Solarflare Solarstorm® 10 Gigabit Ethernet Controller

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

TABLE OF CONTENTS

LIST 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.

SOLARFLARE Solarstorm: EXECUTIVE SUMMARY

- Within the scope of this test we set out to see how the Solarflare's Solarstorm SFC4000E vNIC 10 Gigabit Ethernet controller would perform in a variety of user environments. These scenarios focused around performing at line-rate with uni and bi-directional traffic flows, then transferring that testing environment into a virtualised environment with VMware and Xen, and finally seeing how it could optimise an iSCSI environment.
- Using Ixia's Chariot v6.5 application to generate traffic at 10/20 Gbps (uni/bi-directional) we ran a series of tests over a number of different traffic stream variants on different server OS platforms – Windows 2003, Windows 2008 and Linux Redhat. With a standard combination of 64 KB data files and 1500 bytes Ethernet frame size, we found that we were able to hit and sustain line-rate, while still witnessing significant idle time on the server CPUs, meaning we were achieving line-rate without getting anywhere near maximizing the serverperformance capacity.
- We repeated these line-rate tests in two virtualised environments, using VMware and Xen v3.3 respectively as the virtual machine environments. Testing VMware ESX v3.5 we were able to achieve 9.3 Gbps, considered line-rate given the maximum recordable figures using the 1500 byte frame size (a purely arithmetic limitation) - an excellent result compared with any previously published performance figures we've seen in a virtualised environment. In a Xen test we also achieved line rate, given the parameters of the test setup, which was arguably even more impressive, given Xen's relatively limited benchmarking to date.
- Testing in the iSCSI environment focused upon how the Solarstorm controller could enhance performance with the iSCSI data digest feature enabled. This provides data corruption protection but also imparts a very significant performance overhead, and is usually disabled for this reason. The ideal solution then, is to have the digest enabled to protect the data, but also to remove that performance overhead as much as possible. The Solarstorm controller managed to improve digest-enabled performance by over 300%, giving us near 10 Gigabit line-rate over iSCSI, with full data corruption protection.
- We also compared Solarflare Solarstorm performance with that of an Intel® 82598 10 GbE based adapter, using Ixia's IxChariot and the line-rate tests again, with a 1500 byte frame size. On average we achieved almost double the performance with the Solarstorm controller, over that of the Intel equivalent – a very impressive achievement.
- Throughout the testing, minimal tuning was required on the Solarflare Solarstorm, so in a purely vanilla configuration – a real-world scenario – we proved that the Solarstorm controller, not only achieves line-rate working in a variety of server OS environments, but is also able to optimise other environments, such as virtualised deployments or with iSCSI subsystems.

INTRODUCTION: DO WE NOW NEED 10 GIG?

It's not long ago since we were asking – do we need Gigabit Ethernet.

Now we are asking: do we need 10 Gigabit Ethernet? Well, there's a very obvious answer to that question. If you now have lots of gigabit connections to your desktops and laptops alike on your network, then simple arithmetic dictates that, in order to accommodate all of these gigabit connections, you need something larger than a gigabit core backbone and server connections. Hence the 10 Gigabit argument and the emergence of products such as Solarflare's Solarstorm 10 GbE controller that we are testing here.

Affordability is a key accelerator here. Thanks to stupendously low costs per port now for Gigabit Ethernet copper connections, people are simply throwing 10/100/1000 into their networks by default. While this, at one point, would have been irrelevant, due to the inability of the client machines, and even the servers, being unable to generate a gigabit of traffic, thanks to the enormous recent advances in core processor and serial I/O PCI Express replacing PCI-X bus technology, relatively low-end servers are now fully capable of supporting very high bandwidths.

What were single core, single CPU systems are commonly now dual CPU with dual or even quad core architectures. Perhaps more importantly, the Operating Systems (OSs) and applications are now more able to take advantage of these multi-core, multi-processor architectures. And we haven't even started to talk about virtualised environments, which form a significant element of the testing in the report. More on this later…

With top research companies such as Dell'Oro Group forecasting a major uptake of 10 Gigabit Ethernet switch ports in 2008, the requirement for 10 Gigabit at the server will be magnified even more. This provides us with a valid case for widespread adoption of 10 GbE, both in the core of the network and the server. Whereas throwing bandwidth at the problem is never a sensible option on the WAN/Internet, due to the inherent latency issues that need to be resolved. In the core of a well-designed LAN, that delay problem is typically far less of an issue.

However, the need to guarantee that bandwidth availability is very much a requirement, given the increasing use of real-time applications such as VoIP and video (in its various formats). In these environments it is essential that spare bandwidth is *always* available both in the core and at the server, from where most of this traffic originates.

Moreover, no network manager in their right mind would typically run a network at near capacity and herein lies another potential problem at the server, namely saturating the CPU and hence effectively killing the server. We therefore need high-bandwidth support, but in a way that doesn't over-stress the server, giving sufficient CPU time to carry out the many and various tasks these servers are actually being used for.

Over the years in the Broadband-Testing labs we've seen many times just how easy it is to overwhelm a server. While the latest Intel and AMD technology has certainly moved the game on here, any way that CPU time can be saved is a huge bonus. In this report, then, not only are we looking at what levels of performance the Solarflare Solarstorm can achieve in a variety of different test scenarios, but also how efficient it is in its ability to provide significant levels of idle time at the CPU.

SOLARFLARE Solarstorm – PRODUCT OVERVIEW

The Hardware

The Solarstorm Ethernet controller is available in three variants: 10GBASE-CX4, 10GBASE-SR/LR, and 10GBASE-T.

