VIAVI Solutions

This Former Spirent Business is Now Part of VIAVI

Contact Us +1 844 GO VIAVI | (+1 844 468 4284)

To learn more about VIAVI, visit viavisolutions.com/en-us/spirent-acquisition

DATASHEET

Spirent TestCenter

AI Data Center Test Solution

Application Overview

In today's rapidly evolving technological landscape, Enterprise and data center networking has undergone significant advancements. With the widespread adoption of AI/ML, edge computing, and high-performance computing (HPC) systems, modern data centers demand a network that delivers exceptional performance and minimal latency. This network must also seamlessly adapt to accommodate the tremendous surge in extremely low latency and lossless high-volume data transfer demands driven by data-intensive and compute-intensive applications.

Remote Direct Memory Access technology (RDMA) technology offers efficient and faster ways to move data between xPUs within AI clusters while boosting network and host performance with lower latency, reduced CPU load, and higher bandwidth.

RDMA over Converged Ethernet (RoCE) delivers all the advantages of RDMA in existing Ethernet networks, resulting in significant capital expenditure savings. With the heavy compute requirements associated with intensive AI/ML & storage workloads, RoCE is a fundamental component in ensuring a reliable, end-to-end transport solution for modern data centers. This technology further enhances performance in search, storage, database, financial and high transaction rate applications.

RDMA over Converged Ethernet version 2 (RoCEv2) extends the reach of RDMA technology to be used across both L2 and L3 networks with multiple subnets. This capability allows efficient clustering for elastic and scalable deployments.

Since its inception in 2014, many organizations have been deploying RoCEv2 in their data centers. However, in these large-scale environments, incorrect or non-optimized network settings can lead to poor application performance. Therefore, it is critical to validate switch fabric performance, optimize configurations and ensure network stability and resiliency under congestion conditions.

The pivotal congestion control mechanisms facilitated by switches and NICs encompass Data Center Quantized Congestion Notification (DCQCN) and Priority Flow Control (PFC). While testing with homegrown (with real servers) and open-source test solutions can be valuable, they often fall short of addressing the scale and efficiency requirements.

Spirent offers the industry's most comprehensive solutions, with hardware options supporting 100G to 800G, by accurately emulating Collective Communication Library (CCL) traffic patterns using RoCEv2 as a transport on Ethernet and measuring KPIs for AI Ethernet fabric. It effectively validates a data center's Ethernet infrastructure's ability to support massive AI workloads without the deployment and management challenges of using xPU-based server farms as testing frameworks.

spirent™

Features and Benefits Overview

- **Test coverage options from 100G to 800G connectivity**
- **Realistic AI traffic on Ethernet using RoCEv2 as a transport between xPUs within AI Clusters—**simulate continuous or burst traffic patterns per test port at line rate and support per-flow congestion control using DCQCN and port-level congestion control using PFC
- **Accelerated time-to-market—** validate switches and networks in AI environments and bring products to market faster
- **Lower Cost of Ownership—**replace racks of real xPU-based servers used for testing and save on CAPEX, complexity, and ongoing maintenance
- **Investment Protection—**leverage A1/FX3/MX3 high-density multi-speed test modules in multiple ways:
  - UPGRADE existing B3 & A1 Appliances and FX3 & MX3 high density test modules to RoCEv2 testing via simple licensing
  - REUSE test modules for regular L2-3 testing improving your testing ROI
  - LEVERAGE additional speeds for RoCEv2 testing in future releases
- **Integrated Collective Communication Library—**configure various AI topologies and generate multiple AI traffic patterns through an easy-to-use GUI-based wizard. No additional software or hardware required
- **Leverage Expertise—**simplify the complexity of testing AI data center network fabrics, as a purpose built testing solution

Collective Communication Library (CCL)

Due to the large amount of data communication between xPUs in an AI Data Center, processing this data cannot be handled by a single xPU and must therefore be spread across different xPUs. CCL, built on open MPI, enables high-performance computing and parallel processing in the AI data center. Spirent's CCL library supports the most popular CCL traffic patterns like AlltoAll with reference to NCCL. Through the wizards, the AI traffic patterns are generated easily. When network congestion occurs, Spirent's solution supports the DCQCN and PFC flow control, mirroring real AI Networks. The Job Completion Time (JCT) presented at the end of the test helps evaluate the performance of the network fabric in the AI data center.

