

[0095] Each Trending Server updates its Trending Record selectively based on the received response.

[0096] The local Trending Server 904 (Sydney Trending Server) processes the trending record updates from the trending cloud.

[0097] Optionally the web page update request can be made again after a set time by triggering the same request again and determining a change in the Trending Record.

[0098] If the message 912 is the first request from the user device 902, or an unexpected change is detected, the listening clients are updated via the WebSocket 932 by reporting the changes.

[0099] In some embodiments, the local trending server sends a first type of alert message (e.g., a distinct audio tone or visual cue on the user device) to highlight unexpected change from Sydney, but not from elsewhere. In some embodiments, the local trending server sends a second type of alert message, with its own audio-visual cues, to indicate detection of abnormality from another trending server in the cloud. The user can further investigate the abnormality based on the detail provided in the report 932.

[0100] FIG. 10 depicts an example web page monitored by the trending server 202. This web page includes a number of fields and the corresponding values (alphanumeric strings) that are used to provide information for a movie available for rental from an application server in the service provider's network.

[0101] FIG. 11 depicts examples of some content of messages 912 and 914.

[0102] FIG. 12 depicts examples of some content of messages 928, 930 and 931. Each Trending Server may identify its identity using a “trendingserver” parameter. Furthermore, trending servers may annotate the observed parameter values with trends seen for the values and explicitly list any deviations from normal.

[0103] FIG. 13 depicts an example of content of message 932. In this example, the local trending server indicates that a web page appears to be normal from its vantage point, but some other trending server in the trending cloud has detected a change in the web page. For example, a web page being monitored by a trending server may have been hacked using a DNS exploit at an internet service provider’s caching server, but the original web page (e.g., at Nagra CMS) may still be intact. In such a case, FIG. 13 may represent a message from a trending server monitoring the CMS (and not seeing any irregular changes to the web page) which becomes aware through syncing with the trending server cloud of a trend abnormality noticed by a trending server that is receiving its webpage status responses (e.g., REST responses) from the ISP of the hacked web page.

[0104] FIG. 14 may represent a trend abnormality noticed by a trending server that is receiving its webpage status responses (e.g., REST responses) from the ISP of the hacked web page and therefore notices an irregular change in the monitored web page.

[0105] FIG. 15 depicts another example of message 932 in which the user is alerted with changes in certain fields of the web page monitored.

[0106] It will be appreciated that techniques for monitoring normal operation of a web services system such as a content service provider’s network are disclosed. In some embodiments, the network is monitored from multiple locations in different geographic areas. In some embodiments, a report is generated for the current status of a web page by comparing with a previous trend of the web page and alerting a user of any unexpected or abnormal changes in the web page. The monitoring can be performed remotely from a thin user device, such as from an application running on a mobile device.

[0107] It will further be appreciated that several techniques are disclosed for monitoring trends in changes to a web page by monitoring different data types, e.g., data, audio, video, etc., in a variety of ways.

[0108] The disclosed and other embodiments, modules and the functional operations and modules described in this document can be implemented in digital electronic circuitry, or in computer software, firmware, or hardware, including the structures disclosed in this document and their structural equivalents, or in combinations of one or more of them. The disclosed and other embodiments can be implemented as one or more computer program products, i.e., one or more modules of computer program instructions encoded on a computer readable medium for execution by, or to control the operation of, data processing apparatus. The computer readable medium can be a machine-readable storage device, a machine-readable storage substrate, a memory device, a composition of matter effecting a machine-readable propagated signal, or a combination of one or more them. The term “data processing apparatus” encompasses all apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, or multiple processors or computers. The apparatus can include, in addition to hardware, code that creates an execution environment for the computer program in question, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them. A propagated signal is an artificially generated signal, e.g., a machine-generated electrical, optical, or electromagnetic signal that is generated to encode information for transmission to suitable receiver apparatus.

[0109] A computer program (also known as a program, software, software application, script, or code) can be written in any form of programming language, including compiled or interpreted languages, and it can be deployed in any form, including as a standalone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program does not necessarily correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data (e.g., one or more scripts stored in a markup language document), in a single file dedicated to the program in question, or in multiple coordinated files (e.g., files that store one or more modules, subprograms, or portions of code). A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a communication network.

[0110] The processes and logic flows described in this document can be performed by one or more programmable processors executing one or more computer programs to perform functions by operating on input data and generating output. The processes and logic flows

can also be performed by, and apparatus can also be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application specific integrated circuit).

