Lab Test to Service Management ™

White Paper

The Telco Triple-Play:

Scoring a Successful Deployment with Comprehensive Lab Testing and Service Management

May 2005

By Marc Robins, President, Robins Consulting Group

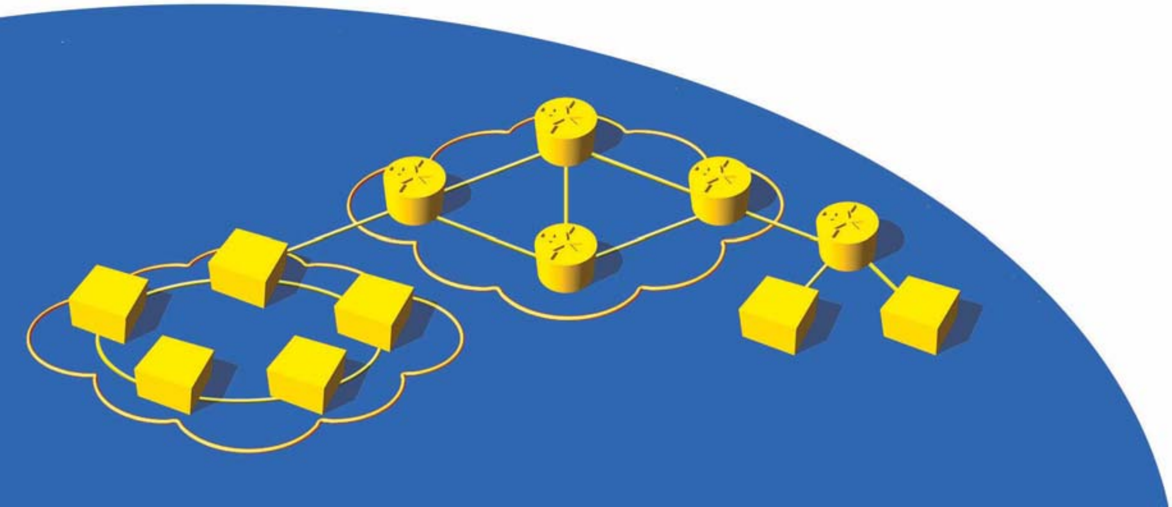

Spirent Communications, Inc.
26750 Agoura Road
Calabasas, CA
91302 USA
Email: productinfo@spirentcom.com
Web: http://www.spirentcom.com/

North America
+1-800-927-2660

Europe, Middle East, Africa
+33-1-6137-2250

Asia Pacific
+852-2166-8382

All Other Regions
+1-818-676-2683

About Robins Consulting Group

Robins Consulting Group (RCG) is a leading marketing communications consultancy providing an array of marketing and other services to the IP communications industry. Marc Robins, an internationally recognized authority in the field of IP communications and emerging new communications technologies, founded RCG in 2003. Prior to RCG, Mr. Robins served as vice president of publications and trade shows, associate group publisher and group editorial director at TMC, publisher of the trade magazine ***Internet Telephony*** and producer of the ***Internet Telephony* Conference & EXPO** trade shows, for which he also served as chief architect and conference co-chairman. For more information about Robins Consulting Group services, call 718-548-7245 or e-mail robinsconsult@optonline.net. This white paper was commissioned by Spirent Communications.

The Telco Triple-Play:

Scoring a Successful Deployment with Comprehensive Lab Testing and Service Management

Contents

Executive Summary

Telcos and service providers the world over are racing to deploy networks that will support the triple-play service bundle – the delivery of voice, video and data services provided on a common network infrastructure. And although there is no agreement among providers in terms of a single plan for such an infrastructure, providers agree that the triple-play represents a major opportunity to forestall the growth of cable MSOs and other competitors.

The availability of new, carrier-grade IP-based technology is also contributing to a major shift in the way providers are designing their networks to provide new video and enhanced telephony services. Indeed, the reality of a true converged, IP-based network architecture opens the door to an array of enhanced applications and services, including IPTV, video telephony, VoIP, and Video on Demand.

Like all major advances in technology, a converged, common network infrastructure for triple-play brings with it a unique set of challenges as well as opportunities. Careful planning and testing throughout the equipment purchase and implementation processes can help telcos meet these challenges and prepare for them before they turn into quality of service issues for their subscribers. Indeed, the most effective way to prepare for and combat the quality issues inherent in a converged triple-play network is to test that network thoroughly before rolling it out.

This paper outlines both the benefits and challenges of deploying such a common network for triple-play. It also discusses the advantages of adopting new IP-based technology for the triple-play network, and explores key quality of service (QoS) issues that can impact a triple-play network. Perhaps most importantly, this paper examines in detail the rationale for the testing of various triple-play network infrastructure, applications and service management operations, and how the development of a comprehensive lab test and service management program can ensure that a telco is able to reap all the benefits of a successful triple-play deployment.

1.0 Telco Triple-Play: The New Mandate for Today's Service Providers

The telco triple-play – the delivery of voice, data and video services over a common network infrastructure – has become the rallying cry of an industry under assault from multiple fronts. Indeed, with their still sizable legacy POTS revenue at risk from upstart VoIP providers and Cable MSOs with their VoCable offerings, and DSL uptake significantly trailing cable broadband, the RBOCs and the majority of local, independent telcos are rapidly facing their moment of truth: either upgrade their infrastructure and expand service offerings or endure a slow and steady contraction of their subscriber base and revenue streams, ending in their eventual demise.

