Performance vs. Resilience – Can You Have Both?

Force10 Networks
C-Series C300
Resilient Switch

A Broadband-Testing Report

First published April 2008 (V1.0)

Published by Broadband-Testing
A division of Connexio-Informatica 2007, Andorra La Vella

Tel : +376 633010
E-mail : info@broadband-testing.co.uk
Internet : http://www.broadband-testing.co.uk

Please note that access to or use of this Report is conditioned on the following:

1. The information in this Report is subject to change by Broadband-Testing without notice.

2. The information in this Report, at publication date, is believed by Broadband-Testing to be accurate and reliable, but is not guaranteed. All use of and reliance on this Report are at your sole risk. Broadband-Testing is not liable or responsible for any damages, losses or expenses arising from any error or omission in this Report.

3. *NO WARRANTIES, EXPRESS OR IMPLIED ARE GIVEN BY Broadband-Testing. ALL IMPLIED WARRANTIES, INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT ARE DISCLAIMED AND EXCLUDED BY Broadband-Testing. IN NO EVENT SHALL Broadband-Testing BE LIABLE FOR ANY CONSEQUENTIAL, INCIDENTAL OR INDIRECT DAMAGES, OR FOR ANY LOSS OF PROFIT, REVENUE, DATA, COMPUTER PROGRAMS, OR OTHER ASSETS, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.*

4. This Report does not constitute an endorsement, recommendation or guarantee of any of the products (hardware or software) tested or the hardware and software used in testing the products. The testing does not guarantee that there are no errors or defects in the products, or that the products will meet your expectations, requirements, needs or specifications, or that they will operate without interruption.

5. This Report does not imply any endorsement, sponsorship, affiliation or verification by or with any companies mentioned in this report.

6. All trademarks, service marks, and trade names used in this Report are the trademarks, service marks, and trade names of their respective owners, and no endorsement of, sponsorship of, affiliation with, or involvement in, any of the testing, this Report or Broadband-Testing is implied, nor should it be inferred.

TABLE OF CONTENTS

TABLE OF FIGURES

BROADBAND-TESTING

Broadband-Testing is Europe’s foremost independent network testing facility and consultancy organisation for broadband and network infrastructure products.

Based in Andorra, Broadband-Testing provides extensive test demo facilities. From this base, Broadband-Testing provides a range of specialist IT, networking and development services to vendors and end-user organisations throughout Europe, SEAP and the United States.

Broadband-Testing is an associate of the following:

NSS Labs (specialising in security product testing)
Limbo Creatives (bespoke software development)

Broadband-Testing Laboratories are available to vendors and end-users for fully independent testing of networking, communications and security hardware and software.

Broadband-Testing Laboratories operates an **Approval** scheme which enables products to be short-listed for purchase by end-users, based on their successful approval.

Output from the labs, including detailed research reports, articles and white papers on the latest network-related technologies, are made available free of charge on our web site at HTTP://www.broadband-testing.co.uk

Broadband-Testing Consultancy Services offers a range of network consultancy services including network design, strategy planning, Internet connectivity and product development assistance.

EXECUTIVE SUMMARY

- Within the confines of this test we looked to prove that it *is* possible to achieve very high levels of performance along with very high availability without having to spend huge amounts of IT budget on multiple redundant devices, as has been common in the past.

- In order to test the Force10 C-Series C300 we created a very powerful testbed based on two, fully-loaded Ixia XM12 chassis' capable of delivering very high levels of real traffic from Layer 2 to Layer 7. We configured a C300 with a full population of 384 10/100/1000 Ethernet ports and redundant Route Processor Modules (RPMs) and Power Supply Modules (PSMs).

- Using the testbed over several weeks of testing, we found that the Force10 Networks C300 chassis-based Ethernet switch is totally capable of running line-rate performance when fully loaded (1.536 Tbps throughput), port-configuration wise, while retaining high-availability.

- We created a series of tests that both proved the C300 capable of sustaining line-rate performance across all 384 ports *and* its ability to support a wide range of application types, notably real-time applications, through delivering very low latency figures. This is a perfect combination for almost any user scenario.

- During simulated failures (pulling out RPMs) we found zero or minimal packet loss.

- Latency figures recorded were extremely low and very impressive.

- These low latency figures would explain why we were able to deliver Converged Services (data, video, VoIP) over 286 ports concurrently, supporting a mix of users (12,954 HTTP, 30,000 video, 12,000 VoIP) that totaled 54,954. We have never seen any previously published figures by any test house showing this level of capability from a single device. Moreover, it proves that for Converged Services or Triple Play environments there are cost-effective routes for deployment.

- Overall, we were very impressed with every aspect of the performance of the Force10 Networks C300 and feel that it boxes well above its weight from a price performance perspective.

- It is the most outstanding chassis-based switch we have tested to date and therefore recommend it fully to anyone in the market for a high-capacity core switch for server aggregation or next generation wiring closet switch.

- It confirms Force10's statement that its C-Series is designed for resilient, critical application support requiring very low latency and cost effective, high-density port population for a wide range of applications, including VoIP, video and wireless – a contemporary solution for a very real-world scenario, in other words.

INTRODUCTION: PERFORMANCE VS RESILIENCE – CAN YOU HAVE BOTH?

The IT world is full of trade-offs.

For example, security and performance are directly at odds with each other. Performance is all about clearing the path; security is about blocking it. And much the same belief exists regarding performance versus resilience. If you want to focus on maintaining the fabled “five nines” uptime, then surely there is some overhead that impacts on performance; that or you simply have to spend enormous amounts of money on multiple, redundant units in order to achieve both targets.

So what about the idea of creating a product that gives you the best of both worlds – a realistic possibility or a pipe dream? That is what we’re here to find out with Force10’s C-Series, chassis-based switch. Is it really possible to offer line-rate performance *and* achieve maximum uptime in a single product?

The Times They Are A’ Changing (and so is the traffic profile)

Another point is that applications are changing. So while, for years, IT designed the enterprise network architecture to emphasise different characteristics at each tier, so high performance and high availability in the data centre, scalability in the LAN core and low cost and basic functionality in the wiring closet, these design principles no longer apply.

Why? Because the traffic mix is far more varied than it once was, with both traditional data (database, office applications) and new generation traffic such as block graphic/bulk data, P2P and other realtime applications making use of voice and video all running concurrently and largely critical to the business.

All these application types therefore have to be catered for, in terms of prioritisation, guaranteed bandwidth and absolute availability – not a trivial requirement, even in a core LAN environment.

