telesoft
everyday. moving forward

Segmented monitoring of 100Gbps data containing CDN video

Telesoft White Papers

Steve Patton
Senior Product Manager
23rd April 2015

Telesoft White Paper:

Segmented monitoring of 100Gbps data containing CDN video

IP Video 'The Challenge'

The growth in internet traffic caused by increasing use of high bandwidth services including video causes problems for systems that need to monitor, filter or process traffic within the network before it reaches the end-user or consumer of the data.

For systems that need to monitor, route or process IP traffic in the network, it is becoming increasingly difficult and commercially unsustainable to continue to purchase additional capacity in both hardware and software.

IP Traffic to 2018

Global IP traffic will reach 1.1 zettabytes (1000 exabytes) per year or 91.3 exabytes (one billion gigabytes) per month in 2016. By 2018, global IP traffic will reach 1.6 zettabytes per month.

Content delivery networks will carry over half of Internet traffic by 2018. Fifty-five percent of all Internet traffic will cross content delivery networks by 2018 globally.

It would take an individual over 5 million years to watch the amount of video that will cross global IP networks each month in 2018. Every second, nearly a million minutes of video content will cross the network by 2018.

Globally, IP video traffic will be 79 percent of all consumer Internet traffic in 2018, up from 66 percent in 2013, not including video exchanged through peer-to-peer (P2P) filesharing. The sum of all forms of video (TV, video on demand [VoD], Internet, and P2P) will be in the range of 80 to 90 percent of global consumer traffic by 2018.

Internet video to TV traffic will be 14 percent of consumer Internet video traffic by 2018, up from 11 percent in 2013.

Consumer VoD traffic will double by 2018. The amount of VoD traffic by 2018 will be equivalent to 6 billion DVDs per month.

Content delivery network traffic will deliver over half of all internet video traffic by 2018. By 2018, 67 percent of all Internet video traffic will cross content delivery networks, up from 53 percent in 2013.

- Cisco© Networking Index (VNI) -10/06/14

Segmented Monitoring

Increase the effectiveness of your monitoring infrastructure by identifying traffic types that are of high value to an end user, or a network operator and then treat each differently.

This can be achieved a number of different ways, such as;

- Traffic can be identified by a user so that all traffic for a specific individual is identified and then grouped together for treatment.
- Traffic for a particular service is identified and grouped together, such as high definition paid for video content for a number of users.

Treatment of the traffic can include routing to specific processing systems for analysis or discarding. Discarding certain unwanted traffic types reduces the traffic volume that is routed to monitoring and analytical tools, by routing and discarding certain traffic types; monitoring and analysis hardware and software can be optimised and reduced.

Since video is forecasted to utilise the largest bandwidth, for systems that need to process video separately, there is a real need to identify video streams quickly for separate treatment.

One example is quality assurance for the delivery of high definition paid for video content, to maintain consumer confidence it is vitally important for a content provider and network operator to know that the quality of the delivered video stream meets with the required paid for quality standard. In this scenario the video content is identified and then routed to quality analysis tools.

However, individuals may use video to communicate. Thus, identifying video content alone may not be sufficient. It may also be necessary to distinguish between video content served as a Video on Demand (VoD) service and video content served as a peer to peer video communication e.g. peer to peer video chat and in application video chat etc.

A content delivery network or content distribution network (CDN) is a large distributed system of servers deployed in multiple data centres across the Internet. The goal of a CDN is to serve content to end-users with high availability and high performance. CDNs serve a large fraction of the Internet content today, including web objects (text, graphics and scripts), downloadable objects (media files, software, documents), applications (e-commerce, portals), live streaming media, on-demand streaming media and social networks.

Content providers such as media companies and e-commerce vendors pay CDN operators to deliver their content to their audience of end-users. In turn, a CDN pays ISPs, carriers and network operators for hosting its servers in their data centres.

Identifying domain names of well-known content providers (for example netflix.com, amazonprime.com, hulu.com…) is insufficient to identify video streams, since the use of CDNs means that the actual source of the video data is likely to be on a server that may not be easily identifiable with known VoD provider identities.

Consequently, there is a need for improved systems to monitor, filter or process traffic at the high speed data rates encountered in today's networks where existing techniques are becoming increasingly inadequate.

