A Spirent magazine

TESTING WORLD

March 2012

100G Ethernet: the Speed of Now

Is it true that you can never be too rich, too thin, or have too much bandwidth? The ever-increasing demand for content seems to validate that claim as far as bandwidth goes. There are now 7 billion potential subscribers in the world, tracking toward 8 billion, possibly as soon as the next decade.

The implications for bandwidth (100G Ethernet) are obvious. Not so obvious is how to assure that reliable and efficient implementations and deployments of these technologies occur on or ahead of schedule. This unprecedented scalability of subscribers and bandwidth, and the emergence of the mobile internet, demand an unprecedented level of performance from the test platforms that make sure the always-on, always-connected generation stays connected.

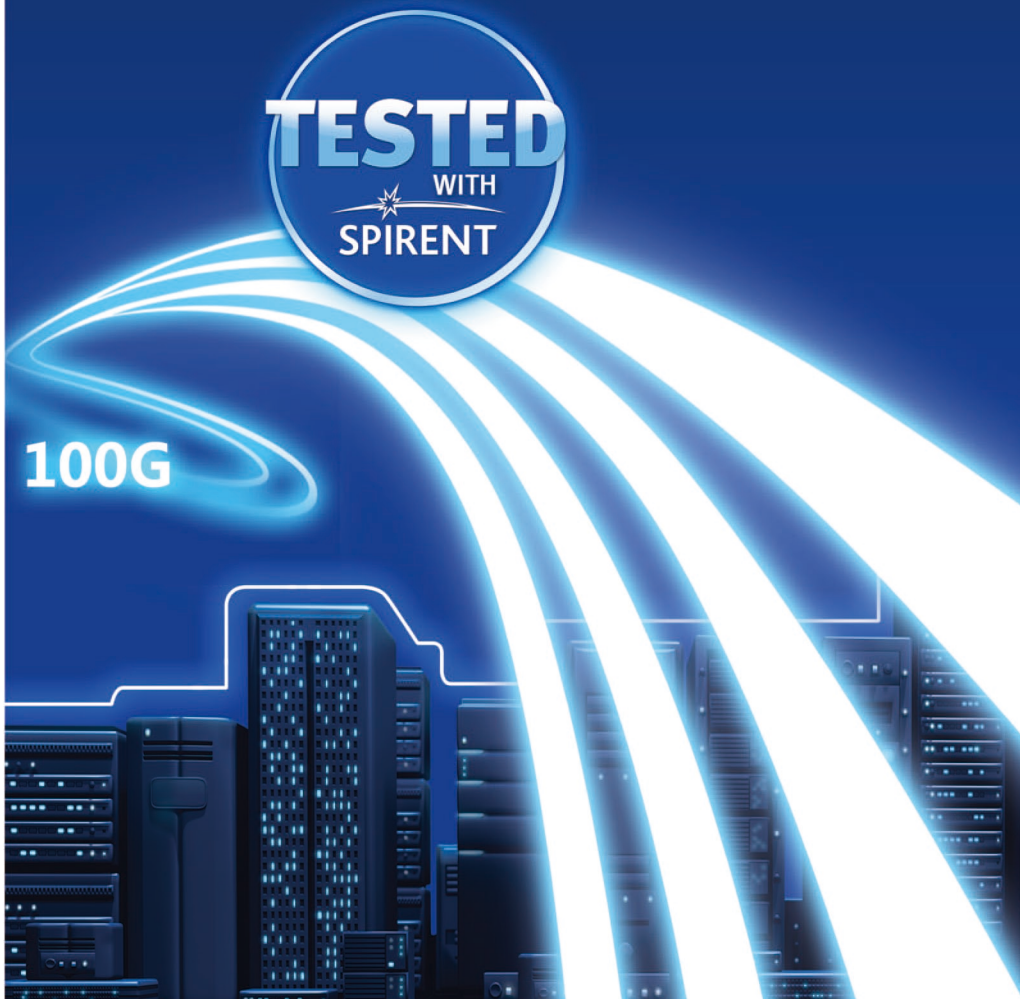

SPIRENT

Embrace Mobile Internet

As you've been aware, Mobile is not just driving significant changes occurring in Network Technologies but also lifestyle you're getting more addicted to...Mobile Internet is converging everything! We will "Embrace Mobile Internet" in 2012!

Never like today, Mobile Internet demands so relentlessly on Scale and Convergence. Technologies we talked yesterday, relatively independent, nowadays got to be treated as a whole. These include LTE & LTE-A, MIMO, LBS, Timing over Packet, MPLS-TP, Mobility, High Speed Ethernet, Data Center, Security, Cloud & Virtualization, etc., just to name a few. Impetus roots in exploding number of new Smartphone Subscribers (iPhone, Android, WM, RIM, etc.), Tablet Users (iPAD, Galaxy Tab, Kindle Fire, etc.) and various form of mobile devices – so-called "Post-PC" era defined by Steve Jobs. All stress added to the network and device to be deployed, we can feel the same in Test and Measurement world. Spirent keeps its leading position in testing the latest challenges by embracing Mobile Internet trend.

- Spirent commits in delivering best converged testing platform, on both hardware and software, to maximize our customer's investment, while still meeting increasing demands in testing high Scale and Convergence. HyperMetricsne Xt test modules, powered by Spirent Cloud Core™ technology, are great success in 2011. Spirent will release more hardware modules based on the same architecture in 2012, enhancing the integrated testing platform for L2-7 and Mobility.
- Spirent will keep consolidating its solution portfolio to test different segments in Mobile Internet, ranging from Mobile Handset, Mobile Radio Access Network, to Mobile Backhaul Network and Mobile Packet Core Network, and further to Core Routing and Data Center, Cloud and Virtualization, Mobile Security and Applications.
- Testing complexity today poses big challenge to testing productivity and quality. Spirent categorizes this market segment as Infrastructure Testing Optimization (ITO) and actively works with industry partners to create standard, friendly, interoperable testing environment. Spirent TestCenter and Spirent iTest are flagship products in addressing ITO needs. Spirent keeps making significant contributions in NTAF (Network Test Automation Forum) where Spirent holds a founder and VP role.
- Spirent provides one-stop shop solution for LTE terminal side, from conformance to performance, physical layer to application layer, single session to service interaction, R&D to carrier acceptance, and commits to bring the real end user experience to lab to speed up smartphone diffusing to mobile internet.
- Aligned with Mobile Internet strategy, Spirent offers best location test solution among the world to address the key factor in Mobile network---location. No matter if it's in house or on high-way, or in mountains or on the sea, with Spirent location test system, best LBS service could be achieved.

In summary, Spirent is well postured in embracing new Mobile Internet era. We dedicate ourselves to provide state-of-art testing solutions to our customers.

Content

Written by Patrick Barry

P07-13

100G Ethernet: the Speed of Now

Is it true that you can never be too rich, too thin, or have too much bandwidth? The ever-increasing demand for content seems to validate that claim as far as bandwidth goes. There are now 7 billion potential subscribers in the world, tracking toward 8 billion, possibly as soon as the next decade.

The implications for bandwidth (100G Ethernet) are obvious. Not so obvious is how to assure that reliable and efficient implementations and deployments of these technologies occur on or ahead of schedule. This unprecedented scalability of subscribers and bandwidth, and the emergence of the mobile internet, demand an unprecedented level of performance from the test platforms that make sure the always-on, always-connected generation stays connected.

Editor Email: janet.peng@spirent.com

Cloud Computing and Standardization: Technical Reports Published

ITU-T's Focus Group on Cloud Computing has completed its preliminary study into cloud computing's standardization ecosystem and has released its FG Cloud Technical Report (Parts 1 to 7). The reports signal the conclusion of the Focus Group's study period and its findings come to form input for the cloud computing work taking place across the ITU-T under the leadership of Study Group 13 (Future Networks), overseen by the Joint Coordination Activity on Cloud Computing.

Clear industry demand for the technology and the promise of new revenues to ICT players has led to great market optimism, with one forecast predicting that global cloud IP traffic will account for more than one-third of total data center traffic by 2015. [Cisco Global Cloud Index: Forecast and Methodology, 2010–2015 http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns1175/Cloud_Index_White_Paper.pdf].

International standards will be key to ensuring a competitive market for cloud service provision, one where services are affordable and the clouds offered by different vendors interoperate. FG Cloud's Technical Reports provide a detailed overview of cloud computing standardization and benefits from telecommunication and ICT perspectives. Two new cloud services categories are identified, communication and network as a service (CaaS and NaaS). In addition, cloud ecosystems and a first ICT cloud reference architecture are defined. These outputs will form a strong reference base for the ITU-T Study Groups engaged with cloud standards development.

