TEST SOLUTIONS TO VALIDATE 400GE DEVICES AND NETWORKS

GET 400GE PRODUCTS TO MARKET FAST

To meet the growing bandwidth requirements of data networks, network equipment manufacturers (NEMs) and their customers are looking to 400GE networking solutions. Ixia's 400GE Higher Speed Ethernet (HSE) test solutions are integral to ensuring they meet performance, quality, and conformance requirements.

NEMs around the globe are racing to bring 400GE products to market. Ixia's QSFP-DD and CFP8-400GE test systems shorten development and test times for these NEMs—accelerating critical time to market. Our 400GE load module validate the development of IEEE 802.3bs-compliant 400GE networking systems, helping critical-path, development teams create innovative networking technologies of the future.

REDUCE 400GE TEST TIME

Ixia-developed 400GE MAC, PCS, and FEC engines for true interoperability with other IEEE 802.3bs vendor IP, and to provide customers with quick response time.

Now, customers can create, test, and verify the interoperability of their next-generation HSE technologies, while maximizing their 400GE test investment. Ixia's QSFP-DD and CFP8-400GE test systems deliver cutting-edge, wire-rate packet generation at 400GE below the minimum and above the maximum Ethernet frame sizes.

Ixia's offers both full-feature and reduced-feature products to meet different test requirements and budget sizes from the hardware/ASIC bring-up teams that need L1 testing and line-rate packet blasting, to those QA, software integration, and protocol teams requiring high-performance L2-7 product performance and protocol analysis.

KEY FEATURES

- Line-rate 400Gbps packet generation, capture, and analysis of received traffic
- Line-rate per-port and per-flow statistics
- High latency measurement resolution at 0.625ns
- RS-544 (KP4) Forward Error Correction (FEC)
- FEC error injection with a comprehensive set of statistics
- Inject packet errors: CRC, runts, giants, alignments, and Out of Sequence
- Standard Ixia instrumentation including timestamp, sequence number and flow identification, and data integrity
- Layer 1 BERT: 8 independent lanes of 56Gb/s PRBS pattern generation and error checking over QSFP-DD; 16 independent lanes of 25Gb/s PRBS pattern generation and error checking over CFP8
- 400G PCS lanes Transmit, Error Injection testing and Receive measurement: Per-lane controls and status, FEC and error monitoring, error insertion, lane skew and swapping
- +/-100 PPM adjustment

LOOKING TO TEST 400GE DEVICES OR NETWORKS? CHECK OUT THESE DOCS.

White Paper: Validating the New World of 400 Gigabit Ethernet

Data Sheet: K400 QSFP-DD 400GE Load Modules

Data Sheet: Ixia CFP8-400GE Test System

Validating the New World of 400 Gigabit Ethernet

THE RETURN OF INDUSTRY BENCHMARKING

WHITE PAPER

ixia
A Keysight Business

915-3767-01-1171 Rev A

CONTENTS

The Leap from 100GE to 400GE Brings New Perils **3**

The Forecast for Bandwidth Growth **3**

More Ethernet Speeds **5**

400GE Brings Big Challenges and New Technologies **7**

Forward Error Correction 7

Transmit Transcoding and Receive Decoding 8

New Optical Transceivers 10

More Silicon Vendors and More Device Diversity **11**

More Silicon Implementation Options 11

Merchant Silicon Market Expands 12

Merchant Silicon Diversity 13

The Need for Performance Benchmarking **15**

A Lesson from History 15

Testing is Key to Overcome 400GE Challenges 16

400GE Benchmark and Performance Testing **18**

RFC Performance Measurement Metrics 19

Summary **20**

THE LEAP FROM 100GE TO 400GE BRINGS NEW PERILS

The networking industry is experiencing unprecedented global demand to deliver more network bandwidth to meet the needs of data, video, mobile, and storage applications. Megatrends within the industry like moving to the cloud and big data, plus those driven by end-users and subscribers such as going mobile, and the proliferation of the Internet of Things (IoT) are pushing industry standards' bodies to develop more and faster Ethernet speeds.

As mega-trends fuel the need for more and faster Ethernet speeds, new silicon chip maker startups are competing for a share of the market.

