



IPTV - Video Server Testing

IxLoad



Copyright © 2006 by Ixia

All rights reserved

Ixia
26601 West Agoura Road, Calabasas, CA 91302
(877) FOR-IXIA


This Test Plan contains a general outline for testing a particular technology.
Not all the capabilities of Ixia technology have been exposed in this document.
Please feel free to contact us if additional capabilities are required.



IPTV – Video Server Test Plan

Contents

| | |
|--|-----------|
| Overview | 1 |
| Background | 2 |
| Generic Video Server Model | 3 |
| Video-on-Demand (VoD) Services Delivery .. | 4 |
| 1. Performance Metrics | 6 |
| 2. Test 1 – Maximum Throughput and Optimal Video Stream Delivery. | 8 |
| Objective | 8 |
| Setup | 9 |
| Input Parameters | 10 |
| Methodology | 11 |
| Results | 12 |
| 3. Test 2 – Cross-Traffic Evaluation | 15 |
| Objective | 15 |
| Setup | 16 |
| Input Parameters | 17 |
| Methodology | 18 |
| Results | 19 |
| 4. Conclusion | 20 |





IPTV – Video Server Test Plan

Overview

In the context of Triple Play services delivery, IPTV represents the most significant challenge and often the most complex technology to test and verify. This test plan focuses on testing and qualifying an IPTV video server to deliver heterogeneous video streams to a large number of subscribers. The goal of the test plan is to examine the quality of the video streams received by the end user under various load conditions.

Ixia's IxLoad application is used to execute all the test plans. IxLoad application uses Ixia's purpose built hardware to generate stateful traffic and provides real-time statistics including the triple play services as outlined in this test plan.

The test plan examines several key video performance metrics that quantify the performance of the video server. Aggregate and per-stream performance metrics are used to determine the operating limits of the video server. To test the video server, unicast streams of video traffic is sourced from the server. The service delivery model closely resembles a video-on-demand (VoD) service that uses IP as the transport for video on converged data networks.

System engineers and testers are the intended audience of this test plan. The test plan begins by testing the video server with multiple single constant-bitrate video streams, using standard video streams to determine the maximum throughput without packet loss. The optimal client capacity is also determined using a similar setup.

A video server must also be able to support various clients with varying bandwidth requirements. For this reason, the video server is tested under cross traffic conditions that are indicative of real-world scenarios with standard-definition (3.75Mbps) and high-definition (19+ Mbps) video streams of unequal usage. The effects of varying video bitrates, traffic burstiness experienced by server, number of active (successful) video streams and correlation between streams are examined to determine the performance and quality of experience (QoS) possible by the video server.

The recommended testing methodology presented here is also meant to be used as a guideline for creating more case-specific testing scenarios to further characterize the performance of an IPTV infrastructure.

Background

Video on demand (VoD) is poised to become one of the most important video services offered over a converged IP network. VoD refers to a highly interactive television experience where the subscribers are able to engage directly with the video stream. For example, subscribers are able to order a movie or an event “on-demand,” and it will be delivered “right away.” There are no scheduled times for the on-demand content. Additionally, individual users are able to interact with the video at will by performing VCR-like functions. It is possible to pause/resume, fast-forward or rewind or jump between chapters (depending on the content being delivered).

To deliver an interactive service with VoD, there are three basic required components: the video server, the transport network, and the video client or set-top box (STB). The requirement of this video service model is to provide an acceptable quality of service (QoS) to the paying subscriber. Specifically, the QoS seen at the video client reflects how well the content was delivered from the video server. Therefore, it is imperative that each of the three components of the VoD service model function optimally to provide the best QoS possible for the video client.

The focus of this test plan is the video server, the first of the three components that make up the VOD interactive service. Since the video server is the source of all video content, it plays an especially critical role. It can be modeled as the first switch in a series of switches that allows the transport of video content to the subscriber. As such, its performance must be fully understood and its limits known. The generic model of a video server is presented later in this test plan to model its performance characteristics.

Generic Video Server Model

To characterize the performance of the video server requires a basic conceptual model of the video server itself. This section provides such a video model; however, it does not preclude more novel video server models nor does it attempt to create a “one-size-fits-all” model that is applicable to all video server architectures.