100% Uptime? Well, it’s not as common as you night think, given that we are in 2008 now, long after so-called “non-stop computing” was introduced. Unplanned down time still costs businesses thousands to tens of thousands of dollars per minute. Even companies whose networks have been designed with best-in-class availability suffer an average of five hours of unplanned down time annually, according to analysts.

There is, then, a general feeling still that you can’t have performance and resilience – in that there must be some overhead or trade-off associated with achieving either. What we are trying to prove in this report, where the focus is on testing Force10’s C300 switch, is that it **is** possible to achieve both, right here, right now. So read on…

C-SERIES – PRODUCT OVERVIEW

The Hardware

The Force10 Networks C-Series has been designed as a resilient, chassis-based switch range that delivers reliability, network control and scalability – all in line with our earlier observations. So the C-Series is designed to support vital applications with very low latency across converged networks.

Figure 1 – C-Series C300 Switch

Up to 384 line-rate 10/100/1000Base-T ports (with full 15.4 W Class 3 PoE support) are available in a 13-RU chassis. Up to 64 line-rate 10 GbE ports with pluggable XFP modules, are also supported. Force10 claims five microsecond switching latency under full load for 64 byte frames with a switch fabric capacity of up to 1.536 Tbps and up to 952 Mbps L2/L3 packet forwarding capacity.

A 1+1 Route Processor Module (RPM + redundant switch fabric) design is aimed at providing maximum resilience and redundancy, along with continuous, runtime data plane monitoring and advanced in-service CLI diagnostic functions.

Add to this, 2+1 system and 4+1 PoE power supply redundancy (eight module slots) with load sharing power bus, enabling uninterrupted VoIP calls during a power supply failure and you have a pretty resilient infrastructure to work with. All components can be inserted and removed online and the chassis features environmental self-monitoring.

Jumbo frame support (9252 bytes) is included, as is link aggregation (8 links per group & 128 groups per chassis), four queues per port and a maximum of 1024 VLANs (tag values can range from 1 - 4094). IPv6 is supported at Layer 2 and the switch can load-balance based on an IPv6 header.

There are several line card options including:

- 4-port 10 Gigabit Ethernet Line Card: With pluggable XFP modules, these line cards supports distances of up to 80 km, thereby spanning LAN, MAN, and WAN options. This line card provides a density of up to 32 line-rate, non-blocking 10 Gigabit Ethernet ports in a single chassis. The C-Series architecture enables claimed line-rate performance with QoS and Access Control Lists (ACLs) enabled, jumbo frame support, and simultaneous L2 switching and IPv4 routing.

- 8-port 10 Gigabit Ethernet Line Card: With a similar specification to its 4-port relation, this line card provides up to 64 line-rate, non-blocking 10 Gigabit Ethernet ports in a single chassis – something that Force10 claims to be industry-leading.

- 48-port 10/100/1000Base-T Line Card: This provides up to 384 line-rate, non-blocking 10/100/1000Base-T Ethernet ports with PoE in a single chassis. Additionally, the C-Series ASICs and architecture enables line-rate performance with QoS and Access Control Lists (ACLs) enabled, jumbo frame support, and simultaneous L2 switching and IPv4 routing.

Figure 2 – 48-Port Copper Ethernet Line Card

- 48-port 10/100/1000Base-T PoE: Providing IEEE 802.3af PoE support, each chassis can provide up to 15.4 W of power per port, while maintaining full system and PoE power supply redundancy. Up to 384 line-rate, non-blocking 10/100/1000Base-T Ethernet ports with PoE in a single chassis are supported, with the same functionality and features of its non PoE relation.

- 48-port Gigabit Ethernet Line Card with Pluggable SFP Modules: With pluggable SFP modules this line card supports distances up to 80 km over fibre, and, with pluggable 1000Base-T modules supports distances up to 100 m over Cat5 UTP. 100Base-FX SFP modules support distances up to 2 km, providing a flexible solution in mixed fibre/copper and mixed speed installations.

Each line card has its own CPU dedicated to protocol housekeeping functions such as sFlow or BFD. Local processing reduces messaging between line cards and preserves Out-Of-Band bandwidth.

The FTOS Operating System

At the heart of every Force10 switch product is the FTOS operating system.

This has been primarily designed to:

- Eliminate single points of failure in as many system components as possible, including both hardware and software.
- Constrain any failures that do occur to only one system component or subsystem.
- Ensure that, when a subsystem fails, becomes compromised, or needs updating, the recovery can be accomplished quickly without disrupting the continued operation of the overall system.
- Minimise the number of operating systems, code branches and releases in development across product lines.

These principles are the basis for the Force10 Networks Resiliency Architecture, which allows the FTOS switch/routers to continue network operations during a number of fault conditions. These include the failure of hardware components, software faults and restarts, link failures, protocol restarts, and attempts by intruders to disrupt normal traffic flow.

The company says it adopted the multi-layered resiliency architecture model after a comprehensive analysis of the causes of system crashes and of all potential sources of catastrophic system failure, which showed that the majority of network disruptions are due to software failures and user mistakes. This conclusion is supported by all the major IT analyst groups.

FTOS is based on NetBSD, which Force10 claims has three unique advantages:

- An unparalleled degree of protocol maturity and stability derived from its roots in BSD Unix.
- A modern kernel architecture, featuring process modularisation and memory protection.
- A high degree of portability across multiple hardware architectures allowing a wide selection of control processors.

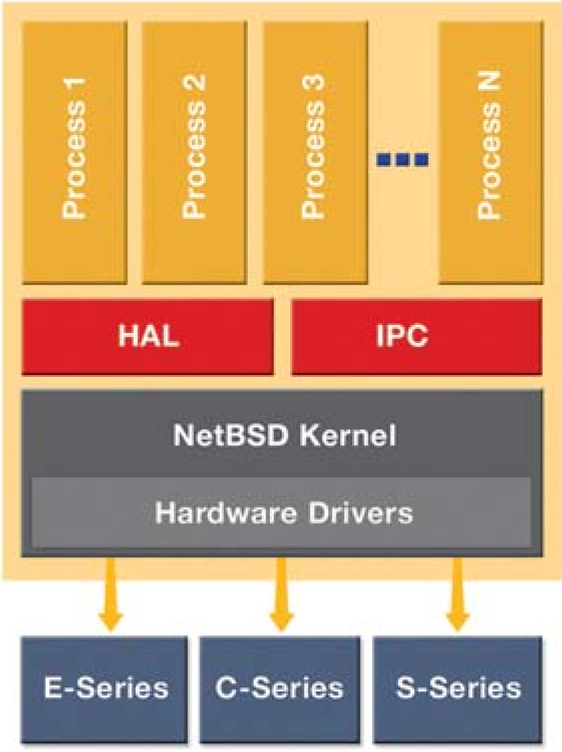

Figure 3 – System Architecture

The basis of the system is the kernel whose functions are limited to providing a stable operating system and performing tasks such as memory allocation and scheduling. Other applications that traditionally execute in the kernel in a monolithic network operating system have been moved to user space where they run as independent, modular processes. Separate processes have been created for:

- Layer 2 functions and protocols (STP, LACP, etc.)
- Layer 3 functions and protocols (ARP, OSPF, BGP, static routing)
- Various system services and management functions (SNMP, CLI, etc)
- Security services and protocols (SSH, TACACS+/RADIUS, ACLs)

When these processes need to share information, the exchanges are all channelled through an inter-process communication (IPC) mechanism, which allows processes to communicate while maintaining memory protection between them.

The hardware abstraction layer (HAL) is a layer of software that decouples FTOS from the specific details of the underlying hardware. The primary function of HAL is to isolate the modular processes from a major rewrite whenever the FTOS is ported to a different hardware platform.

A Modular Approach

Protected Memory: The FTOS memory architecture allocates a separate protected address space for each independent process and its associated subsystems.

Fault Containment: The combination of modularisation and protected memory space ensures that faults in one process are limited only that process (plus any other processes that require current information from the failing process via IPC).

Pre-emptive Process Scheduling: Pre-emptive process scheduling prevents a single process from monopolising one of the control plane CPUs.

Portability: FTOS is readily portable to new hardware platforms due to the abstraction and code efficiencies provided by HAL.

Cross-Platform Portability and Scalability: FTOS can readily scale to support all of Force10 Networks switch/router product families from high-end core and data centre switches (the E-Series) to aggregation switches (the C-Series) to access switches (the S-Series).

Distributed Processing in the FTOS Control Plane: In order to achieve massive scalability together with unparalleled resiliency, the Force10 Networks switch/routers use a fully distributed architecture with all packet forwarding decisions made by the data plane ASICs, and all control functions performed by the switch/router CPU(s) running FTOS.

In the C300, FTOS processing is distributed across a processor on each route processor module (RPM), plus an additional processor on each line card. The modularity of FTOS allows the individual control plane processes to be readily partitioned among the available processors in the chassis. Resilience is significantly improved because failure of a process on a CPU will not affect other functions.

Management & Configuration

OpEx (Operating Expenses) are a major, but misunderstood, cost of ownership. A key contribution to OpEx are ongoing device management costs. So, having common management functionality and a common user interface across the Force10 product range is an obvious benefit – reduced training, reduced issues, reduced costs.

At the CLI (Command Line Interface) anyone familiar with the classic “IOS style” command set will be instantly familiar here. Configuration is therefore straightforward for anyone familiar with a classic CLI. Force10 does also offer a GUI-based management platform, the Force10 Management System (FTMS).

In terms of automated management, the FTOS kernel monitors all processes to ensure operations are within normal limits of resource utilisation. FTOS also provides system-wide monitoring for out-of-range environmental conditions and other fault conditions, such as unsynchronised configurations of line cards. Fault reporting and automated fault correction are included minimise system interruption.

A modular approach meant that the tracing of software errors to specific processes is made easy, providing for quick resolution. Importantly, from a resilience perspective, out-of-band (OOB) switched Ethernet is used for communications within the control plane and for communication between the control plane and the line cards.

Figure 4 – FTMS Management GUI

FORCE10 C300: PUT TO THE TEST

The Testing Outlined

As we outlined during the introduction to the report, the emphasis of the testing is to prove that it is possible to achieve high performance *and* availability with a thoroughly contemporary mix of traffic types.

Our testing therefore encompasses line-rate performance at switch capacity, availability in the event of simulated line card failure, latency monitoring to ensure support for real-time applications, even in a stressed environment, and genuine Converged Services testing using the highest port count we have yet encountered, or seen published to date.