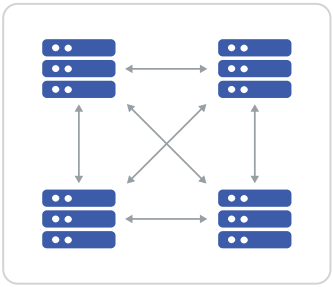

AlltoAll

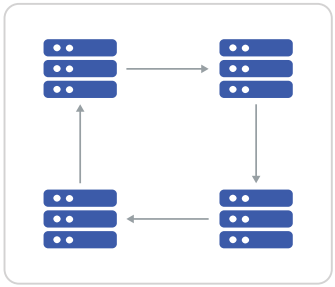

RingAllReduce

Spirent's **AI test solution** offers a high density, multi-speed testing environment that is readily available, cost-effective and easy to deploy. Built on open and transparent architecture, it ensures consistent and repeatable results. Through CCL, it generates realistic AI traffic patterns at line rate while supporting congestion control. Users can measure key network KPIs, such as JCT, throughput, tail latency, packet latency, packet loss, re-ordered packet count, late packet count, to evaluate an AI fabric. These KPIs help ascertain the resiliency of AI fabric under congestion conditions and link flap scenarios.

It enables users to characterize the performance of AI fabric and optimize the relevant network configurations, including buffer size, ECN, load-balance algorithm, and QoS settings.

Spirent's **AI test solution** is an innovative offering that runs on four of the existing best-in-class high-density, multi-speed Ethernet test modules.

- The B3 4-port and 8-port QSFP-DD and OSFP 800G appliances with industry-leading port density advantage for QSFP-DD and OSFP.
- The A1 400G 16-port appliance is the industry's highest density 400GE platform.

Both platforms support emulation of AI workloads and are multi-purpose to test both AI and routing/switching use cases concurrently.

FX3-QSFP28-6 offers 3x 100G ports or fan out to 12x 25G ports. In addition, both MX3-QSFP28-4 and FX3-QSFP28-4 offer 2x 100G ports or fan out to 8x 25G ports, making them ideal for RoCEv2 testing.

Multi-user support is available with per-port user reservation. This allows ports from a single test appliance/module to be used in either single or multiple concurrent test sessions. Additionally, these test appliances/modules can be used for regular L2-3 tests on a per port basis (outside RoCEv2 profile) for 800/400/200/100/50/40/25/10G speeds, maximizing your investment.

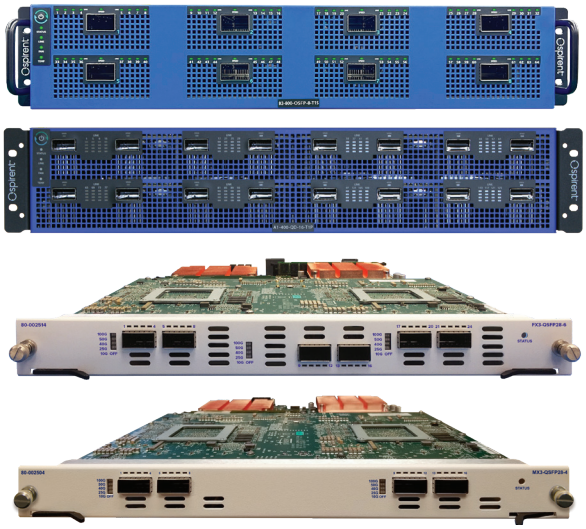

This solution is an integrated component of Spirent TestCenter and works with other Spirent TestCenter components to deliver easy, consistent TCL and REST API for scripting and Command Sequencer NoCode automation.