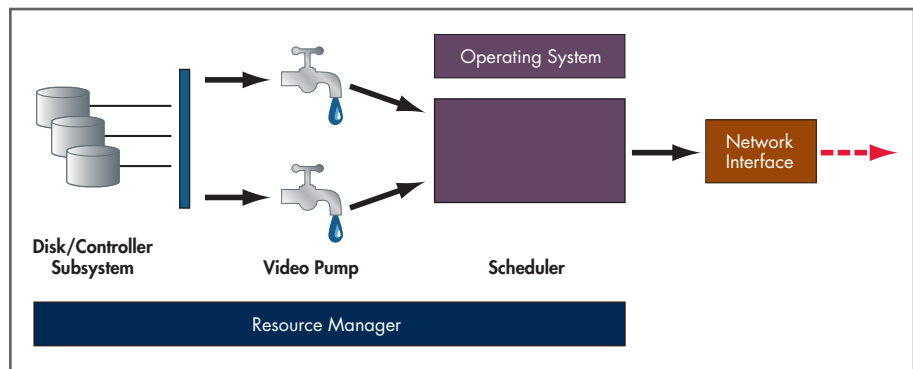


Figure 1. Generic Video Server Model

The video server architecture comprises of the following:

- Highly available (RAID) disk subsystem for storage of several videos.
- **Video pump** – Reads data from the disk and writes it to the network interface. The video pump preserves the timing information in the stream. For example, for a constant bit rate (CBR) video, it must deliver packets to the network interface at the rate of the video stream, which is available in the stream data. There is a video pump for every video stream.
- **Resource Scheduler/Manager** – Responsible for controlling how the video pumps are allocating the video streams from the disk subsystems, and provides access to the network interface to stream the content.
- **Network Interface** – One or more interfaces that all video pumps share for network access. The network interface resolves network resource contention among the many video streams by the means of a resource manager.

Keep in mind that there may be other components that comprise a specific video server implementation or deployment.

Video-on-Demand (VoD) Services Delivery

Having established a generic model of a video server, the VoD services delivery model can now be briefly discussed. VoD service typically uses unicast traffic flows to deliver movies/ events that are interactive. The movie/event (media) is delivered via a media transport protocol and the user interacts with the media stream using a control protocol.

VoD is similar to renting a movie. An individual “copy” of the movie is delivered to the subscriber “on-demand.” The subscriber has control over the playback using VCR-like commands: play, pause, resume, fast-forward, rewind, and stop (among others depending on the type of content).




Figure 2. High-Level Video-On-Demand Services Delivery

Some of the technologies and protocols used in a VoD IPTV deployment are listed below.

- **Real Time Streaming Protocol (RTSP)** – A session control protocol used to deliver interactive content such as VoD in an IPTV network. It uses TCP for reliable transport.
- **Real Time Protocol (RTP)** – A stateless media transport protocol used to deliver multimedia content. It uses UDP or TCP as the transport protocol.
- **MPEG Transport Stream** – Contains multiplexed video/audio streams that are delivered over a media transport protocol such as RTP/UDP or raw UDP.

Knowing how well a video server performs helps guide resource expenditures for the right type and number of devices that are able to scale to support a growing subscriber base. For example, to deliver 100,000 standard-definition videos (3.75Mbps) on-demand subscribers at a 5% peak use would require close to 19Gbps of capacity. To service this subscriber base would require (aside from considering the transport network requirements), a cluster of video servers (at least 20+ servers serving content at near line-rate).

1. Performance Metrics

To determine the performance of video servers, several performance metrics are used. The following terms are defined and used to provide objective performance. Note that the terms defined here comprise generally accepted industry terms and protocol-specific terms that help characterize the performance.

Connection – A single TCP connection between two end hosts, using connection establishment (3-way handshake).

Concurrent Connections – Multiple TCP connections established between two end hosts.

Throughput – The rate at which the DUT sends or receives data, measured in bits per second.

Video Bitrate – Measure of the rate of information content in a video stream. It is measured in bits per second or Megabits per second. A higher bitrate video generally allows better quality.

Effective Bandwidth – Total rate of information including protocol overhead. The effective bandwidth is always greater than the video bitrate and is dependent on transport protocols and link-layer technologies.

Client Capacity – Number of active video clients that a video server can support. The maximum client capacity of a server is the threshold where the server can no longer realize any additional clients without degradation of quality in media streams.

Success Rate – The ratio of active video streams calculated over the total attempted video streams at the video clients.

Latency – The time (delay) it takes for a packet to get from one designated point to another. It is generally measured either one-way (the time it takes to receive a packet from a sender) or round-trip (the time it takes a source to send a packet and receive a response back). Latency is generally measured end-to-end for a given data stream such as one video stream. It can also be measured for a given video stream as inter-packet latency. It is measured in milliseconds.