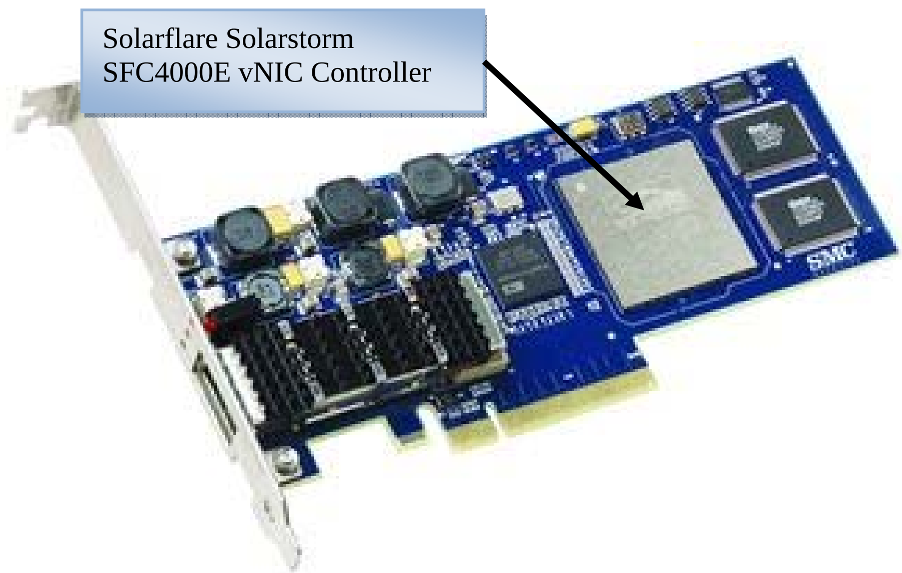

Figure 1 – Solarflare Solarstorm Controller

These controllers are also available through 3rd parties, such as SMC, under its' TigerCard brand. They are based on a high-speed PCIe x8 host interface, designed to combine high performance with low latency, something we have proved in past Broadband-Testing reports on the Solarflare card. The XFP variant with pluggable modules for SR or LR means that 10 GbE performance can be delivered over distances from 300m (SR) to 10km (LR) between ports.

The Solarstorm controller includes the basic MAC logic required for a 10 GbE connection plus several basic hardware functions that support what Solarflare calls stateless offload. In this arrangement the hardware that performs wire-speed checksums for IDP, UDP and TCP transactions is completely transparent to both the upper-layer hardware and the TCP/IP stack. It also performs a header and data digest (a CRC operation) function which is often disabled to speed throughput because it is so compute-intensive (see iSCSI testing later).

The controller also accelerates I/O virtualisation by providing a bypass stage that handles hypervisor forwarding operations between the host system and its virtual guest operations. Its on-chip hash table does not divert all flows passing through the device for offload, but is structured to identify and deal with the most frequently-occurring operations for offload. This allows it to resolve and accelerate enough flows to greatly relieve host processing of interrupts. It supports 4096 virtual systems via an on-chip lookup table and vNIC interface logic.

A transmit pacing mechanism keeps each vNIC separate from the others and ensures they maintain bandwidth allocations. Priority is assigned to provide an upper boundary on the number of packets/s each vNIC can pass and guarantees that packets are transmitted in the correct order. QoS and bandwidth allocation to each virtual system is handled by a multi-level scheduler which uses a weighted round-robin algorithm. Each of its queues can have one of 20 priority levels. This vNIC or Virtual NIC technology is the primary building block of the controller architecture.

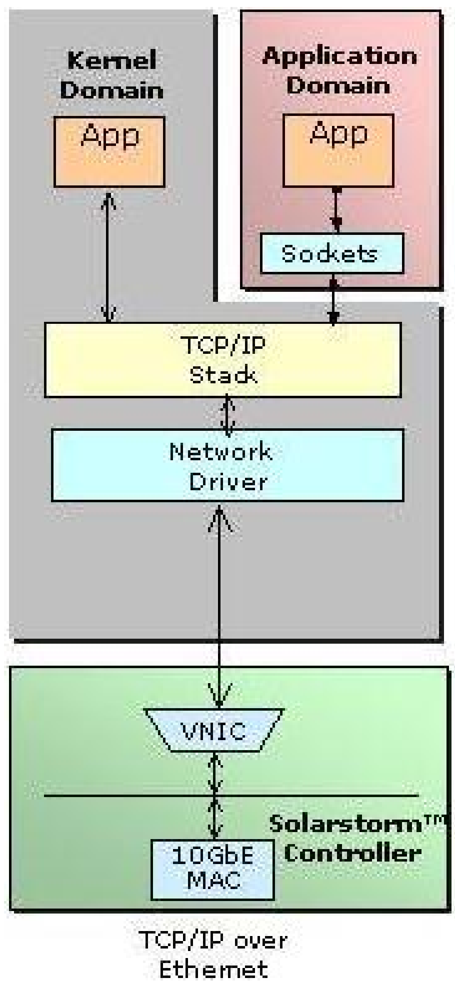

Figure 2 – The vNIC Architecture

This enables virtual, protected interfaces to be created for each running guest operating system or application, giving it a direct pipeline to the Ethernet network. This connection can be made through either a traditional virtualisation OS or directly through a hardware interface that handles MAC address mapping and traffic scheduling. In bypassing the virtualisation OS, the controller reduces processing power significantly.

Solarflare claims this is far more effective than the alternative TOE (TCP Offload Engine) approach favoured by some vendors, not least because the advances in the standard server hardware and operating systems is arguably making this technology redundant anyway. In this way it provides the most efficient way of using both the network and CPU. It is important to reiterate that the controller is completely industry standard, with plug and play installation, providing binary compatibility with all data centre and HPC applications.

Other key features include:

- RSS: RSS (receive side scaling) distributes I/O load across all CPUs and cores, allowing for greater use of servers, especially in virtualised environments.
- IP/UDP/TCP checksum offload: This provides calculation and validation for the checksums found in IP, TCP and UDP headers, saving a significant number of valuable CPU cycles.
- Transmit rate pacing (per queue): This provides a mechanism for enforcing bandwidth quotas across all guest operating systems. The software is re-programmable on the fly to allow for adjustment as congestion increases on the network.
- iSCSI acceleration: iSCSI Header and data digest CRC provide greater levels of data integrity without consuming additional, expensive CPU cycles.
- Jumbo Frame Support: The control provides support for up to 9000 byte jumbo frames.
- MSI/X Support: Supporting both MSI and MSI-X enables higher levels of performance on both legacy and contemporary systems.
- Remote Boot: The controller includes support for etherboot, PXE boot, iSCSI Boot and Linux BIOS, thereby providing flexibility in cluster design and diskless servers. A small programmable on-board ROM contains the boot code which provides this functionality.
- XAUI interface: This provides flexibility on the 10 GbE physical layer interface, implementing four lanes at 3.125 GHz full duplex.
- IP flow filtering: This enables the hardware to steer packets based on IP, TCP and UDP header contents.