With the lines drawn in the battlefield, the strategy is clear: telcos must add their own video services -- including video telephony, broadcast video and video-on-demand -- to their current voice and data offerings in order to effectively compete in today's marketplace. To highlight the growing threat, consider that broadband cable providers have taken more than 38% of the broadband Internet market and 2.2% of the local telco market, according to Heavy Reading, and this trend is expected to grow dramatically if left unchecked. In fact, in Omaha, Nebraska, the number one telephony provider today is Cox Communications, the city's leading MSO! And by 2008, the market for cable telephony services is projected to reach 14.7 million users, according to research firm ABI.

Clearly, there are significant business opportunities for the telco offering a triple-play service bundle. Aside from protecting their subscriber base from churn, video service has the potential to grow a telco's ARPU (annual revenue per user) by 100%, something that no other service offering can currently provide. But there are also significant challenges ahead for the RBOCs: key strategic issues include how to cost-effectively deploy triple-play infrastructure in their extensive networks, and how to best deliver triple-play services to millions of subscribers' homes in a reliable manner.

1.1 Key Elements of Triple-Play Deployment

This document explores three key issues necessary for the successful understanding, development and deployment of a triple-play network.

a) Every implementation is different. The variety of possible video deployments and the need to interwork with legacy networks means that there is no "one size fits all" solution.

b) The user experience is based on proper operation of (i) the triple-play applications, (ii) the services provided by the service provider and (iii) the underlying infrastructure.

c) A clear understanding of both the individual point elements in a network and their function as a dynamic interactive system are required to operate a successful triple-play network and pinpoint diagnosis of problems as they arise.

1.2 The Different Approaches to Achieving the Telco Triple-Play

While the destination is universally agreed upon – namely the ability for telcos to offer a robust service bundle of voice, data and video -- there is no universal agreement among telcos on the path that will take them there. Currently, there are a number of different plans for triple-play infrastructure, each with their own set of challenges. Additionally, there is currently no common approach to defining network requirements or agreement regarding the appropriate level of service for voice vs. data vs. video.

The big four RBOCs -- Verizon, SBC, BellSouth and Qwest -- are all devising different network architectures to deliver the triple-play service bundle. For example, while Verizon seems to be favoring Fiber to the Home (FTTH), the other three RBOCs are focusing primarily on Fiber to the Node (FTTN), with new DSL-based delivery mechanisms that will allow them to leverage their existing DSLAMs and copper-based local loops, while reserving FTTH for mostly Greenfield deployments.

A FTTH Triple-Play architecture, while still costly to implement over long distances, offers virtually unlimited bandwidth – currently 50,000 times the capacity of coax cable, or 50 Terahertz vs. 1 Ghz. Verizon's plan is to offer Voice over IP on their fiber access connections, currently running to 1 million homes in nine states and eventually to all 27 million households throughout their territory. For video, Verizon can employ the same type of multicast video distribution architecture as a cable MSO.

On an interesting note, the digital cable service offered by MSOs today is actually transmitted in an analog format, in such a way that enables cable companies to increase the channel capacity of their systems. In current systems, up to five video channels are converted to a digital format, multiplexed, and then converted into an analog format and sent in a single 6-MHz channel slot for transmission. This way, digital channels can be combined with traditional analog NTSC channels on the system. Cable systems currently support more than 100 6-MHz channel slots, which allows them to offer the 500 or so channels currently in the channel lineups today. While a TV tuner can access and select analog channels, a set-top box is required to select and decode digital channels.

In contrast, FTTN architecture employs fiber up to the local distribution cable. From that point on, the service would run over copper loops to the subscriber using DSL. The advantages of this approach are that it leverages the existing twisted pair infrastructure – something the telcos own a huge amount of – as well as their sizeable investment in DSL technology. The downside is that the copper local loop doesn't have nearly the same capacity as fiber.

In an effort to create the necessary capacity, a number of key developments are in the works that makes an FTTN Triple-Play deployment feasible. The first is MPEG-4 (H.264), a new digital video-encoding algorithm that reduces the bandwidth of a standard-definition video channel from 2.5-3 Mbps to 1.5 Mbps, and a high-definition channel from 15-19 Mbps to around 10 Mbps. In addition, higher capacity flavors of DSL have been developed, including bonded ADSL2+ (favored by BellSouth) and VDSL (for very high bit-rate DSL, favored by SBC and Qwest) that are based on the same transmission method employed by current generation DSL modems, and will allow telcos to deliver a 25 Mbps digital signal over the local loop at distances up to 3,000 feet. A newer, more robust version favored by Japanese service providers, VDSL2, promises to support a 1000 Mbps data rate over 1,000 feet and at least 25 Mbps over 4,000 feet – hopefully enough to cover the majority of local loops.

The challenge for the telcos, even with these new developments, is that there will still be a need for even greater capacity, considering that the infrastructure will have to support not only a multitude of standard and high definition video channels, but also broadband Internet and voice. Some options include aggregating links for more bandwidth, and creating new, more efficient error correction and encoding algorithms to squeeze more capacity out of VDSL2, but even these methods won't be enough to solve the bandwidth dilemma.

With respect to voice, current options are to transmit telephone service the old analog way, on the low-frequency portion of the local loop as it's currently performed with today's ADSL services, or to migrate to a VoIP offering. The analog method reserves the available 25 Mbps for video and broadband Internet, but leaves subscribers with the same old POTS service they've always had. A migration to VoIP, on the other hand, allows telcos to effectively counter the threat from the cable companies and upstart broadband telephony providers such as Vonage, and offer a far richer application/feature set in the process.