IEEE Announces Five New Standards and a Standards-Development Project for the Smart Grid

IEEE recently announced five new standards, as well as a modified standards-development project, that are all intended to aid the efficient rollout of the smart grid worldwide. The new standards and standards projects recently approved by the IEEE Standards Association (IEEE-SA) Standards Board carve critical new dimensions into the IEEE portfolio of more than 100 active standards or standards in development relevant to the smart grid.

Smart-grid standards newly published by IEEE-SA include the following:

- IEEE C37.118.1™-2011 – Standard for Synchrophasor Measurements for Power Systems – is intended to define synchronized phasors and frequency measurements in substations, along with methods and requirements for verifying such measurements in power system analysis and operations under both static and dynamic conditions.
- IEEE C37.118.2™-2011 – Standard for Synchrophasor Data Transfer for Power Systems – is intended to specify a method (including messaging types, use, contents and data formats) for real-time communications among phasor measurement units (PMUs), phasor data concentrators (PDCs) and other power-system applications.
- IEEE C37.238™-2011 – Standard Profile for Use of IEEE Std. 1588 Precision Time Protocol in Power System Applications – is designed to provide precise time synchronization within and among substations across wide geographic areas via Ethernet communications networks. The standard is intended to extend proven techniques for precise time distribution to applications such as mission-critical power-system protection, control, automation and data communication.
- IEEE C37.232™-2011 – Standard for Common Format for Naming Time Sequence Data Files (COMNAME) – is designed to define the naming of time sequence data (TSD) files that originate from digital-protection and measurement devices. The standard procedure—gaining in popularity among major utilities, independent system operators and manufacturers and recommended for use by North American Electric Reliability Corporation (NERC) and the Northeast Power Coordinating Council (NPCC)—helps resolve problems associated with reporting, saving, exchanging, archiving and retrieving large numbers of files.
- IEEE 1020™-2011 – Guide for Control of Small (100 kVA to 5 MVA) Hydroelectric Power Plants – updates an existing IEEE standard to address significant technology changes impacting small hydro-plant control issues and monitoring requirements that have emerged since the guide's original publication.

ITU Establishes Focus Group on Machine-to-Machine Service Layer

The January meeting of the Telecommunication Standardization Advisory Group (TSAG) has established a new Focus Group on Machine-to-Machine (M2M) Service Layer.

M2M refers to the ability of a machine to sense and measure certain variables, and communicate this information to other machines in a network. Included under the larger umbrella of the "Internet of Things" (IoT), M2M technologies have applications in a number of industries – e-health, fleet management, sales and payment, security and surveillance, intelligent transport systems (ITS) etc.

The group will study and evaluate the M2M landscape and M2M work currently being undertaken by regional and national standards development organizations (SDOs), with a view to identifying a common set of requirements.

The Focus Group will initially focus on the APIs and protocols to support e-health applications and services, and develop technical reports in these areas. It is suggested that the Focus Group establish three sub-working groups: "M2M use cases and service models", "M2M service layer requirements" and "M2M APIs and protocols."

The group is expected to carry out the following specific tasks:

- Research, collect and analyze the vertical market M2M service layer needs, initially focusing on e-health.
- Identify a minimum common set of M2M service layer requirements and capabilities, initially focusing on e-health applications and services.
- Study APIs and relevant protocols that satisfy the above requirements and capabilities to support the communications between the M2M applications and the telecom networks.
- Develop technical reports to address the identified gaps and propose future standardization work for ITU-T developments on M2M.
- Support global harmonization and consolidation by inputting its final deliverables to the parent Study Group and other relevant Study Groups as appropriate.

IPTV Standards the Basis for Transcontinental IPTV Experiment

An international experiment deploying ITU-standardized IPTV technologies has taken place 6-8 February 2012. IPTV services were used to live-stream scenes from the Sapporo Snow Festival in Japan and to provide supporting Video-on-Demand (VoD) segments.

ITU Headquarters in Geneva received the stream from the head-end server in Japan, participating alongside organizations from Japan, Singapore and Thailand in what is the first transcontinental broadcast of a live event using IPTV technology standardized end-to-end by ITU. The connection uses native IPv6 from ITU Headquarters to Japan.

Proprietary IPTV services have hampered the growth of this exciting new market, and such experiments - together with ITU IPTV Interoperability events – are important steps towards broadening the IPTV market through globally-interoperable services. Standardized IPTV will lead to a whole new market for innovation, and ITU standards will ensure this market remains open, competitive and accessible to all.

First approved in 2009, Recommendation ITU-T H.762, a "Lightweight Interactive Multimedia Environment" (LIME) for IPTV services, is the standard with which Sapporo's live-stream IPTV application complies. Hokkaido Television Broadcasting (HTB) developed this application, and is one of many broadcasters, manufacturers and research institutes involved in the IPTV experiments. The experiments have been organized by Japan's National Institute of Information and Communications (NICT) and are being conducted over its IPv6 research network, Japan Gigabit Network-eXtreme (JGN-X).

Testing Data Center Fabrics for High Performance at Massive Scale

By Jurrie van den Breekel

How do you test a data center switching fabric that scales exponentially beyond legacy platforms?

The implications go beyond the QA department. Of course, system testing is required as part of the product lifecycle and a large-scale system requires large-scale testing. But such a system poses an equally daunting challenge to marketing and sales—demonstrating to customers that high performance and high scalability can be achieved in their environment.

Independent test lab Network Test recently used Spirent's data center testing solution to demonstrate that Juniper Networks' QFabric™ can deliver unprecedented performance at scale.

Legacy switch fabrics scale to hundreds of ports. The QFabric solution scales to thousands. A fully loaded QFabric system supports 6,144 ports configurable as a single device in an any-to-any port topology with low latency and the appearance of unlimited bandwidth.

In addition to high speed, high-port density, and the ability to measure latency with precision, a proper test of such a fabric requires a full-mesh network topology traffic pattern. It is the most demanding possible configuration both for the test platform and the device under test, but it more accurately reflects the ultimate limits of system performance, an important consideration in a highly-dynamic environment where the demands of the next killer app are not known.

The test bed for this industry-first test comprised 1,536 x 10 GbE ports connected to the same number of Spirent TestCenter test ports and 128 redundant 40 Gbps fabric uplinks, a test configuration four to five times larger than previous high-port density Ethernet switch tests. Juniper QFabric System components took two 42U racks. Spirent TestCenter chassis loaded with 32-port HyperMetrics DX modules filled the other two racks.

The test methodology included standards-based benchmarking tests for unicast throughput, unicast latency, and multicast performance. Because of the full-mesh topology, Spirent TestCenter generated and analyzed throughput and latency over 2.3 million streams in real time and over 15 Tbps of traffic at line rate. Network Test used Spirent TestCenter's analysis of Juniper Network's QFabric to validate that:

- Throughput for Layer 2 traffic was virtually identical in store-and-forward and cut-through modes, with rates approaching the maximum channel capacity for most frame sizes
- QFabric forwarding delay is low and consistent across all tests (less than 5 microseconds for all frames sizes up to 512 bytes) when offered loads below the throughput rate
- Average latency for Layer 3 traffic is 10 microseconds or less for most frame sizes tested
- Multicast throughput was close to line rate in all tests, regardless of frame length, with the system moving traffic at speeds of up to 15.3 terabits per second
- Multicast average latency was low and consistent in all tests, never exceeding 4 microseconds
- The QFabric System successfully interoperated with Cisco Nexus 7010 and Cisco Catalyst 6506-E switch/routers when using common data center protocols such as link aggregation, OSPF equal cost multipath, and BGP

For more information on the individual tests methodologies, configurations, and detailed test results, see the test report from Network Test.

100G Ethernet: the Speed of Now

Is it true that you can never be too rich, too thin, or have too much bandwidth? The ever-increasing demand for content seems to validate that claim as far as bandwidth goes. There are now 7 billion potential subscribers in the world, tracking toward 8 billion, possibly as soon as the next decade.

The implications for bandwidth (100G Ethernet) are obvious. Not so obvious is how to assure that reliable and efficient implementations and deployments of these technologies occur on or ahead of schedule. This unprecedented scalability of subscribers and bandwidth, and the emergence of the mobile internet, demand an unprecedented level of performance from the test platforms that make sure the always-on, always-connected generation stays connected.

Market Drivers: Truths and Consequences

Three Truths

As the always-on, always-connected lifestyle sweeps across the globe, three truths of networking have emerged.

There will always be:
1. More subscribers
2. More demand for content
3. More addressable devices per person

More Subscribers

According to Miniwatts Marketing Group, in March 2011 there were over two billion internet users, representing 30 percent penetration worldwide. While North America approaches saturation at 78 percent, other areas, such as Asia (23 percent) and Africa (11 percent), are huge growth markets.