Delivery of Content Sequence

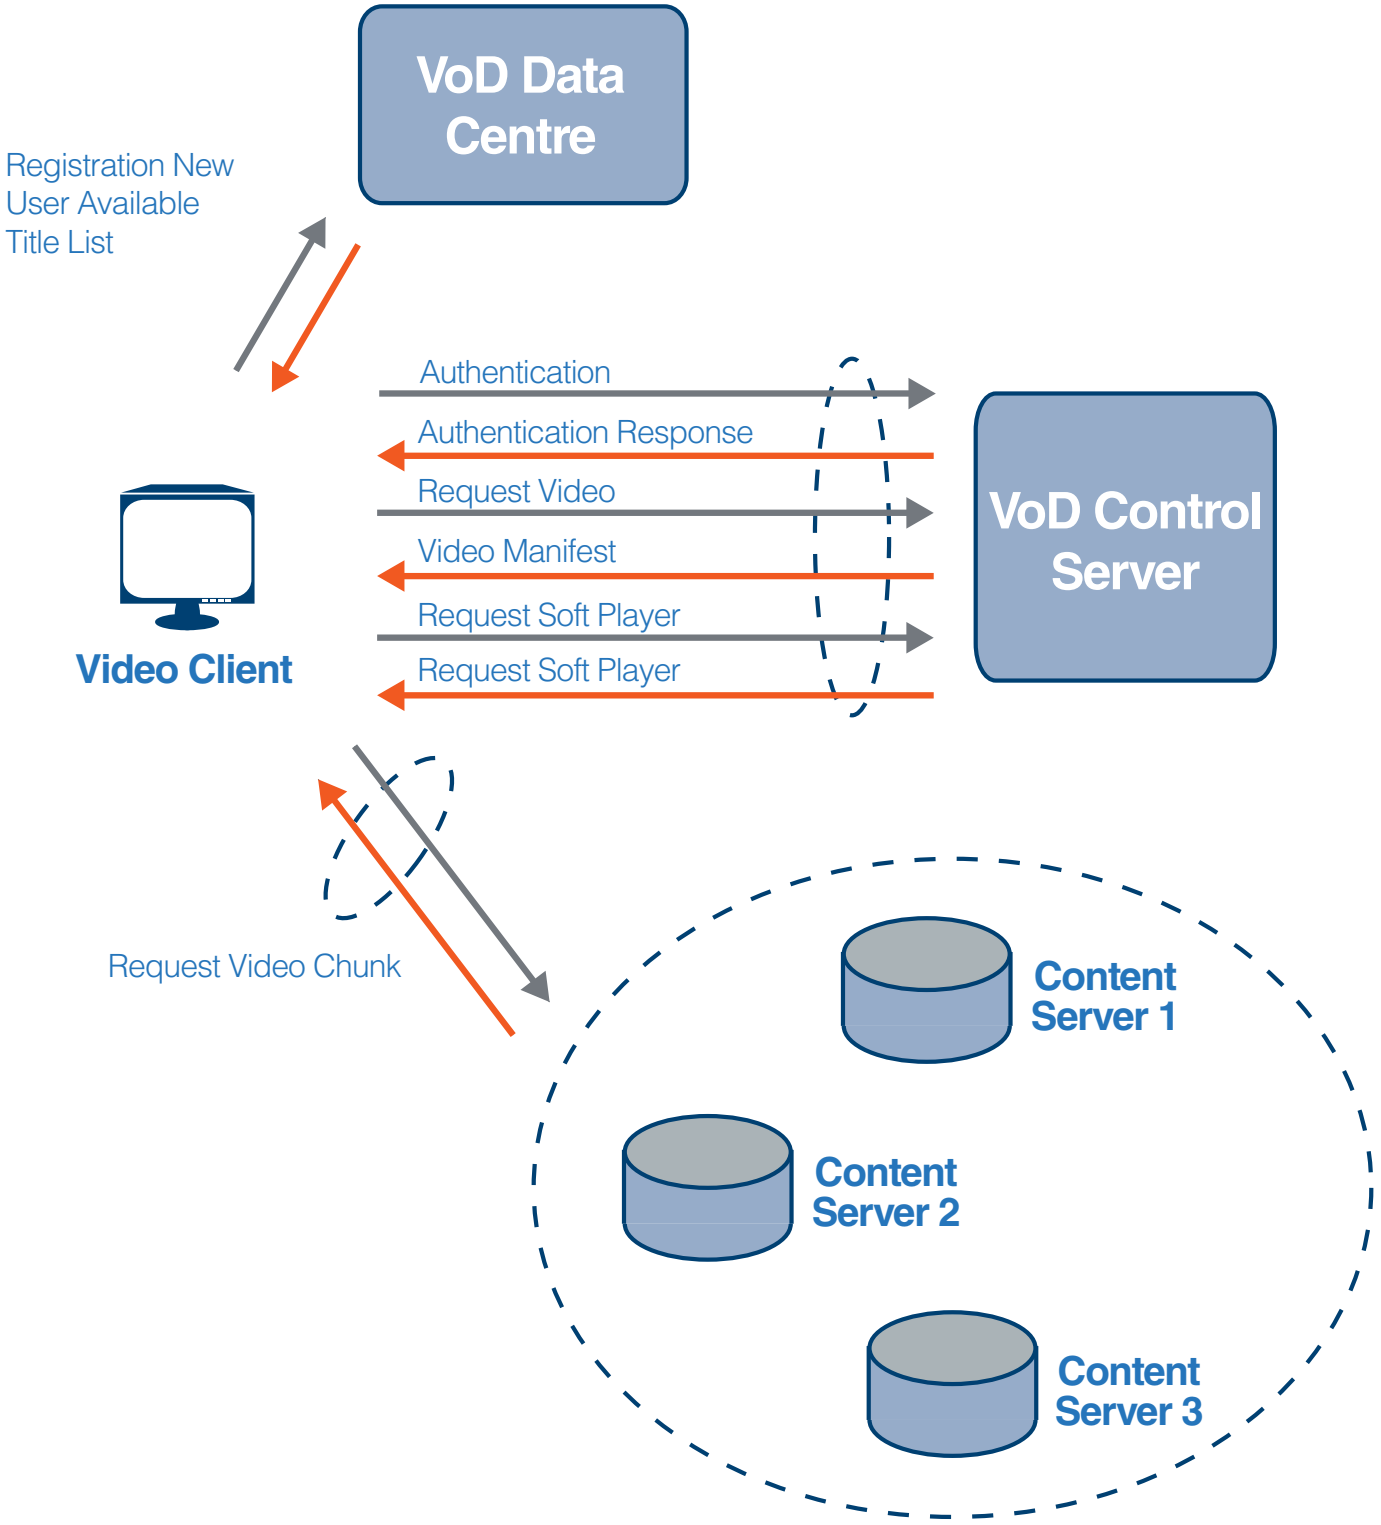

Basic Principles of Video on Demand (VoD)

1. Registration and payment

A subscriber (subscriber 100) to a VoD service registers an account and provides payment information to a server located in a VoD Data Centre 101. Once authorised, subscriber 100 will be given access to lists of available content. Registration, payment and browsing lists of hosted content is not bandwidth intensive, it is therefore of little interest to identify such activity in order to relieve traffic loading on monitoring systems.

2. Selection of video playback

For some VoD services, when a title is selected for play the first operation is to download a soft (software) based player appropriate to the viewing device (phone, table, PC, set-top box, TV…) from a VoD Control Server 102. Although downloaded for each different title played, the soft player is small, which is a low bandwidth activity compared to streaming and viewing actual video content, making identification unnecessary.

3. The user is then authenticated

Authentication takes place between subscriber 100 and VoD control server 102, ensuring that the subscriber is permitted to access the requested content.

4. Manifest file fetch

After authentication, the player on subscriber 100's device fetches a manifest file from the control server 102. This manifest file contains information telling the video client (subscriber 100) where to fetch the video content, yielding a low bandwidth level which is of little interest to most applications.

Steps 3 and 4 are delivered over encrypted sessions, such as SSL, making extracting useful information to identify video streams difficult and time consuming, especially at very high data rates e.g. at 100Gbps and higher.

5. Video streaming starts according to the manifest file