1.3 The Promise of IP: Delivering a Packet-based Voice, Video and Data Service Bundle

As we've discussed above, even with more efficient coding algorithms and new versions of DSL, telcos that follow the FTTN path to triple-play deployment will still hit a wall in terms of capacity where a 25 Mbps connection isn't enough to support the services required. As a result, these telcos are eyeing a new IP-based network architecture employing gigabit Ethernet technology.

In fact, IP is emerging as the telcos' most potent weapon in their technology arsenal. The migration to a converged, IP-centric network infrastructure enables all service providers to become full-fledged triple-play providers, allowing them to leverage multiservice IP access solutions to provide video, data and voice over a common network platform. In fact, the deployment of switched digital video services could potentially favor the telco in the battle with cable: most large telcos have an innate understanding of switched services, and have the operational infrastructure to effectively support massively scaleable switched services to boot.

By leveraging existing Ethernet and packet-based technologies, telcos can efficiently deliver a triple-play service bundle over their existing DSL connections in a highly flexible and cost-effective manner. Indeed, IP has now become the network-layer and service protocol of choice, and is widely deployed in the core of service-provider networks as well as in corporate and residential networks, putting it on both sides of the access network.

From a network equipment perspective, a host of new, IP-centric products are on the market that form the puzzle pieces for a migration from a circuit switched to packet-based network infrastructure. Products such as softswitches for voice services, digital head-ends for IP video, and broadband loop carrier systems are all serving to enable the convergence of voice, data and video on a single access network infrastructure. From a services perspective, an array of IP-based applications are available, including VoIP, Internet access, virtual private networks (VPNs) and other data services, video telephony and conferencing, and broadcast video (IP TV) and video-on-demand.

In addition, IP networking is being continually enhanced, and today encompasses a broad set of network services. For example, the differentiated services (DiffServ) architecture defines a set of IP-based quality of service (QoS) mechanisms based on the classification and marking of packets for special handling within the network – something that comes in quite handy for prioritizing one type of service over another. In addition, IP-based multicasting protocols also allow for the efficient delivery of a video stream to multiple subscribers.

1.4 The Challenges of Developing, Managing and Maintaining Service on a Converged Network Infrastructure

Even with all the technological resources at their disposal, telcos entering the triple-play arena with video offerings face some daunting obstacles. For one, they are the last in the ring after the cable and satellite service providers, making it imperative that they offer services that deliver an experience equal to or better than that of the competition.

In addition, since it's still not clear which delivery and pricing models and applications will gain market acceptance, the availability of a high speed broadband connection is simply not enough. Telcos must design their broadband delivery infrastructure to support the delivery of high-speed, QoS-sensitive applications, regardless of what they are or how they are billed. Given the economic necessity to maximize investment in the existing infrastructure, service provider network evolution must tie into existing Internet backbones while adding the service intelligence required to tie into a variety of access networks and new service delivery networks.

One such service delivery challenge involves IP TV. In an FTTN deployment, IP multicasting will be employed with channels terminating at a local neighborhood concentrator. Up to five TV channels and Internet service will be delivered to a subscriber – a model that is radically different than the delivery method employed at cable MSOs. With only five channels at a time sent to a subscriber's set-top box, a signaling protocol is now needed to send an alert to the concentrator when a subscriber requests a new channel.

A big unknown with this mechanism is what the response time of the network video server will be. From a competitive standpoint, the channel change (join and leave times) must be as quick as it is with today's digital cable service. Actual performance, however, will depend on the ability of the telco to effectively develop, manage, and maintain the network elements for optimal performance during periods of a large number of concurrent channel requests (such as during the Super Bowl or the final episode of the latest *American Idol*).

2.0 Stress in the Network: Quality of Service (QoS) Issues Related to Triple-Play Deployment

So far, we've covered some of the major challenges surrounding the deployment of a triple-play service bundle. At this point, it would be helpful to review what might happen when converged IP-based voice and video services go live without proper network assurance and testing. The quality of service (QoS) issues below represent some of the most common problems that might be experienced by subscribers.

2.1 Data Packet Loss

For email and HTTP traffic, a few lost packets are no big deal. As long as packet header information is in order, the packets will re-assemble themselves at the destination without any apparent change or perceptible loss to the end user. But an IP-based triple-play network deployment is an entirely different beast, as even a low level of packet loss can create unacceptable degradation of voice and video traffic quality.

Packet loss can result in everything from garbled voices on a phone call to graininess and dropouts in a video signal, where the screen during a video conference or TV show loses clarity or momentarily goes blank. In the worst case, severe packet loss can cause an entire communications session to be dropped by the network, meaning that a phone call gets disconnected or video session abruptly ends.

2.2 Voice and Video Jitter

The time it takes for packets to travel from point to point within a converged network is measured in delay or latency levels. Jitter is a variation of delay/latency levels over time, caused by actions like queuing and routing that affect the path of packets as they travel through the network. A congested network will generally have higher levels of jitter, but proper quality of service controls like queuing and bandwidth allocation can control the problem.

Testing is essential for determining how much jitter a network will experience and how it should be handled. Most quality of service solutions include a jitter buffer, which adds small amounts of delay to packets received so they all appear to have equal and acceptable amounts of latency. This also ensures that packets are transmitted in the correct order.