Communications chip manufacturers, network equipment manufacturers (NEMs), cloud service providers (CSPs), carriers and service providers, hyperscale and large enterprise data centers, financial, and storage network organizations must adopt and deploy new Ethernet speeds like the fast-emerging 400 gigabit Ethernet (GE) and 200GE, plus new implementations of 100GE and 50GE. There is a tremendous amount of new product development occurring between 2015 and 2020 to meet the bandwidth demand. Now, we are right in the middle of this period, and new chip makers are coming to market, introducing high-performance products. There is no longer just one or two primary suppliers of communications chips and devices.

As a result, this is a critical time to renew efforts to compare and benchmark the performance of these new network products and devices.

THE FORECAST FOR BANDWIDTH GROWTH

Bandwidth growth is the driving factor for the need for 400GE. Past studies have defined the growth in bandwidth. The best-fit rules for bandwidth growth prediction has been:

- Core networking doubles its bandwidth demand every 18 months
- Server I/O doubles its bandwidth demand every 24 months

However, recent application growth and a shift by large enterprises to move their IT services to the cloud has tilted the scales a bit. Now, bandwidth demand is defined by traffic that is measured relative to whether it is non-data center IP traffic or global data center IP traffic. It is no longer about the core network growth, it is more about data center and server I/O computing bandwidth. Note in Figure 1, that global data center IP traffic has a growth rate 4x larger than the previous traditional non-data-center IP traffic.

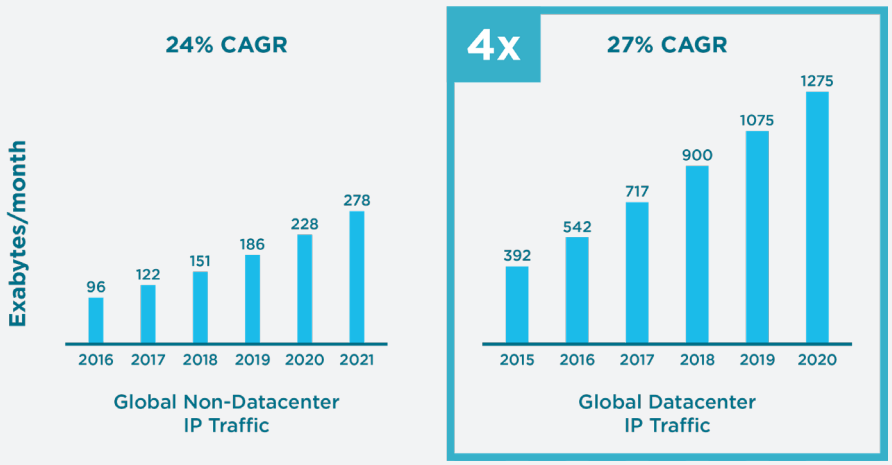

Figure 1. Non-data-center versus data-center bandwidth growth comparison, (Source: Cisco Global Cloud Index and Cisco Visual Networking Index).

Today, hyperscale data center companies establish the bandwidth growth baseline. These companies drive the networking industry now, rather than the carriers who did in the past. It is data center interconnects that drive the market. Data center bandwidth requirements are doubling every 12-15 months. NEMs must build extremely high throughput and density switches, designed for data centers to accommodate the bandwidth that the application servers require. Figure 2 shows the distribution of traffic based on applications and where the majority of the bandwidth demand is—within the data center.