More Bandwidth, Lower Power Consumption

Solarflare claims its Solarstorm vNIC is the industry's lowest-power 10 Gigabit Ethernet controller. The Solarstorm controller operates at just 2.2 watts, setting the bar for low power, whereas other vendor controllers have been reported as operating at as high as 15+ watts – a very significant difference.

The reason it is so significant is because increasing power, cooling and footprint costs, combined with a growing number of applications, users and transactions are acting as drivers for a number of new developments, not least server and storage virtualisation. This brings more energy considerations with it, so power optimisation becomes almost as important as performance optimisation.

The company believes that energy and power efficiency will continue to be top-of-mind for CIOs and IT administrators in 2008, a belief backed up by market analyst reports.

SOLARFLARE Solarstorm: 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 Solarflare's Solarstorm 10 Gigabit controller can perform in a variety of user environments, as follows:

- Performing at line-rate with uni and bi-directional traffic flows.
- Performing at as close to line-rate as possible in a virtualised environment with VMware and Xen respectively.
- Optimising an iSCSI subsystem environment.

The Testbed

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 Chariot software, in v6.5 format. Its important to mention here that, in keeping with Broadband-Testing's requirement to provide as close to real world test conditions as possible, the Ixia test bed produces real traffic – http, ftp, VoD etc – not just meaningless, generic traffic generation.

For the throughput-related tests we used the standard Chariot file size of 64 KB with a frame size of 1518 bytes, and the sendw32 test script. These tests were performed on an Intel 5400 dual-core, dual-processor platform with 8 GB RAM. For the iSCSI optimisation testing we used two dual-core, single processor Dell 2950 servers running the DiskTest tool with 64 KB files and jumbo frame size (9216 KB) as we were simulating bulk data transfer.

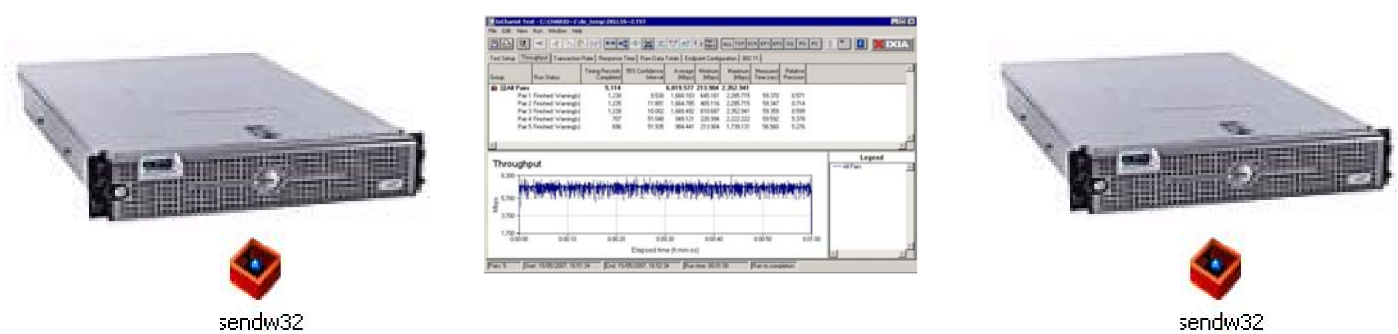

Figure 3 – The IxChariot-based Test Bed

The Test Results

Line-Rate Performance

Our first test focused on line-rate performance. Using IxChariot we ran a series of tests across both Windows (2003 and 2008 Server) and Linux (Redhat Enterprise Linux 5) environments, using a mix of Chariot tests based around five and 32-stream uni-directional and 10/64 stream bi-directional tests. We ran several iterations of the tests to ensure consistency and tested with both 1500 byte and 9000 byte (jumbo) frame sizes.

We also compared the Solarflare Solarstorm performance with that of an Intel-branded controller to give us an idea of the relative level of performance achieved by the Solarflare product. All tests were rerun several times, including on long, overnight runs in order to validate the consistency of the tests as well as the reliability of the Solarstorm controller.