More Demand for Content

The mobile internet is upon us, along with interactive mobile applications, hi-def video conferencing, high-speed file transfer, mobile clouds, and a host of others.
Real-time entertainment has moved to the forefront of online content. In general, for decades the demand for content has increased and is now concentrating on high-bandwidth, latency-sensitive traffic.

More Addressable Devices per Person

IP-addressable devices have expanded beyond traditional computing and network devices to include smartphones, mobile entertainment centers, vehicle systems, surveillance systems.

Three Consequences

The inevitability of growth in number of subscribers, number of devices per subscriber, and demand for content has three implications for those building networks.

Networks will continue to require:
1. More bandwidth.
2. More addresses.
3. Better Quality of Experience (QoE)

Ethernet dealt with bandwidth concerns by going from 1GE to 10GE to 40GE and now 100GE in a decade. In the process, the changes in the technology have gone far beyond bit rate to multi-lane paths and issues of lane swapping and skew.

IPv6 resolved the address issue. The depletion of IPv4 addresses in February 2011 highlighted the urgency of IPv6 deployment.

And because of increasing expectations of QoE from the expanding subscriber base, networks must deliver high-bandwidth, low-latency, real-time entertainment with unprecedented levels of quality.

A Quick Technology Tour

While it is true that the future is bigger (IPv6) and faster (100GE), the implications for network applications go beyond questions of address length and bit rate.

IPv6

The address space of IPv6 is large enough to give each person on the planet a personal address space larger than the entire IPv4 address space. How much larger? Over 11,000,000,000,000,000,000(1.13E+19) times larger.

IPv6 poses challenges to existing network solutions.

Switch Fabric and Queuing

With an address length of 128bits as opposed to 32 bits, IPv6 requires rewriting and integrating into hardware decades-old algorithms that are hardware-optimized for IPv4. Solutions that soft switch packets with the CPU will not scale, causing capacity and performance issues.

In addition to length, IPv6 address formats differ from IPv4 formats. Legacy algorithms that optimize performance based on a fixed length format must be supplemented with more flexible routines. IPv6 solutions must be evaluated through capacity testing to reveal how the switch fabric will scale with IPv6 traffic and with rate testing to determine how fast the system can add addresses to the Forwarding Information Base.

Frame Size and Frame Rate

Frame sizes for IPv6 are much more likely to reach 16K than in IPv4. This can affect switch performance, especially in dual-stack implementations, which must handle packets of both versions with the wide range of frame sizes found on the internet and the sudden changes in frame rate that happen during microbursts.

Peaceful Coexistence

The transition from pure IPv4 networks to pure IPv6 networks will be a long journey. During the interim, there are technologies in place to allow both to coexist on the same network. These solutions are workarounds that require translation or tunneling, which introduce extra processing steps in the path and the latency that goes along with it.
For some end users this delay is just an inconvenience. For others, such as financial services, delay is potentially disastrous for the top and bottom lines. For service providers, low latency is a competitive differentiator.

IPv6 Deployment Challenges

IPv6 is substantially different from IPv4. CPUs, ASICs, queues, buffers, RIBs, microcode, drivers, and OS Code are at high risk for:

- Unexpected data plane performance / forwarding optimization
- Inefficient queuing algorithms and un-optimized queue hash management
- Decreased stability and reliability and increased security breaches
- Loss of quality of service (QoS),quality of experience (QoE), and the ability to sell service level agreements (SLAs)
- Poor performance, compliance, and scale of control plane protocols

100GE

It's More Than Just Moving Bits Faster

While it's true that 100G Ethernet is still Ethernet, just faster, there's a lot packed into that phrase "just faster." In some ways it's more of the same, in others it changes everything.

At the upper layers, that means each component in a device or step in a process must accomplish the same thing it currently does in one tenth of the time. Consider a router, for example, which strips lower-layer information from an incoming packet, queues it, performs a route lookup, and sends it to the proper outbound queue to be packetized, while performing filtering, SLA monitoring and policing, and CoS/QoS prioritization. In addition, a router also exchanges per-VPN MPLS label information, builds multicast routing trees, performs routing table updates for multiple protocols, maintains statistics and logs (performance, alarm, event and failure) and performs firewall and security functions, such as key exchanges, attack detection and prevention.

A router with 100G interfaces must do all this at ten times 10G speeds without dropping packets, introducing excessive jitter, compromising VPN boundaries, or reordering packets, which is especially disruptive for storage and real-time entertainment.

Clock Precision

Imprecise clocking between systems at 100G can increase latency and packet loss. At 10G, especially at high densities, the specification allows for a little variance for clocks. As you aggregate traffic into 10G ports, small differences between clocks can cause high latency and packet loss. At 100 GB, the tolerances are even smaller.

100GE Deployment Challenges

Key test metrics for IEEE 802.3ba 100GE are:

- **Layer 1 skew performance.** Lane skew was a contributor in older Ethernet roll out issues, being able to add skew and to measure effects are a critical part of the physical layer setup.
- **Layer 1 lane swapping.** Measuring the ability of the DUT to manage the virtual to physical translation is critical because lane swapping errors can lead to interface link-down problems that are difficult to debug in deployed systems.
- **Layer 1 per-lane unique BERT (PRBS).** The ability to test the physical pathways using unique PRBS BERT patterns reveals physical lane stability and crosstalk issues.
- **Latency and jitter.** Testing 100 GB to 100 GB and 100 GB to multiple 10 GB ports is a critical test of backwards compatibility. The ability to measure jitter to 2.5 ns in the core is essential.
- **RFC2544 100GE throughput.** Measuring the forwarding rate at 100G is critical to prove forwarding performance.

- **Correct sequencing detection.** Validating the buffering storage-and-reassembly algorithm while under load and across different Ethernet technologies is essential. Reported metrics include lost, duplicate, late, and reordered packets.

Test Realism: Emulation versus Simulation

The ability to accurately evaluate a solution against the PASS parameters is based on the simple but essential element of test realism. Realistic testing means re-creating the environment that the application lives in, from the provider to the customer, in all its dynamic and daunting complexity.

Test realism comes down to the difference between simulation and emulation. Simulation attempts to predict the behavior of a system by creating an approximate mathematical model based on a set of assumptions. Emulation replaces a part of the system and performs in exactly the same way as the element it replaces. It adjusts dynamically to a changing environment and responds to actual stimuli from the system it interacts with.

Simulation is appropriate for analyzing a system to infer predictions about how it might behave, but you need emulation to see what will actually happen in a given situation. You need emulation to know how your solution will respond under real-world conditions.

There are three essential elements of test realism.

1) Real user behavior

Users are as unique as their fingerprints. They vary significantly in how long and in what manner they navigate through an application and how they respond to sluggish performance, picture and voice dropouts, dropped calls, and other problems. They violate usage and security policies in different ways. They find unique ways to break your system.

Realistic testing means the flexibility and sophistication to emulate a wide range of user behavior, both benign and malicious.

Hitting a device with a mix of traffic types (say, mixing HTTP requests/responses with IPTV channel changes and P2P file sharing or gaming traffic) isn't emulating user behavior. It's a simulation, just hitting the system with a mix of traffic, static packets that don't respond statefully to the incoming messages. Emulating real user behavior means supporting stateful traffic that emulates how a user operates, including think time, click-through, abandon, channel surfing, etc. It means good users and malicious users simultaneously attempting to achieve their good and bad goals.

User-centric traffic reveals the performance of the device in a real environment. Queues, buffering, and other mechanisms behave differently depending on the order and the nature of the transactions.

2) Real converged traffic

The consolidation of networks onto Carrier Ethernet and MPLS enables mobile and fixed-line voice and data, residential video, and MPLS-based VPNs. The different types of traffic carried on the converged network have different characteristics and requirements, but they all travel on the same path, dependent on CoS and differentiated traffic rules to keep everyone happy.

Realistic testing means the power to create line-rate, fully-emulated, stateful traffic across hundreds of ports.

Real converged traffic not only means a realistic mix of traffic types, but also realistic traffic encapsulation. For example, if the deployed system tunnels user PPP sessions over MPLS, then testing PPP setup-teardown rate and throughput performance without MPLS is not real converged traffic. It's a dangerous shortcut that will mask problems your users will discover after deployment. Real converged traffic means emulating the actual deployed topology, regardless of how complex or simple, including all encapsulations.

3) Real network conditions

The network creates time-varying conditions that are linked to a complex set of conditions, influenced by routing table updates, signaling protocols, queuing algorithms, buffering, traffic management and policing policies, malicious attacks, EMI and other environmental factors.