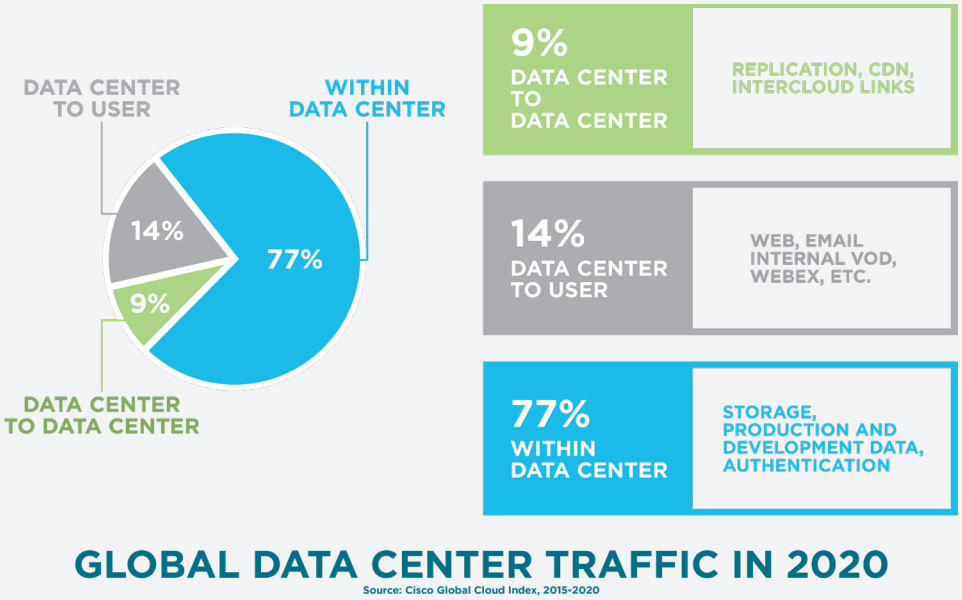

Figure 2. Global Data Center Traffic by Destination in 2020 (Source: Cisco Global Cloud Index, 2015-2020).

The vast number of servers deployed by just dozens of global companies drives bandwidth needs. Figure 3 shows the massive scale of the number of servers deployed by recognized brand-name companies.

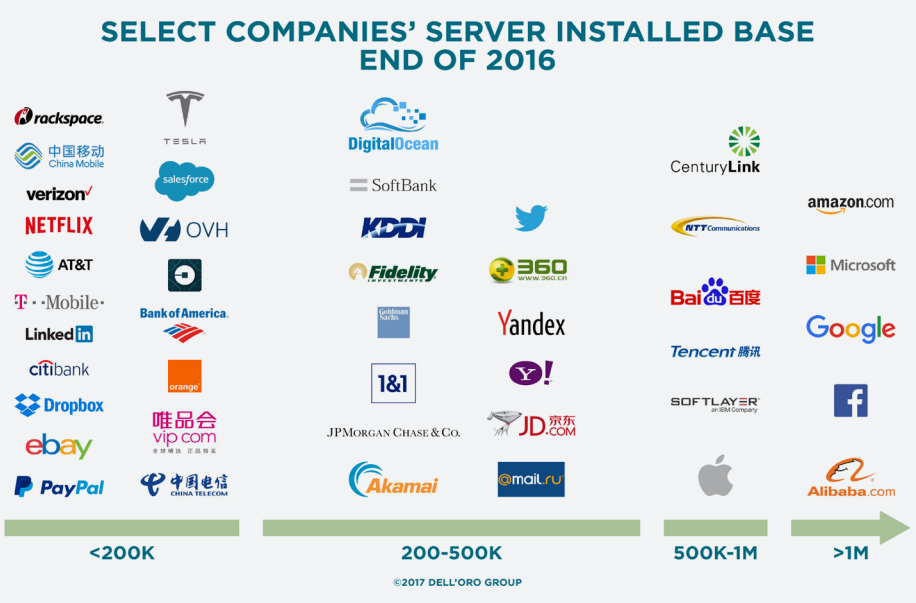

Figure 3. The massive scale of deployed servers by brands we all know (source: Dell'Oro Group).

In 2016 Amazon, Alphabet, and Microsoft together racked-up nearly $32bn in capital expenditure and capital leases, up by 22% from the previous year, according to the Wall Street Journal.

The need for more bandwidth to support this kind of growth is established. The question is how does the industry evolve Ethernet, the primary networking interface technology to provide more speed and subsequently, greater bandwidth.

MORE ETHERNET SPEEDS

Standard's bodies have provided the industry with a diverse range of Ethernet speeds to address a broad range of application requirements. Figure 4 briefly describes the well-known efforts to deliver more Ethernet speeds.