Features

- RoCEv2 support* for 800G, 400G, 200G, 100G, 25G
- RoCEv2 traffic over Lossy and Lossless fabric using hardware-based, per-flow congestion control triggered by CNP/ECN and hardware-based PFC congestion control at the port level
- Integrated CCL traffic patterns with an easy-to-use wizard
- Configuration to optimize congestion control mechanisms DCQCN and PFC according to network requirements
- Background traffic sharing the same port as CCL traffic emulates complicated real-world traffic
- Automatic Benchmark Tool offers convenient and simplified test configuration. A single wizard run can conduct a series of use cases. Test report is available in either TestCenter IQ or Results Reporter
- Facilitation of different traffic scheduling mechanisms and topology options to simulate both xPU and storage data transfers
- Multiple xPU and RoCEv2 server emulation per port and High scale with extensive Queue Pairs per port
- QoS settings for L2 (VLAN) and L3 (DSCP) mode
- Per QP/Flow congestion control support
- JCT, tail latency, and bandwidth reporting
- Simulated One-to-One data transfer between Queue Pairs
- By enabling multiple traffic patterns in single test, it is possible to comprehensively evaluate the performance in heterogeneous scenarios, effectively overcoming the limitations of single-traffic pattern verification
- Traffic Chart visually depicts data flow directions and facilitates clear identification of problematic streams
- Emulate Queue Pair Establishment process and support for Reliable Connection
- Spirent TestCenter+ offers a user-friendly web interface for streamlined test configuration, execution, and results analysis

Benefits

- **Assess & Validate AI Network infrastructures** at massive scale
- **Emulate xPUs** for large scale testing, complexity of in building CCL topology and traffic patterns masked by the RoCEv2 & CCL Wizard
- **Emulate servers** in a test bed, thereby reducing inventory cost and maintenance significantly
- **Reduce time-to-test** with easy configuration of realistic test topology
- **Scale up easily** with extensive Queue Pairs in a few clicks using a wizard
- **Collect key statistics** and generate intuitive reports with TestCenter IQ
- **Expand set of use case coverage** of AI Data Center test functionality in the Spirent TestCenter+ web-based user interface

Technical Specifications	
Parameters	
Application Support	
Application CCL Configuration	Spirent TestCenter • AlltoAll • RingAllReduce • Halving and Doubling • Dual Tree • RingReduceScatter • RingAllGather • Alternating Ring • Mixed Traffic Patterns
RoCEv2 Configuration	• RoCEv2 profile applied per test module • Multiple server emulations per port • 1K Queue Pairs per port supported
RoCEv2 Throughput	800, 400, 200, 100 and 25Gbps
QoS	• L2 - VLAN based 8 priority values • L3 – DSCP based 64 priority values
Congestion Control	ECN, PFC combination for lossy and lossless fabric • ECN detection, CNP notification • CNP interval • QoS priority for CNP notification • DCQCN rate control settings • PFC with eight queues
RoCEv2 Traffic Configuration	• Traffic endpoints – queue pairs • Traffic mode – continuous and bursty • Full mesh, one-to-one, many-to-many traffic between servers • Bidirectional and unidirectional traffic
Key Results (TestCenter IQ)	• JCT, Ideal JCT • Total Transmission Time, Total Training Time • Average Iteration Time, Max Iteration Time, Min Iteration Time • Tail latency • Throughput, packet loss, re-ordered count • ECN, CNP TX/RX • Algorithm bandwidth (algbw) and Bus bandwidth (busbw) based on JCT • CCL Traffic Chart • algbw and busbw based on transmission time
Test Appliance Support	
A1-400-QD-16	**Port Density** for AI testing – 16x 400G ports / 32x 200G ports / 32x 100G PAM4 Ports
RoCEv2 Supported	Yes
CCL Configuration	Yes
User Reservation	Per-port user reservation
B3-800-OSFP, B3-800-QD, B2-800-OSFP, or B2-800-QD	Port Density for AI testing for up to 8 x 400G* ports and up to 4 x 800G ports
RoCEv2 Supported	Yes
CCL Configuration	Yes
User Reservation	Per-port user reservation
Test Module Support	
FX3-QSFP28-6	**Port Density** for AI testing – 3x 100G -Or- 12 x 25G ports (L2-3 Testing Datasheet)
FX3-QSFP28-4	**Port Density** for AI testing – 2x 100G -Or- 8 x 25G ports (L2-3 Testing Datasheet)
MX3-QSFP28-4	**Port Density** for AI testing – 2x 100G -Or- 8 x 25G ports (L2-3 Testing Datasheet)
RoCEv2 Supported	Yes
CCL Configuration	No
User Reservation	Per-port user reservation
Chassis Support	
SPT-N12U	**Slots per chassis:** 12 **Port Density** for AI testing: • 36 x 100G -Or- 144 x 25G ports (**FX3-QSFP28-6**) • 24 x 100G -Or- 96 x 25G ports (**MX3-QSFP28-4**) • 24 x 100G -Or- 96 x 25G ports (**FX3-QSFP28-4**)
SPT-N4U	**Slots per chassis:** 2 **Port Density** for AI testing: • 4 x 100G -Or- 16 x 25G ports (**MX3-QSFP28-4**) • 4 x 100G -Or- 16 x 25G ports (**FX3-QSFP28-4**)

