



- (51) **International Patent Classification:**
H04N 9/80 (2006.01)
- (21) **International Application Number:**
PCT/US2019/014218
- (22) **International Filing Date:**
18 January 2019 (18.01.2019)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
15/875,796 19 January 2018 (19.01.2018) US
- (71) **Applicant: MUX, INC.** [US/US]; 1182 Market St., Suite 425, San Francisco, CA 94102 (US).
- (72) **Inventors: BROWN, Adam;** 1182 Market St., Suite 425, San Francisco, CA 94102 (US). **DAHL, Jonathan;** 1182 Market St., Suite 425, San Francisco, CA 94102 (US). **HEFFERNAN, Steven;** 1182 Market St., Suite 425, San Francisco, CA 94102 (US). **SANFORD, Justin;** 1182 Market St., Suite 425, San Francisco, CA 94102 (US). **WARD, Matthew;** 1182 Market St., Suite 425, San Francisco, CA

94102 (US). **KIDDER, Scott;** 1182 Market St., Suite 425, San Francisco, CA 94102 (US). **DODSON, Benjamin;** 1182 Market St., Suite 425, San Francisco, CA 94102 (US). **DIEHL, Alex;** 1182 Market St., Suite 425, San Francisco, CA 94102 (US).

(74) **Agent: STEVENS, David, R.;** Stevens Law Group, 1754 Technology Drive, Suite 226, San Jose, CA 95110 (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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, 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, ST, SZ, TZ,

(54) Title: VIDEO ANALYTICS SYSTEM




Fig. 2

(57) **Abstract:** A video monitoring system can include multiple collectors to receive video beacon data from multiple video monitoring interface modules. At least one beacon stream is connected to receive data from multiple collectors. A processing module receives the beacon stream and provides a real-time event stream used for real-time data analysis and a video view stream used for long-term data analysis.



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))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

VIDEO ANALYTICS SYSTEM

TECHNICAL FIELD

[0001] The present disclosure relates to a video monitoring and analytics systems, and more specifically to a system for real-time data capture of information relating to video playback on video playback platforms.

BACKGROUND AND SUMMARY

[0002] Streaming video, movies, music, news, and business or personal communications are increasingly provided to users by a wide range of devices that are connected to a complex range of networks. As compared to broadcast cable or similar dedicated media distribution system, ensuring that this content is timely provided, with needed quality, can be challenging. For example, a poor-quality viewing experience for typical social media user on a smartphone, tablet, or personal computer can result from a number of factors beyond the control of a content provider, including inadequate hardware, operating systems or players having legacy software that do not have necessary updates, or poor wired or wireless network topology. However, in many situations, such hardware, software, or network limitations can be compensated for by appropriate quality adjustments to provided bitrate, video resolution, frame rates, or buffering times. In such an environment, it is valuable to have the capability to quickly determine whether a viewing experience is adequate for users, and make necessary adjustments to video content or players to provide the best possible viewing.

[0003] In some embodiments, such a novel video monitoring system can include multiple collectors to receive video beacon data from multiple video monitoring interface modules. A

beacon stream can be connected to receive data from multiple collectors and a processing module receives the beacon stream and provides a real-time event stream used for real-time data analysis. A separate video view stream also derived from the beacon stream is used for long-term data storage and deeper data analysis. In some embodiments an alert module is connected to the real-time event stream. The alert module can be adaptively configured to modify alert types and statistical thresholds. In other embodiments, a bot or crawler-detection module can be used to determine errors due to bot or crawler activity.

[0004] In other embodiments, the video view stream can be connected to a database writer. A connected database system (including database writer) can include a slim database supporting data required to respond to API requests from a customer interface module, as well as a view database supporting data collected from each available data view. Metadata information about individual events that define the available data view. The database system can also include an aggregation processing cluster for pre-computing aggregated data. A hashed combination of internal customer ID and names of the columns that compose an aggregate are used to partition a dataset into shards.

[0005] In another novel described embodiment, a video monitoring system includes a local video monitoring interface module to provide local video data. The video data can be sent to a remote video monitor system connectable to receive data from the local video monitoring interface. An automated video request identification module included in the remote video monitor system is configured to identify automated video requests that can generate video errors. Identification can be based on IP address, geographic origin, software version, vendor information, or characteristic viewing request patterns.

[0006] Another embodiment of a video monitoring system includes a local video monitoring interface module to provide local video error data and a remote video monitor system connectable to receive error data from the local video monitoring interface. An error history module is included in the remote video monitor system and is configured to store error history. A statistical anomaly module is used to compare received error data and error history to determine statistical anomalies. A machine learning module can be used to receive information from the statistical anomaly module and determine whether to send an error alert. In some embodiments, the machine learning module is able to adaptively set error rate standards for sending an error alert. The machine learning module can be a binary or other statistical classifier, as well as a deep convolutional or neural network.

[0007] In another embodiment a video monitoring system includes a local video monitoring interface module to provide local video data from a video player, the interface module having an associated viewer experience rating system a remote video monitor system is connectable to receive data from the local video monitoring interface and an error tracking system is used to adjust tracked error categories based on information provided by the viewer experience rating system. In some embodiments, errors are deliberately injected into the video player, while in others a placebo quality improvement button is associated with the video player. Video attributes or video player operating parameters can be adjusted at least in part due to viewer experience rating.

[0008] BRIEF DESCRIPTION OF THE DRAWINGS

[0009] Non-limiting and non-exhaustive embodiments of the present disclosure are described with reference to the following figures, wherein like reference numerals refer to like parts throughout the various figures unless otherwise specified.

[0010] FIG. 1 generally illustrates one embodiment of a video analytics system;

[0011] FIG. 2 illustrates in more detail an embodiment of a video analytics system;

[0012] FIG. 3 illustrates one embodiment of a method for processing data by a video analytics system;

[0013] FIG. 4 illustrates one example of a data payload for processing by a video analytics system;

[0014] FIG. 5 illustrates a bot/crawler detection procedure;

[0015] FIG. 6 illustrates a procedure for adaptive error rate handling;

[0016] FIG. 7 illustrates a procedure for viewer experience tracking;

[0017] FIG. 8 illustrates one embodiment of a database system for a real-time and near real-time video analytics system; and

[0018] FIG. 9A-C illustrate various exemplary screen shots for a customer user interface to a video analytics system.

DETAILED DESCRIPTION

[0020] FIG. 1 generally illustrates one embodiment of a real-time video monitoring and analytics system 100. This system 100 can evaluate completed video views in real or near real-time, comparing selected performance metrics against those of other video views across numerous breakdown classifications, including video title, browser, country, and more. In effect, system 100 is stream-processing application that applies a real-time map-reduce operation across multiple breakdown dimensions. Once a sufficient number of views have been collected (e.g. 100 views), a summary metric (e.g. error-rate) can be calculated for the views. The summary metric is compared against previous values for that customer account to determine whether it suggests a problem. Advantageously, system 100 does not require completion of a video before data analysis. Instead of requiring, for example, that data analytics for a thousand or so hour long content views must wait for an hour of viewing time to expire, system 100 can begin immediately providing useful information, leading to much lower latency in detecting problems and providing needed alerts.

[0021] This system is illustrated in FIG. 1, which shows a beacon stream 110 connected to a processor system 120. The beacon stream includes data payloads from large numbers of video clients that each support a video analytics plugin or module. This data can be collected as a real-time event stream 130 that can be stored in database 140, immediately viewed by partners 150 interested in the information, or cause automated creation of alerts in an alerting system 160. Alerts can be stored in the database 140 for later retrieval, or immediately provided to partners 150. A wide range of content can have data metrics tracked by the described system. For example, movies, commercial or user generated videos, music, sports, games or the like can have associated data analytics. This content can be viewed on smartphones, tablets, personal

computers, televisions, in-home or business projection systems, movie theaters, and augmented or virtual reality systems. Players can be dedicated apps or applications, browser plugins or addons, set top boxes or view cast systems, or any other suitable video viewing system that can support collection and transmission of content related data.