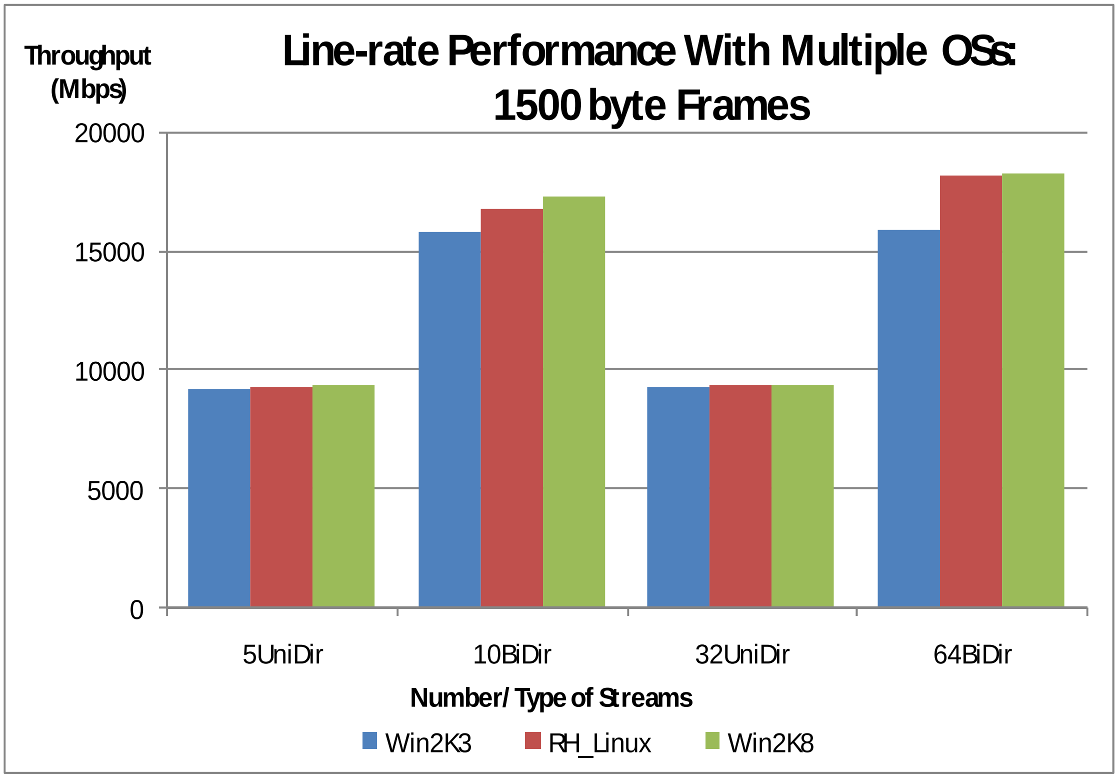

Figure 4 – Line-rate Performance With Multiple OSs 1500 byte Frames

Testing with 1500 byte frames, we were able to achieve in excess of 9.3 Gbps uni-directionally and almost 18 Gbps in bi-directional mode, giving us effective line-rate allowing for the arithmetical allowance required when using the 1500 byte frame size.

Note: It is important to note that, where the line-rate appears to fall below 100% we worked out that this was a function of the configuration and the way Ixia has to calculate throughput based on 100% divided by the frame size used (1500 bytes in this case) so "line-rate" appears less than 10 Gbps for our tests.

Overall we found performance between the different OSs to be relatively consistent meaning, in the real-world, that users would not see significant differences, regardless of which OS they chose to run. Of the three, Windows Server 2003 was the worst performing of the three, while Server 2008 and Redhat Linux (release 2.6.18-53.1.4.el5PAE kernel, running in 32 bit mode) were very closely matched.

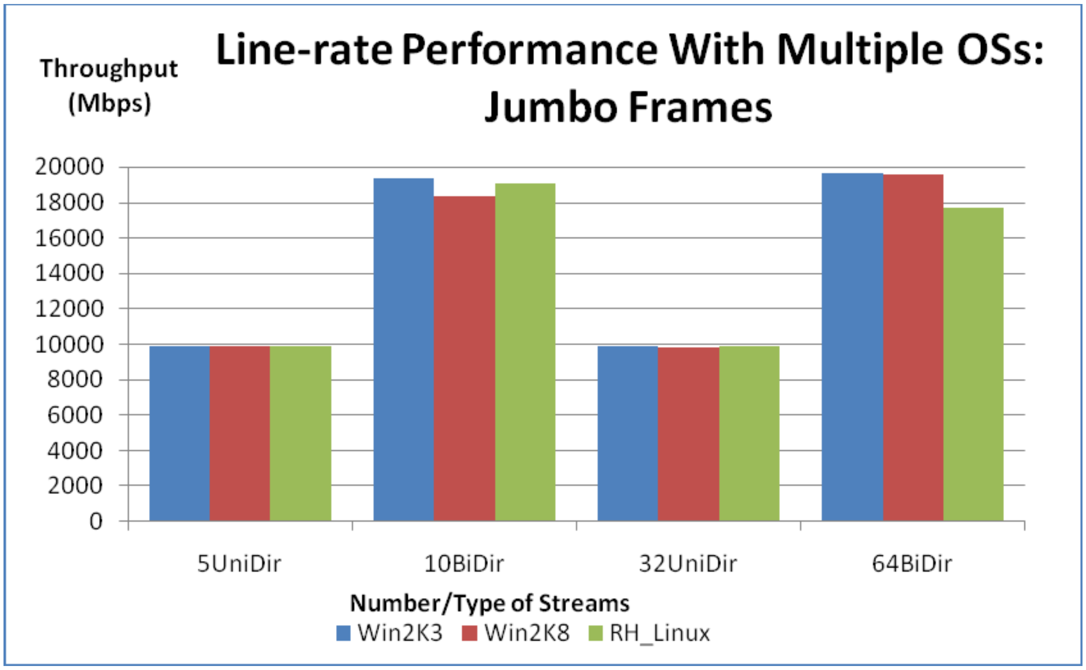

Figure 5 – Line-rate Performance With Multiple OSs Jumbo Frames

With the jumbo frame size we again effectively hit line-rate, in this case peaking at almost 9.9 Gbps in uni-directional mode and over 19.6 Gbps in bi-directional mode. With the jumbo frames, the OSs were more evenly matched still, with very little real-world difference in performance across the three – good news for users with a mixed OS environment.

Despite effectively achieving 10 Gigabit line-rate with the Solarstorm controller, the IxChariot client and server CPU utilisation figures were exceptionally low, rarely exceeding 30% and generally below 25%, regardless of the OS under test. This is an excellent result and shows the Solarflare controller to be extremely efficient in operation.

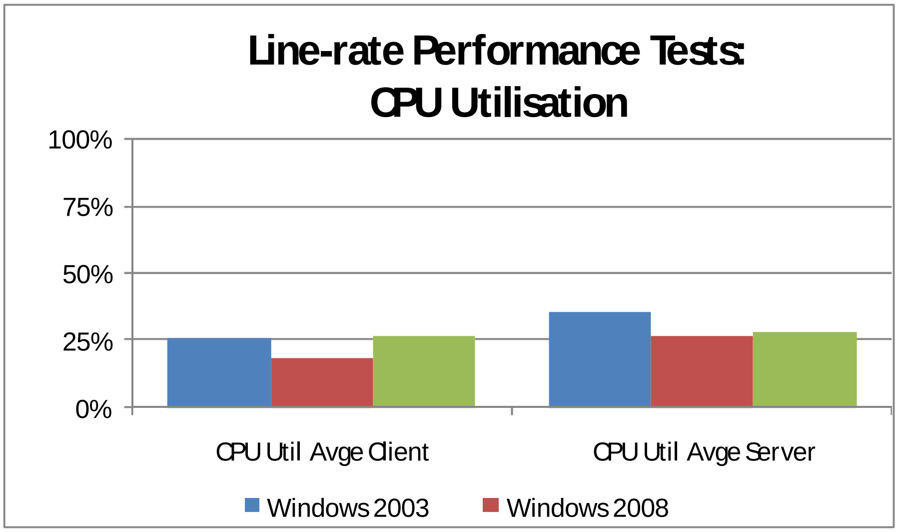

Figure 6 – Line-rate Performance: CPU Utilisation

We also compared the Solarstorm with an Intel® 82598 10 GbE controller to see if the latter could match the Solarflare product's performance, as a means of establishing whether we were witnessing exceptional performance here, or merely typical.