The Testbed

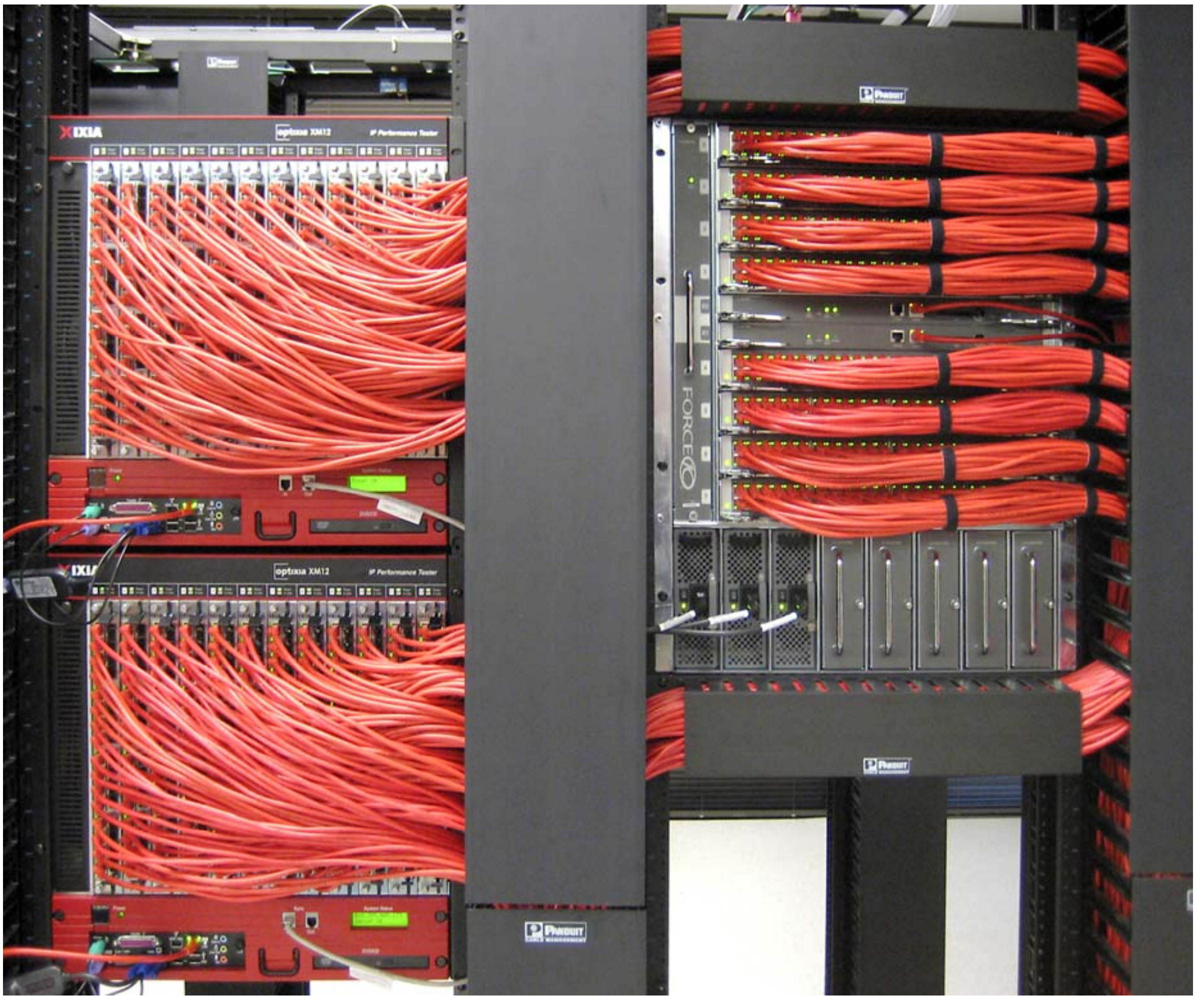

Figure 5 – The Testbed – Ixia XM12s Plus Force10 C300 at Ixia's iSimCity

In order to generate both sufficient levels of data traffic and the right combination of traffic types, we turned to Ixia (www.ixiacom.com) and its XM12, chassis-based traffic generation technology – two fully-loaded XM12s in fact. Coincidentially Force10's C300 switch is the first product to be tested in Ixia's new proof-of-concept lab Santa Clara, California called iSimCity (Figure 5).

A key point here is that, in keeping with Broadband-Testing's requirement to provide as close to real world test conditions as possible, the Ixia testbed produces real traffic – HTTP, FTP, VoD etc – not just meaningless, generic traffic generation. Using the Ixia LSM10GXM3 load modules we were able to generate real Layer 2 through Layer 7 traffic at line-rates.

Controlling the test hardware were Ixia's IxAutomate and IxLoad software applications. A combination of the software modules enabled us to both fully-load the switch at the lower layers, but equally run very high levels of Layer 7 application traffic, providing us with a completely realistic test.

Throughput tests followed the RFC2544 guidelines for testing across a number of different frame sizes – from 64 bytes to 1518 bytes – across different switch port topologies including a fully meshed test across all ports. We also tested with a 9216 byte Jumbo frame sizes as a matter of course, since the C300 supports frame sizes up to 9252 bytes. However, this is not something we see as being that important in comparison with the other tests. The latency tests were also based around the same guidelines.

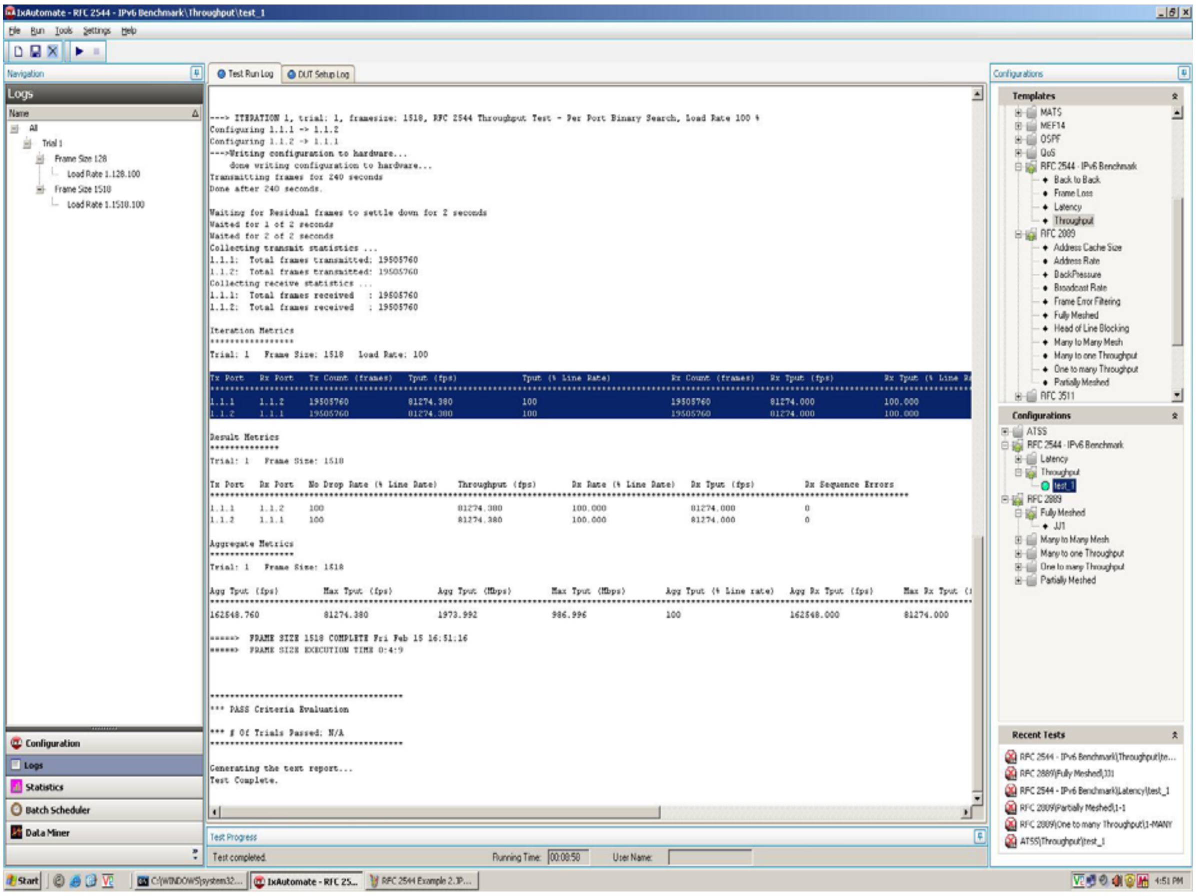

Figure 6 – IxAutomate – RFC2544 Test Example

For the device under test (DUT) configuration, the C300 came fully loaded with 384 10/100/1000 (8x48-port line cards) Ethernet ports, dual (redundant) Route Processor Modules, and three PSMs for a 2+1 redundancy configuration, though this can be extended significantly with eight PSU slots actually available, designed for PoE use. An interesting alternative configuration, perhaps for a future test, is the 8x8 10 GbE line-rate card option, giving 64 ports of 10 GbE, but that's for another day…

It is important to note that the C300 chassis has a totally passive copper backplane with no active components, so there is no obvious single point of failure. Total switch capacity in this configuration is 1.536 Tbps.