[0111] Processors suitable for the execution of a computer program include, by way of example, both general and special purpose microprocessors, and any one or more processors of any kind of digital computer. Generally, a processor will receive instructions and data from a read only memory or a random access memory or both. The essential elements of a computer are a processor for performing instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto optical disks, or optical disks. However, a computer need not have such devices. Computer readable media suitable for storing computer program instructions and data include all forms of non-volatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto optical disks; and CD ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

[0112] While this document contains many specifics, these should not be construed as limitations on the scope of an invention that is claimed or of what may be claimed, but rather as descriptions of features specific to particular embodiments. Certain features that are described in this document in the context of separate embodiments can also be implemented in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable sub-combination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a sub-combination or a variation of a sub-combination. Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results.

[0113] Only a few examples and implementations are disclosed. Variations, modifications, and enhancements to the described examples and implementations and other implementations can be made based on what is disclosed.

CLAIMS

What is claimed is:

1. A computer-implemented method to monitor status of a content provider network, comprising:
 - sending, from a trending server, a first request message to the content provider network, wherein the first request message has an associated type value;
 - receiving, at the trending server, a first response message from the content provider network;
 - determining, by the trending server, whether or not the first response message is normal by comparing with a trend record for the type value associated with the first request message;
 - selectively updating the trend record based on the first response message; and
 - generating an alert when it is determined that the first response message is abnormal.
2. The method of claim 1, further comprising:
 - receiving, at the trending server, a second request message from a user device; and
 - generating the first request message based on the second request message.
3. The method of claim 2, wherein the trend record is based on the second request message.
4. The method of claim 1, wherein the trend record accumulates a predetermined number of previously received response messages having the type value.
5. The method of claim 1, wherein the operation of determining whether or not the first response message is normal comprises:
 - comparing a parameter included in the first response message with a range of normal values for that parameter.
6. The method of claim 1, further comprising:
 - providing a second response message, based on the first response message, to a user device originating the first request message.

7. The method of claim 6, wherein the second response message comprises the first response message and the second response message is provided irrespective of whether the first response message is determined to be normal or not.

8. The method of claim 6, wherein the second response message comprises the first response message and further includes an alert message when the first response message is determined to be not normal.

9. The method of claim 1, wherein the trending server is located external to the content provider network.

10. The method of claim 10, further comprising:
configuring the trending server to be a part of a plurality of geographically separated servers configured to perform trending and monitoring operations and wherein at least one of the plurality of geographically separated server performs trending and monitoring without receiving the first request message from the user device.

11. An apparatus for monitoring operation of a content provider's network, comprising:
a status request transmission module that transmits a plurality of status requests to the content provider's network, each status request having an associated request type;
a status response reception module that receives status responses from the content provider's network;
a response sorting module that assigns, for every received status response, a corresponding request type based on a status request for which the status response was received;
a trending record-keeping module that maintains trend records based on a predetermined number of prior status responses received for each request type;
a trending estimation module that estimates a trend for each request type, based on the status response and a user setting for that request type;
a trend checker module that checks, for each received status response, whether or not the received status response is within the trend for the corresponding request type; and
a trend alert module that issues an alert when the received status response is not within the trend for the corresponding request type.

12. The apparatus of claim 11, further including:
a trend user interface module that receives an input for the user setting, wherein the input provides an input request type and an associated trend to use for the input request type.
13. The apparatus of claim 11, wherein:
the status request transmission module transmits the plurality of status request messages using a stateless protocol.
14. The apparatus of claim 11, wherein a first status request comprises a request to purchase a multimedia program and a corresponding first status response comprises a value representing cost of purchase of the multimedia program.
15. The apparatus of claim 11, wherein a first status request comprises a request to receive a web page and a corresponding first status response comprises contents of the web page.
16. A computer program product comprising a computer-readable storage medium having code stored thereupon, the code, when executed by a processor, causing the processor to implement a method of monitoring operation of a content provider's network, the method comprising:
receiving a first message from the content provider's network;
determining whether the first message is normal by comparing with a trend of previously received messages having a same message type as the first message;
selectively updating, based on the first message, the trend for the message type of the first message; and
generating an alert when the first message is determined to be not normal.
17. The computer program product of claim 16, wherein the method further comprises:
transmitting a status request message to the content provider's network and wherein the first message is received in response to the status request message.

18. The computer program product of claim 17, wherein the method further comprises:

- receiving a second message from the user device; and
- transmitting the status request message based on the second message.

19. A system for monitoring operation of a web service provider, the system comprising:

- one or more web servers under management of the web service provider, the one or more web servers located inside the web service provider's network;

- a plurality of monitoring servers positioned external to the web service provider's network and geographically separated from each other; wherein each monitoring server is configured to:

- transmit sync messages to the one or more web servers;
- maintain a record of expected responses to the sync messages;
- receive a status response messages from the one or more web servers;
- issue an alert when the status response message is not an expected response.

20. The system of claim 19, wherein each monitoring server is further configured to:

- update the record when a new status response message is received.