2.3 Delay/Latency

Exceeding the acceptable delay/latency levels for a voice application can make for a voice call in which parties are forced to pause each time they speak and wait for the other party to hear what they've said. For video and IP TV applications, likely impairments can include "freezed" frames and corrupted images displayed on monitors and TVs.

These problems can be solved by quality of service measures that monitor and buffer packets as they travel across the network, ensuring that they are transmitted with acceptable amounts of delay and in the proper order. The network must be thoroughly tested under varying traffic conditions to accurately determine the maximum level of latency that can occur, and to properly configure quality of service solutions to accommodate that level.

2.4 Voice Echo

In a hybrid network, analog phones can be the source of echo, but a pure IP telephony configuration is also susceptible, since some calls will inevitably interface with TDM network elements at some point when they go "off net." Voice echo occurs when participants can hear their own voices coming out of a phone's earpiece, and if extreme, can make a voice call completely intolerable.

Factors such as the volume and length of the echo and the impact of hardware like handsets and routers can contribute to the echo "level" or measurement of the echo. There are generally acceptable levels of echo for IP telephony calls, and echo cancellation and suppression are the most common ways to keep echo at an appropriate level within a network. Echo cancellers (digital signal processing solutions) are generally used in media gateways and other equipment to control echo levels.

2.5 Clipping

Clipping within VoIP calls occurs when either the beginning or end of words, or whole words, seem to be cut off during a conversation. This can occur when voice activity detectors and other solutions that work in tandem with echo cancellers are thrown out of sync. Echo cancellers deal with background noise and "double talk" in addition to echo. For instance, an absence of background noise during a call can confuse users into thinking a call has been dropped. But improper levels of background noise can result in annoying voice clipping. A careful balance must be struck, ensuring that proper levels of latency, jitter, packet loss and echo are maintained throughout the network at all times.

2.6 The Ramifications of Deploying Triple-Play Services on an Untested Network Infrastructure

Each potential quality of service issue detailed above has the potential to not only negatively impact a service provider's overall network performance, but can also lead to massive subscriber unrest and revolt, where users cancel their services and leap into the waiting arms of the local cable MSO.

Indeed, if QoS issues aren't uncovered during comprehensive testing in the lab or in pre-deployment network testing, network operators will undoubtedly be in for a shock when services go live. At the very least, dropped calls, unacceptable voice and video quality, bandwidth bottlenecks, and general chaos will ensue. The net result for a telco's business operation is a negative impact on the company's bottom line.

3.0 Testing the Triple-Play Network:

Ensuring the Optimal Performance of Network Infrastructure, Applications, and Service Management Operations

Of course, a totally different outcome is possible than the one described just above. By adopting a comprehensive testing program that stresses the network infrastructure, triple-play applications and end-user processes under controlled conditions, and details the appropriate measures needed to ensure optimal levels of service, the results can be just the opposite. In this case, happy subscribers will eagerly signup for the service bundle, show a propensity to trial and signup for new applications and services, and spread positive word-of-mouth to friends, family and co-workers. – all of which serve to dramatically grow marketshare, revenues and profits.

When it comes to devising such a testing program for triple-play, it's important to recognize that while there are many different types of "point test" solutions, where individual network elements or functions are tested individually, it's also essential to test network elements and the network together, as a whole entity. Such a holistic testing methodology will uncover the full impact a particular element has on other areas of the network, and greatly aid in the measurement and management of network services, and the diagnosis of any problems.

However, ensuring that triple-play services function in the lab and operate in the network leads to complex testing scenarios. Such scenarios run the gamut of measuring the speed of the network to the buffering capacities of end devices; diagnosing data and control plane issues; managing subscribers ability to join multicast groups; monitoring issues relating to large MPLS networks and the performance of the systems; and determining the level of protection required, in terms of what happens when an element in the network goes down and whether the network can reconfigure itself within a reasonable time frame.

Finally, it's important to recognize that a triple-play deployment involves the coordinated operation of a series of related network architectures, including a head-end network, aggregation network, service delivery network, access network, and finally, the CPE at the subscriber household, and that each architecture has unique and specialized testing requirements.

3.1 Infrastructure Testing

Infrastructure testing is performed to ensure that the hardware and equipment to be deployed delivers optimal performance, and involves tests that run all the equipment through their paces to ensure that they are performing as expected. Areas of testing include throughput capability and bandwidth utilization, packet loss, and measurements of latency created by each piece of equipment. Also vital are tests to ensure that all the required protocol stacks and video and voice coding algorithms employed by the equipment are operating correctly, and that the specific interfaces are working as expected.

As a first step in testing the overall performance of the network, there is a need to know the performance characteristics of each piece of equipment. This information will be invaluable at some point, such as when video service is added to the network. If the service isn't working at all, or there are certain impairments in the picture generated, an operator will need to know that each individual piece of equipment is operating properly in order to narrow down the causes of the problems.

3.1.1 Ethernet Aggregation Router Scalability

Testing Ethernet aggregation router (EAR) scalability involves testing the capability of the routers to receive the traffic that a service provider is trying to pass through their network, and ensure that the traffic is passing through it successfully. Specific tests include throughput, amount of latency and packet loss, protocol stack performance related to PPP, L2TP, IPTV, DHCP, Ping, 802.1x, and L4-7 HTTP, and interface support for 10/100 Ethernet, Gigabit Ethernet, ATM and DSL.