Realistic testing means the power and complexity to create the dynamic, time-varying conditions found on deployed, production networks.

Real network conditions can't be emulated through static rates of delay/loss or distribution-based mathematical models of impairment. Real networks don't introduce impairments at fixed rates or follow neat curves. They behave in seemingly non-deterministic ways due to the number of factors affecting them. Testing under real network conditions means emulating this complexity to discover issues before the real network finds them for you.

Beyond QoS: Testing and Delivering QoE

Realistic testing goes beyond simple, deterministic QoS test methodologies. While real-time services are built on a foundation of a QoS-enabled network, the real deliverable is not a set of priority rules and metrics. It's a high quality of experience for the subscriber.

Accelerated Testing: Product Lifecycles at the Speed of Now

Another aspect of the faster, bigger world of the mobile internet is the acceleration of the technology life cycle. Earlier iterations of Ethernet may have taken a decade to transition through cycle – introduction, growth, maturity, and decline – but successive leaps in bandwidth have cycled through increasingly shorter timeframes.

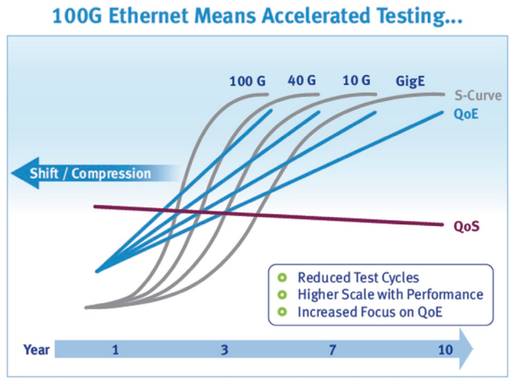

For each successive generation, the S-curve from introduction to maturity gets steeper as schedule expectations are collapsed.

Testing 100GE at the Speed of Now

The 100GE delivery test requirements are different from 1GE and 10GE because of the strong emphasis placed on QoS/CoS, realism, stacked protocols, and multiplay services over 100GE.

Validate Data Transmission in Loopback Mode

IEEE 802.3ba.

This test determines whether the system under test (SUT) works error free in loopback. This test places the DUT in loopback mode and verifies that traffic is received correctly without errors. In this mode, each lane maps traffic to its respective virtual lane. Then, virtual lanes are mapped to the physical layer. On the receive side, the procedure is reversed.

This test provides long duration verification and determines basic pathway and switching robustness. It creates a baseline of expected results before involving the 100 GB module, provides long duration verification of the SUT, and determines that the basic pathway and switching functions work correctly on the SUT.

Validate Performance of the Optics

IEEE 802.3ba.

By measuring potential errors over a long duration, this test determines whether the optics function correctly. Numerous traffic patterns are generated across the optics. The optics are then measured for traffic rate and error conditions. This test characterizes the optics, determines whether the optics can pass traffic successfully, determines whether the optics can run error free, provides long duration verification of the optics, and determines the skew introduced while running traffic.

Validate Short and Long Term Performance of the PMA Layer

802.3ba

This test case determines whether the SUT generates traffic without errors at the PMA layer. The primary function of the PMA is to multiplex M input lanes to N output lanes where needed. The PMA also performs clock recovery, clock conversion, test pattern generation and detection

and loopbacks where applicable. This test case validates multiplexing validity and clocking, recovery mechanisms, and determines the interoperability between the SUT and the tester at the PMA layer.

Validate PCS Marker Alignment

IEEE 802.3ba

This test case verifies that the virtual lane markers are properly implemented in the 100G SUT design by generating large amounts of traffic and determining whether physical-to-virtual translation occurs correctly. Validating PCS lane markers enhances interoperability and reduces the possibility of random errors in the field.

Assess Maximum Lane Skew Limits

IEEE 802.3ba

This test case measures the range and correctness of individual lane skew inside the DUT HSE implementation. Lane skew is a difference in timing delay across lanes. Skew must be compensated for on the receive time before the data stream may be reconstituted by the RX port. Maximum skew compensation is an important metric in interoperability.

Standards-Based Benchmarks

RFC 2544: Benchmarking Methodology for Network Interconnect Devices

Industry benchmarks provide a standard method to measure and evaluate a solution through a uniform procedure for the generation of traffic to and from the system under test supported by a normalized procedure for analysis and reporting. The goal of the benchmark is to generate metrics in a reproducible and unbiased fashion for comparability.

Testing 100G IPv6 at the Speed of Now

IPv6 and 100GE are pushing networking to new levels of performance. The massive scalability required has implications not only for the devices used to build the network, but the systems that test the devices and networks.

There are three key capabilities required of a test system for 100GE and IPv6.

1) Real-Time IPv4/v6 Routing at 100GB with Real-World Scenarios

Dual-stack implementation is the foundation of the foreseeable future, along with the processing overhead and potential performance impact that goes with it.

Currently the internet routes 380,000 IPv4 prefixes aggregated to 220,000+ table entries, compared to 7,500+ IPv6 prefixes aggregated to 6,400+ table entries. However, for IPv6 that represents an over 100% increase year-on-year, compared to the 12% increase for IPv4.

Your test system must have the power and sophistication to emulate a dual-stack, MPLS-enabled, real-world routing environment at line rate.

2) L2-L7 True Multi-Protocol Testing at 100GB

To deliver the QoE the customer expects, QoS must function without errors under line-rate stateful traffic. That means your test system must deliver the realism of true emulation, not just simulation, at 100GB to go beyond QoS to deliver QoE.

3) High Port-Density Testing at 100GB

Depending on the platform, a router chassis can have up to 32 x 100GE ports. While initial deployments may not require 3 terabits of throughput per chassis, current growth rates indicate that the requirement is not that far in the future. Your test system needs the power to match the scale of the production network.

Testing the Tester

Testing goes beyond the ability to throw traffic at a box or system. Testing the unprecedented scale of speed and users involved with 100GE and IPv6 with realism requires a test platform that can not only scale to stress the system, but can insulate your workflow from the negative effects of high-scale testing and results.

A test platform is a competitive differentiator that can mean the difference between accelerated time to market or costly schedule delays. Don't hesitate to ask the tough questions about your existing system or a system under consideration. Make sure it supports:

- **Detailed statistics with targeted analysis.** The power of testing is the ability to gather comprehensive results and interpret them. Before selecting a system, look at the results it produces from real tests and ask pointed questions about using them to quickly identify root problems. A test system that streamlines troubleshooting through powerful results can significantly accelerate design and development. Good results can help you stay on or ahead of schedule.
- **Scalability.** Numbers matter. You need a good understanding of your target environment, including plans for growth. Armed with this information, verify that the test platform can scale to your requirements in terms of numbers of users, clients, subscribers, application services, VMs, and ports.
- **Ease of use.** There should be one GUI that controls all test scenarios. Your day is complicated enough without having to learn multiple interfaces for different elements of a single test system.
- **Automation.** The ability to automate individual tests and suites of tests is a fundamental element in test lab productivity. In these days of reduced resources and increased expectations, simple, manageable automation is not a luxury, it's essential.

Conclusion: Testing at the Speed of Now

The combination of more subscribers than ever demanding more high-bandwidth, delay-sensitive, real-time content than ever requires equipment manufacturers and service providers to deliver more powerful solutions and higher QoE than ever.

The scale of IPv6 and 100GE brings new challenges to the market and specifically to the test lab. Test realism, scalability (whether in ports, subscribers, sessions, VPNs, routes, tunnels, or VLANs), automation and intelligent results that help you quickly pinpoint root causes are more important than ever.

Make sure your test platform can deliver what you need to stay ahead of schedule, ahead of demand, and ahead of the competition.

For more information see www.spirent.com/ethernet.

Spirent TestCenter 3.90

Spirent TestCenter 3.90 release adds many new features and functions including significant improvements in Spirent TestCenter time-to-test.

Spirent TestCenter 3.90 includes improvements in usability, enhancements to the Dynamic Protocol Generator functionality, and new protocol support in the Mobile Backhaul, Routing, Cloud and Data Center segments. This release also introduces a new 3U chassis, as well as HyperMetricsneXtfX and mX test modules designed for 40G and 100G testing.