[0022] The beacon stream 110 can include collective information from video client players or viewers. For example, plugins, addons, or integrated software can be provided for players such as HTML5 Video Element, Video.js, Brightcove, Ooyala, Bitmovin, Chromecast, THEOplayer, JW Player, iOS, and Android. The players can provide technical and content data such as browser type, browser version, country, operating system, OS version, player name, player version, player software, player source type, source hostname, stream type, video title, video series, and content source. Hardware, network, and site or owner-specific data such as CPU and motherboard identity, operating frequency, connection type (e.g. wired or wireless), connection speed, packet drop or loss metrics, IP address, ISP, geographic location, and owner identity can be provided. Tracked data metrics can include an overall viewing performance score, number of playback failures, time to first frame, number of exits before video start, player load time, page load time, rebuffer percentage, rebuffer frequency, seek time, upscale percentage, and downscale percentage. In one embodiment, such second-scale beacon data is provided every 1 to 59 seconds. In other embodiments, minute-scale beacon data (e.g. 1 to 59 minutes) can be provided. In still other embodiments, frequency of beacon data transfer can be adaptively increased or decreased to match available content, player, or network resources. For example, a 1-5 second beacon might be provided for tablet player predominantly providing social media clip views, while a set-top box player (e.g. Roku or Amazon Fire) may only need to provide beacon data every 5-30 seconds. Advantageously, this procedure for beacon streaming

can be used to quickly identify a spike in error-rates occurring in the first minute of playing video content. This speed favorably compares to a system that waits until video completion and might need more than an hour to provide an alert, a latency that is not suitable for a real-time alerting system.

[0023] FIG. 2 illustrates in greater detail a system 200 similar to that discussed with respect to FIG. 1. Multiple video clients 202 are connected via the Internet or other suitable network connection to a load balancer 204. The load balancer 204 feeds into multiple collectors 206 which provide limited categorization and group the received data in one or more beacon streams 210. The beacon streams 210 are connected to a processor system 220, with all beacons from a single local module being directed to the same respective processors of the processor system 220. The processor system 220 further filters, organizes, begin statistical compilation, and directs the organized data into a real-time event stream 230 and a video view stream 232. The processor system 220 is supported by a view cache 234.

[0024] Real-time event stream 230 can be stored in a real-time database 240, immediately viewed by partners 250 interested in the information, or cause automated creation of alerts in an alerting system 260. Alerts can be stored in the database 240 for later retrieval, or immediately provided to partners 250 via calls to a customer interface module 270.

[0025] Longer term and more complex information can be stored and processed in database system 280. Database system 280 includes a master database 290 connected to customer interface module 270. Data is provided via the video view stream 232 connection to a database writer 281. The database writer 281 provides data to both slim view databases 282 and video view databases 283. The slim view databases 282 contain the data required to quickly respond to certain types of API requests (via customer interface module 270) for data about

views, before data has been aggregated and fully processed. The video view databases 283 instances contain all data collected about each view, as well as metadata information about the individual events that define that view. For commonly needed data, the master database 290 can receive information from aggregation workers 284 directed by a directed acyclic graph (DAG) workflow system 286 and supporting aggregation processors 288.

[0026] FIG. 3 illustrates operation of certain aspects of a real-time video monitoring system 300. A video viewer 310 with associated video data beacon sends data upon playback start. Additional status data is sent at predetermined or adaptively determined time intervals (typical 1-10 seconds) until video playback ends, causing playback finished beacon data to be sent. The beacon data is collected by beacon collection system 320, which can include dedicated local or remote servers, virtual server systems, cloud systems or the like.

[0027] When the playback finished beacon data is received, a real-time error alert system is invoked to provide preliminary data concerning error events and other view attributes. In operation, received video views can be continuously mapped to various easily determined view attributes such as browser type, country of viewing, or title of video. At determined intervals, or when the accumulator window is full, various summary metrics can be calculated such as number of particular browser types used in each country, or number of video titles completely watched. If no alert incident is currently open, summary metrics can be compared against previous history to see if any statistically unusual situations have developed. If there is statistically unusual behavior, an alert incident can be opened and preserved in an alert incident database 340. A customer 350 can be alerted, allowing them to take appropriate manually or automatically determined actions in near real-time. If an alert incident is already open, summary metrics can be compared against previous history to see if the identified statistically unusual

situations has returned to historically normal levels. If it has, the alert incident can be closed, with the alert incident database updated and notification of alert closure provided to customer 350. Otherwise, the summary metrics are added to the history of the error-rate alert system 330 and processing continues.

[0028] Beacon derived data can be grouped into various payload data packages as illustrated with respect to FIG. 4. For example, depending on designated data to be tracked, customer interest, or bandwidth or processing limitations, various levels of data detail can be provided. For example, data and metadata for a particular customer can require video view ID, IP address, IP mask, customer ID, and video source (as identified through hostname, IP address, and host ID). One or more data records can be included, with a timestamp, duration, and type and value of a customer requested metric being included.

[0029] Providing accurate real-time or near real-time statistics can involve a need for identifying outliers, bad data, or systemic errors such as might be caused by interaction of video players and systems with robots (bots) or crawlers. Web crawlers and scrapers are a class of software designed to follow links in web content and scan large sections of the Internet. Crawlers are often used to extract metadata from web content to support useful activities like web searches with search engines at Google, Facebook, or Microsoft often relying on data gathered by proprietary crawlers. Scrapers can be less benign, often downloading entire web pages and their resources including videos, images, and CSS.

[0030] Video error-rates and other performance metrics can be influenced by bots that crawl the web. Bots can trigger video playback in the browser engine, leading to video playback metrics being sent even though no user has requested video playback. Bots often have a higher error-rate compared to human views, and have the potential to inflate aggregate error-rates or

provide misleading data. For example, web crawlers and scrapers often use browser user-agent strings with software versions that lag behind versions of software they are emulating, which in effect misidentifies browser version. As another example, JavaScript crawlers often depend on a headless browser like PhantomJS and the Chromium Embedded Framework (CEF). Such crawlers using PhantomJS can trigger the creation of video-views in real-time video monitoring system such as described herein, with those views will likely registering as an error. This is extremely problematic because a wave of crawlers accessing the website of a customer can trigger an unwarranted alert or a flood of error-views. As another example of issues relating to crawlers, many crawlers often run in environments with self-imposed CPU and network throughput constraints. This is mutually beneficial to the crawler operator and the site being crawled. It enables the crawler to simultaneously access multiple sites while not overwhelming the site network and server resources. However, these constraints have the potential to distort performance metrics for video playback if the crawler happens to trigger the loading of an embedded video, while still limiting processor or network availability. Network throughput limits could lead to extremely poor measurements for key metrics like time-to-first-frame (TTFF). In effect, automated web monitoring or data capture bots and crawlers have the potential to trigger video playback and influence performance metrics unless steps are taken to exclude them.

[0031] As seen with respect to FIG. 5, in one embodiment a bot or crawler-detection procedure/module 500 can be used in a processing flow that inspects several metrics to determine whether the view came from a bot or crawler. If the view is suspected as coming from a bot or crawler, then special steps are taken to exclude it from alerting calculations and aggregate metrics. Bot and crawler detection begins with receipt of a new video beacon by a

processor 510. Beacon information is checked against an IP lookup service 520, geographic IP database 530, and an Autonomous System Network (ASN) database 540 that maps IPv4 and IPv6 addresses to ASN. The ASN database can be also be updated with Border Gateway Protocol prefix information. Further information can be provided by a crawler model 550 and historical dataset 550 that help in determining identifying signatures for bots or crawlers, along with frequency and IP or geographic associations. This information can be collected together and used to provide a basis for classification of crawler type, and likely associated errors. As new crawlers are found, or error associations change, the classification information used by the processor 510 can be updated, added, or otherwise modified.

[0032] In some embodiments, machine learning (ML) can be used to classify bots or crawlers. ML methods can include, but are not limited to neural networks, deep neural networks, convolutional narrow networks, recurrent neural networks, boosted tree networks or other suitable machine learning modules or system. Machine learning can be supervised or unsupervised. Multiple systems or types of systems can be used. Well-defined learning tasks can be developed to precisely define a video related attribute or function to be learned. Data can be automatically cleaned and put in a suitable form for usability by the ML algorithm, and initial and periodic updates to training data can be necessary.

[0033] In one preferred embodiment, multiple models for identifying bots or crawlers can be used, with the best performing model being selectable as a consensus model. The consensus model can be periodically updated as new bots or crawlers are found, or parameters or training data for available models are adjusted. The models can be used to identify bot or crawler intent, and data on bot or crawler view can be segregated from user viewing data.