Ordering Information

Test Appliance

Part Number	Description
A1-400-QD-16	A1 16-Port QSFP-DD 400G/200G/100G/50G T1 APPLIANCE
HWO-A1-16-ROCE-ECN-400-16	Spirent RoCE ECN-400G FOR A1-400-QD-16. 16 X 400G PAM4 ports total per appliance.
HWO-A1-16-ROCE-ECN-200-32	Spirent RoCE ECN-200G FOR A1-400-QD-16. 32 X 200G PAM4 ports total per appliance.
HWO-A1-16-ROCE-ECN-100-32	Spirent RoCE ECN-100G FOR A1-400-QD-16. 32 X 100G PAM4 ports total per appliance.
HWO-A1-16-ROCE-ECN-500A	Spirent RoCE ECN-400G/100G SUITE FOR A1-400-QD-16. 16 X 400G PAM4 or 32 X 100G PAM4 ports total per appliance.
HWO-A1-16-ROCE-ECN-600A	Spirent RoCE ECN-400G/200G SUITE FOR A1-400-QD-16. 16 X 400G PAM4 or 32 X 200G PAM4 ports total per appliance.
HWO-A1-16-ROCE-ECN-700A	Spirent RoCE ECN-400G/200G/100G SUITE FOR A1-400-QD-16. 16 X 400G PAM4 or 32 X 200G PAM4 or 32 X 100G PAM4 ports total per appliance.
HWO-A1-16-CCL-16	Spirent CCL FOR A1-400-QD-16. All speeds are included. 16 X 400G PAM4 or 32 X 200G PAM4 or 32 X 100G PAM4 ports total per appliance.
HWO-A1-8-ROCE-ECN-400-8	Spirent RoCE ECN-400G FOR A1-400-QD-16. 8 X 400G PAM4 ports total per appliance.
HWO-A1-8-ROCE-ECN-200-16	Spirent RoCE ECN-200G FOR A1-400-QD-16. 16 X 200G PAM4 ports total per appliance.
HWO-A1-8-ROCE-ECN-100-16	Spirent RoCE ECN-100G FOR A1-400-QD-16. 16 X 100G PAM4 ports total per appliance.
HWO-A1-8-ROCE-ECN-500A	Spirent RoCE ECN-400G/100G SUITE FOR A1-400-QD-16. 8 X 400G PAM4 or 16 X 100G PAM4 ports total per appliance.
HWO-A1-8-ROCE-ECN-600A	Spirent RoCE ECN-400G/200G SUITE FOR A1-400-QD-16. 8 X 400G PAM4 or 16 X 200G PAM4 ports total per appliance.
HWO-A1-8-ROCE-ECN-700A	Spirent RoCE ECN-400G/200G/100G SUITE FOR A1-400-QD-16. 8 X 400G PAM4 or 16 X 200G PAM4 or 16 X 100G PAM4 ports total per appliance.
HWO-A1-8-CCL-8	Spirent CCL FOR A1-400-QD-16. All speeds are included. 8 X 400G PAM4 or 16 X 200G PAM4 or 16 X 100G PAM4 ports total per appliance.
AON-A1-AI-16-700A	A1 16-PORT ROCE ECN AND CCL 400G/200G/100G BUNDLE (AION).
AON-A1-AI-8-700A	A1 8-PORT ROCE ECN AND CCL 400G/200G/100G BUNDLE (AION).
B3-800-OSFP & B3-800-QD	B3 800G Appliances available in 4-Port and 8-Port configurations for OSFP or QSFP-DD.
HWO-B3-ROCE-ECN-400-8*	Spirent RoCE ECN-400G FOR B3-800-OSFP or B3-800-QD for 8 X 400G PAM4 ports total per appliance.
HWO-B3-ROCE-ECN-800-4	Spirent RoCE ECN-800G FOR B3-800-OSFP or B3-800-QD for 4 X 800G PAM4 ports total per appliance.
HWO-B3-ROCE-ECN-1200A	SPIRENT ROCE ECN-800G/400G SUITE FOR B3-800-QD or B3-800-OSFP. 4-port B3 and 8-port B3: 4 X 800G PAM4 ports or 8 X 400G* PAM4 ports total per appliance.
HWO-B3-CCL-4	Spirent CCL FOR B3-800-OSFP or B3-800-QD. Speeds are included: 4 X 400G PAM4 or 4 X 800G PAM4 ports per appliance.
HWO-B2-ROCE-ECN-800-4	SPIRENT ROCE ECN-800G FOR B2-800-QD or B2-800-OSFP. 4-port B2: 4 X 800G PAM4 ports. 8-port B2: 4 X 800G PAM4 ports total per test module.
HWO-B2-ROCE-ECN-400-8	SPIRENT ROCE ECN-400G FOR B2-800-QD or B2-800-OSFP. 4-port B2: 8 X 400G PAM4 ports. 8-port B2: 8 X 400G PAM4 ports total per test module.
HWO-B2-ROCE-ECN-1200A	SPIRENT ROCE ECN-800G/400G SUITE FOR B2-800-QD or B2-800-OSFP. 4-port B2: 4 X 800G PAM4 ports or 8 X 400G PAM4 ports. 8-port B2: 4 X 800G PAM4 ports or 8 X 400G PAM4 ports total per test module.
HWO-B2-CCL-4	SPIRENT CCL FOR B2-800-QD or B2-800-OSFP. All speeds are included. 4-port B2: 4 X 800G PAM4 or 8 X 400G PAM4 or 8 X 200G PAM4 or 8 X 100G PAM4 ports. 8-port B2: 4 X 800G PAM4 or 8 X 400G PAM4 or 16 X 200G PAM4 ports or 16 X 100G PAM4 ports total per test module.
TPK-1106	CCL AUTOMATION BENCHMARK TEST PACKAGE. Requires RoCEv2 and CCL license. Works with A1, B2, and B3 ONLY.
AON-PB-AI	AION AI BASIC Bundle. Include TPK-1106 CCL Automation benchmark test package. Requires AON-A1-AI, AON-B2-AI or AON-B3-AI which are sold separately.