Inter-Packet Latency – The amount of time measured between packets in a single video stream. It is measured as the inter-packet arrival time between packets for a single video stream. It is measured in milliseconds.

Jitter – The variation in time between packet arrivals. Network jitter is inherently introduced by network elements such as switches/routers. Source jitter can be introduced by video servers due to poor implementation of protocol or resource management. Jitter can cause buffer underruns and overflows in a video client leading to poor picture quality.

Burstiness – The tendency of a device/network to experience sudden increase in bandwidth utilization.

Video Transport Quality Metrics – A scoring mechanism that combines packet delay variation (jitter) and media packet loss to determine the ability of the network to transport good quality video.

Other performance metrics may be of interest to the tester to characterize the video server performance and tests can be designed to meet these objectives.

2. Test 1 – Maximum Throughput and Optimal Video Stream Delivery

Objective

Based on the generic video server model presented earlier in the document, one of the most basic requirements for testing a video server is determining its raw performance capability. The raw performance of the video server can be qualified in terms of the effective use of network bandwidth and optimal video delivery.

Knowing the maximum performance for delivering video is very desirable. It ensures that the server is able to perform under high network utilization. With the increasing number of active video streams comes the challenge of ensuring that the server is able to resolve resource contention so that such contention does not translate into poor performance. Resource contention on the video server can be caused by processor and memory limitations, disk access, and network interface access (including buffer limitations) for multiplexing an increasing number of video streams. Such contentions can produce jitter, which usually impacts video quality experienced by the user.

Optimal video delivery refers to the nominal rate of video delivery that is perceived to provide the best user experience under expected load conditions. This is generally achieved at a lower threshold than the maximum throughput limit. Optimal video delivery is a complex measurement because it involves the video client and how well it is able to receive the video streams and play the video. This measurement begins with determining the video server's delay and jitter characteristics and then correlating these characteristics as part of a complete end-to-end video-on-demand solution.

Setup

The test setup requires at least one Ixia test ports to generate stateful VoD traffic. The video server can be directly connected to the test port if there is one video server under test. If there is a cluster of video servers to be tested, the test results will inherently include the aggregate effects of the IP network (and any content balancers) in the transit path of the video client and servers. In this case, the test is not designed to specifically qualify the video server performance but the network/system-under-test.

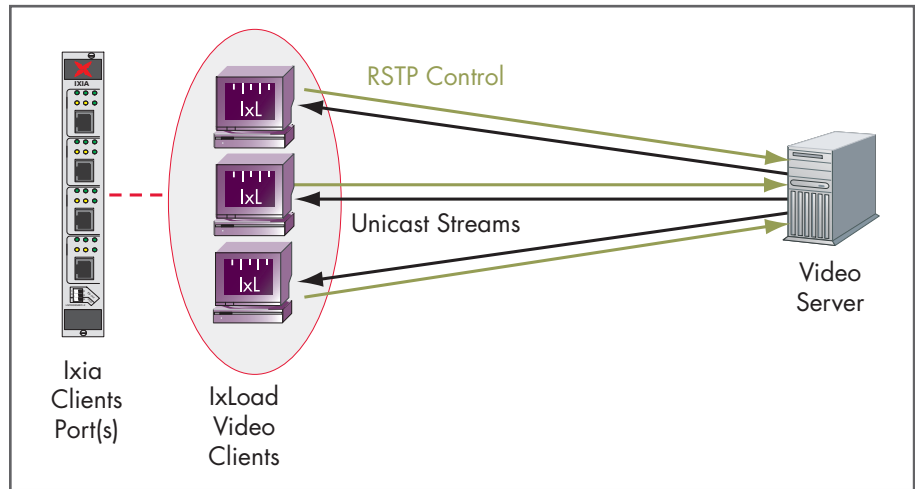


Figure 3. Topology for Test 1 – Maximum Throughput and Optimal Video Stream Delivery