[0034] FIG. 6 illustrates another embodiment 600 of a real-time video monitoring system that supports automated adaptive alerts. Defining and maintaining customer relevant error rates is tedious and prone to obsolescence as hardware, software, and network connectivity change. For error rate timeseries with evolving patterns, thresholds should be calibrated to reflect their most recent state. Data-driven thresholds require periodic readjustment. This approach emphasizes the importance of monitoring as a process and yields high rates of accuracy.

[0035] In the embodiment illustrated in FIG. 6, error data 610 is collected as historical data 620. For example, each error-rate summary metric can be saved with up to 20,000 error-rate calculations that preceded it. The error data 610 and historical data 620 is passed to a statistical anomaly module 630 that identifies potential statistical anomalies and passes the relevant information a machine learning system 640. If the alert is for an error class that customers can benefit from having an alert, and the error rate statistically significant, an alert 650 can be sent to a customer.

[0036] Similar to machine learning systems discussed in connection with bot and crawler identification, ML methods for providing adaptive alerts can include, but are not limited to neural networks, deep neural networks, convolutional narrow networks, recurrent neural networks, boosted tree networks or other suitable machine learning modules or system. Machine learning can be supervised or unsupervised. Multiple systems or types of systems can be used. One embodiment uses a recurrent neural network (RNN) to identify anomalies in error-rate. The RNN can be created with an autoencoder that identifies key features in predicting error-rates. Data would move to the autoencoder offline, which encodes the key features and moves it an online model RNN model using Long Short Term Memory (LSTM) cells. The RNN model can predict the number of errors, and if the divergence between prediction and observed is large

enough, an alert is triggered. In another embodiment, a statistical classifier is used to identify important alert conditions that match characteristics of historical alerts that affected large numbers of viewers. A statistical classifier uses predefined categories, and is used to categorize new probabilistic observations into said categories. When there are only two categories the problem is known as statistical binary classification. Statistical binary classification can include decision trees, random forests, Bayesian networks, support vector machines, neural networks, and systems that use logistic regression. Typically, the type of classifier used depends on the number of observations, the dimensionality of the feature vector, the expected noise in the data, and other factors.

[0037] FIG. 7 illustrates another embodiment 700 of a real-time video monitoring system that supports viewer experience tracking. Quality of Service (QoS) errors and/or deliberately injected errors 710 act to mar or reduce a user's viewer experience. Such artificial impairment by injecting errors can allow experiments to test importance of single errors, types of errors, or error groupings. Using after viewing polls, online questionnaires, surveys, or the like, viewer experience data 720 is gathered. In some embodiments features that ostensibly improve quality can be provided to indirectly give useful feedback. For example, use of placebo, non-operational, or placebo buttons, toggles, triggers, or the like can provide non-written feedback. In effect, the more often a placebo "quality improvement" button is pressed, the worse the viewing experience. Based on viewing experience information, tracked error categories (730) can be deleted, added, or adjusted to provide better matching between viewer experience and QoS errors that can tracked in real-time. In effect, viewing satisfaction can be more accurately correlated with trackable errors. Ultimately changes to video distribution or attributes can be made, as well as changes or adjustments to a video player, to improve viewer experience.

[0038] In one embodiment a user's viewing experience can fall into into three categories:

[0039] Satisfied (1): the user was satisfied with the experience

[0040] Tolerating (0.5): the user tolerated the experience, and was neither satisfied nor frustrated

[0041] Frustrated (0): the user was frustrated with the experience.

[0042] Various scores can be determined, including:

[0043] Overall Score: the min of scores 2-5

[0044] Playback Success: 1 if playback succeeds; 0 if playback fails before the user quits watching video.

[0045] Startup Time: 1 if time to first frame is under 2 seconds; 0.5 if time to first frame is under 8 seconds; 0 time to first frame is over 8 seconds.

[0046] Rebuffering: 1 if the view had no rebuffering. 0.5 if rebuffering percentage is 0%-12% AND rebuffering count is 1-4. 0 if rebuffering percentage of more than 12% or rebuffering count is 5+.

[0047] Playback Quality: 1 if number of source pixels is at least 50% of the number of display pixels. 0.5 if the number of source pixels 25%-50% of the number of display pixels. 0 if <25%.

[0048] An aggregate score can be determined by taking all of the individual view scores and dividing it by the number of views. In many situations, a single aggregate score is all that is required for most users of a real-time monitoring systems.

[0049] In one alternative embodiment illustrated with respect to FIG. 8, a real-time video monitoring system can be supported by a database system 880 that includes an aggregate master database 890 connected to customer interface module 870 having a dashboard API access.

Processed data identified with unique view IDs is supplied to database writer 881. The database writer 881 provides data to both slim view databases 882 and video view databases 883. The slim view databases 882 contain the data required to respond to API requests (via customer interface module 870) for data about views. The video view databases 883 instances contain all data collected about each view, as well as metadata information about the individual events that define that view. Typically, all of these servers are memory optimized cloud instances with, for example, hardware support by 4 CPU cores and about 30 GB of available memory.

[0050] On an hourly cadence (for example, 1-2 hours behind current time), a scheduled process copies data from the video view databases 883 onto a set of database workers referred to as the aggregation processing cluster 888. The aggregation processing cluster 888 is responsible for computing the pre-aggregated set of data that allow quick response to API requests via customer interface module 870. CPU intensive instances through cloud providers (each instance has 8 cores and ~15 GB of memory) can support the aggregation processing cluster 888. A process is followed to generate a list of all combinations of dimensional column combinations that need to be pre-aggregated, and then splits that list evenly across the available database instances. An example of a pre-aggregated combination is country code, operating system, and browser. Once all pre-aggregated data has been generated on the aggregation processing cluster 888, the resulting dataset can be transferred onto several instances of airflow workers 886. These instances run a set of scripts that sort the data set into thousands of individual files. The content of these files is determined by the partitioning scheme of the aggregation data database cluster. A hashed combination of internal customer ID and the names of the columns that compose the aggregate (e.g. country code, operating system, browser) can be used to partition the dataset into shards. The shards are then further partitioned into tables by customer ID

(allowing enforcement of retention durations on a per customer basis). Once all data has been sorted into the appropriate files, these files are loaded into the aggregation master database 890 directly from the airflow workers 886. When a user interacts with the customer interface module 870 having a dashboard API access, requests for data are run. Stored procedures that live on the aggregate master database 890 are callable to access both the pre-aggregated data living on the aggregate master database 890 as well as the raw data located on the slim view databases 882. The aggregate master database 890 can combine the datasets from the two locations into a single result set, which is in turn passed back to a user through a suitable user interface.

[0051] FIGS. 9A-C illustrate some example screenshots of graphical data presentation of generated result sets. For example, FIG. 9A illustrates various actions such as average start up time, play time, rebuffering time, and advertising play time. An x-axis is provided to indicate when events occurred in the video playback timeline.

[0052] FIG. 9B illustrates a mechanism for identifying errors or problems. A specific error that has occurred can be identified and support a callout with additional information. A triangle symbol on the video playback timeline will mark where and when it happened. Hovering over a bar in the video playback timeline will display its event category type, duration, and playback time. Also included is a list of events that make up the event category block. The first event listed is the start event, and the last ends the current event category block.

[0053] FIG. 9C illustrates a graphical presentation of user initiated events such as pause and seek explicitly on the video playback timeline since such events can explain non-user initiated event behavior on the video playback timeline. For example, the video playback timeline may contain a high amount of startup events, but that may be caused by a user seeking

or pausing frequently. Seeking and pause event category blocks are now marked with symbols above the main chart timeline.

[0054] In the foregoing description, reference is made to the accompanying drawings that form a part thereof, and in which is shown by way of illustration specific exemplary embodiments in which the disclosure may be practiced. These embodiments are described in sufficient detail to enable those skilled in the art to practice the concepts disclosed herein, and it is to be understood that modifications to the various disclosed embodiments may be made, and other embodiments may be utilized, without departing from the scope of the present disclosure. The foregoing detailed description is, therefore, not to be taken in a limiting sense.