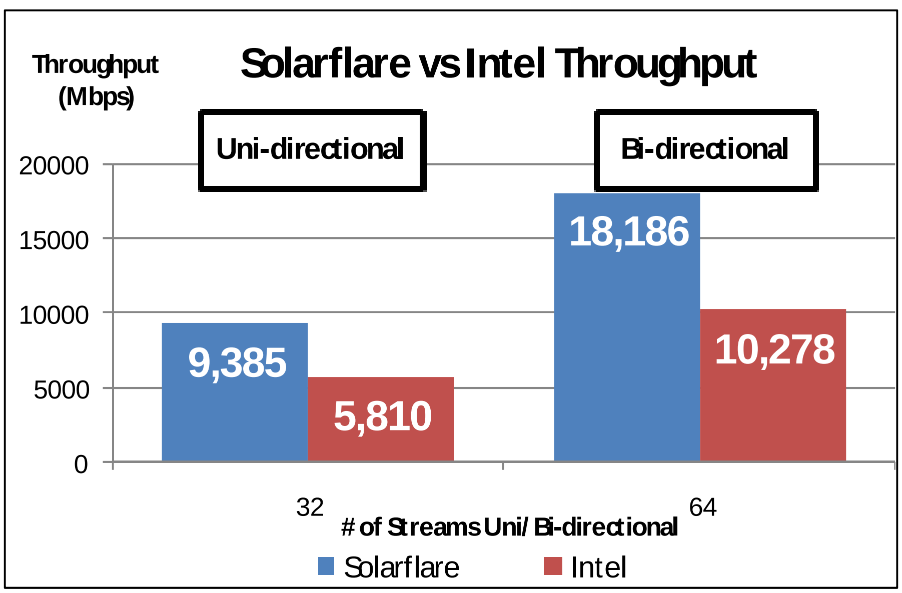

Figure 7 – Solarflare versus Intel Performance

Taking a subset of our IxChariot tests – the 32-stream unidirectional and the 64-stream bi-directional – we can see that the Solarflare controller provided almost double the performance of the Intel equivalent testing with a 1500 byte frame size using Redhat Linux as our OS in this example. And this is without any kind of non-standard enhancements, just pure and simple, native Ethernet support.

Virtualised Environment Performance

Line-rate performance in a native OS environment is one thing, but we were interested to see just how close we could get to this level of performance in a virtualised environment, given the dramatic increase in virtualisation adoption right now.

Using the same basic test configuration as for the OS performance tests, running IxChariot with 64 KB file size and 1500 byte frame size, we created two virtualised environments; one running VMware ESX v3.5 and a second running the latest version of Xen from Xensource in the form of v3.3. We then ran a series of tests where we increased the virtual machine count over a multi-step test, to a maximum of 16 virtual machines – a significant test environment.

Starting with VMware, historically, performance in virtualised environments has been less than impressive, but the likes of VMware have been busy trying to resolve this issue, with the introduction of features like its NetQueue support. NetQueue is a performance technology that is claimed to significantly improve performance in 10 GbE virtualised environments so we put it to the test here (note – it is effective only when receiving, not transmitting, traffic).

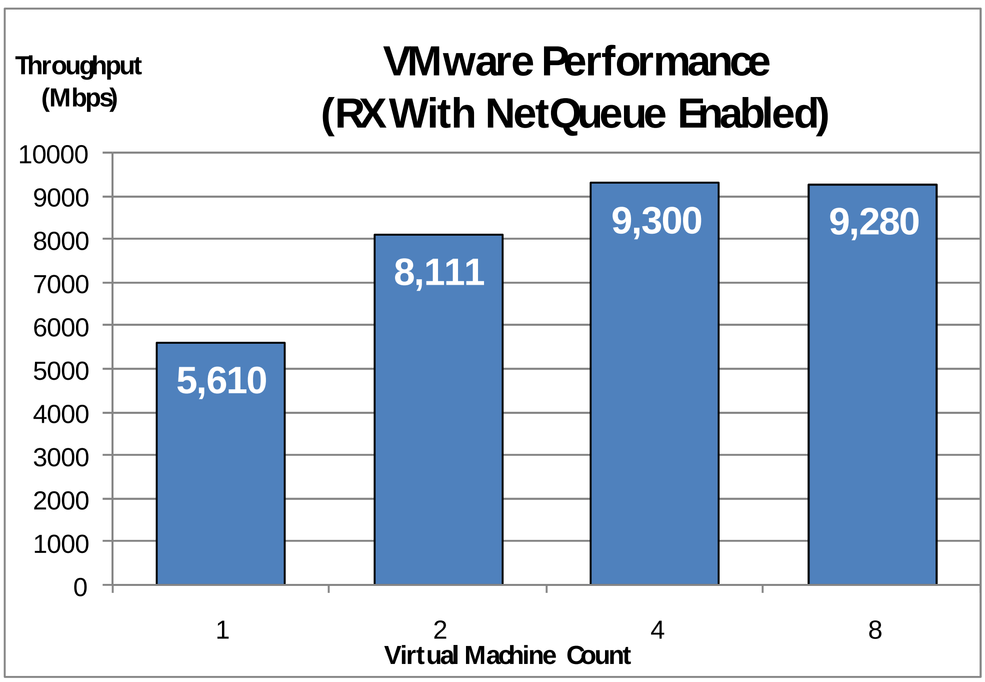

Figure 8 – VMware Performance With NetQueue Enabled

With NetQueue enabled we recorded a very impressive series of results, achieving near line-rate performance from two to 8 virtual machines. This is a dramatic improvement on performance in virtualised environments of old.