Input Parameters

| Parameter | Description |
|---|---|
| Number of Emulated Clients | 50 – 300 users per test port |
| VLAN Configuration | One or more VLAN with untagged and/or 802.1q tagged packets per VLAN |
| Client Viewing Profiles | Request channels as per viewing profile selected per run |
| | PLAY – for 10 sec – 10 minutes |
| | STOP after a fixed duration |
| | PLAY (again) – for 10 sec – 10 minutes |
| | PAUSE/RESUME – not used for this test |
| Ramp-up rate – configured as part of Test Objective to simulate load conditions | |
| Client RTSP Parameters | Simulate different video clients types by choosing Client Emulation as required |
| Test Objective | Iterative method to determine the <i>maximum simulated</i> users. |
| | Ramp-up rate must be selectable to create bursts of traffic to the video server |

Table 1. Input Parameters for Test 1 – Maximum Throughput and Optimal Video Stream Delivery

Methodology

1. Before starting the test, determine the baseline performance characteristics of the test equipment. This will ensure that the upper limits of the test equipment are known and that the performance per-port or per-system is well defined.
2. Once the baseline is known, configure the video server as outlined in the *Input Parameters* table.

For a single video server, configure the maximum number of VoD streams possible for the server. This may be provided by the vendor and this test will validate this based on criteria that is deemed relevant by the tester. Configure all streams to use the same video payload – effectively serving the same content at the same bitrate (use 3.75Mbps for SD streams).

3. Set up the video clients on the test equipment as outlined in the *Input Parameters* table.

Note the rtsp:// path required to access the video streams. The command list (viewing profile) for each subscriber can be a simple profile (PLAY, STOP, etc.) or can include an incremental method of requesting the streams. The viewing profile is used to drive the Test Objective.

4. Configure the Test Objective to achieve the maximum Simulated Users.

An iterative process is required to ensure that the video server can be tested for different levels of burstiness (using ramp-up rate and several viewing profiles) and to examine the effects of any sudden load experienced across the active streams.

To determine how the server is performing under various load conditions, monitor the CPU and memory usage, Disk, and Network Utilization of the server.

The performance of the server will translate directly into the quality of experience for the subscribers.

5. Once the maximum throughput has been determined (considering no packet loss is seen at the video clients, not necessarily 100% request success rate alone), set the Test Objective to 75% of the maximum throughput or the desired video clients per server and examine the video client statistics. The video quality metrics per-stream will characterize the typical delay and jitter contribution of the video server.

Results

This objective of this test was to determine the maximum throughput and the optimal (nominal) video client capacity of a video server.

Several iterations of the tests were conducted to determine how the video server’s performance affected the user’s experience. Several viewing profiles were tested to simulate varying levels of burstiness experienced by the server. For example, when the number of video clients was increased by 160 clients (to 230), the server was able to accommodate all the clients, indicating a success rate of 100%. However, there were packet losses observed on several existing and new streams.

Below, is a sample of the video streams affected.

| Stat Name | Stream Bit Rate (bps) | MDI-MLR | Min MDI-DF (us) | Max MDI-DF (us) | MPEG2 TS Loss | Jitter (ns) | Min Inter Pkt Arrival Time (ns) | Max Inter Pkt Arrival Time (ns) | Min Packet Latency (ns) | Max Packet Latency (ns) |
|-------------------|-----------------------|---------|-----------------|-----------------|---------------|-------------|---------------------------------|---------------------------------|-------------------------|-------------------------|
| user0_201.100.0.1 | 3750001 | 1 | 2841 | 5649 | 7 | 1505 | 2671020 | 16856040 | 37060 | 40080 |
| user1_201.100.0.1 | 3750000 | 2 | 2841 | 5685 | 14 | 1724 | 2671280 | 5616600 | 37040 | 40060 |
| user2_201.100.0.1 | 3750001 | 2 | 2841 | 5650 | 14 | 1645 | 2671580 | 16860220 | 37020 | 40100 |
| user3_201.100.0.1 | 3750001 | 0 | 2840 | 3136 | 0 | 1465 | 2671780 | 30594380 | 37220 | 40260 |
| user4_201.100.0.1 | 3750003 | 0 | 2841 | 3136 | 0 | 1425 | 2671540 | 30594440 | 37180 | 40000 |

Table 2. Per-Stream Statistics for Test 1 – Packet Losses Observed

The result of the bursty traffic that was injected is significant. Table 2 does not show all streams; however, the effects of the new streams were correlated with existing ones. There was variable packet loss seen across multiple existing and new streams. The server’s inability to protect existing streams from being affected by new streams identifies a design limitation of the server, namely, its inability to control stream admission well, or effectively perform resource management.