[0055] Reference throughout this specification to “one embodiment,” “an embodiment,” “one example,” or “an example” means that a particular feature, structure, or characteristic described in connection with the embodiment or example is included in at least one embodiment of the present disclosure. Thus, appearances of the phrases “in one embodiment,” “in an embodiment,” “one example,” or “an example” in various places throughout this specification are not necessarily all referring to the same embodiment or example. Furthermore, the particular features, structures, databases, or characteristics may be combined in any suitable combinations and/or sub-combinations in one or more embodiments or examples. In addition, it should be appreciated that the figures provided herewith are for explanation purposes to persons ordinarily skilled in the art and that the drawings are not necessarily drawn to scale.

[0056] Embodiments in accordance with the present disclosure may be embodied as an apparatus, method, or computer program product. Accordingly, the present disclosure may take the form of an entirely hardware-comprised embodiment, an entirely software-comprised embodiment (including firmware, resident software, micro-code, etc.), or an embodiment

combining software and hardware aspects that may all generally be referred to herein as a “circuit,” “module,” or “system.” Furthermore, embodiments of the present disclosure may take the form of a computer program product embodied in any tangible medium of expression having computer-usable program code embodied in the medium.

[0057] Any combination of one or more computer-usable or computer-readable media may be utilized. For example, a computer-readable medium may include one or more of a portable computer diskette, a hard disk, a random access memory (RAM) device, a read-only memory (ROM) device, an erasable programmable read-only memory (EPROM or Flash memory) device, a portable compact disc read-only memory (CDROM), an optical storage device, and a magnetic storage device. Computer program code for carrying out operations of the present disclosure may be written in any combination of one or more programming languages. Such code may be compiled from source code to computer-readable assembly language or machine code suitable for the device or computer on which the code will be executed.

[0058] Embodiments may also be implemented in cloud computing environments. In this description and the following claims, “cloud computing” may be defined as a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned via virtualization and released with minimal management effort or service provider interaction and then scaled accordingly. A cloud model can be composed of various characteristics (e.g., on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service), service models (e.g., Software as a Service (“SaaS”), Platform

as a Service (“PaaS”), and Infrastructure as a Service (“IaaS”)), and deployment models (e.g., private cloud, community cloud, public cloud, and hybrid cloud).

[0059] The flow diagrams and block diagrams in the attached figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present disclosure. In this regard, each block in the flow diagrams or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It will also be noted that each block of the block diagrams and/or flow diagrams, and combinations of blocks in the block diagrams and/or flow diagrams, may be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions. These computer program instructions may also be stored in a computer-readable medium that can direct a computer or other programmable data processing apparatus to function in a particular manner, such that the instructions stored in the computer-readable medium produce an article of manufacture including instruction means which implement the function/act specified in the flow diagram and/or block diagram block or blocks. Many modifications and other embodiments of the invention will come to the mind of one skilled in the art having the benefit of the teachings presented in the foregoing descriptions and the associated drawings. Therefore, it is understood that the invention is not to be limited to the specific embodiments disclosed, and that modifications and embodiments are intended to be included within the scope of the appended claims. It is also understood that other embodiments of this invention may be practiced in the absence of an element/step not specifically disclosed herein.

CLAIMS:

1. A video monitoring system comprising:
 - multiple collectors to receive video beacon data from multiple video monitoring interface modules;
 - a beacon stream connected to receive data from multiple collectors;
 - a processing module to receive the beacon stream and provide a real-time event stream used for real-time data analysis and a video view stream used for long-term data analysis; and
 - an alert module connected to the real-time event stream that can be adaptively configured to modify alert types and statistical thresholds.
2. The video monitoring system of claim 1, wherein beacon data from each respective video monitoring interface modules is sent to a same processor in the processing module.
3. The video monitoring system of claim 1, further comprising connection of the video view stream to a database writer.
4. The video monitoring system of claim 1, further comprising a database system having a slim database supporting data required to respond to API requests from a customer interface module.

5. The video monitoring system of claim 1, further comprising a database system having a view database supporting data collected from each available data view, and metadata information about individual events that define the available data view.
7. The video monitoring system of claim 1, further comprising a database system having a master database supporting a slim database and a view database.
8. The video monitoring system of claim 1, further comprising a database system having an aggregation processing cluster for pre-computing aggregated data.
9. The video monitoring system of claim 1, wherein a hashed combination of internal customer ID and names of the columns that compose an aggregate are used to partition a dataset into shards.
10. The video monitoring system of claim 1, further comprising a bot or crawler-detection module to determine errors due to bot or crawler activity.
11. A video monitoring system comprising:
 - a local video monitoring interface module to provide local video data;
 - a remote video monitor system connectable to receive data from the local video monitoring interface; and

an automated video request identification module included in the remote video monitor system and configured to identify automated video requests.

12. The video monitoring system of claim 11, wherein identification of automated video requests is based on detected IP.

13. The video monitoring system of claim 11, wherein identification of automated video requests is based on determined geographic origin of video requests.

14. The video monitoring system of claim 11, wherein identification of automated video requests is based on identified viewing request patterns.

15. A video monitoring system comprising:

a local video monitoring interface module to provide local video error data;

a remote video monitor system connectable to receive error data from the local video monitoring interface;

an error history module included in the remote video monitor system and configured to store error history; and

a statistical anomaly module to compare received error data and error history to determine statistical anomalies.

16. The video monitoring system of claim 15, further comprising a machine learning module to receive information from the statistical anomaly module and determine whether to send an error alert.
17. The video monitoring system of claim 15, further comprising a machine learning module to receive information from the statistical anomaly module, with the machine learning module able to adaptively set error rate standards for sending an error alert.
18. The video monitoring system of claim 15, further comprising a machine learning module that is a statistical classifier.
19. The video monitoring system of claim 15, further comprising
a beacon stream connected to receive data from multiple collectors; and
a processing module to receive the beacon stream and provide a real-time event stream used for real-time data analysis and monitoring by the statistical anomaly module.
20. A video monitoring system comprising:
 - multiple collectors to receive video beacon data from multiple video monitoring interface modules;
 - a beacon stream connected to receive data from multiple collectors;

a processing module to receive the beacon stream and provide a real-time event stream used for real-time data analysis and a video view stream used for long-term data analysis;

an error history module configured to store error history from the real-time data analysis and a video view stream used for long-term data analysis; and

a statistical anomaly module to compare received error data and error history to determine statistical anomalies.

21. The video monitoring system of claim 20, further comprising a machine learning module to receive information from the statistical anomaly module and determine whether to send an error alert.
22. The video monitoring system of claim 20, further comprising a machine learning module to receive information from the statistical anomaly module, with the machine learning module able to adaptively set error rate standards for sending an error alert.
23. The video monitoring system of claim 20, further comprising a machine learning module that is a statistical classifier.
24. The video monitoring system of claim 20, wherein beacon data from each respective video monitoring interface modules is sent to a same processor in the processing module.

25. A video monitoring system comprising:

a local video monitoring interface module to provide local video data from a video player, the interface module having an associated viewer experience rating system;

a remote video monitor system connectable to receive data from the local video monitoring interface; and

an error tracking system able to adjust tracked error categories based on information provided by the viewer experience rating system.

26. The video monitoring system of claim 25, wherein errors are deliberately injected into the video player.

27. The video monitoring system of claim 25, wherein a placebo quality improvement button is associated with the video player.

28. The video monitoring system of claim 25, further comprising a system for video attribute adjustment at least in part due to viewer experience rating.

29. The video monitoring system of claim 25, further comprising a system for video player adjustment at least in part due to viewer experience rating.