The major concern with respect to any EAR performance issues is that if traffic is not being passed through at the required speeds, and if packet loss is too severe, then service interruption or service failure is a likely outcome. The latency issue is not an enormous issue for data, but is a huge issue for voice and video.

3.1.2 Broadband Remote Access Server (BRAS)/Broadband Network Gateway (BBNG) Capacity

First generation broadband remote access servers provided session management to route traffic over different ISP networks. Today, the BRAS serves as the critical connection point between the access network and service delivery networks. As such, it must be able to recognize and deliver voice, video and data applications in a secure, reliable fashion – no matter where those applications originate or what billing model is used.

Testing is required to ensure these components are providing sufficient throughput, as well as to measure the latency through each device. And because the BRAS/BBNG is the point in the network where subscriber services will be handled, testing is also needed to determine that the devices can support the number of subscribers expected, that the methods of authenticating subscribers are appropriate, and that they are working correctly. In addition, necessary protocol stack tests and interface support testing are similar to that of the router testing described above.

If testing is not conducted, and the devices are not up to snuff, then when a provider moves to deploy a service, customers won't be able to gain access, or service may be dropped on an intermittent or extended basis.

3.1.3 DSLAM QoS and MultiService Performance

With an FTTN triple-play offering, there will be many different types of services that are going to be run through the DSLAM. Therefore, it's essential to know that the DSLAM will be able to handle all services with the required QoS for each service. Testing should include load testing with traffic generation and analysis, as well as multiservice performance testing that involves setting up multiple traffic groups and determining whether or not the DSLAM is able to support them without causing an impact to any one given service.

3.1.4 IPTV End-to-End Testing

For comprehensive testing of the performance of IPTV services, it's vital to not just test the services across the DSLAM, but to also test across the entire network as well, including Ethernet aggregation routers and BRAS/BBNG devices. Tests should be designed to measure channel change time (the time difference between changing from a current channel to a new channel), average join and leave latencies, and for channel validation, channel zapping modes (sequential or random), and the overall quality of the video signal.

In addition, tests should also include a number of enhancements to provide additional realism to the testing process. Such enhancements might include allowing joins and leaves when multicast data is being received, the ability to correlate data timestamps from an analysis and report back all latency measurements to the application, provide an emulation of independent subscriber viewing duration, and the ability to increase the number of subscribers emulated during a channel zapping test.

3.1.5 DSLAM IPTV

In addition to testing the QoS of the DSLAM, and determining how many subscribers can be supported and the throughput it can support for that given number of subscribers, it's also essential to test the join and leave request capabilities of the DSLAM itself. Tests types are similar to those in the section above, and can help uncover what will happen when multiple subscribers are connected to an IPTV service and perform multiple join and leave requests in a very short period of time. Additional tests can also include the emulation of different types of set-top box behaviors, as well as the testing of actual video streams and L4-7 applications.

If a DSLAM has trouble with this "channel zapping" scenario, then subscribers could start to see delays in channel changing of more than one or two seconds – a situation that is of great interest to DSLAM vendors and service providers alike. Remember: in order for a telco's video services to be successful in the marketplace, it is imperative for them to ensure that the services they provide are at least equivalent to or better than a satellite or cable system's. Being able to accurately measure a DSLAM's response to join and leave requests is therefore critical.

3.1.6 Simulating the Local Loop and CPE Access Link Testing

Simulating a variety of local loop scenarios, and testing the various types of CPE that will reside at the subscriber household, such as TV set-top boxes and IADs for broadband data and VoIP services, is essential in order to provide a real-world scenario for testing the overall triple-play service delivery. Remember that in the real world, customers will be connected at the end of 5,000 or more feet of copper loop. It's critical, then, to be able to simulate the real impact of service on that loop in order for the service provider to know the quality of the signal that is going to reach the customer.

Testing will help ensure that any type of noise or service impairments that might be on the loop can be accommodated for, and to ensure that, depending on the type of customer loop, sufficient signal is still getting through to the end points in order to provide acceptable service.

3.2 Applications Testing

Applications testing is performed to ensure that various applications and end-user processes support the triple-play services to be deployed. In this way, telcos can have the confidence that all data, voice and video services are offered seamlessly and without any interruption.

A comprehensive program of applications testing must involve the testing of all the equipment and devices that enable and support applications such as streaming video, video telephony, video-on-demand (unicast video), and IP TV (multicast video). This includes a host of video telephony equipment including media gateways and videophones, unicast streaming servers, streaming aware devices such as firewalls, switches, routers, and server load balancers, and multicast IP TV servers. Key areas of testing and measurement include bandwidth utilization, capacity and scalability, interoperability tests including protocol stacks and supported interfaces, and overall voice, data, and video service quality.

3.2.1 Video Telephony Distributed Testing

The idea of a videophone is not new: AT&T first introduced a Picturephone in 1956, which sent an image every two seconds over a standard POTS connection, and another attempt to popularize the service was made in the late 1980s. Although these early videophone services never took off, it appears that the growth of high-speed broadband connections and the advent of IP telephony could very well be setting the stage for such services to finally stick in the marketplace.

Today, in order for telcos to offer dependable video telephony services, it's important for them to conduct video telephony testing in a distributed fashion – end-to-end on the network. In this way, they can observe and measure all the different types of telephony signals encountered in the network, and to ensure that the video and voice quality a customer experiences with his video phone or video conferencing application/system is acceptable. Equipment that is part of a video telephony distributed test includes not only video phones and video conferencing systems, but also IP-based and legacy PBXs, IP phones, call managers, gateways/switches/hubs, voice messaging servers, and meeting hubs and servers.