After viewing the traffic profile that led to packet losses, the rate of ramp-up was reduced to determine if the server was able to sustain all 230 video clients at a lower rate. Unfortunately, packet losses were still observed. After a few iterative test runs, it was determined that the maximum number of clients reliably supported by this video server was 225 video clients. The video stream bitrate was at a CBR of 3.75Mbps for all streams (SD).

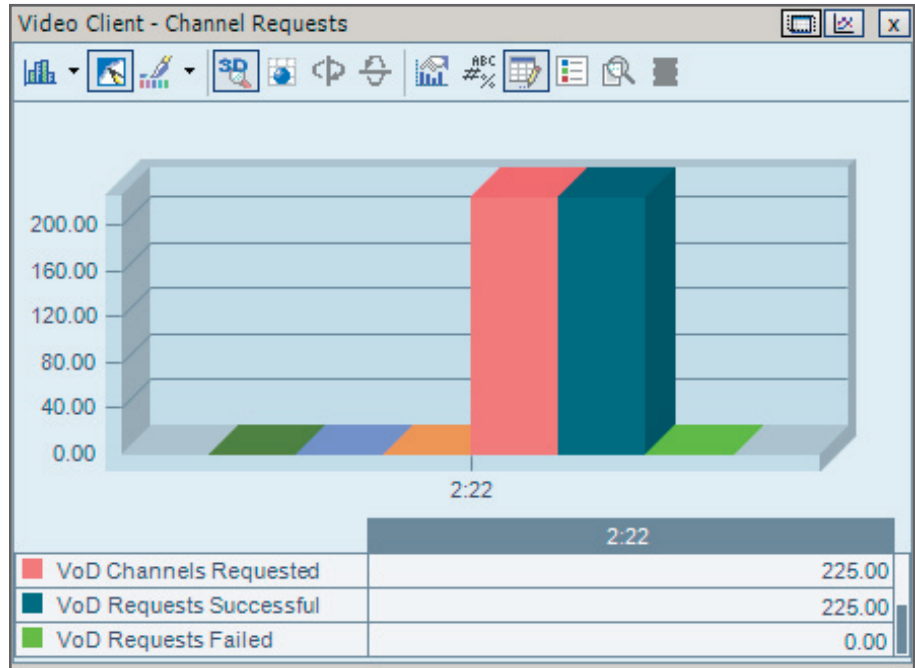


Figure 4. Maximum Video Client Handling (With 3.75Mbps CBR Video Streams)




Figure 5. Maximum Video Server Throughput – 225 X 3.75 Mbps = 845.75 Mbps

IPTV – Video Server Test Plan

The optimal video stream delivery was determined to be to 200 simulated users. This was slightly higher than the 75% of the 225 maximum supported video streams; however, the video server was verified to perform reliably under this load condition for a significant, sustained duration.

| Stat Name | Stream Bit Rate (bps) | MDI-MLR | Min MDI-DF (us) | Max MDI-DF (us) | MPEG2 TS Loss | Jitter (ns) | Min Inter Pkt Arrival Time (ns) | Max Inter Pkt Arrival Time (ns) | Min Packet Latency (ns) | Max Packet Latency (ns) | Join Latency (ms) |
|-------------------|-----------------------|---------|-----------------|-----------------|---------------|-------------|---------------------------------|---------------------------------|-------------------------|-------------------------|-------------------|
| user0_201.100.0.1 | 3750002 | 0 | 2832 | 2841 | 0 | 564 | 2790180 | 2825400 | 37120 | 40240 | 1103 |
| user1_201.100.0.1 | 3750002 | 0 | 2832 | 2843 | 0 | 516 | 2786460 | 2828920 | 37160 | 40200 | 1103 |
| user2_201.100.0.1 | 3750002 | 0 | 2831 | 2845 | 0 | 484 | 2788740 | 2826460 | 37140 | 40180 | 1103 |
| user3_201.100.0.1 | 3750002 | 0 | 2831 | 2847 | 0 | 876 | 2786000 | 2829200 | 36980 | 39940 | 1103 |
| user4_201.100.0.1 | 3750002 | 0 | 2832 | 2850 | 0 | 964 | 2786300 | 2830220 | 37120 | 40100 | 1103 |

Table 3. Per-Stream Statistics for Test 1 – Optimal Video Stream Delivery

Some further observations can be made referring to Table 3. When comparing the reference video stream to the received video stream, the stream bitrate is slightly higher. The actual difference was insignificant for this test; however, this difference may become an issue if the test is run for a long duration, and the video server keeps increasing the sending rate. Drift in streaming bit rates due to clocking issues can also lead to lossy behavior – as the relation between streaming bit rate and playout bit rate translates into the length of the input buffer needed on the client side. For this reason, drift in streaming rates is not desirable. At some point, the video client's playout buffer may overflow and packet loss will occur. The Join (RTSP PLAY) latency is also seen to be even across several streams. This evenly distributed inter-arrival packet time confirms that the video server is delivering all the video streams consistently well.