IEEE Draft Standards	Ethernet Speeds	Applications	Projected Ratification Date
IEEE 802.3bs	400GE, 200GE	Data center, core/edge routing	End of 2017
IEEE 802.3cd	200GE, 100GE, 50GE	Data center for higher bandwidth switches and servers	End of 2018
Consortium 25G/50G	25GE, 50GE	Data center for higher bandwidth switches and servers	V1.6 is released

Figure 4. Summary of industry standards for 400/200/100/50/25 Ethernet speeds.

The timelines for the standards are not gating any effort by NEMs to deliver 400GE, 200GE, 100GE, and 50GE prior to the ratification of the standards. Companies around the globe are racing to be first to market with 400GE.

Momentum in the development of 400GE equipment is in late 2017 through 2018. Deployment is geared toward 2019 and 2020 as shown below in figure 5.

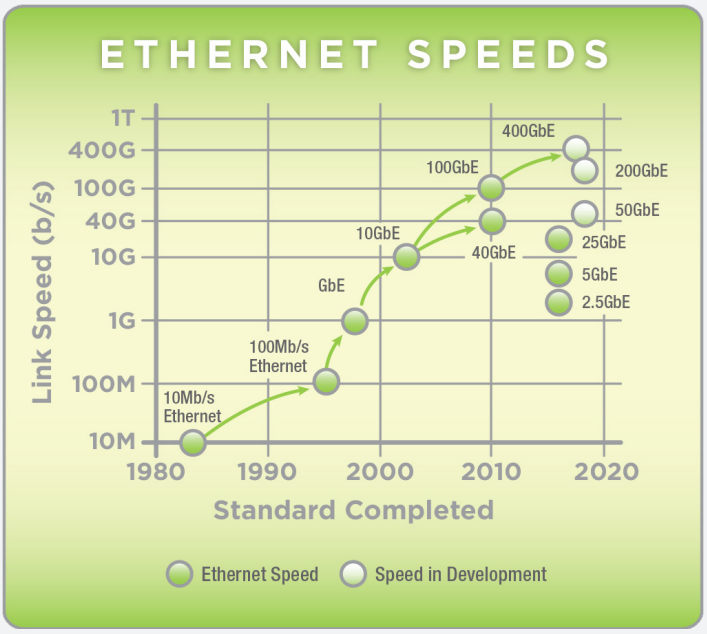

Figure 5. Ethernet Alliance industry timeline for 400GE development.

Earlier this year, Alan Weckel, Founder and Technology Analyst at 650 Group, discussed 400GE growth. He indicated that in 2019 there will be nearly 1 million port shipments of 400GE with a ramp up to 5 million port shipments in 2021. He continued by saying, "Clouds and telco service providers look towards higher speeds both within the data center and interconnecting geographically dispersed data center locations. Semiconductor designs for data center networking systems are expected to undergo a rapid pace of advancement over the next five years in order to keep up with strong demand from the cloud, especially Amazon, Facebook, Google, and Microsoft."

"Clouds and telco service providers look towards higher speeds both within the data center and interconnecting geographically dispersed data center locations." —Alan Weckel, Founder and Technology Analyst at 650 Group

Communications chip manufacturers and NEMS all want a fair share of the emerging 400GE market. The key takeaways for understanding the state of 400GE today are summarized below:

- 400GE is undergoing rapid product development right now
- Early interoperability tests are occurring in late 2017-1H-2018
- Performance and benchmark testing will occur throughout 2018 and 2019
- New high-speed serializer/deserializer (SerDes) devices with 56Gb/s per lane signaling will implement full-feature sets in 2018
- Data center applications are largely driving the need for new technology to meet fast-emerging application requirements

400GE BRINGS BIG CHALLENGES AND NEW TECHNOLOGIES

400GE has brought several new items to the formula for successful Ethernet packet transmission and reception. Notice that “error-free” is not a phrase used when discussing successful Ethernet packet transmission and reception. What the industry must accept is that 400GE has inherent errors, enough to warrant the need for a highly complex, yet efficient, error-correction mechanism that is now a permanent part of the 200GE and 400GE link.

For the first time in the evolution of Ethernet, forward error correction (FEC) is mandatory when deploying 200GE and 400GE.