Equally importantly, CPU utilisation was extremely low, typically below 20% at the client machine and around 50% at the server – a very impressive set of results.

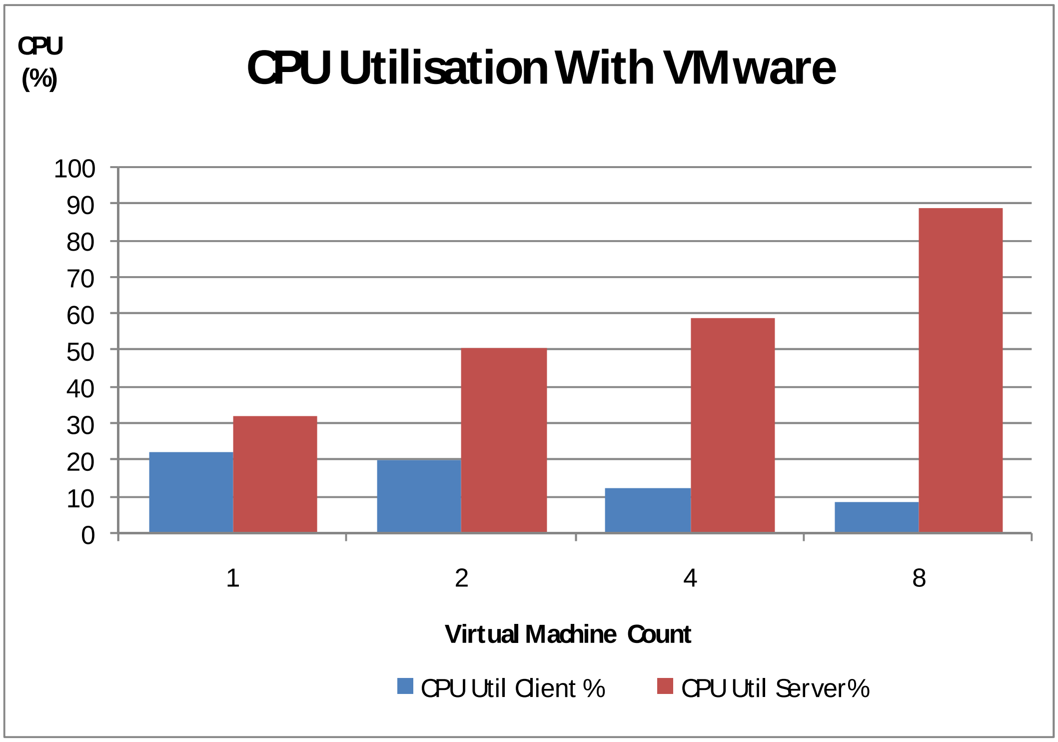

Figure 9 – VMware CPU Utilisation

We also tested in a Xen environment. Here we also put Solarflare's netfront acceleration for Xen to the test. This module provides a fast path for network traffic where there is hardware support available for the netfront driver to send and receive packets directly to a NIC (such as that from Solarflare).

For TX traffic, netfront gives each accelerated front end the option of sending each packet, which it can accept (if it wants to send it directly to the hardware) or decline (if it thinks this is more appropriate to send via the normal network path).

As the graph (next page) shows, performance before enabling the Solarflare acceleration was very poor, falling well short of 3 Gbps. With acceleration peaking in excess of 9.3 Gbps which, given the arithmetic allowances required when working with a 1500 byte frame size, means we essentially achieved line-rate once more.

What we felt was equally significant, was any impact (of a positive kind) the Solarflare acceleration could have on CPU utilisation, so we put this to the test.

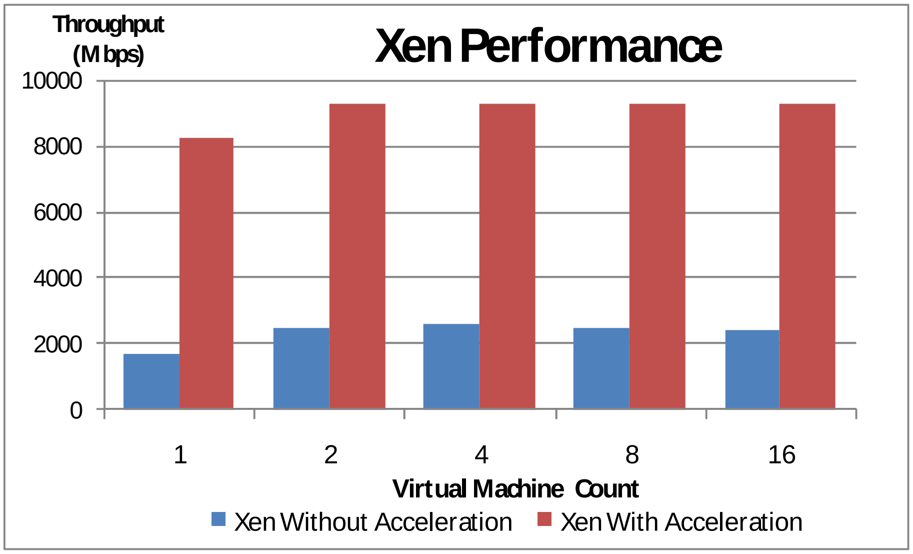

Figure 10 – Xen Performance: With and Without Solarflare Acceleration

The CPU utilisation was measured using Xentop in batch mode, reporting on per domain statistics, including the number of seconds of physical CPU time spent by the domain.