3. Test 2 – Cross-Traffic Evaluation

Objective

The previous test focused on assessing the video server performance empirically, i.e., the test used several streams of the same bitrate to determine the performance limits of the video server, and then used these performance metrics to infer what their impact would be on the quality of experience for the user.

This test extends to include several video streams of varying bitrate to model a production network. For example, standard-definition (SDTV) streams have a bitrate of 3.75Mbps, DVD-quality video streams run at 9.8Mbps, and HDTV content have a bitrate of 19Mbps or more. Such increased bandwidth requirements of differentiated video streams impose more stringent limits on the video client buffer lengths based on latency and jitter tolerances. Note that high definition content can be delivered using codecs that significantly reduce the required bandwidth.

The goal of this test is to determine how the quality of the subscriber's experience is affected by varying load conditions on the video server.

By introducing cross traffic that includes video streams of varying bandwidth requirements, the resource contention on the video server becomes more complex. It changes the way the video server services the traffic requests. For example, the resource scheduler would give the HDTV video pumps more frequent access to the network interface so that the SDTV streams receive a more constant flow of data at a higher bitrate. The correlation between different streams of different bandwidth will also not be the same as for streams of the same bandwidth.

Evaluating the server under such real-world conditions is critical to understanding the performance dynamics of the video server and how this translates into the quality of experience for the subscribers.

IPTV – Video Server Test Plan

Setup

The setup requires at least one Ixia test port to generate stateful VoD traffic. The video server can be directly connected to the test port if there is only one video server under test. If there is a cluster of video servers to be tested, the test results will inherently include the aggregate effects of the IP network (and any content balancers) in the transit path of the video client and servers. In such as case, the test is not specifically qualifying the video server performance but the network/system-under-test.

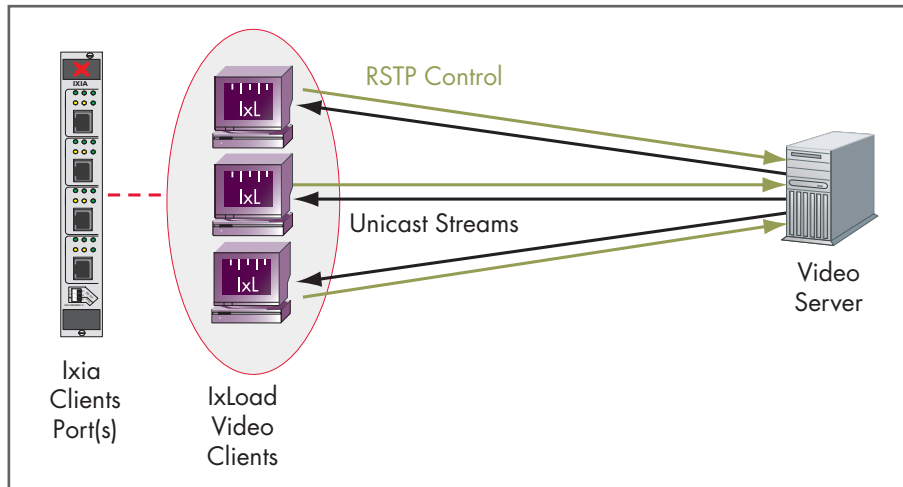


Figure 6. Topology for Test 2 – Cross Traffic Evaluation

Input Parameters

| Parameter | Description |
|-----------------------------------|---|
| Number of Emulated Clients | 50 – 300 users per test port |
| VLAN Configuration | One or more VLAN with untagged and/or 802.1q tagged packets per VLAN |
| Client Viewing Profiles | Request channels as per viewing profile selected per run |
| | Create several viewing profiles that uses all the commands below. |
| | PLAY – for 10 sec – 10 minutes |
| | STOP, PAUSE, RESUME – required to create varying traffic profiles (load conditions) |
| Client RTSP Parameters | Simulate different video clients types by choosing Client Emulation as required |
| Test Objective | Iterative method to determine the <i>optimal simulated</i> users. |
| | Ramp-up rate must be selectable to create bursts of traffic to the video server |

Table 4. Input Parameters for Test 1 – Maximum Throughput and Optimal Video Stream Delivery

Methodology

1. For a single video server, configure a mix of VoD streams. Refer to the *Input Parameters* table.

The VoD streams will be of various CBR so that the video clients request different types of video streams as part of their viewing profile.