Hardware

- SPT-3U Chassis
- HyperMetricsneXt mX-100G-F1/2 Modules (40/100G)
- HyperMetricsneXt mX-100GO-F1/2 Modules (100G only)
- HyperMetricsneXt mX-40G-F1/2 Modules (40G only)
- HyperMetricsneXt fX-100G-F1/2 Modules (40/100G)
- HyperMetricsneXt fX-100GO-F1/2 Modules (100G only)
- HyperMetricsneXt fX-40G-F1/2 Modules (40G only)

Software

Expanded application support for DX-10G-S32
Usability Improvements

- Summarized port results counters
- Validate Dynamic Result View counts in Command Sequencer
- Application now retains user settings on upgrades
- Filterable Device Grid
- Select Statistics to Save
- Histogram Enhancements

Cloud and Data Center

- New wizard for large scale control plane configurations for Shortest Path Bridging
- Enhanced FCoE Wizard, setup enhancements to STC Virtual
- New protocol support for LISP and VEPA (IEEE 802.1Qbg)
- Native protocol support for CIFS has also been added in 3.90

Mobile Backhaul

- MPLS-TP Y.1731 OAM Performance Monitoring (delay measurement and loss measurement) support
- GMPLS support
- Improvements to the RSVP point-to-multipoint and FRR (fast re-route) wizards
- New templates for MPLS-TP BFD for easy negative testing of RFC 6428
- Conformance Test Suites for MPLS-TP OAM and 1588

PPPoE

- Independent control of PPPoE ICMPv4 and ICMPv6
- Result states for both are displayed

Cisco TIP Support

ABR (Apple HLS) Support
DPG Designer
Automation

- Native Python API support
- LINGO support

Avalanche 3.90

Avalanche 3.90 includes the Turbo HTTP performance feature; enhancements to SAPEE, the Action List Generator, and EZ Test; plus test start time improvements

Software

- Support for simultaneous 1G and 10G operation on Avalanche 3100
- Turbo HTTP
- SAPEE Import "Anything"
- Port Reservation and Start Time Improvements
- Action List Generator
- EZ Test and SAPEE
- SNMP Real Time Statistics

Spirent TestCenter 3U Chassis

Spirent TestCenter 3U chassis (SPT-3U) lays the foundation for testing converged multi-service devices using the HyperMetricsneXt™ family of test modules with Cloud Core™ processing.

The SPT-3U chassis is the next generation in mobile chassis architecture, designed to handle tomorrow's most complex multiprotocol scale and cloud virtualization testing needs. The SPT-3U chassis supports all Spirent TestCenter HyperMetrics modules up to 100 Gbps. With ultra high-scale, green low-power per port, intelligent fan control, and full support for IPv4 and IPv6 control, the SPT-3U chassis provides the foundation for the highest performance mobile test system available today.

Because there is no central operating system or slow software proxy server bottleneck, the SPT-3U switched-Ethernet backplane design optimizes efficiency and performance when configuring tests and gathering results across multiple users or automation processes. This architecture approach yields the fastest time to test and eliminates concern with managing a virus prone operating system in the chassis. The SPT-3U innovative Android™ based touch screen allows easy chassis status monitoring and configuration.

Applications

- Mobile chassis solution for Spirent HyperMetrics neXt test modules
- Fully backward compatible with current and older generation cards
- Portable high-density 10G and 40/100G Ethernet
- Stratum-3 clock solution for ultra-precise PTP Grandmaster Clock emulation (Upgrade Stratum-1 with GPS Option)
- Full support for Routing Topology emulation and full Layer 2-7 Traffic
- Multi-user capability

Spirent TestCenter 3U Chassis

HyperMetrics fX 40/100G Ethernet Test Module

Spirent TestCenter HyperMetrics fX 40/100G Ethernet test module with Cloud Core™ processing enables performance and scalability testing of high-speed Ethernet networks. Targeting testing of multi-terabit routers and high-speed Ethernet cloud infrastructure, the HyperMetrics fX ensures dataplane QoS performance over realistic routing and cloud infrastructure topologies. With four 40G ports and two 100G ports per module, the HyperMetrics fX 40/100G delivers the most comprehensive set of features and performance at a very competitive price.

The Spirent TestCenter HyperMetrics fX 40/100G Ethernet modules are available in 2-port 40G/1-port 100G, 4-port 40G/2-port 100G variants. Also available are versions that support only 100G operation and only 40G operation for those applications that do not require dual speed capability. With the combination of Cloud Core processing and the deep real-time analysis that Spirent TestCenter is known for, these modules deliver on realistic testing of complex multi-protocol topologies.

Applications

- High Scale Terabit Routers—Test 40G and 100G Ethernet core routers with multi-protocol topologies and line rate traffic
- Data Center Fabrics—Validate the forwarding performance and functional capabilities of ultra high-scale, next-generation multi-terabit cloud data center fabrics

- Enterprise Switches—Validate forwarding performance and functional capabilities of large, next-generation enterprise campus and data center switches with ultra low-latency, high port density and FCoE capabilities

HyperMetrics mX 40/100G Ethernet Modules

Spirent TestCenter HyperMetrics mX 40/100G Ethernet modules are available in 2-port 40G/1-port 100G, 4-port 40G/2-port 100G variants. Also available are versions that support only 100G operation and only 40G operation for those applications that do not require dual speed capability. With the combination of Cloud Core processing and the deep real-time analysis that Spirent TestCenter is known for, these modules deliver enhanced realism with scale and performance. By combining the HyperMetrics mX 40/100G module with the mX or dX 10G, Spirent TestCenter is ideally suited for testing multi-terabit cloud and carrier solutions.

Applications

- High Scale Terabit Routers—Test 40G and 100G Ethernet core routers with high scale, multi-protocol topologies
- High Capacity Multiservice Routers—Validate IP throughput and Any G mobility with millions of subscribers and per-port line-rate data with minimum-sized packets and detailed per- mobile statistics

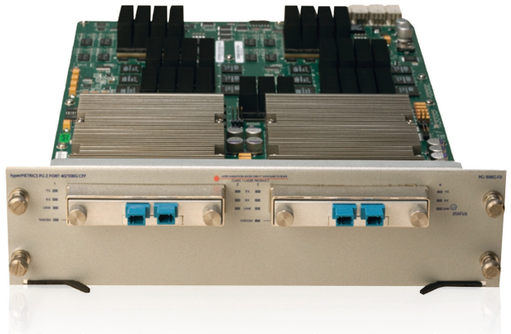

- Data Center Fabrics—Validate the forwarding performance and functional capabilities of ultra-high scale, next-generation multi-terabit cloud data center fabrics

CS8 Wireless Device Design Tester

For years, network operators and test labs have enjoyed the benefits of Spirent expertise—test equipment that provides both efficient ease of use and the most realistic replication of the live cellular network. Now these same Spirent innovations are available in the mobile device design lab for design verification testing (DVT). CS8 is the perfect testing solution for all stages of the device design cycle: RF & baseband development, radio protocol, and system or platform-level development testing. It is also the perfect testing platform for developers of LTE, UMTS, CDMA and multi-mode UEs and MS's. A single hardware platform can be used in every stage of mobile device development.

Spirent's CS8 is a state-of-the-art platform that provides a wide range of multi-technology testing scenarios for the device development engineer. Multiple test modes tailor the interfaces and functionality to most efficiently address the task at hand, making this single platform the ideal solution at all stages of the design cycle.

CS8 introduces the world's most advanced Evolved Packet Core (EPC) emulation. This state-machine-based EPC brings realistic IPv4/v6 network-side message responses and timing to the lab without requiring a single line of script. CS8 boasts multi-cell, multi-RAT capability, integrated MIMO

support and purpose-specific interfaces designed to ensure efficiency and optimal testing workflow.

Key Features

Powerful purpose-specific applications support device testing at every stage of the design cycle, supporting below testing mode:

- RF/baseband development mode
- Radio protocol development mode
- Platform validation mode
  - Protocol testing
  - Feature Testing
- Inter-RAT testing
- System test mode

Key facts – Platform

- LTE:
  - Multi-cell support
  - Support for all 3GPP LTE bands and bandwidths
  - UL signal capture and in-band measurement capabilities
  - Integrated SISO, SIMO, MIMO 2x2, MIMO 4x2
  - Integrated fading
  - L1 - L3 protocol stacks comply with 3GPP Release 9
- UMTS:
  - Multi-cell support
  - Support for all 3GPP UMTS bands and bandwidths
  - UE Categories 1-14 (HSDPA) and 1-5 (HSUPA)
  - Uplink Power Control algorithms 1 and 2

CS8 Testing Modes

CS8 is a single platform designed to address all stages of the mobile device design cycle. Separate testing modes combine to customize CS8 functionality to address each specific stage of device development. Once you realize how different development testing tasks can be performed by a single test stand, CS8 may quickly become the most valuable part of your development lab.

For RF and baseband developers the CS8 provides a touchscreen interface to a complete set of TX and RX measurements involving all protocol layers. For radio protocol development, the system provides a TTCN-3 programming environment to develop customized protocol stack testing.