In order to test and stress the chassis in as many different ways as possible, in some case we tested against three specific topologies:

- Ports on the same MAC on the same line card
- Ports on different MACs on the same line card
- Ports on different MACs on different line cards

The Test Results

Line-Rate Performance

Our first test focused on line-rate performance. Using Ixia IxAutomate and following the RFC2544 guideline spec, we started at a relatively low port-count, working our way up to 384 ports to identify the point at which the C300 was unable to function at 100% line-rate across all ports. We never did reach that point because the chassis performed flawlessly and this in a fully-meshed, bi-directional test.

It is important to note that, where the line-rate appears to fall very marginally below 100% we worked out that this was a function of the configuration and the way Ixia has to calculate throughput so, sometimes, "100%" is actually 99.99% 99.98% because that is the maximum that can be tested/calculated.

384 Ports Full Mesh Throughput							
Frame Size	Agg Tx Tput %	Agg Tx Rate %	Agg Rx Tput (fps)	Agg Rx Rate %	Agg Tx Count (frames)	Agg Rx Count (frames)	Agg Frame Loss
64	571,418,842	100	571,418,842	100	5,714,188,416	5,714,188,416	0
128	324,323,174	100	324,323,174	100	3,243,231,744	3,243,231,744	0
256	173,912,640	100	173,912,640	100	1,739,126,400	1,739,126,400	0
512	90,213,965	99.99	90,213,965	99.99	902,139,648	902,139,648	0
1024	45,974,707	100	45,974,707	100	459,747,072	459,747,072	0
1280	36,915,072	99.98	36,915,072	99.98	369,150,720	369,150,720	0
1518	31,208,678	100	31,208,678	100	312,086,784	312,086,784	0

That the C300 was able to sustain line-rate at all frame sizes is especially outstanding. We have observed in the past, for example, that smaller frame sizes such as the 64 byte frame, can be difficult to sustain at line-rate, so well done to Force10. A complete absence of frame loss through the testing confirmed that line-rate had been achieved with zero data loss. Even with Jumbo frames, we observed a transmit line-rate of 99.9%

and a receive rate of 99.03% indicating an aggregate frame loss of only 0.87% - minimal in the real world. Bear in mind, however, that the Jumbo testing is less applicable to majority real-world use than the smaller frame sizes tested, but revealed impressive performance nonetheless.

Latency

Of course, line-rate performance is just one measure of switch performance, albeit a very important one in high-traffic environments (and who doesn't want to maximise their investment in hardware nowadays?). For certain application types, such as real-time applications, latency, or delay, are a vital issue. Even in more traditional application areas, poor RTR (Round-Trip Response) times, as a result of latency in the network, can be – at best – frustrating and – at worst – can make an application effectively unusable.

For the latency tests, we especially needed to consider the three port combinations defined in the testbed definition section, since we expected different latencies in each case because of the different topologies. Again, we used Ixia IxAutomate. Since the C300 uses store and forward switching, we configured the Ixia software accordingly.

From an IxAutomate perspective, the Store-and-Forward script measures the raw time the frame is gone from the tester, then subtracts the frame duration to give the forwarding decision time of the DUT. A DUT operating in Store-and-Forward mode, such as the C300, knows what egress port to forward to early in the frame, but waits for the whole frame to come into the ingress port so that it can check the CRC and drop it if it's bad, which is good switch behaviour. Larger frames take longer to forward, but if the DUT is keeping up, the decision time should be constant, so we were looking to achieve low latency results throughout the test, though we naturally expected a gradual increase as frame size increased, because of the nature of store and forward switching. Consistency was therefore a key requirement to look for here. We started with a configuration where all ports were on the same MAC on the same line card, then where ports were on different MACs on the same line card and finally where ports on different MACs were on different line cards.