Video streaming is controlled by instructions in the manifest file that the video client (subscriber 100) downloads. The manifest file provides the player with information to conduct adaptive video streaming from content server 105.

Manifest files are client-specific and are generated according to each client's playback capability. For instance, if the user player indicates it is capable of rendering h.264 encoded video, h.264 format video is included in the manifest file. If the player indicates that it can only playback .wmv format, then only .wmv format video will be included.

The manifest file contains several key pieces of information including the list and priority/rank of the available CDNs content servers, location of trickplay data, video/audio chunk URLs for multiple quality levels and timing parameters such as time-out interval and polling interval.

6. Audio and video chunk downloading

Audio and video contents are downloaded in chunks (also known as segments). Download sessions are more frequent at the beginning, allowing build up to the player buffer. Once the buffer is sufficiently filled, downloads become periodic. The manifest file contains multiple audio and video quality levels. For each quality level, it contains the URLs for individual CDNs.

7. Trickplay

Most players support pause, rewind, forward and random seek, collectively referred to as "trickplay". This is achieved by downloading a set of thumbnail images for periodic snapshots. The thumbnail resolution, pixel aspect, trickplay interval and CDN from where to download the trickplay file are described in the manifest file.

8. Dynamic Adaptive Streaming (DAS)

Most players support fetching and playing video and audio content at variable quality and bit rate allowing continuous viewing as network conditions (latency, congestion and available bandwidth) change e.g. DASH (DAS over HTTP) and Apple's HTTP Live Streaming (HLS) solution.