FORWARD ERROR CORRECTION

For the first time in the evolution of Ethernet, forward error correction (FEC) is mandatory when deploying 200GE and 400GE (i.e., 200GBASE-R and 400GBASE-R). FEC is defined in many different forms and is used in a variety of applications such as radio transmission, disk storage (e.g., CD, Blue-Ray), high-speed computer memory I/O, and in various forms of data communication transmission applications (e.g., xDSL, WiMAx, Ethernet, OTN).

FEC has been used in networking standards such as synchronous optical networking (SONET) and optical transport network (OTN), which use stronger FEC than the one called out in the IEEE 802.3bs 400GE draft standard. FEC was introduced into 10GE and 100GE initially for backplane uses—then for a few of the front-panel use cases, such as copper cabling for 100GE, 50GE, and 25GE over the QSFP28 physical interface, to deliver more economical optical transceiver and cable technologies.

FEC adds a pre-determined number of redundant bits into a data transmission that are error-checking bits (encoding these with the data). The error-checking bits are then used by the receiver of the data transmission to decode and correct errored-bits. The purpose of FEC is to provide a way to send/receive data in extremely noisy signaling environments. The benefit is error-free data transmissions. However, the penalty is carrying the additional bits of overhead for the FEC mechanism to encode and transmit, and the time required to decode and correct the errors on the receive side. The overall plus for Ethernet is a reduction in data packet re-transmissions.

Let’s look at why FEC is required for 200GE and 400GE communication systems. The speed of the signaling for the electrical interfaces that host the subsequent optical transceivers and cables used for high-speed transmission are fundamentally noisy signaling environments. Both forms for electrical interfaces defined by the IEEE, either the 16 electrical lanes of 28Gb/s per lane signaling with non-return-to-zero (NRZ) encoding, or the newer 4- or 8-lanes of 56Gb/s per lane signaling with 4-level pulse amplitude modulation (PAM-4) require FEC. In other words, the bit error rate (BER) that these electronics generate on their own require an error-correction method.

The IEEE802.3bs Task Force selected an error-correction transmission encoding and reception decoding method know as Reed-Solomon FEC (RS-FEC). The well-known and efficient RS-FEC uses an approach where it operates on a finite-sized block of bits known as block code. In Ethernet, this is called a codeword. RS-FEC is a cyclic type of FEC method. However, it works in a linear code (fixed block of bits) fashion, using a concept of redundant checking symbols and messages. For 200GE and 400GE applications, RS-FEC is also known as KP4 FEC and RS 544 FEC. A FEC symbol is 10-bits in size. Symbols in RS 544 FEC are critical, so let's review this in a bit more detail.

RS-FEC (544, 514) Symbols Break Down
544 = the total number of symbols in a codeword; 544 x 10-bit symbols = 5,440 bits in the codeword or block
544 - 514 = 30 symbols, which is the number of check symbols per codeword or block
514 x 10-bit symbol = 5,140 bits in total, which is the size of the information bits per block
(544 - 514)/2 = 15 symbols, which is the maximum number of symbols that can be corrected in a codeword or block

Figure 6. Understanding how RS 544 FEC symbols are applied sheds light on how 400GE is transmitted.

TRANSMIT TRANSCODING AND RECEIVE DECODING

In addition to RS 544 FEC, there are other new processes that make up a 400GE transmission and reception. Transmit transcoding and receive decoding are new additions to Ethernet. These take the 66-bit encoded data blocks and compresses them to 257-bit blocks to reduce line encoding overhead (i.e., better efficiency). The concept of interleaving FEC codewords using an A/B interleave helps to improve the FEC error-correction in the presence of naturally bursty Ethernet traffic. Shown separately below are the transmit and the receive processes for 400GE. FEC is shown where it fits into the entire communication process.

Transmit transcoding and receive decoding are new additions to Ethernet.