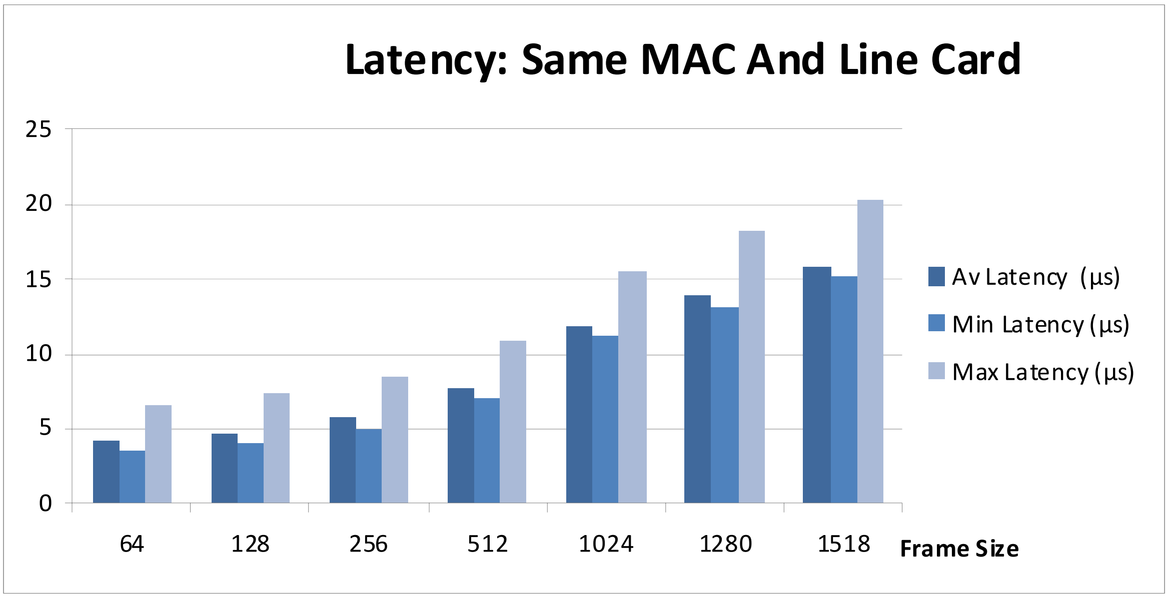

Figure 7 – Latency Test: Same MAC and Line Card

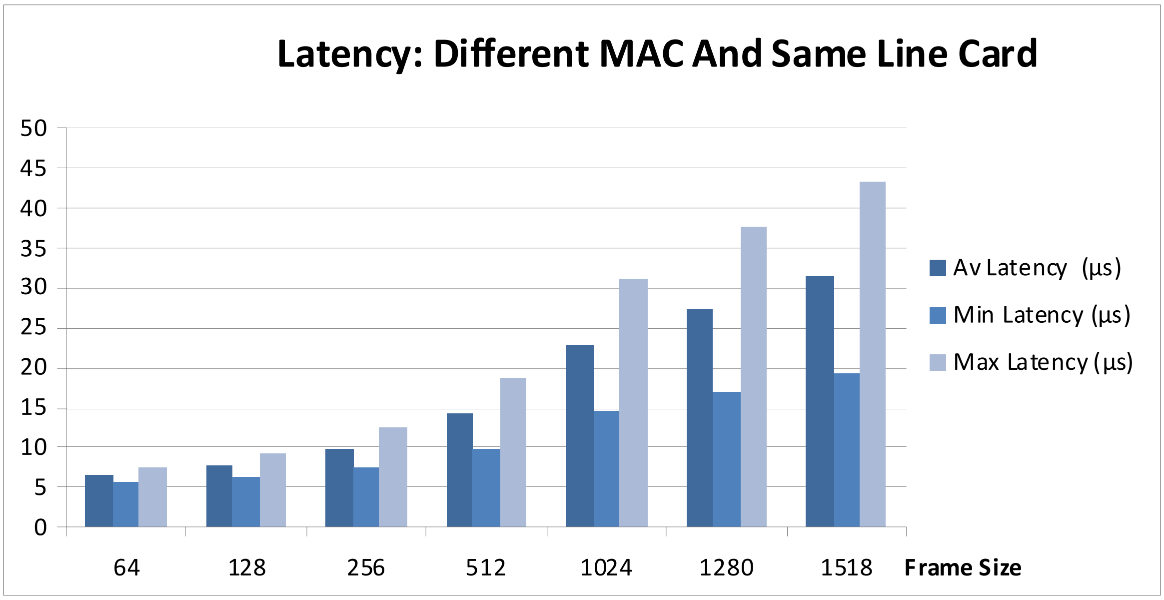

Figure 8 – Latency Test: Different MAC and Same Line Card

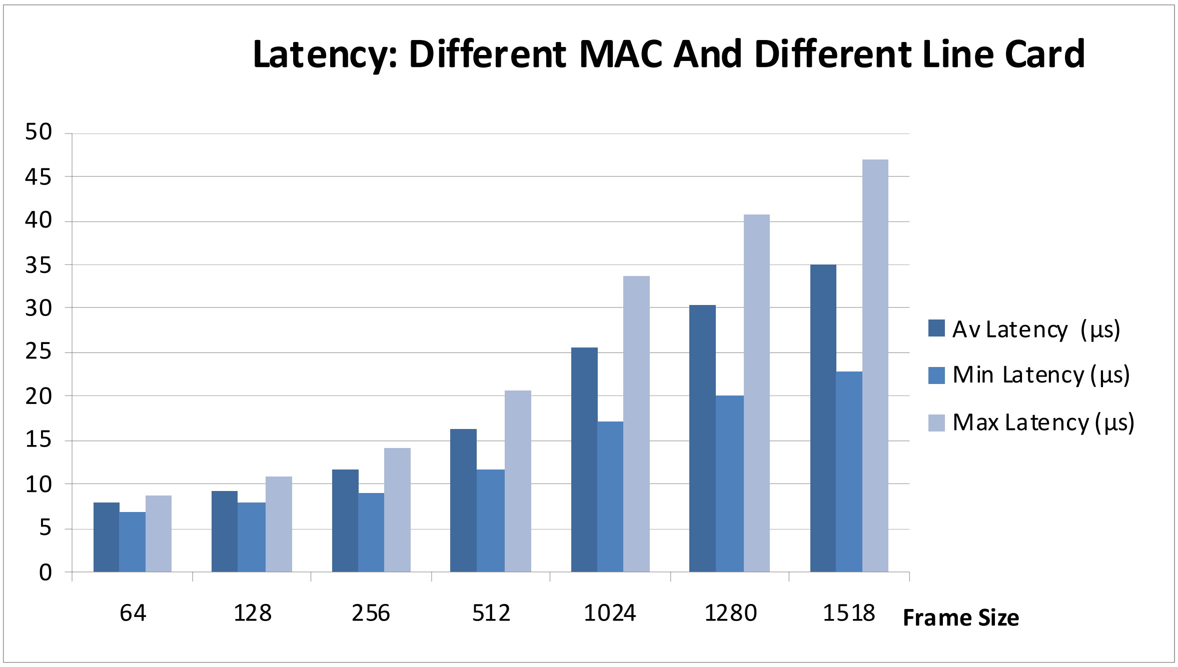

Figure 9 – Latency Test: Different MAC and Different Line Card

As we can see from the set of three results, in all cases latency figures recorded were *extremely* low. Bear in mind that these are in microseconds, not milliseconds (the figures were too small to graph) and the results speak for themselves.

The slight increase in latency as we added complexity to the test network port topology in each of the three stages was exactly in line with expectations. That the "worst case" scenario above shows very low latency indeed still, is a truly excellent result.

High Availability

To reiterate the point we made earlier – ultra-high performance without ultra-high availability and reliability is largely irrelevant. So we put the C300 to the test with a series of simulated line card failures while running traffic across, first 48 ports, then all 384 ports.