2. Set up the video clients on the test equipment as outlined in the *Input Parameters* table.

Note the rtsp:// path required to access the video streams. The command list (viewing profile) for each subscriber can be a simple profile (e.g., PLAY, STOP) or include an incremental method of requesting the streams. Setup the viewing profile according to the Input Parameters table to ensure that different video clients request different types (SD, DVD, HD) of video streams in the viewing profile.

3. Configure the Test Objective according to the viewing profile to ensure that the test equipment (or the video server) bandwidth limits are not reached by oversubscribing the link.

For example, a constant ramp-up of HD clients will oversubscribe the link capacity and cause undesired load conditions and packet loss.

4. This test must be iteratively run to simulate all the viewing profiles to accurately assess the various load conditions on the video server.
5. Ensure that there are no packet losses and that the video quality metrics are within acceptable boundaries.

Results

The objective of this test was to profile the video server characteristics under realistic load conditions by simulating various cross traffic conditions. The cross load conditions included HD and SD video streams.

The optimal setup details are given below.

- 15 HD streams – 20Mbps CBR
- 160 SD streams – 3.75Mbps CBR
- Viewing profile – varying profiles
 - Example: Simulate 15 HD video clients continuously.
 - Introduce 160 SD video clients with a ramp-up of 100 clients/sec after 60 seconds.
 - Perform various PAUSE and RESUME functions to simulate load conditions throughout the test duration.

Several other viewing profiles were iteratively tested. The test profile above was used for analysis here.

| Stat Name | Stream Bit Rate (bps) | Avg MDI-DF (us) | MDI-MLR | MPEG2 TS Loss | Jitter (ns) | Inter Pkt Arrival Time (ns) | Packet Latency (ns) | Join Latency (ms) | Leave Latency (ms) |
|--------------------|-----------------------|-----------------|---------|---------------|-------------|-----------------------------|---------------------|-------------------|--------------------|
| user0_201.100.0.1 | 20000109 | 690 | 0 | 0 | 117 | 526280 | 38160 | 8 | 0 |
| user1_201.100.0.1 | 20000109 | 690 | 0 | 0 | 1097 | 525300 | 38280 | 9 | 0 |
| user13_201.100.0.1 | 20000140 | 690 | 0 | 0 | 2184 | 528580 | 38420 | 16 | 0 |
| user14_201.100.0.1 | 20000140 | 690 | 0 | 0 | 844 | 527240 | 37900 | 12 | 0 |
| user0_201.100.0.1 | 3750002 | 2836 | 0 | 0 | 464 | 2807000 | 38460 | 5 | 3 |
| user1_201.100.0.1 | 3750002 | 2836 | 0 | 0 | 1384 | 2806080 | 38960 | 7 | 10 |
| user2_201.100.0.1 | 3750002 | 2835 | 0 | 0 | 556 | 2808020 | 38640 | 8 | 5 |

Table 5. Per-Stream Statistics for Test 2

The statistics above confirm no packet losses and the expected MDI-DF values for the bitrates used (i.e., 690us for 20Mbps and 2836us for 3.75Mbps video streams). The Join latency was very acceptable. The Leave latency refers to the PAUSE command as this per-stream statistic was captured during a time when a PAUSE was initiated by some of the video clients.

4. Conclusion

A video server is one of the three most important components of an IPTV deployment. Benchmarking its performance provides several benefits. It provides confidence that the video server will perform well in a production deployment. It also ensures that the server is able to work under heavy/loaded conditions. Finally, knowing the specific performance characteristics of the video server, such as delay and jitter helps ensure that downstream devices are within acceptable tolerances for the end-to-end requirements of a deployed Triple Play network.

A general goal of the kind of testing presented in this plan is to study the video output pattern to determine if and how it deviates from an ideal or expected output. Some of the key metrics used to help characterize this include active streams capacity, packet loss, delay, and jitter contributions (i.e., the delay distribution associated with the inter-arrival of video streams) under normal and loaded conditions. To model this kind of behavior in a test requires that all other network components be removed from the setup. This ensures that the assessed performance represents the effective performance of the video server alone.