TRANSMIT PATH

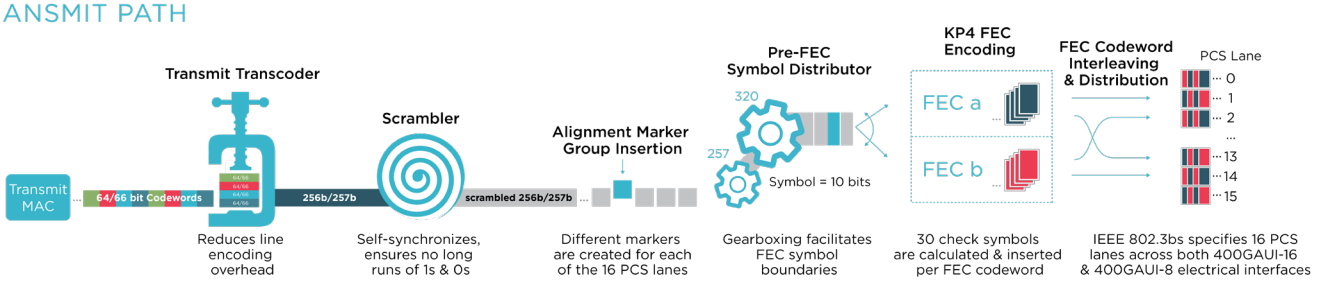

Figure 7. Transmit path processes for a 400GE link

RECEIVE PATH

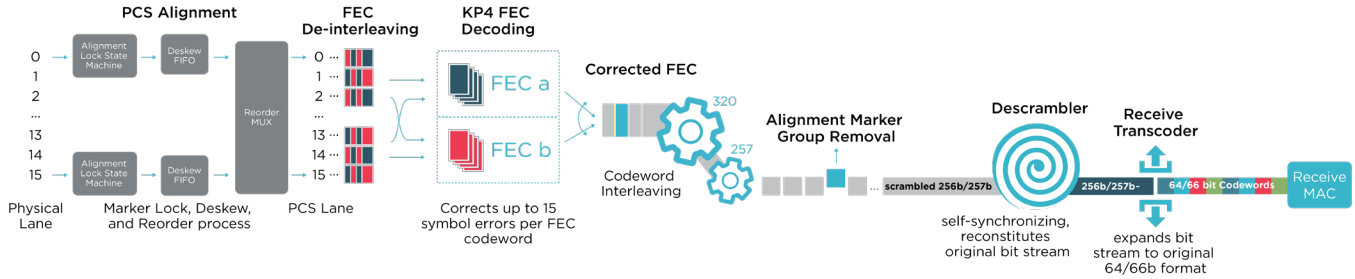

Figure 8. Receive path processes for a 400GE link

For a 200GE or 400GE data transmission, if the number of corrected symbols exceeds 15, then the link goes into an error condition. This is called an uncorrectable codeword or block. Once there are too many uncorrectable symbols in a codeword, it means the Ethernet link now exhibits errors that are shown in the physical coding sublayer (PCS). There is now potential for loss of synchronization of alignment markers in the encoded data blocks. This may lead to a link failure. The FEC symbol errors are shown over one or more PCS lanes and may include one or more actual data packets. The FEC codeword, due to its size, can contain data information bits from smaller packets, and/or from fragments of larger-sized packets. In the presence of uncorrectable FEC symbol errors, there may be one or more dropped packets.

The FEC symbol errors on the PCS lanes are the naturally occurring errors shown in the PCS that are caused by the electrical lane signaling discussed previously. It is important at some point in evaluating a 200GE or 400GE link, to know what causes the errors.

It is important at some point in evaluating a 200GE or 400GE link, to know what causes the errors—whether FEC symbol errors, and/or the naturally occurring PCS lanes errors from the port electronics, optical transceivers, and/or cables.

Since FEC is always on, the PCS is always running, and there can be errors in the FEC transmission and uncorrectable FEC errors in reception—it begins to become a bit complicated to evaluate the health of two link partners. The full view of 200/400GE transmission errors only comes by understanding:

- Pre-FEC BER
- Mapping of errors across PCS lanes:
  - Symbol error distribution and density
  - Number of uncorrectable code words
  - Frame loss ratio
- Optical transceivers and/or cables errors

Pre-FEC BER, FEC symbol error density (i.e., the number of symbol errors per code word), the number of uncorrectable FEC codewords, and the frame loss ratio are key metrics in determining the health of the link.

- **Pre-FEC BER—**A measure of the bit error rate before there is any error correction performed by the FEC checking symbols. It shows the errors between the transmitter of one host to the receiver of the other host.