Adaptive streaming works by breaking the content into a sequence of small HTTP-based file segments, each segment contains a short interval of playback time of content, which potentially may be many hours in duration, such as a movie or the live broadcast of a sports event. The content is typically made available at a variety of different bit rates. As the content is played back, the client automatically selects from the alternative next segments to download and playback based on current network conditions. The client selects the segment with the highest bit rate possible that can be downloaded in time for playback without causing stalls or buffering events in the playback. Thus, an adaptive bit rate client can seamlessly adapt to changing network conditions and provide high quality playback without stalls or buffering events.

The Telesoft STR (Segmented Traffic Router)

The Telesoft STR uses latest generation FPGAs to select CDN video media sources from within 100Gbps data, connected as either 100GbE (LR4 or SR10) or from 10x10GbE. Data streams sourced from these identified CDN media servers can then be selectively monitored or discarded, allowing video Quality of Experience (vQoE) systems to only process video traffic and other monitoring and analysis equipment to only process the data that they require e.g. voice traffic.

As traffic rates rise (mainly due to video), existing monitoring equipment requires little if no modification to cope with the changes in bandwidth and traffic throughput. Our specialist equipment can be seamlessly added to selectively process video traffic.

Meeting the Challenge

One method to reduce the load on monitoring equipment is to identify and select video traffic streams that originate from a CDN and then process or route this separately, either to be processed for analysis e.g. video Quality of Experience (vQoE) and Quality of Experience (QoE), or to be discarded to allow probes that are not interested in video to process the data they need to.

In the case of streaming video (broadcast video or VoD), the system is only concerned with the content delivery itself – the pre-authentication, authentication and handover steps are of no interest largely because the final handover is opaque.

To detect the video content delivery streams a number of detection stages can be employed, including identification of HTTP GET requests to fetch video chunks, examination of HTTP GET header fields, uri analysis and analysis of DASH requests if present.

A matching combination of these factors indicates that the traffic in the flow being analysed is streaming video. If the HTTP referrer field is present, then this can confirm the match.

As an additional and optional stage, it may be advantageous to identify the CDN provider (e.g. Level 3, Akamai, Amazon …) to allow all traffic from a specific CDN provider to be treated in the same way. In some circumstances, a Fully Qualified Domain Name (FQDN) may be present or this may be done by matching addressing to a known allocated IP address range.

Once a CDN source of VoD is identified, all future communication sessions from that IP address can be treated in the same way and can be routed for specific monitoring processing or discarded.

Conclusion

Rising use of high bandwidth services such as VoD is already starting to cause throughput problems for network monitoring systems. This is only set to become worse, with the volume of VoD traffic predicted to be equivalent to 6 billion DVDs per month by 2018. One solution is to segment traffic by its value to the monitoring application or the content provider, this can be done by subscriber or data type. Since VoD is one of the largest classes of data to be carried across networks and is often provided by a CDN, identification of CDN traffic and the CDN itself can allow certain high bandwidth VoD traffic to be segmented and treated differently from other traffic for the purposes of monitoring and analysis.

Telesoft has a patent applied for a system (STR) that identifies and routes VoD and CDN traffic to allow monitoring infrastructure to scale up with this rise in traffic. To find out more, call us today, or visit our website www.telesoft-technologies.com

Talk to Telesoft today, to discuss monitoring at 100Gbps

About the Author

Steve Patton B.Sc (Eng), C.Eng, MIET is a leading authority on test, monitoring and analysis of IP data networks. He began his career in electronics and software design of encryption systems, subsequently building on this to work in a broader field of data networking and communications.

Steve has worked as a design engineer, business development manager, sales director and is currently Senior Product Manager at Telesoft Technologies, a UK based datacoms infrastructure vendor specialising in extreme line rate, low latency packet processing systems built around FPGA.

Headquarters

Telesoft Technologies Ltd,
Observatory House, Stour Park
Blandford DT11 9LQ UK

t. +44 (0)1258 480880

f. +44 (0)1258 486598

e. sales@telesoft-technologies.com

Americas

Telesoft Technologies Inc
125 Townpark Drive, Suite 300
Kennesaw, Georgia, GA 30144 USA

t. +1 770 454 6001

e. salesusa@telesoft-technologies.com

India

Telesoft Technologies Ltd (Branch Office)
Building FC-24 Sector-16A, Noida 201301
Uttar Pradesh, INDIA

t. +91 120 466 0300

f. +91 120 466 0301

e. salesindia@telesoft-technologies.com

www.telesoft-technologies.com