For chipset-platform validation, the CS8 emulates multiple technology networks with a fully developed real-time IPv6 Evolved Packet Core (EPC), providing a multi-RAT system with realistic connection anchoring points. This is the same EPC used by network equipment manufacturers to ensure proper operation of network products.

The CS8 further accelerates LTE device development and debugging with the CS8 Software Development Kit (SDK). The SDK provides deep control of LTE Network Emulation functions and events. With the hundreds of functions and commands available,users can create and run custom tests to meet R&D needs and internal test plans.

Finally, the CS8 Development Library UI offers a clean and intuitive interface to help you develop custom test cases quickly and efficiently.

Spirent Test Expert 1.30

Spirent Test Expert 1.30 is the third major version of Spirent Test Expert. Major highlights of the release include:

- Highly anticipated test cases for EtherSAM, or Y.1564, which addresses the requirements of major service providers
- Test cases written for Avalanche to measure bandwidth performance now take advantage of multiple cores, meaning that much higher bandwidth of application traffic can be achieved.

Additional test cases are also provided for applications & security, and data center. Several new usability enhancements have also been made.

New Test Cases

- EtherSAM (Y.1564)
  - EtherSAM Service Performance Test with Turbo QoS
  - EtherSAM Service Configuration Ramp Test
- Data Center
  - Load Validation Test
  - Maximum Forwarding Rate Test
  - DCBX Feature Negotiation

- Multicast
  - RFC 3918 Multicast Join/Leave Latency Test
- Applications and Security
  - Applications
- HTTP Adaptive Streaming Video Bitrate Validation under Scale
  - RFC 3511
- HTTP Transfer Rate
- Maximum HTTP Transfer Rate

New Infrastructure Features

- Usability enhancements:
  - Copper ports, when available, are now the default
  - STE automatically finds the installation of Spirent TestCenter and Avalanche when installed on the same machine
  - Upon installation, STE prompts the user to set the appropriate paths within the preferences, with an option to select default paths
  - STE provides an option to launch or not launch after installation
  - A basic progress indicator is now provided
- Licensing: Spirent Test Expert now supports the ability for test cases to be licensed. The licensing mechanism is not applied to any test cases in STE 1.30.
- Authoring Enhancement: Authors creating test cases using the custom tool to export test cases to Spirent Test Expert may now add diagrams to their test cases

Features and Benefits

- Features
  - Built-in library of test cases
  - Framework for developing your own test cases
- Benefits
  - Leverage existing expertise to accelerate development of test libraries
  - Speedy adoption of technology
  - Customize to meet your specific testing needs

Spirent TestCenter C1

Spirent TestCenter C1 offers the power of Spirent's award-winning Layer 2 through Layer 7 router, switch and application optimization and security device test solutions in a portable form factor. With support for two line-rate 10 Gig test ports and the complete suite of Spirent test solutions, the C1 offers the power of a professional test tool used by the world's top network equipment manufacturers, service providers and Fortune 100 at an affordable, entry level price point. Users looking for a professional, converged test tool instead of relying on multiple open source utilities or ad-hoc testing methods can now improve their product realization efforts by moving up to Spirent TestCenter C1.

Solution Overview

Test a network with a solution that reproduces anything less than a realistic environment and you are risking slower product development and deployment and lower network service quality. Spirent TestCenter minimizes that risk by mirroring actual network scenarios and traffic patterns so networks, services and individual network elements can be validated quickly.

The C1 lowers the barrier to entry-enabling companies of all sizes to realize the higher ROI. That Spirent's industry-leading test capabilities offers. The Spirent TestCenter C1 supports Spirent TestCenter, Avalanche, and Spirent Test Expert software packages.

Whether you are an Enterprise or Service Provider network engineer, manufacturing production line technician, product developer, or systems engineer, Spirent will empower you to better manage your solutions and deliver on the promise of next-generation services.

Applications

The C1 is an ideal fit for:

- Network element engineering development, design, and test groups requiring physical access to a test port on the workbench
- Network equipment manufacturers doing burn-in and production line tests requiring low port count
- Technical and field marketing groups needing a low density portable solution
- Enterprise network engineers and technicians performing pre-deployment testing and service rollout testing
- University Computer Science departments and technical training organizations

Features & Benefits

Spirent sets the standard for testing the network environment with unparalleled realism and test packages for:

- Mobile and Fixed Infrastructures
- Cloud Infrastructures
- Cloud Applications
- Enterprise Campus and WAN Networks
- Data Centers

Realism

- Realistic Layer 4-7 user emulation to test applications and application infrastructure
- Realistic Layer 2-3 traffic to test Quality of Service (QoS) mechanisms

Productivity

- Intelligent results allow users to quickly confirm positive results and identify problematic areas
- Real-time traffic and protocol controls enable the tester to validate and troubleshoot problems by altering the test while it is running
- Real-time results views allow the user to see how the network responds to changes in specific test conditions without having to stop the test and save the results
- Built-in wizards and automated test scenarios reduce test setup and execution times
- Easy Automation for novice programmers

Spotlight on Data Center Testing from Professional Services

By Patrick Barry

Facing the Challenges of Data Center Deployments

Costly business disruption and remedial work can be attributed to performance problems in data centers, operational network systems and software applications. These issues often go undetected prior to launch due to the difficulty of conducting realistic performance testing. Quantifiable consequences of undetected performance problems include: revenue loss, customer base erosion, and brand damage. All of these issues call into question a business's long-term viability.

Since the advent of network testing, Spirent test engineers have helped hundreds of clients save thousands of testing hours, which resulted in bringing products and services to market faster and with higher quality. Spirent's industry-certified test experts offering packaged test plans, and employing a world-class service delivery process, have delivered reliable test results on time, on budget, worldwide. Spirent Professional Services has developed a range of test services, with special expertise in data center testing, which have been designed to detect and mitigate the full range of deployment challenges before customers launch new infrastructure and applications.

Using Spirent's Europe region as a sample for an overview of Professional Services' testing success in the data center space, three brief customer case studies of data center testing challenges, and the solutions provided by Spirent, are presented below.

Enterprise — Pharmaceutical

Challenge: The system integrator for this large French multi-national enterprise was responsible for building two new data centers and connecting them to the customer's existing center. The compatibility of the new devices with the original data center needed to be assessed before cut over. The customer needed to select the devices that delivered the right performance and high availability. Solutions from each vendor needed to be verified with the live data center for performance, QoS and failover scenarios.

Solution: Employing Spirent TestCenter, Spirent Professional Services engineers implemented a rigorous test plan for Throughput, Latency and Failover Time in various scenarios.

In addition, they tested the impact of network congestion on critical traffic, especially voice.

Benefit and Outcome: Spirent testing revealed firmware and configuration defects in all vendor solutions, especially when under high load. Test reports that included vendor performance data, provided both objective and critical comparisons of the different vendor solutions. With this report, the system integrator was able to recommend the solutions that best fulfilled the customer's requirements for the data center deployment. After selection, the customer requested two phases of re-testing to assure that all defects had been addressed. As a result, the systems integrator was fully confident that the network architecture would work as designed. When the customer went live they had no critical loss of service and are very confident that their data center solution meets their future requirements.

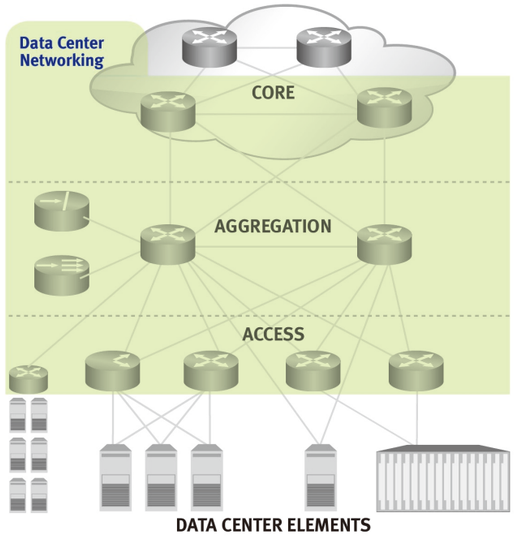

Enterprise — Construction Group

Challenge: With 70,000 employees, the French multi-national company's new data center had critical business applications, including Microsoft SharePoint, which had to support 32,000 local and remote users. Application performance and availability in the new data center infrastructure needed to be verified for local and remote users in reading and writing modes. The new data center needed to be launched without interruption to workforce productivity or the company's tight construction schedules, which were intensely budget-sensitive.