As demonstrated throughout this test plan, testing for video quality is extremely complex because of the close interplay between the various elements of the VoD interactive service. To simulate various load conditions realistically, video clients must be used as part of the test bed. This requires understanding how a video client is designed. A video client is designed to conceal these errors and mask acceptable levels of delay and jitter present in the video stream. It is important, therefore, that the video clients used in the test topology to not perform error concealment. If this is not possible, then the behavior of the video clients must be taken into consideration when determining how well the video server is actually performing during the test.



About Ixia

Ixia is a leading provider of performance test systems for IP-based infrastructure and services. Its highly scalable solutions generate, capture, characterize, and emulate network and application traffic, establishing definitive performance and conformance metrics of network devices or systems under test. Ixia's test systems are used by Network and Telephony Equipment Manufacturers, Semiconductor Manufacturers, Service Providers, Governments, and Enterprises to validate the functionality and reliability of complex IP networks, devices, and applications. Ixia's Triple Play test systems address the growing need to test voice, video, and data services and network capability under real-world conditions. Ixia's vision is to be the world's pre-eminent provider of solutions to enable testing of next generation IP Triple Play networks. Ixia's test systems utilize a wide range of industry-standard interfaces, including Ethernet, SONET, ATM, and wireless connectivity, and are distinguished by their performance, accuracy, reliability, and adaptability to the industry's constant evolution.

Contact Ixia

For more information, contact Ixia or visit our Web Site at <http://www.ixiacom.com>.

Ixia Worldwide Headquarters

Corporate Center

26601 W. Agoura Rd.
Calabasas, CA 91302

(Toll Free North America)
1.877.367.4942
(Outside North America)
+1.818.871.1800
(Fax) 818.871.1805

www.ixiacom.com

Info: info@ixiacom.com

Investors: ir@ixiacom.com

Renewals: renewals@ixiacom.com

Sales: sales@ixiacom.com

Support: support@ixiacom.com

Training: training@ixiacom.com

Ixia USA Sales

Phone: 1.866.355.4942

Email: sales@ixiacom.com

Ixia Canada Sales

Phone: 1.877.367.4942

Email: salescanada@ixiacom.com

Ixia China Sales

Phone: +86.10.84549199

Email: saleschina@ixiacom.com

Ixia Europe, Middle East, & Africa Sales

Phone: +44.1753.722056

Email: salesemea@ixiacom.com

Ixia India Sales

Phone: +91.80.25633570

Email: salesindia@ixiacom.com

Ixia Japan Sales

Phone: +81.3.5365.4690

Email: salesjapan@ixiacom.com

Ixia Oceania Sales

Phone: 1.818.292.1561

Email: salesoceania@ixiacom.com

Ixia South Korea

Phone: +82.11.897.1326

Email: salessouthkorea@ixiacom.com

Ixia Federal Sales

Phone: 1.703.822.7527

Email: salesfederal@ixiacom.com

© 1998-2006 Ixia. All rights reserved.

This publication may not be copied, in whole or in part, without Ixia's consent.

Ixia and its licensors retain all intellectual property rights in all products identified in this publication. Such products may be covered by one or more patents and/or pending patent applications, including but not limited to the following U.S. patents: 6,717,917; 6,408,335; 6,397,359; 6,061,725; 5,937,165; 5,881,237; and 5,838,919. All software and related documentation identified in this publication is licensed, not sold, pursuant to a separate license agreement between Ixia and the recipient. The recipient's use of such software and documentation is subject to the terms of that agreement.

Restricted Rights Legend

Use, duplication, or disclosure by the U.S. Government is subject to the restrictions set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

THIS PUBLICATION IS PROVIDED "AS IS" AND WITHOUT ANY WARRANTIES OF ANY KIND, EITHER EXPRESS OR IMPLIED. IXIA SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. THE INFORMATION HEREIN IS FURNISHED FOR INFORMATIONAL USE ONLY, IS SUBJECT TO CHANGE BY IXIA WITHOUT NOTICE, AND SHOULD NOT BE CONSTRUED AS A COMMITMENT BY IXIA. IXIA ASSUMES NO RESPONSIBILITY OR LIABILITY FOR ANY ERRORS OR INACCURACIES CONTAINED IN THIS PUBLICATION.

Ixia, the Ixia four petal logo, and IxLoad are either trademarks or registered trademarks of Ixia in the United States and/or other countries. All other trademarks belong to their respective owners.