The type of tests that should be performed in a comprehensive video telephony distributed testing program includes tests of video quality, voice quality, and application functionality. Scalability tests need to be run to measure how many concurrent video telephony sessions the network can support before service is impacted. It's also vital that tests are performed to uncover any interoperability issues among devices that could potentially lead to poor performance or service failure. For example, if tests expose excessive jitter and latency, these conditions could result in video quality impairments that include "blockiness" of the video image, blurring and various other video problems, as well as poor voice quality.

In addition, tests need to be run that employ all the various protocols that would be in use on the network, including IP-related protocols such as SIP, H.323, MGCP, Megaco/H.248, and Skinny; and PSTN related protocols such as CAS, PRI, SS7, NFAS, V5 and GR303. In addition, testing should also encompass all relevant network interfaces, such as Gigabit Ethernet 1000Base-SX, 1000Base-LX, 10/100/1000Base-T, Analog, T1/E1 and T3/E3.

3.2.2 Video Telephony Network Infrastructure and Media Gateway Testing

Rather than testing the network end-to-end, video telephony network infrastructure testing involves the testing of individual components that pertain to video telephony applications. Tests that should be performed include those that simulate video traffic in order to understand how a variety of CPE responds. Specific types of CPE include IP phones, video conferencing devices, and media gateways.

The types of tests that need to be performed on each device include video quality, voice quality, and functionality tests, which include tests employing common IP telephony protocols such as SIP and H.323. Necessary speech quality measurements include PESQ, PSQM, MOS, and R Factor, while video quality measurements include MOS, blockiness, blur, and PSNR.

3.2.3 Video-On-Demand Infrastructure Testing

According to research firm Forrester Research, video-on-demand services will attract 7.5 million users spending approximately $700 million annually by the year 2007, and the total market for on-demand television will grow to about $6 billion. This growth is expected to come from a fundamental shift in consumer behavior involving a transition from the $25 billion per year movie rental business to an on-demand distribution model that involves a much wider library of content and concurrent higher usage rate. It's no surprise then, that telcos see an enormous opportunity ahead. The challenge, of course, is to deploy the necessary infrastructure and VoD applications on a scale and in a fashion that enables them to effectively to provide robust, competitive services.

Indeed, the creation of new on-demand video services will have significant impacts on the distribution networks of service providers. It's expected that each VoD session per subscriber will consume the bandwidth equivalent of a broadcast video channel. Consider the traditional broadcast video distribution model, where providers offer 500 or so channels to hundreds of thousands to millions of subscribers. With this model, the equivalent bandwidth to deliver a VoD service could potentially be consumed by a mere 500 concurrent VoD sessions! The challenge – and ultimate solution – is to develop a localized unicast video delivery infrastructure and VoD applications that prevent the erosion of multicasting services from unicast VoD sessions on the wide network.

Testing this VoD infrastructure involves testing the entire video streaming infrastructure, as well as individual network components, including routers, switches, and DSLAMs. Tests related to performance, scale and capacity are essential, as operators will need to know how many concurrent active users the infrastructure will support, how much bandwidth the system can handle, and how the network will respond to times of peak demand. Specific tests should center on the overall VoD quality of service, and include measuring the amount of packet loss, latency and jitter, as well as the number of joins and leaves supported, and the average join and leave times.

From the standpoint of the user's experience, an untested VoD infrastructure can cause a number of frustrating service impairments, including video image quality problems such as graininess, blurring, banding, and freezed and dropped frames. In addition, if the infrastructure isn't tested to validate the necessary capacities of various equipment, service problems such as the inability of the network to serve up requested VoD programming could occur – a situation that could likely infuriate a subscriber.

3.2.4 Unicast Streaming Servers

Conducting tests on unicast streaming servers and other Internet data-related servers, such as SMTP servers, is also important in terms of gauging the ability of the these devices to scale appropriately to meet expected demands of the applications. Tests should center on the number of active users supported, the number of sessions per second supported, and the amount of bandwidth consumed.

3.2.5 Streaming Test- Aware Devices

Likewise, traffic aware devices such as firewalls, server load balancers, switches and routers need to be tested to ensure that they are handling streaming video traffic appropriately in order to provide quality video service. One key area for testing aware devices is to check that the policies in place that give priority to voice and video over data traffic are performing effectively. In essence, network operators need to make sure that these devices are able to actually implement established policies regarding how different types of traffic will be routed through the network.

Let's put this in terms of the user's experience: It is much less onerous to make a customer wait four seconds for an email to come across the network than it is to make him wait an additional four seconds for his Super Bowl video to be refreshed. In short, subscribers are much more sensitive to video latency than to data-related latency, and operators must ensure that policies are in place and functioning properly.

3.2.6 Multicast Streaming Testing -DSLAM IPTV

According the latest market data from research firm Multimedia Research Group, new competition, trials and deployment in Europe, Asia and North America are resulting in accelerated growth rates for telco-related IPTV services. It's predicted that the market will grow from around two million IPTV subscribers in 2004 to more than 26 million by the year 2008.

And when it comes to the deployment of IPTV, the DSLAM is in many respects the point in the network where "the rubber meets the road." In an FTTN triple-play network architecture, it's vital to be able to test the multicast streaming performance of DSLAMs under real world loads to effectively gauge true IPTV service quality.