100
↙

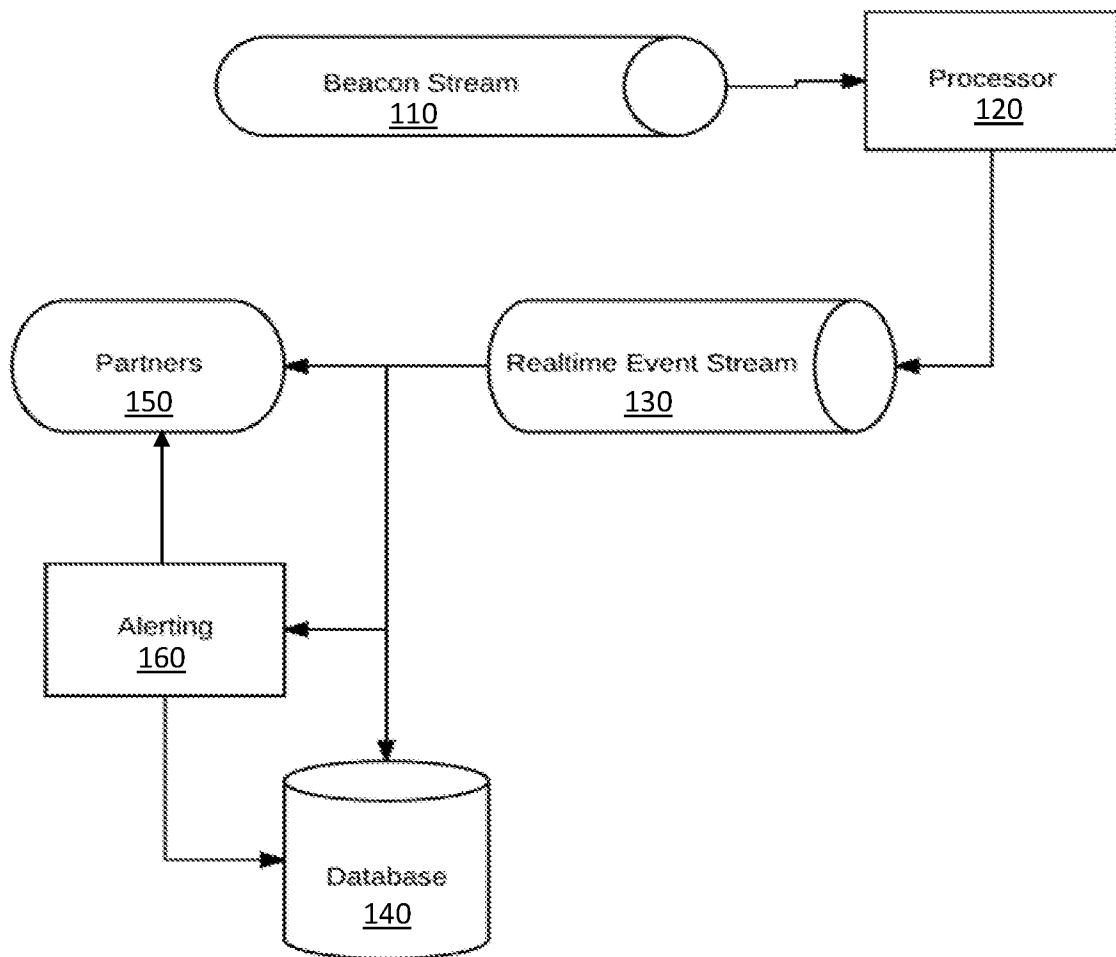


Fig. 1




Fig. 2

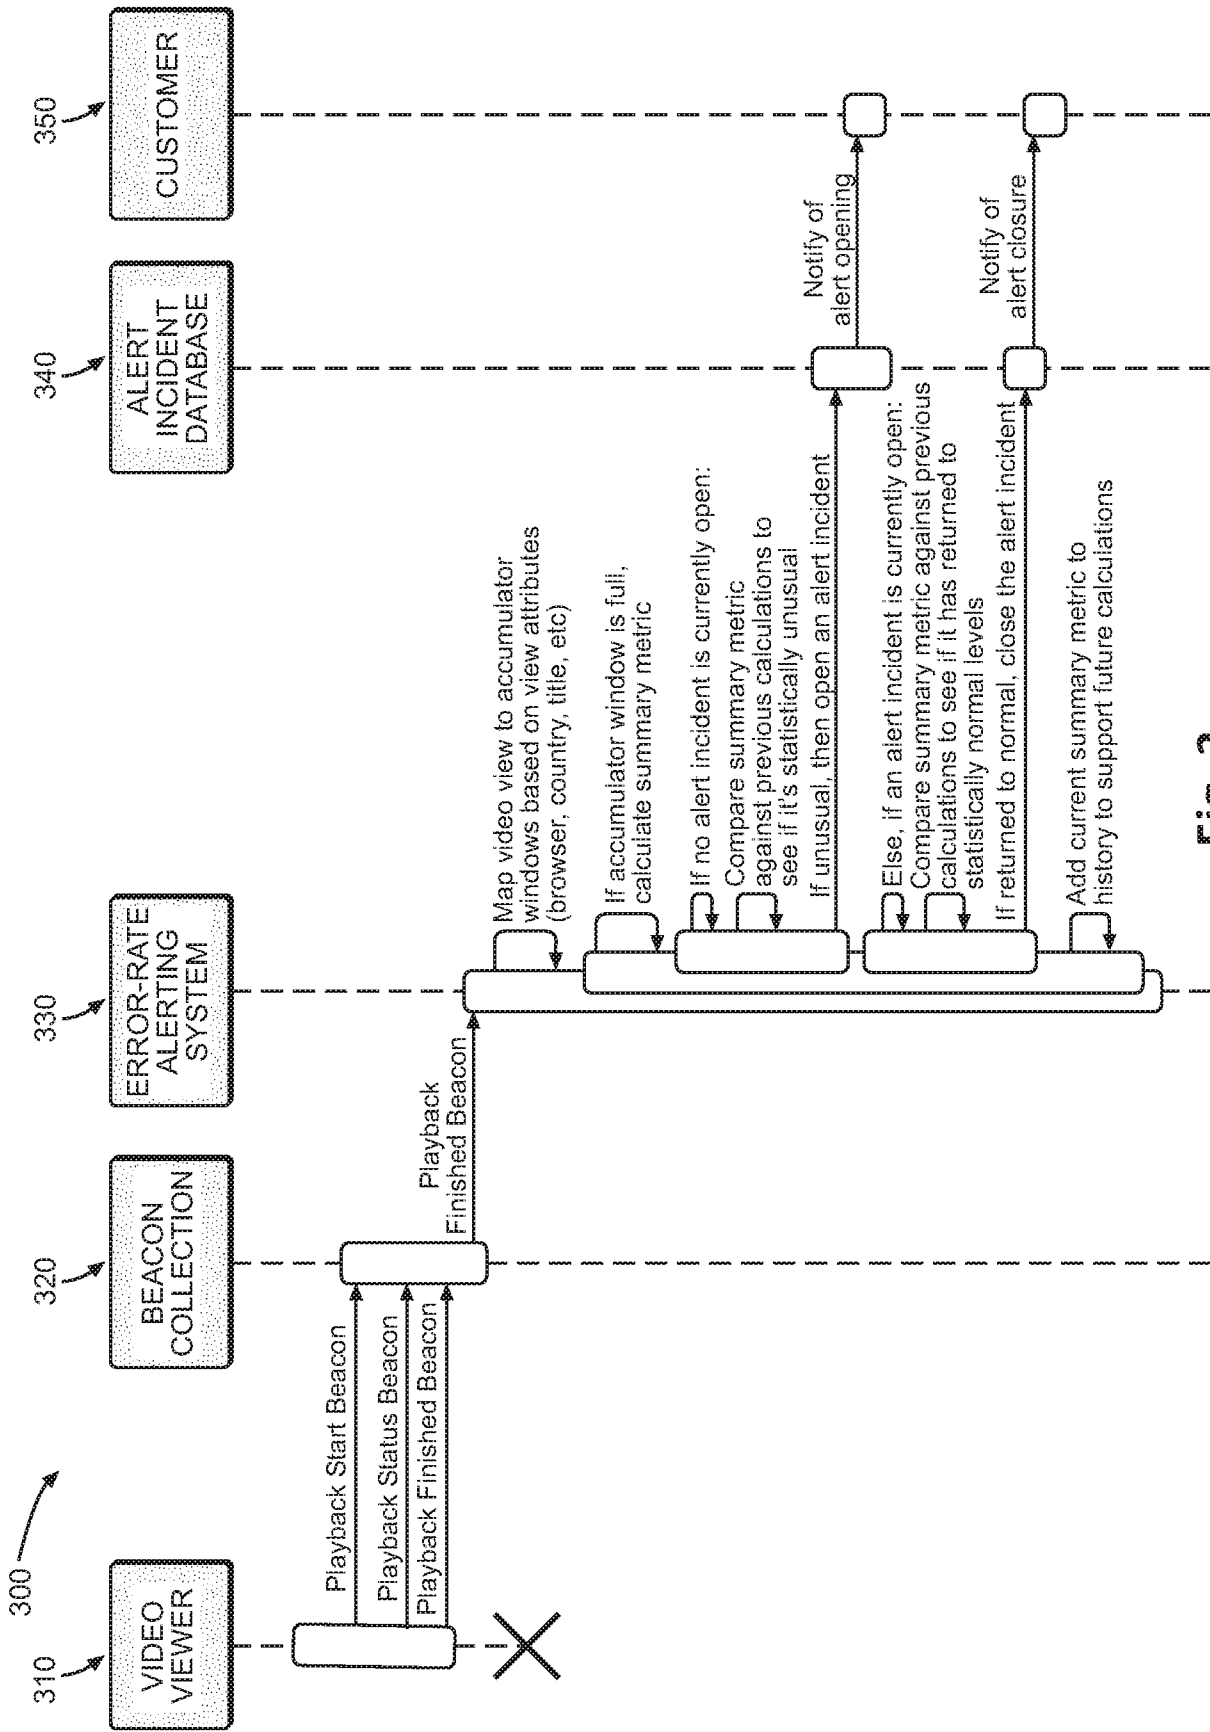


FIG. 3

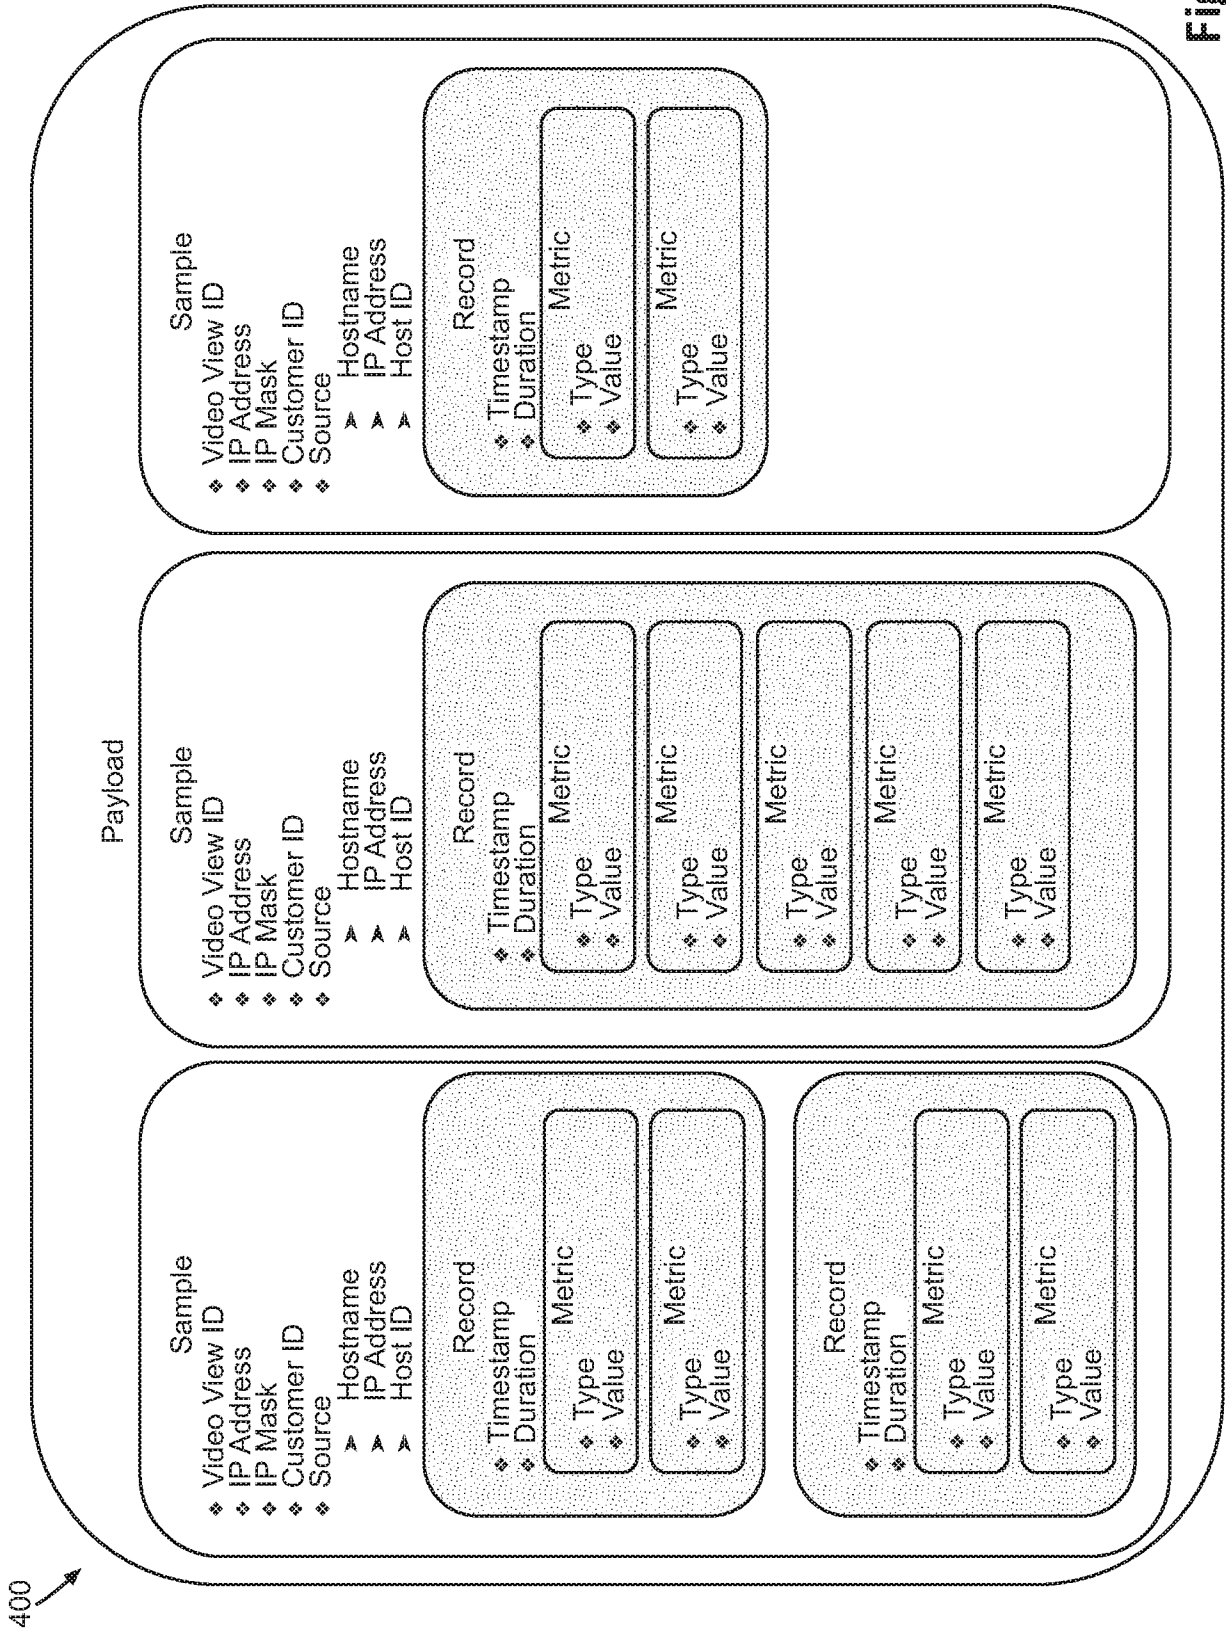


Fig. 4

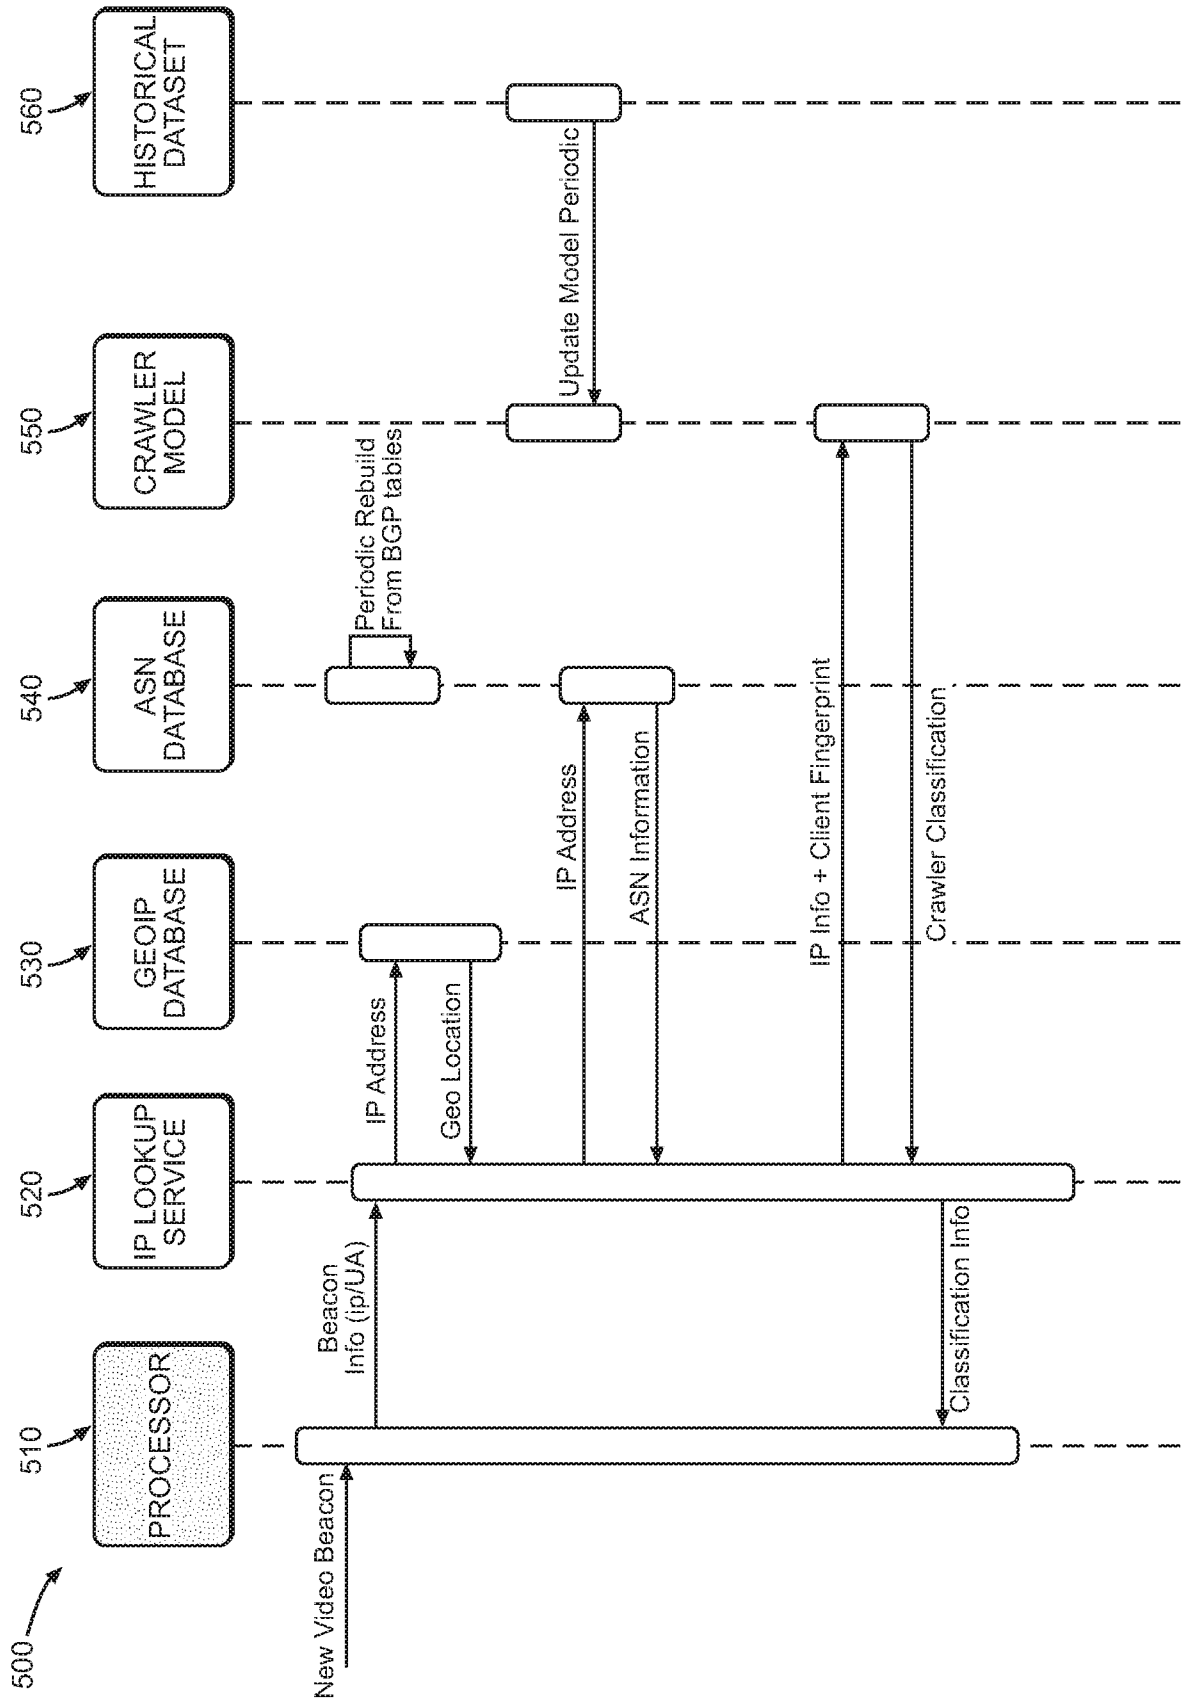


Fig. 5

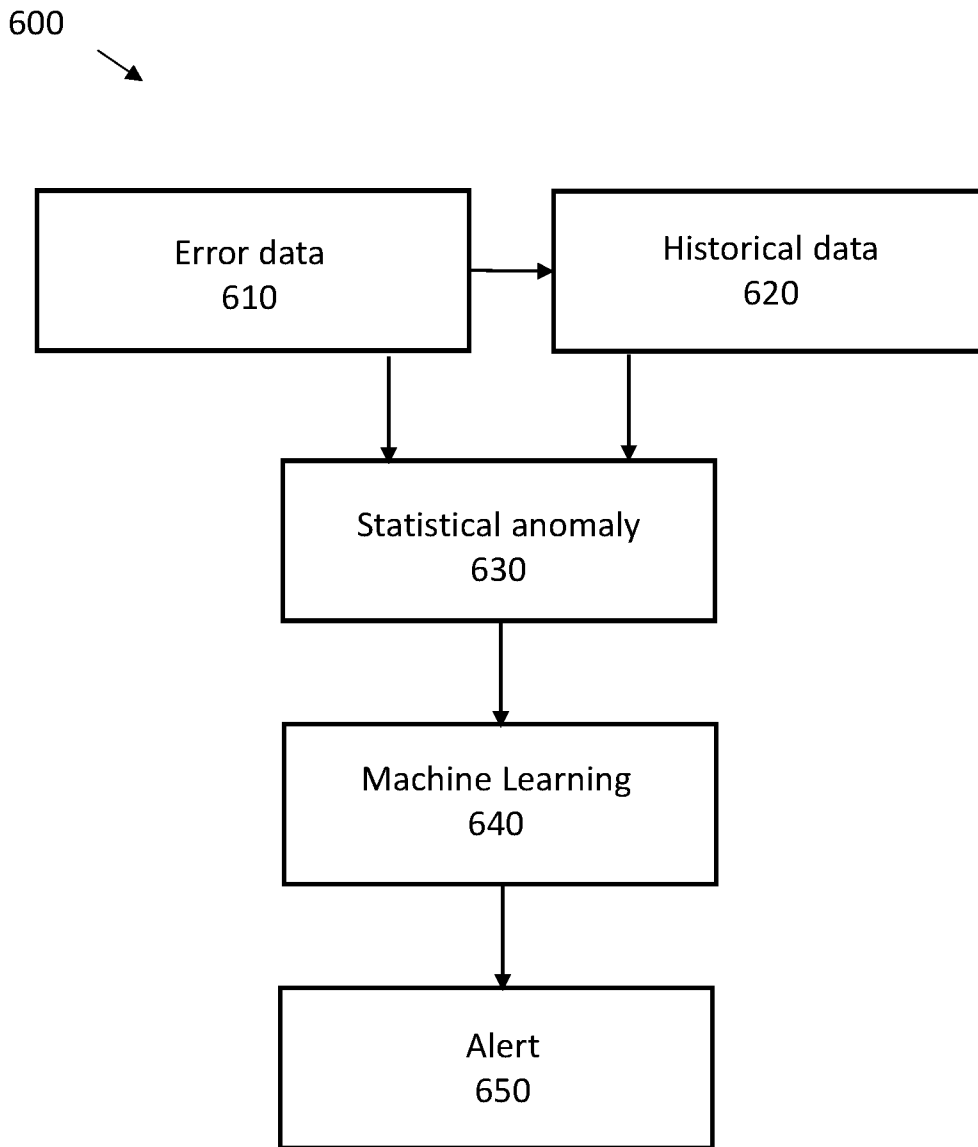


Fig. 6

700 ↘

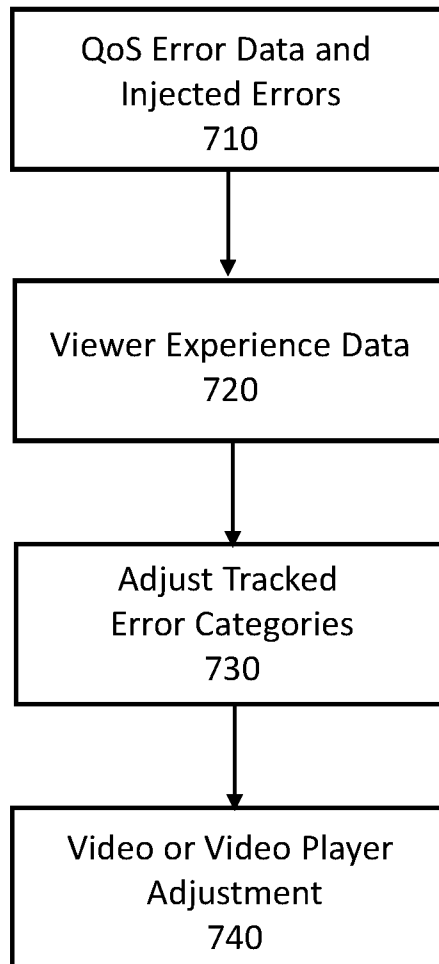


Fig. 7




Fig. 8




Fig. 9A

910

Code 4
① MEDIA_ERR_SRC_NOT_SUPPORTED
[view this error](#)




Fig. 9B

930 ↗




FIG. 9C

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US 19/14218

<p>A. CLASSIFICATION OF SUBJECT MATTER IPC(8) - H04N 9/80 (2019.01) CPC - G11B 27/105, G11B 2220/2562, H04N 5/85, H04N 9/8042, G11B 27/329, H04N 9/8042, G11B 27/105, G11B 27/3027, G11B 2220/20, H04N 5/775, G11B 27/34, H04N 9/8042</p>		
<p>According to International Patent Classification (IPC) or to both national classification and IPC</p>		
<p>B. FIELDS SEARCHED</p>		
<p>Minimum documentation searched (classification system followed by classification symbols) See Search History Document</p>		
<p>Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched See Search History Document</p>		
<p>Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) See Search History Document</p>		
<p>C. DOCUMENTS CONSIDERED TO BE RELEVANT</p>		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X -- Y	US 2015/0067185 A1 (AKAMAI TECHNOLOGIES, INC.) 05 March 2015 (05.03.2015), entire document especially abstract; para [0006], [0007], [0013], [0014], [0038]-[0042], [0057], [0064], [0071]-[0089]	1-5, 8 ----- 7, 9, 10
Y	US 2017/0085634 A1 (WAL-MART STORES, INC.) 23 March 2017 (23.03.2017), para [0047], [0050], [0052], [0062]	7, 9
Y	US 2014/0040015 A1 (Haley et al.) 06 February 2014 (06.02.2014), para [0025], [0042]	10
<p><input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/> See patent family annex.</p>		