Test Modules

Part Number	Description
FX3-QSFP28-4	Spirent FX3 QSFP28 4-Port
FX3-QSFP28-6	Spirent FX3 QSFP28 6-Port
MX3-QSFP28-4	Spirent MX3 QSFP28 4-Port
Chassis	SPT-N12U Chassis, SPT-N4U Chassis
HWO-FX3-6-ROCE-ECN-100-3	Spirent RoCE ECN-100G Upgrade for FX3 QSFP28-6 - 3 PORTS - 3 X 100G ports total per test module. For non-ECN operation, additional normal speed license required.
HWO-FX3-4-ROCE-ECN-100-2	Spirent RoCE ECN-100G Upgrade for FX3 QSFP28-4 - 2 PORTS - 2 X 100G ports total per test module. For non-ECN operation, additional normal speed license required.
HWO-MX3-4-ROCE-ECN-100-2	Spirent RoCE ECN-100G Upgrade for MX3 QSFP28-4 - 2 PORTS - 2 X 100G ports total per test module. For non-ECN operation, additional normal speed license required.
HWO-FX3-6-ROCE-ECN-25-12	Spirent RoCE ECN-25G Upgrade for FX3 QSFP28-6 - 12 Ports -128 X 25G ports total per test module. For non-ECN operation, additional normal speed license required.
HWO-FX3-4-ROCE-ECN-25-8	Spirent RoCE ECN-25G Upgrade for FX3 QSFP28-4 - 8 PORTS - 8 X 25G ports total per test module. For non-ECN operation, additional normal speed license required.
HWO-MX3-4-ROCE-ECN-25-8	Spirent RoCE ECN-25G Upgrade for MX3 QSFP28-4 - 8 PORTS - 8 X 25G ports total per test module. For non-ECN operation, additional normal speed license required.

Note: corresponding hardware speed license is required. Some RoCEv2 and CCL features will be available in future software releases on specific platforms.
* Ordering availability / and support will be available in upcoming software releases. RoCEv2 speed support varies by platform.