Solution: Employing Spirent Professional Services, remote sites were simulated and 20 test scenarios were executed to determine the maximum load supported by the system. This test plan included integrating existing and new user behavior scenarios with the applications as defined by the customer.

Benefit and Outcome: Spirent testing validated the data center infrastructure design as being in accordance with the exacting customer requirements. As a result, no additional upgrades were required and the customer launched their data center on schedule and on budget, maintaining uninterrupted productivity of their workforce. As a result of the confidence inspired by working with Spirent Professional Services, the customer expanded their testing strategy to encompass web portal and firewall assessments, to assure that these IT initiatives were also as ready for launch as was their data center initiative, ensuring budgets, operations and schedules were not adversely affected.

Enterprise — Financial Institution

Challenge: A major bank in the United Kingdom, in conjunction with a key global network equipment manufacturer, planned to implement a new data center network infrastructure. Network performance, availability and functionality of other network features needed to be assured before launch. In addition, issues had been observed with ARP performance and traffic distribution under the QoS configuration in the financial institution's live network. How network failure affected live applications running on the network needed to be identified before launch.

Solution: Spirent developed a test plan using Spirent best practices and test methodology and employing Spirent TestCenter. A series of core network tests were executed. These included testing the server farm (both stand-alone and virtualized), ARP performance, load balancing module performance and failover tests.

Benefit and Outcome: Spirent testing identified configuration and firmware issues in the pre-production network. It also revealed the strengths and weaknesses of the various configurations tested. This evaluation allowed an objective decision to be made between the virtualized or stand-alone servers. The availability of servers with various hardware and operating systems combinations was successfully evaluated

facilitating successful deployment of the data center and meeting the launch deadline.

Observations on Trends with Virtual Data Centers

In the boom years companies tended to add new servers each time they bought new business applications, rather than risk compromising the performance of the already installed system. As a result of this proliferation, most servers in a typical data center run at just 5-10% of capacity. Given the power of today's servers, this approach doesn't make sense. Virtualization makes it possible to unite that processing power into one massive server, which then gets broken down by virtualization software into many smaller virtual servers working at a healthier 60-70% capacity, and the remaining processing power can be switched off and power saved. Even though the virtual servers run on shared hardware, virtualization software preserves their independence so they retain all the safety features of running applications on separate servers.

Staffing costs are saved by concentrating a highly skilled team into one central data center instead of scattering them across the enterprise. This trend toward data center virtualization is an important solution for the numerous enterprises that are running out of infrastructure as this outdated and inefficient model of limited processing capacity utilization remains operative in their existing networks.

This is especially true for banks and other financial institutions who cannot afford to take chances. With the process of planning a new data center taking as long as 18 months, recognizing the benefits of exercising full due diligence in that process is essential, and testing is a logical part of any important IT initiative.

The Quantitative Importance of Testing

Nemertes Research reports in their issue paper *Strategic IT Initiatives Need Strategic Testing:* "Without proper testing, such strategic initiatives can fail, with serious unforeseen consequences, including significant hard-dollar and opportunity costs." The Nemertes findings also state that gains from strategic initiatives can be minimized or erased if testing is not implemented before an initiative goes live, that testing should be performed throughout the lifecycle of a strategic initiative, and that budgeting in advance for testing should be standard practice, allocating 2%-5% of the overall budget, including capital expenditures and operational costs.

The customers above clearly recognized the benefits of testing their initiatives before launch to assure the quality of their solution. Indeed a number of our customers continue to partner with Spirent to assure that their evolving solutions meet their and their customers' exacting standards of excellence.

Why Test Your 40/100G Network?

10 Things You Need to Know

The exponential growth, expanding bandwidth demand, and delivery requirements in global networking for video, voice and data applications are accelerating at a speed never before seen. For many, the original transport backbone architecture built on 10G wavelengths is at the threshold of capacity. Higher-speed 40G and 100G are the next generation transport solutions that enable and define future success for service providers, network equipment manufacturers and a growing number of enterprises. Recognizing the right steps of due diligence, which include testing, can assure the delivery of all the benefits expected of this game-changing solution.

Q: On a high level, I understand the growth of demand that pushes the need for 40/100G. Can you give me a deeper insight into the market trends?

A: As Telco's compete with cable providers, with IPTV and on demand use increasing, the bandwidth requirements are also exploding. HDTV uses over five times the bandwidth of SDTV (19 Mbps vs. 3.5 Mbps). YouTube has experienced a 20 percent to 50 percent monthly growth rate. In wireless, video and other applications are constantly pushing the limits of the backhaul network. Social networking is responsible for spikes of usage as people increasingly upload full resolution images from low-cost digital cameras. Storage and data centers continue to grow and data center virtualization represents a critical increase in inter-server traffic as processes are spread across multiple physical machines. The equities trading industry has increasing demand for massive volumes of data (billions of transactions and terabytes of information per day) where low latencies are a matter of survival and surplus bandwidth is essential. 1

Q: I need more details about what aspects of scalability and performance are the critical points of 40/100G testing.

A: Scalability – scale to the highest physical (ports) and/or virtual (emulated subscribers and devices) density; Performance – full line rate traffic generation with the highest connection rates for application layer traffic and emulated protocols. 2

Q: We are planning to manually test our 100G solution before launch with our small team of experts who have studied the associated technology and methodologies. We have the right up-to-date test equipment for the job. Are we missing anything?

A: Due to the significant increase in bandwidth associated with the move to 100G, there is a corresponding increase in test results to be analyzed. More of everything must be supported – more routes, more VPNs, more tunnels, more queues. Because of this it becomes more important to keep the number of test runs to a minimum without sacrificing test coverage and product quality. Hence, if manual testing has been the strategy in the past, it may not be suitable to handle the needs of the future. Automation is an essential element in reducing test time, especially with a small test team, and particularly with the greater test demands imposed by 100G. Test case automation, which uses Tcl scripts or saved configurations to quickly configure and run tests without intervention from the lab engineer, is the most common form of automation. The benefits and advantages are critical in that they free the engineer to do other critical work, forgoing the repetitive requirements of manual testing. Another advantage of automation is 24/7 testing, which maximizes test time. Other opportunities for automation also exist, such as test lab automation. All automation solutions help you speed your 100G solution to market, with accurate results that help you assure the quality of service you and your customers require. 3

Q: What network performance tests are recommended for 40/100G?

A: RFC 2544 testing, which measures latency, Jitter, inter-arrival time, frame loss. RFC 3393, which calculates real-time jitter as required by MEF-10 and RFC 5481. In addition, sequencing, as well as PRBS Bit Error Rates should be tested. In lower speeds this is not an issue in testing. In 40G and 100G, however, this is critical. For example, as the length of cable increases we see packet drops, even in a straight piece of cable. In this stratum of testing, bit errors happen much more frequently, and accounting for error rate speed is critical in 40/100G testing. 4

Q: What are quantifiable benefits of performance testing my 40/100G network before launch?

A: Testing can verify network equipment manufacturer product claims, acquiring consistent, comparable statistics to provide objective comparison of each tested vendor. In addition, not all networks will be composed of equipment from a single vendor, so testing allows you to find performance data between different vendors. You are able to establish network/device performance under load, identifying break points and allowing issues to be resolved prior to the subscriber experience.

5

Q: I recognize the need to test my 100G solution deployment for performance and scalability issues and the need for up-to-date equipment. I'll just have my in-house team get up to speed on the technology and methodologies and take over with the testing guidelines I learned about here, right?

A: Best practices clearly indicate that testing before a 100G network deployment goes live is an effective practice to mitigate costly risks. However, relying on in-house engineering resources, which may lack the required up-to-date expertise and hands-on experience, can produce unreliable test results. The choice of a qualified test partner is critical, especially when deadlines are short, budgets tight, and margin for error is zero. Employing a partner with a team of test experts conversant in the challenges in the field of 40/100G testing is essential. Also, since the cost of the required test equipment can be prohibitive, selecting a partner that also provides rental options offers added advantages.

6

Q: What specific qualifications should I look for when choosing a third party test partner?

A: First, the partner must be a neutral party, without a motivation to influence the test results, as may be the case with a system integrator testing their own solution. Some systems integrators who do their testing in house have a vested interest in delivering positive results. Second, testing should be a core competency of the partner, not an ad hoc solution offered on request. This means the partner is an established global name in the test and measurement industry with verifiable experience and expertise and a team of quality test engineers. The testing team should have a holistic understanding of networks, be able to articulate testing benefits and ROI vs. risk, and have an established delivery process. A test partner's qualifications are further strengthened if they have extensive expertise in lab and test automation and can provide references of successful engagements. The most advantageous test partners that provide test rental equipment are those who supply devices manufactured by the same company.