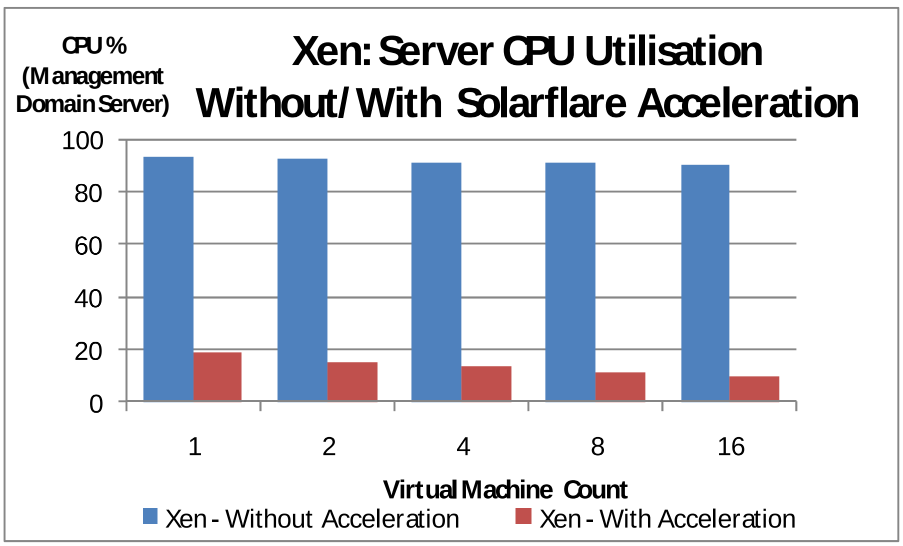

Figure 11 – Xen CPU Utilisation: With and Without Acceleration

The code works out the rate at which CPU is being used while the system is active (e.g. ignoring the times when Chariot is waiting around) and converts that to a percentage, so we get a very accurate measurement.

By comparing the management domain CPU activity we were able to see just how effective or not the Solarflare acceleration proved to be. The answer was, hugely effective, with a dramatic decrease in CPU utilisation with acceleration enabled, despite the equally dramatic increase in performance. This was a truly spectacular result – absolutely the best of both worlds.

Looking at overall CPU figures for the test, outside of the important management domain server activity, we still saw similar or marginally (50% maximum) more CPU utilisation accelerated against unaccelerated performance, despite the near 400% performance increase, with CPU utilisation always staying well below 50%.

iSCSI Optimisation

For our iSCSI testing, the focus was on optimisation within an environment where the digest was enabled.

The digest was created as a means of preventing data corruption. TCP defines a 16-bit checksum to detect corruption of TCP packets. But there are instances where the TCP checksum does not protect iSCSI data, as when data is corrupted while being transferred on a PCI bus or while in memory. The iSCSI protocol therefore defines a 32-bit CRC digest on iSCSI packets in order to detect data corruption on an end-to-end basis. However, this digest creates a very significant performance overhead, so is typically disabled, meaning that data corruption is unprotected – not ideal!

We therefore created a test bed, using Dell 2950, single processor, dual core systems, back to back, running the DiskTest utility with 64 KB writes and reads (bypassing local buffering to ensure true traffic measurement) and using jumbo frame sizes (appropriate for bulk data transfer) in order to assess the impact of the Solarstorm controller in a digest-enabled configuration.

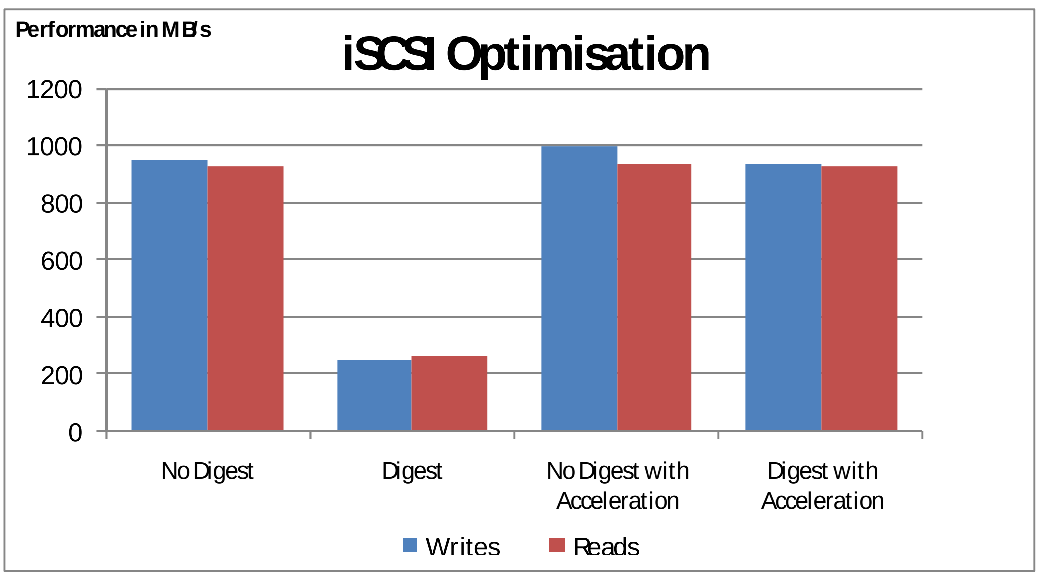

Figure 12 – iSCSI Optimisation

Our starting point was in a “vanilla” iSCSI environment with digest disabled, in order to get a baseline performance figure. In this configuration we achieved write speeds of 949.18 MB/s and read speeds of 925.63 MB/s. At the same time we recorded CPU utilisation as 64% and 63% for writes and reads respectively.

Turning the digest on and rerunning the tests resulted in a huge performance loss. Write speeds went down to 245.13 MB/s while reads max’d out at 260.23 MB/s. Despite the huge performance reduction, CPU utilisation was relatively high at 55% and 54% for writes and reads respectively.

Our final change was to reconfigure the setup with the Solarflare digest offload enabled and digest, first disabled, then enabled. With digest disabled we recorded a write speed of 996.39MB/s and a read speed of 934.41 MB/s – the best recorded figures but still roughly in line with our initial “vanilla” run, which is exactly what we would expect. However, on then enabling the digest, we recorded only marginally reduced performance figures of 933.47 MB/s for writes and 924.88 MB/s for reads.

In other words, the Solarflare acceleration preserved performance and was almost four times better than when running in a standard iSCSI environment with digest enabled. CPU utilisation also remained constant, 53% on writes and 43% on reads, so there was no other significant overhead when running the Solarflare acceleration.