<p>* Special categories of cited documents:</p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier application or patent but published on or after the international filing date</p> <p>"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)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p> <p>"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</p> <p>"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</p> <p>"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</p> <p>"&" document member of the same patent family</p>		
<p>Date of the actual completion of the international search 12 May 2019 (12.05.2019)</p>		<p>Date of mailing of the international search report 24 MAY 2019</p>
<p>Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-8300</p>		<p>Authorized officer: Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774</p>

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 19/14218

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

--PLEASE SEE EXTRA SHEET--

1. As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:
1-5, 7-10

Remark on Protest

- The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- No protest accompanied the payment of additional search fees.

Box No. III Observations where unity of invention is lacking

This application contains the following inventions or groups of inventions which are not so linked as to form a single general inventive concept under PCT Rule 13.1.

Group I: claims 1-5, 7-10 directed to a video monitoring system configured to modify alert types and statistical thresholds based on a real-time event stream.

Group II: claims 11-14 directed to a video monitoring system configured to identify automated video requests.

Group III: claims 15-24 directed to a video monitoring system configured to compare received error data and error history data to determine statistical anomalies.

Group IV: claims 25-29 directed to a video monitoring system configured to adjust tracked error categories based on a viewer experience rating system.

The inventions listed as Groups I-IV do not relate to a single general inventive concept under PCT Rule 13.1 because under PCT Rule 13.2 they lack the same or corresponding technical features for the following reasons:

Groups II-IV do not include the special technical feature of an alert module connected to the real-time event stream that can be adaptively configured to modify alert types and statistical thresholds of Group I.

Groups I and III-IV do not include the special technical feature of an automated video request identification module included in the remote video monitor system and configured to identify automated video requests of Group II.

Groups I-II and IV do not include the special technical features of an error history module included in the remote video monitor system and configured to store error history; and a statistical anomaly module to compare received error data and error history to determine statistical anomalies of Group III.

Groups I-III do not include the special technical feature of an error tracking system able to adjust tracked error categories based on information provided by the viewer experience rating system of Group IV.

None of these technical features are common to the other groups, nor do they correspond to a special technical feature in the other groups. Therefore, unity of invention is lacking.

Shared technical features of the Groups are taught at least by US 2016/0080804 A1 to Comprobo Limited (hereinafter, Comprobo) and US 2015/0067185 A1 to AKAMAI TECHNOLOGIES, INC. (hereinafter, Akamai):