To do this, tests can be conducted that on one side of a DSLAM emulate an IP-based network that is carrying multiple video channels, and on the other side emulate set-tops boxes undergoing channel surfing behavior by a number of subscribers, including varying speeds of both sequential and random channel changing.

The key benefit of this testing is that it provides valuable measurements that illustrate how good or bad the user experience will be, in the form of packet loss, jitter, number and average times of joins and leaves, and channel zapping time.

If proper testing isn't performed to uncover potential QoS issues, a number of service impairments can rear their ugly heads. From the user's perspective, such impairments can manifest themselves as grainy pictures, blurred images, banding across the screen, frozen frames, and painfully slow channel change times. Worst-case user scenarios can involve situations where channel change requests are not performed at all to completely blank screens.

3.3 Network Assurance and Service Testing in a Scaled Deployment

The reality of IP networks is that they are extremely dynamic environments, meaning that as more people come on the network to use a variety of applications, service is inevitably affected. While network architects can design the network with all due respect to delivering optimal performance, and technicians and operators can build it and turn it on, the plain truth is that when services and live users come online, things will change. Live subscribers always create a different result than any emulation.

Therefore, in addition to the testing of specific triple-play infrastructure elements and applications, it’s also critical to perform a variety of proactive network assurance and service tests in a live, scaled deployment to assure that a customer’s services are working correctly. In short, it’s simply not enough to perform tests under a simulated load during pre-deployment -- it’s also essential to perform tests on a continuous basis with live subscribers.

After a telco has rolled out service with actual customers sitting on the ends of its lines, it’s vital that the provider is able to know when the network and its users are having problems, such as grainy pictures, painfully slow joins/leaves, or poor connect times for voice calls. And if there are problems, the provider must be able to figure out what the problems are, and then fix them as quickly as possible. Such live testing helps ensure that customers are enjoying a good experience with the service(s) – a condition that reaps many benefits: an optimal user experience incents positive word of mouth, generating new business, and also serves to reduce or eliminate customer churn – the keys to success in the marketplace. If testing isn’t performed, however, the quality of the network will ultimately degrade and customers will revolt.

Finally, the growing importance of SLAs (Service Level Agreements), where providers promise to provide a certain, measurable level of performance (such as guaranteed levels of bandwidth or service uptime) to their customers, also makes it mandatory that regular assurance and service testing is performed. Although SLAs are common with corporate customers, the inherent flexibility of IP networking means telcos will soon have the ability to offer residential customers different tiers, or classes, of service as well. If testing isn’t performed, and contracted service levels aren’t maintained by the provider, penalties usually apply in terms of rebates or credit applied to a customer’s account.

3.3.1 Triple-Play Provisioning and Change Verification

The first step during a commercial, live deployment of triple-play is to provision service – to put a customer on the network. During the deployment of new services, there’s always some concern about it not working right, even though the standard configuration was followed. To minimize (and hopefully eliminate) the frequency of this occurrence, the best thing to do is for the provider to perform provisioning tests to make sure the service works the way it's expected to before a customer's account goes live.

For example, if a telco sells a customer 5Mbps of data bandwidth, testing is needed to confirm this is what's being delivered before the provider tells the customer the service is active. Another situation that sometimes arises is when marketing claims don't jive with what customers are actually experiencing. Provisioning testing helps ensure that those two things are as close to being the same as possible.

Change verification testing is required to verify when a customer changes from one level of service to another, such as when a customer upgrades from a 3Mbps residential service to the business grade 10Mbps. A benefit of such testing is the ability to generate a report that validates that a change has been made, and provides for valuable backup when a customer calls in to say they're not getting the service they're paying for.

3.3.2 Triple-Play Service Performance Monitoring

One of the big issues with the new IP networks that are being deployed to support triple-play services is that there is typically very little visibility regarding what is going on in the network. Common methods for providing some level of visibility include CISCO SAA, RMON, or performing ping tests every five minutes. However, there are some intrinsic problems with these methods. One problem is that there is very poor granularity in the data collected: for example, ping testing every five minutes can't provide any meaningful data regarding what's happening with jitter and packet loss on a millisecond basis.

Simply increasing the ping rate to improve the granularity of the data doesn't provide a good solution, because the more the ping rate is increased, the more impact there is on the performance of the network itself. And no self-respecting network manager would dream of turning on RMON in a live network, as the act would cause her to bring down the entire network!

However, by putting dedicated performance monitoring hardware in the network, service providers can gain the added visibility required to get an accurate, real-time picture of network performance. Active measurements of performance can be made -- for example, by periodically making active calls across the network and reporting a MOS score back to a network manager. The MOS score can then be used with respect to an SLA. Or if the MOS score falls below a certain threshold, it can serves as an indication that there is a problem in the network that needs to be checked and corrected. Passive monitoring can also be performed to report on the performance of a number of QoS conditions, such as out of order packets, jitter trends, etc.

3.3.3 Triple-Play Diagnostics

Even with the best of intentions, sometimes even the best-laid plans can go awry. Let's say services are provisioned correctly, monitoring is up and running -- and a customer calls in to report a problem. Since a service provider can't monitor everything in the network at all times (if they did, the network couldn't scale), some problems can get past the safeguards set up. In addition, many problems are at the customer premises, past the demarcation point. Issues such as a misconfigured router or IAD, faulty home wiring or a dead mike on a phone fall into this category.