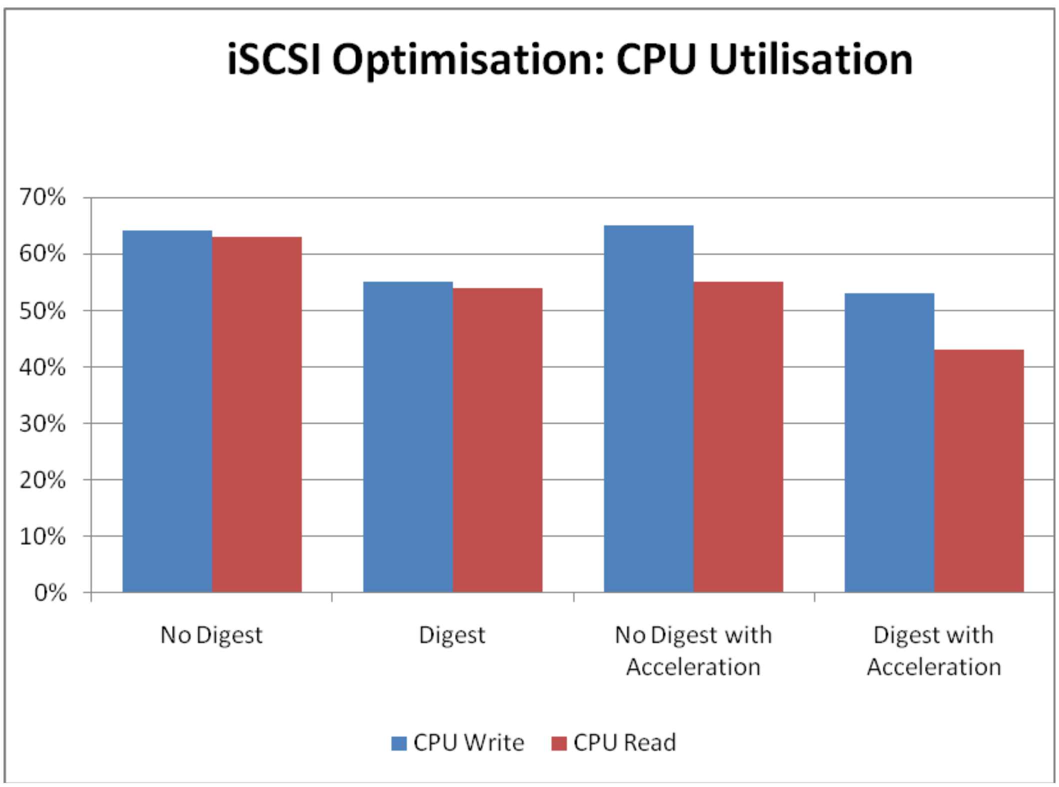

Figure 13 – iSCSI Optimisation: CPU Utilisation

SUMMARY & CONCLUSIONS

For this report we set out to prove or disprove Solarflare's claims that it's Solarstorm SFC4000E vNIC 10 Gigabit Ethernet controller is capable of optimising performance in many different environments.

Moreover, this optimisation was not deemed only to relate to improved throughput but also in terms of reduced CPU utilisation – a vital element of server performance optimisation.

We found that, not only is the controller capable of performing at line-rate in a standard environment, and with low CPU utilisation, but that it was also able to perform this feat in two virtualised environments – VMware and Xen. We also compared Solarflare Solarstorm performance with that of an Intel 82598 10 GbE based adapter and, on average, achieved almost double the performance with the Solarstorm controller, over that of the Intel equivalent.

In all these cases we were able to hit and sustain line-rate, while still witnessing significant idle time on the server CPUs, meaning we were achieving line-rate without getting anywhere near to maximizing the server performance capacity.

Testing in the iSCSI environment we found that the Solarstorm controller could enhance performance with the iSCSI data digest feature enabled, improving digest-enabled performance by over 300%, giving us near 10 Gigabit line-rate over iSCSI, with full data corruption protection.

Throughout the testing, minimal tuning was required on the Solarflare Solarstorm, so in a purely vanilla configuration – a real-world scenario – we proved that the Solarstorm controller can, not only achieve line-rate working in a variety of server OS environments, but is also able to optimise other environments, such as virtualised deployments or with iSCSI subsystems.

We therefore have no hesitation in recommending it to anyone looking to acquire 10 Gigabit controllers. The Solarflare product did everything we asked of it and much more.

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.

A wide array 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 10GE 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 10GE 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.

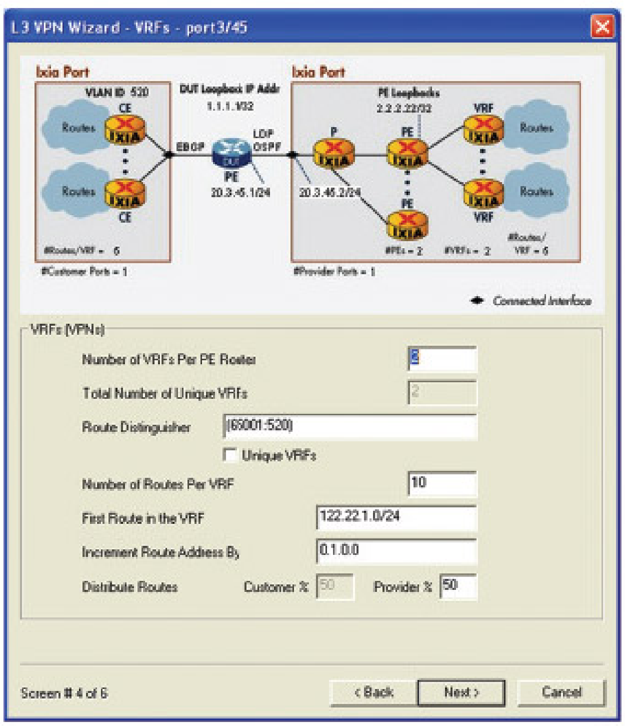

IxNetwork offers users the flexibility to customize the application to meet a wide range of requirements for testing complex network topologies consisting of thousands of network devices. Millions of routes and reachable hosts can be emulated within the topology.

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.