We were looking for minimal packet loss using the 384 GbE ports, despite the complete loss of a loaded 48-port line card in each case. We ran a fully-meshed test in four variations, failing over in turn to each of the processor modules (RPM), as follows:

- Initial baseline test with one RPM, traffic on primary RPM
- Remove primary RPM, traffic on primary fails over to secondary RPM
- Re-insert secondary RPM, traffic on primary RPM
- Remove secondary RPM, traffic on primary RPM

High Availability 384 Ports Full Mesh

	Agg Tx Tput (fps)	Agg Tx Rate %	Agg Rx Tput (fps)	Agg Rx Rate %	Agg Tx Count (frames)	Agg Rx Count (frames)	Agg Frame Loss	Agg Frame Loss %
Run 1	571,418,842	100	571,418,842	100	5,714,188,416	5,714,188,416	0	0
Run 2	571,418,842	100	552,741,709	96.73	5,714,188,416	5,527,417,093	186,771,323	3.27
Run 3	571,418,842	100	571,418,842	100	5,714,188,416	5,714,188,416	0	0
Run 4	571,418,842	100	552,112,909	96.62	5,714,188,416	5,521,129,087	193,059,329	3.38

As we can see, removal of an RPM (primary or secondary) resulted in just over 3% aggregate frame loss on a 384-port test at line-rate. Hot insertion of the secondary RPM resulted in absolutely zero loss. In each case, after failover, the test continued to run perfectly meaning, even in a worst-case scenario, downtime will be minimal (subsecond at most).

Converged Services Scenario

Having proved that the C300 could sustain line-rate at full capacity, with high availability and minimal latency, we then wanted to put the switch into an application scenario and see how far we could push it (and the test equipment!).

We chose Converged Services, given that there is so much noise being made currently about IPTV/Video delivery and related real-time applications and services, and the complexities of delivering these services. Here we simulated the kind of deployment of video over IP, as part of VoIP + data delivery that you would get in an office suite, apartment block or hotel, for example, where video is piped from a basement machine room or equivalent. It is important to note that these results were achieved without the need to add any kind of QoS or CoS bandwidth controls. We combined large-scale HTTP (web browsing) activity, with live video and VoIP (using SIP) to simulate a massive Converged Services or so called 'Triple Play' delivery scenario. Typically these kinds of tests involve the simulation of hundreds of users across tens of actual Ethernet ports. In this case we successfully worked our way out to 286 live Ethernet ports, simulating tens of thousand of users.

Converged Services Test Summary

Protocol	Sim Users	Servers	TCP Connections Established		Avge TCP Conn Req Rates		Client Tput (bytes/sec)
HTTP	12,954	48,000	515,633	526,927	86.1	0	1,479,750
IPTV/Video	30,000	96	38,754	31,119	56.6	0	-
SIP	12,000	48,000	0	0	0	0	-
Totals	54,954	96,096	554,387	558,046	142.7	0	1,479,750

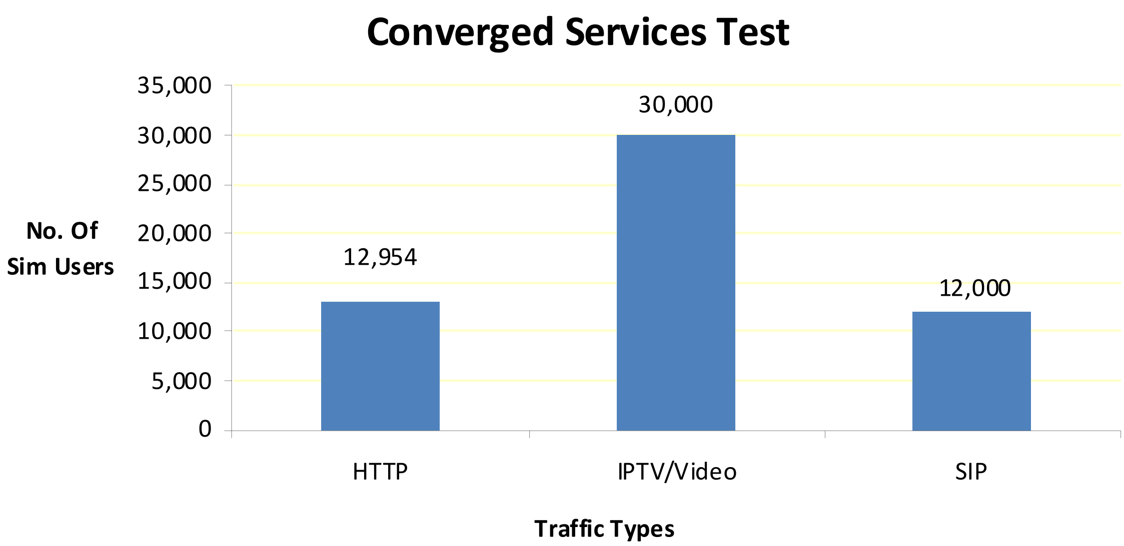

Figure 10 – Converged Services Test: Simulated User Totals

These figures show what can be achieved in terms of concurrent Converged Services delivery – to almost 55,000 users/subscribers – with a single chassis such as the Force10 C300. These are very significant numbers; from a service provider or content provider perspective, it means that deployment costs can be kept in check – vital for profitable deployment of Triple Play services.

SUMMARY & CONCLUSIONS

Within the confines of this test we looked to prove that it *is* possible to achieve very high levels of performance along with very high availability without having to spend huge amounts of IT budget on multiple redundant devices, as has been common in the past.

Using the testbed over several weeks of testing, we found that the Force10 Networks C300 chassis-based Ethernet switch is totally capable of running line-rate performance when fully loaded (1.536 Tbps throughput), port-configuration wise, while retaining high-availability. We created a series of tests that both proved the C300 capable of sustaining line-rate performance across all 384 ports *and* its ability to support a wide range of application types, notably real-time applications, through delivering very low latency figures. This is a perfect combination for almost any user scenario. Moreover, because of the port combinations versus the per-blade bandwidth capacity, with a full 10 GbE card capacity we can run the chassis harder still – watch this space on that one…

Latency figures recorded were extremely low and very impressive. These low latency figures would explain why we were able to deliver full Converged Services (data, video, VoIP) over 286 ports concurrently, supporting a mix of users (12,954 HTTP, 30,000 video, 12,000 VoIP) that totaled 54,954. We have never seen any previously published figures by any test house showing this level of capability from a single device. Moreover, it proves that for Converged Services or Triple Play environments there are cost-effective routes for deployment.

Overall, we were very impressed with every aspect of the performance of the Force10 Networks C300 and feel that it boxes well above its weight from a price-performance perspective. It is the most outstanding chassis-based switch we have tested to date and therefore recommend it fully to anyone in the market for a high-capacity core switch for server aggregation or next generation wiring closet switch. It confirms Force10's statement that its C-Series is designed for resilient, critical application support requiring very low latency and cost effective, high-density port population for a wide range of applications, including VoIP, video and wireless – a contemporary solution for a very real-world scenario, in other words.