7

Q: We have very reliable performance testing gear in our lab which we have used with great satisfaction over the last three years. We intend to employ this platform for testing our new 100G solution. Does this testing strategy make sense?

A: Resolution of legacy test equipment (including those for 10GE testing) is 20ns (nanoseconds) or more, which cannot deliver accurate measurements for 40/100G. Here is why: Assume one frame every 672 bit times, a 12-byte inter-frame gap, an 8-byte preamble, and the smallest frame size of 64 bytes. The time between two frames is 67.2ns for 10G, 16.8ns for 40G and 6.72ns for 100G. Therefore, if the traffic generator has a clock resolution of 20ns it will send two frames at 40G and three frames at 100G with the same Tx or Rx time stamps. One must have at least 16.8ns resolution at 40G or 6.72ns resolution at 100G to uniquely timestamp contiguous packets. Having the right test equipment with the right nanosecond resolution is essential to test properly and get the results you need to assure your 100G solution is ready for launch.

8

Q: We are planning to migrate to a 40G solution. My highly-qualified team of engineers has ample experience with designing and implementing 10GE networks. If we have their stamp of approval, is it really necessary to test our newly-proposed solution any further before launch?

A: First, 40G is very different from 10G technologies. It's not simply a matter of being four times faster than 10GE, but a multi-lane, parallel wavelength system. Jitter and sequencing testing may not be that essential to a 10GE network but they are critical to multi-lane systems. Having a team of engineers conversant with 10G networks does not guarantee their expertise in 40G and testing properly with a suitably qualified team before launch addresses that issue.

9

Q: Our budget is approved for upgrading our testing platform for our planned 100G network. Are there other components I need to keep in mind as well?

A: Not all system clocks are created equal. Thorough research is essential to assure the one you have is right for your testing requirements. A seemingly insignificant detail such as clock resolution can affect the ROI of your test lab investment as well as the TCO (total cost of ownership) for the equipment purchased. The time to ask the hard questions, such as can the system scale to test routing at high speeds, or stream and analyze real video, is before the test platform purchase is made.

10

Spirent Puts Performance and Scale of Juniper Networks Qfabric Data Center Networking Solution to the Test

Any-to-Any Connectivity and Low-Latency Data Center Performance

March 6, 2012 — Spirent announced that Network Test, an independent test lab, selected its solutions to complete the industry's first public 1,536 port 10G Ethernet test of QFabric™, Juniper Networks data center network fabric. Network Test worked with Spirent and Juniper Networks on an innovative methodology to test QFabric's performance at massive scale, under the most stressful yet realistic conditions. The results allow Juniper Networks to confidently address the scalability, performance and latency needs of increasingly large data centers. With a port density more than four times greater than any previous public test of its kind, this effort involved 1,536 10G Ethernet ports, measuring nanosecond latency at terabits per second.

Today's data centers are growing rapidly in scale, with increasing numbers of physical and virtual servers. This results in the need for high-density, high-speed fabric based networks with terabits of any-to-any port traffic. For these increasingly large scale data centers, it is imperative to understand how the network fabric will perform under stress, with latency-sensitive traffic running over it. This independent test of QFabric reflects real-world data center setups, with millions of application flows between any set of ports. It has important implications for cloud computing, storage and financial applications such as high-frequency trading and big data analytics.

Spirent Puts Juniper Networks Mobility Solution to the Test

EANTC Uses Spirent Solutions to Validate Performance, Scale and Reliability of MobileNext Broadband Gateway

March 7, 2012 – Spirent announced that the European Advanced Networking Test Center (EANTC) and Juniper Networks selected its mobile network solutions to test the Juniper Networks® MobileNext™ Broadband Gateway. EANTC's tests used Spirent solutions to test the performance, scalability and reliability of the gateway and its ability to support large metropolitan areas with several million subscribers accessing 3G or LTE applications.

EANTC validated MobileNext Broadband Gateway's ability to scale up to millions of mobile subscriber data sessions, with seamless mobility across 3G and LTE networks. Although service provider networks are not yet required to support this scale, it is critical that converged mobile gateways be prepared to support the exponential growth rates in mobile data traffic forecasted for the next few years.

Spirent Selected to Test Antenna Implementations for LTE-Advanced and TD-LTE in Asia

Spirent's Award Winning VR5 Channel Emulator Aids Implementation of Advanced MIMO Beamforming Services

February 14, 2012 — Spirent announced that several leading companies and government research organizations in the Asia Pacific region are using the Spirent VR5 HD Spatial Channel Emulator to test implementations of advanced MIMO beamforming. The Spirent VR5 emulates real-world radio conditions and identifies potential pitfalls of the antenna design in the mobile device before deployment.

MIMO beamforming combines two separate antenna techniques: MIMO, which offers a significant gain in data rates when compared to traditional wireless technologies; and beamforming, which steers wireless signals towards heavy usage areas, thereby increasing system efficiency. MIMO beamforming is a planned part of TD-LTE technology, expected to be widely deployed in several key markets worldwide, and of LTE-Advanced. However, due to the complex interactions between radiated signals and between the techniques themselves, turning theory into reality involves intensive laboratory work.

Spirent Expands Leadership in Testing E911 and Location Based Services for LTE Networks

Enabling E911 and Moving One Step Closer to Accurate Everywhere Location

February 21, 2012 – The rise in popularity of Location Based Services (LBS) and the FCC mandate for optimal E911 performance are driving the wireless industry to rapidly evolve positioning capabilities on LTE networks and devices. Recognizing the increase in LBS momentum and the need for improved positioning performance on mobile devices, Spirent Communications has expanded its LBS LTE test solution to support LTE Positioning Protocol (LPP) and Observed Time Difference of Arrival (OTDOA). These latest enhancements are critical for meeting E911 requirements on LTE networks, and they bring the mobile industry one step closer to accurate everywhere location, the key enabler for truly mainstream LBS adoption.

Although the need for emergency services (E911) and LBS is nothing new, the arena in which these services need to be delivered has grown significantly more complex. The deployment of LTE enables the introduction of new positioning technologies and positioning protocols for communicating location information between the mobile device and network entities, with exciting consequences for the new services they can enable.

Spirent First in the Industry to Put Performance of Wi-Fi Offload Gateways to the Test

Solution Enables Mobile Operators to Reduce Network Costs, Boost Quality of Experience

February 27, 2012 — Spirent announced the addition of Wi-Fi Offload Gateway testing capability to its Spirent Landslide solution. Landslide is the industry's first and only solution to test the performance of Wi-Fi Offload Gateways which handle offloading of data from sources such as Over-The-Top video, as well as traditional services such as voice calling and SMS, from a 3G/4G/LTE cellular network to a Wi-Fi network.

As mobile operators face overwhelming bandwidth demands from data-hungry subscribers, Wi-Fi offloading is an important approach to helping ease congestion and control capital expenditure. It helps ensure subscriber quality of experience (QoE) by freeing up mobile spectrum for users who are truly mobile and by offering all users higher bandwidths. The Wi-Fi Offload Gateway is a critical element in ensuring seamless and secure service handoff from the mobile network. Unlike competitive products, Spirent's solution tests the performance and scalability of the Wi-Fi Offload Gateway and the mobile core network, enabling authentication and security scenarios to be verified, as well as handoffs.

Spirent Testcenter Voted By Communications Industry Peers as the Best Testing Platform

Spirent Receives Test & Measurement World Magazine's 2012 Test of Time Award, Recognized for Providing Industry's Highest Performance in Testing

February 28, 2012 – Spirent Communications' flagship test platform Spirent TestCenter™ received the test and measurement industry's most prestigious honor – the Test & Measurement World Magazine's 2012 Test of Time award. Selected based on the greatest number of votes from the magazine's readers and editors, the award was presented to Spirent for excellence in testing and commitment to providing the industry with state-of-the-art testing performance over the long term.

Originally introduced in 2005, Spirent TestCenter is an integrated solution that tests the performance, availability, security and scalability of communications and networking technologies. With more than 1600 customers across a wide range of industries, including retail, financial services and networking, Spirent TestCenter offers the broadest coverage in high performance testing of networks and applications. The platform has constantly evolved to meet the testing needs of users in fast-growing technologies areas such as virtualization, cloud computing, backhaul and evolved packet core for 3G/LTE networks, security and high-speed Ethernet.

SPIRENT®

Spirent Communications
1325 Borregas Avenue
Sunnyvale, CA 94089 USA

SALES AND INFORMATION
sales-spirent@spirent.com
www.spirent.com

Americas
T: +1 800.SPIRENT
+818 676.2683

Europe, Middle East, Africa
T: +33 1 6137.2250

Asia Pacific
T: +852 2511.3822