The solution is to make use of diagnostic appliances to help troubleshoot, and eventually pinpoint the problem. Both active and passive monitoring methods can be employed to perform diagnostic tests. With active monitoring, a service technician will try to replicate a problem reported by a customer. A common scenario for this is when a customer says he can't make a call to a particular area code or phone number. With passive monitoring, technicians are able to reference a passive historical record that reveals what happened on the network on any given date or time, showing any anomalies that might be causing service impairments.

With either active or passive diagnostic monitoring, the automation of the troubleshooting process is essential. Monitoring activities typically generate a large amount of data that needs to be turned into actionable, usable information. By employing advanced automation features of network-based diagnostic appliances that effectively automate the process of discovery, technicians can quickly perform automatic analyses of various network elements and make intelligent decisions based on the results of the tests being performed.

Ultimately, the benefit to the telco is the ability to decrease the escalation of service calls to service experts, the mean time to repair, and minimize costly truck rolls out to subscribers' households.

4.0 Reaping the Benefits of Triple-Play Testing

By now, the challenges, as well as the opportunities, of deploying a common network system for triple-play services should be clear. The benefits of such a deployment include infrastructure cost savings created by leveraging existing plant and equipment, fully integrated applications, world-class service offerings and newfound flexibility afforded by IP. Just as important, the capability of service providers to offer their customers seamless voice, video and broadband data services has the potential to deliver a host of additional benefits – not the least of which are new revenue streams and profits.

However, a triple-play deployment is likely to be the source of service disruptions, costly downtime and even network failure if it is rolled out haphazardly, without the necessary testing performed to ensure that all infrastructure elements and triple-play applications and services are working as expected. Indeed, there is no room for compromise when it comes to testing the triple-play network. Customers won't tolerate lesser service than they're getting with existing technologies. Comprehensive lab testing, pre-deployment testing, and performance monitoring and diagnostics must be must be performed in order for any real benefits to materialize.

4.1 Optimizing Network Performance

To be sure, telcos desire nothing less than five-nines uptime for their triple-play services, and they want their networks to be running at peak performance at all times. Deploying a well-built, converged IP-based network architecture goes a long way toward ensuring that services are delivered to subscribers as efficiently as possible. But as we know, converged networks and network overlays come with a unique set of challenges and potential problems. Without proper testing of the types described above, providers run the risk of being inundated by QoS issues and suffering from a range of serious to fatal service impairments. The only real way to achieve optimum network performance is through careful testing.

4.2 Driving Cost Efficiency

Of course, optimizing a network's performance through comprehensive testing comes with a price tag. However, the capital spent on proper testing equipment and solutions is money well spent, as it will surely save a great deal of money in the long run. If testing is performed correctly, it will result in a highly efficient network infrastructure that is properly scaled to accommodate the correct number of customers and traffic load.

Through miscalculation, a provider can easily overspend on equipment and other resources, or fail to provision effectively at the outset – inducing expensive redeployment efforts. Indeed, failure to test or to test appropriately is a major risk that will likely end up costing a telco much more money to fix in the long run than the incremental costs of testing.

4.3 Enabling New Features and Services

Perhaps the greatest benefit of the triple-play for telcos is the ability for them to provide a host of new enhanced services and applications that would have been impossible for them to offer before over their relatively low speed, inflexible legacy network infrastructures. Not only will a true IP triple-play allow telcos to enter the entertainment marketplace with video on-demand and IPTV, but the integration and bandwidth improvements will also help facilitate and promote a host of new and unique services, including video phone and video mail, e-learning, medical alerts, and SMS alerts on mobile devices of incoming calls to the home phone.

Indeed, with appropriate testing and network assurance procedures in place, providers can make the most of their triple-play network investments. The potential even exists for telcos to be able to leapfrog over cable MSOs and satellite providers in the race to provide the most advanced service bundle in the marketplace.

4.4 Satisfying Subscribers and Growing The Bottom Line

As we've discussed above, the benefits of a satisfied subscriber base go right to a telco's bottom line in terms of increased loyalty and uptake of new services, reduced churn, and increased revenues and profits. Therefore, it's of paramount importance that a telco ensures that the services and applications running on its network perform as flawlessly as possible. The only way to do this is to develop and apply a comprehensive program of triple-play network assurance and testing.

5.0 Conclusion

While the telco triple-play presents a huge opportunity for traditional service providers to mount a real offensive against competing cable MSOs -- and a chance to capture a greater "share of wallet" from subscribers who sign up for video, voice and data services -- the challenges of deploying a common network platform that supports such a service bundle, while providing subscribers the highest levels of QoS, are no trivial matter!

As we've seen, there isn't only one, universal method of implementation for triple-play, but rather a number of different network architectures and implementations that are possible. As there is no "one-size fits all" solution, a variety of different video telephony, VoD, and IPTV network strategies exist, and the need for a telco to ensure that the new video service infrastructure seamlessly interworks with legacy networks is of paramount importance.

In the end, the quality of the user experience is the factor that will ultimately determine whether a telco's efforts will succeed or not. And in order for a provider to offer a world-class user experience, the operation of triple-play applications, services and underlying infrastructure must be effectively managed for optimal performance. A solid understanding of the performance of both individual elements in the network and how they function in the holistic network sense – as a part of a dynamic interactive system of many network elements – is essential to maintaining, managing and delivering successful triple-play services. By developing and conducting a comprehensive program of lab testing and service management, telcos can assure this outcome and achieve the success they desire.