APPENDIX: MORE TESTBED DETAILS – THE IXIA XM12 CHASSIS AND COMPONENTS

Ixia test systems claims to deliver the industry's most comprehensive solutions for the performance, functional, and conformance testing of networks and networked applications.

The 12-slot Optixia XM12 modular chassis provides an ultra-high density, highly flexible platform on which an Ixia test system can be built. Operating in conjunction with the Aptixia family of test applications, the Optixia XM12 provides the foundation for a complete, high performance test environment.

Wide arrays of interface modules are available for the Optixia XM12. The chassis supports up to 192 Gigabit Ethernet ports, 36 - 10 Gigabit Ethernet ports, and 24 Packet over SONET (POS) or Asynchronous Transfer Mode (ATM) ports. These modules provide the network interfaces and distributed processing resources needed for executing a broad range of data, signalling, voice, video, and application testing from Layers 2-7.

Each chassis supports an integrated test controller that manages all system and testing resources. Resource ownership down to a per-port level coupled with hot-swappable interface modules ensures a highly flexible, multi-user testing environment. Backward compatibility is maintained with existing Ixia interface modules and test applications to provide seamless migration from and integration with existing Ixia test installations.

Ixia's 10 Gigabit Ethernet XM LAN Services Modules (LSMs) offer unprecedented scalability, performance, and service testing flexibility as part of the Optixia XM test system. The 10 GbE XM modules provide the industry's highest density 10 Gigabit Ethernet test solution with up to 36 - 10 Gigabit Ethernet test ports in a single Optixia

XM12 test system. A broad portfolio of edge/core testing solutions are supported, including performance, scalability, and conformance testing of Layer 2-3 devices at the control and data planes, and high performance Layer 4-7 testing of content-aware devices and networks.

The Ixia 10 GbE XM LSM supports a comprehensive portfolio of service testing solutions for next generation service provider networks, including Metro Ethernet E-LAN and E-LINE services; and MPLS VPNs, such as Layer 2 VPNs, Layer 3 RFC 2547 VPNs, and VPLS.

As networking devices become increasingly complex, so must the analysis equipment designed to assess their performance. Such sophisticated analysis systems must support multiple powerful routing/bridging protocol emulations that are flexible, highly scalable, and easy to use. In addition, the analysis systems must be able to generate wire-rate traffic and automatically analyse thousands of traffic flows with comprehensive QoS analysis. For this reason, Ixia offers a range of applications capable of delivering test solutions at all levels, such as the IxAutomate and IxLoad software applications featured in this test.

As a leader in performance analysis solutions, Ixia's Aptixia IxNetwork meets these requirements and is specifically targeted for the performance and functionality testing of high-speed, high-capacity routers, switches and application servers.

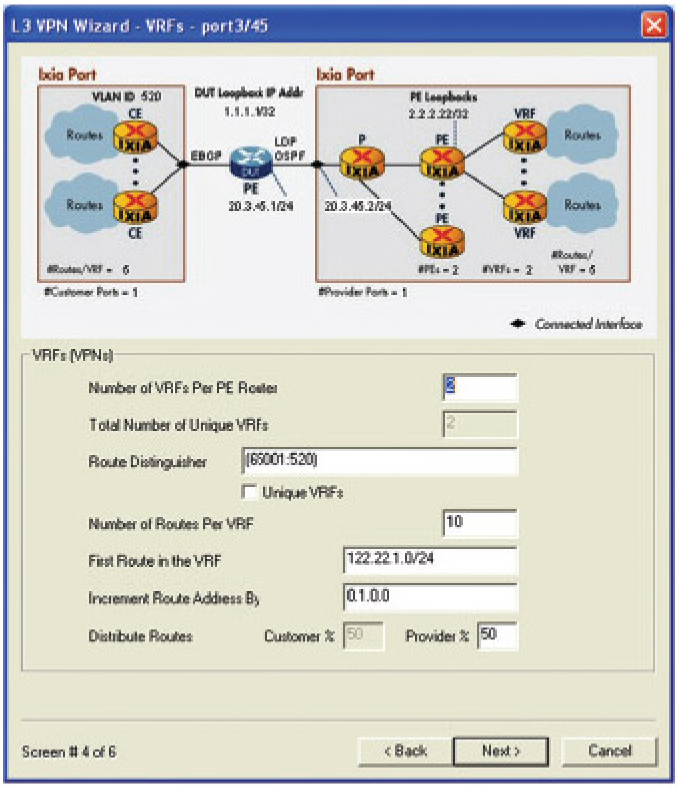

IxNetwork offers users the flexibility to customize the application to meet a wide range of requirements for testing complex network topologies consisting of thousands of network devices. Millions of routes and reachable hosts can be emulated within the topology. IxNetwork also provides users with the ability to customize millions of traffic flows to stress the data plane performance. Sophisticated configurations can be created using powerful wizards and grid controls in the graphical user interface. With its enhanced real-time analysis and statistics, IxNetwork is capable of reporting comprehensive protocol status and detailed per-flow traffic performance metrics.

As network functions continue to be aggregated into devices, it becomes increasingly important to consider security and encapsulation protocols, such as NAC, PPP and L2TP. IxNetwork provides the ability to authenticate emulated clients and to establish broadband sessions. Traffic can then be encapsulated over the tunnelling protocols.

Key Features

- Emulation of Internet-scale routing topologies to determine scalability limits.
- Simulation of network instabilities to characterize the performance of network convergence.
- Easy-to-use Protocol Wizards to quickly and precisely set up complex topologies.
- Powerful Traffic Wizard to create millions of traffic flows for validating emulated networks and hosts.
- Realistic emulation of enterprise application traffic, including stateful TCP, HTTP, E-mail, Video, RTSP, Telnet and FTP. Now with voice and triple play traffic and IPv4/IPv6.
- Setup of and encapsulation through PPP sessions and L2TP tunnels.
- Authentication utilizing 802.1X and Cisco Network Admission Control (NAC).
- Comprehensive protocol and per-flow traffic statistics for detailed troubleshooting and analysis.
- Flow Detective™ to find best and worst performing flows in real-time.
- Built-in data-rate capture and analysis tools.
- Flexible Test Scheduler to dynamically flap emulated topology on the fly.
- Use of industry standard RFC-based data plane tests, along with other automated testing.
- Enhanced Tcl API allows complete automation of IxNetwork functions.