Groups II, III and IV: a local video monitoring interface module to provide local video data. Comprobo, (Fig. 1; para [0035], [0056], [0132], [0134], [0139], [0153]) and Akamai, (Fig. 3, 4, 6; para [0026], [0045], [0058], [0061], [0103]).

Group II, III and IV: a remote video monitor system connectable to receive data from the local video monitoring interface. Comprobo, (Fig. 1; para [0035], [0056], [0132], [0134], [0139], [0153]) and Akamai, (Fig. 3, 4, 6; para [0026], [0045], [0058], [0061], [0103]).

Groups I, multiple collectors to receive video beacon data from multiple video monitoring interface modules. Comprobo, (Fig. 1; para [0035], [0056], [0132], [0134], [0139], [0153]) and Akamai, (Fig. 3, 4, 6; para [0026], [0045], [0058], [0061], [0103]).

Groups I and III: a beacon stream connected to receive data from multiple collectors. Comprobo, (Fig. 1; para [0035], [0056], [0132], [0134], [0139], [0153]) and Akamai, (Fig. 3, 4, 6; para [0026], [0045], [0058], [0061], [0103]).

Groups I and III: a processing module to receive the beacon stream and provide a real-time event stream used for real-time data analysis and a video view stream used for long-term data analysis. Comprobo, (Abstract; para [0024], [0112], [0128], [0129], [0139], [0143], [0153]) and Akamai, (Fig. 3, 4, 6; abstract; para [0018], [0026], [0038], [0045], [0058], [0061], [0086]-[0088], [0103]).

Groups II, III and IV: a remote video monitor system connectable to receive data from the local video monitoring interface. Comprobo, (Fig. 1; para [0035], [0056], [0132], [0134], [0139], [0153]) and Akamai, (Fig. 3, 4, 6; para [0026], [0045], [0058], [0061], [0103]).

Accordingly, unity of invention is lacking as the groups do not share a same or corresponding special technical feature providing a contribution over prior art.

* Note - Claim 6 is missing from the recitation of Claims.