- **Number of uncorrectable FEC codewords—**Once a FEC codeword has >15 symbol errors, RS 544 FEC can no longer correct these errors and the codeword becomes an uncorrectable FEC codeword, and its contents are dropped by the Rx MAC. When there are 3 uncorrectable codewords in a row, the link is dropped. Various PCS-layer processes such as alignment marker lock and PCS synchronization will restart. Once all the conditions are met, the link is re-established.
- **Frame loss ratio (FLR)—**Equivalent 64B packets lost due to uncorrectable FEC codewords over the total number of equivalent 64B packets received. As frames are lost the FLR will become > than zero.

Both 200GE and 400GE present new challenges, compared to traditional Ethernet, in determining if a design or implementation is error-free enough to maintain error-free data transmissions.

NEW OPTICAL TRANSCEIVERS

Both 200 and 400GE bring new optical transceiver physical medium dependent (PMD) architectures into the market. Implementation of these optics have been through multi-source agreements (MSAs), where optics industry leaders define specific form factors, and management and controls specifications. The form factors at present are quad small form-factor pluggable double density (QSFP-DD), octal small format pluggable (OSFP), and 100G (C) form-factor pluggable 8 (CFP8) where each implements the majority of the architectures shown in Figure 9.

Both 200 and 400GE bring new optical transceiver physical medium dependent (PMD) architectures into the market.

Physical Medium Dependent	Pluggable Transceiver Standalone	Host Electrical I/F	Mode	No. of Fibers	Reach Range	Encoding Method
400GBASE-SR16	Yes	16x25Gb/s	Multimode	16	100 meters	NRZ
400GBASE-DR4	Yes	8x50Gb/s	Single mode	4	500 meters	PAM-4
400GBASE-FR8	Yes	8x50Gb/s, 16x25Gb/s	Single mode	8 WDM	2 kilometers	PAM-4
400GBASE-LR8	Yes	8x50Gb/s, 16x25Gb/s	Single mode	8 WDM	10 kilometers	PAM-4
200GBASE-DR4:	Yes	8x50Gb/s	Single mode	4	500 meters	PAM-4
200GBASE-FR4	Yes	8x50Gb/s	Single mode	4 WDM	2 kilometers	PAM-4
200GBASE-LR4	Yes	8x50Gb/s	Single mode	4 WDM	10 kilometers	PAM-4

Figure 9. New optical transceiver PMD architectures

Links to specific MSAs for 400GE optical transceiver information:

- http://www.cfp-msa.org/documents.html
- http://www.qsfp-dd.com/
- http://osfpmsa.org/

The adoption rate of 400GE is expected to be at least 2x faster than was the adoption of 100GE. The diversity of implementation and test challenges, plus the wide range of new optical form factors provides an interesting environment for innovation.

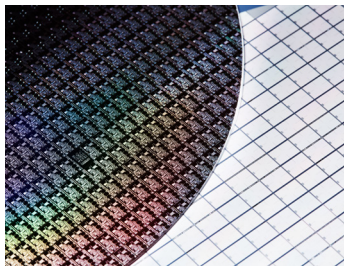

The real shakeup is happening in what was once an oligopoly—the silicon market.

MORE SILICON VENDORS AND MORE DEVICE DIVERSITY

The race to the market by NEMs and optical transceiver manufacturers is fast and furious. But the real shakeup is happening in what was once an oligopoly—the silicon market. New silicon communications chip companies have entered the market, also working at breakneck speeds to bring their 400GE product to market. Let's look at what's happening in the chip market to see what's in store for 200/400GE performance testing.

The companies that must respond first to pressing market demands like more bandwidth are the silicon chip providers. These are the companies that provide the physical-layer silicon chips that go into networking devices. These chips are the underpinnings that support the building of more Ethernet speeds, high-density switches, routers, NICs, servers, and many other networking devices.

As the industry prepares for early deployment of 400GE with high port counts per silicon device, there are several factors that are different compared to the early days of 100GE deployment.

MORE SILICON IMPLEMENTATION OPTIONS

Fundamental approaches on how NEMs implement the 400GE physical-layer silicon for new, high-speed Ethernet network devices include:

- **Merchant Silicon—**A NEM purchases the physical-layer silicon chip and any required software and intellectual property from another company to build their network equipment. The company will still have to integrate the chip with their software and hardware electronics.
- **In-House ASIC (application-specific integrated circuit)—**A NEM builds its own physical-layer silicon chip. Creates its own intellectual property, and writes all required software. Often these companies will collaborate, mixing their own developed portions with a series of vendors for silicon chip intellectual property (IP), and required software and operating systems.
- **Merchant and In-House ASIC—**A NEM may choose to buy merchant silicon for some products, and develop their own ASIC for other products.
- **OEM/ODM (original equipment manufacturer / original design manufacturer)—**A NEM will contract an OEM/ODM to build the network device for them. It will be branded and resold as the company products. The OEM/ODM typically buys merchant silicon that meets the need of the purchaser, and does all the development and often the testing. Some of the OEM / ODMs build their own products and sell them as well.

Although this is a simplified explanation of how NEMs build networking devices, it brings to the forefront a subject of critical importance. With the different combinations and ways of building new, high-speed 400/200/100/50/25GE networking devices, where everyone is diligently striving to meet the industry standards, it is inevitable that all the devices will not be built the same way. Nor will they perform to the same level, even though they may support the same Ethernet industry specifications, speeds, protocols, and feature sets. There are too many variables.

MERCHANT SILICON MARKET EXPANDS

There are two notable trends happening in the switch chip market.

1. In 2016 and 2017, new companies have arisen that are challenging the established market leaders in Ethernet silicon. Rather than one or two primary sources for physical-layer silicon chips, we are seeing more options and vendors.
2. NEMs are using more merchant silicon in their products as opposed to building switch ASICs in-house.

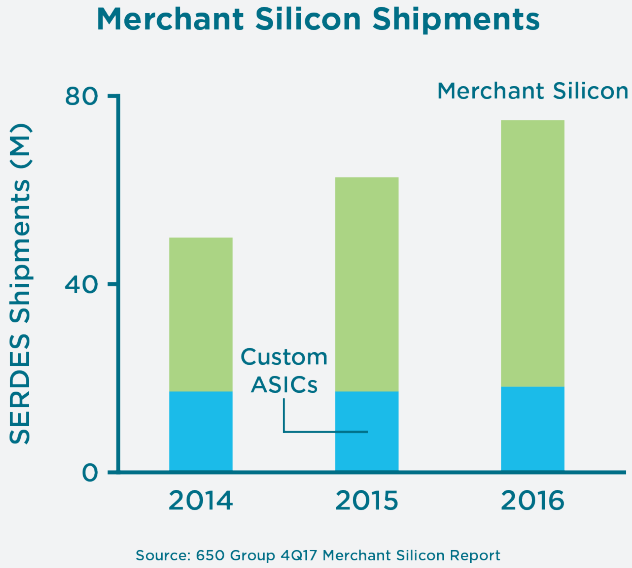

Figure 10. Merchant silicon shipments increasing (source: 650 Group, 4Q17 Merchant Silicon Report)

We know there will still be ASICs developed in-house. However, there are now multiple and different solutions available from the merchant silicon vendor base. For the companies like Apple, Amazon, Facebook, Google, and Microsoft that are driving innovation, merchant silicon is now a good option as it allows companies to select the appropriate merchant silicon that meet their specific application. Additionally, the timelines will be different for both release of merchant silicon solutions and their manufacturing customers to test and deploy them. This will create a great diversity of switch products, features, and capabilities spread out over the next few years.

Once NEMs, ODMs, and white-box manufacturers start to deploy the different switch chip solutions, it is inevitable that there will be differing features and performance levels. In the coming multi-vendor, best-of-breed 400GE marketplace, final end-users of the switches will require independent performance and benchmark testing to validate the equipment before committing capital expenditure (CAPEX).

MERCHANT SILICON DIVERSITY

As mentioned, the merchant silicon chip vendor field has increased by a few new players. Next, a review of their products will reveal their diverse nature. This diversity is exactly why the industry must return to performance benchmark testing.

A snapshot of the leading merchant silicon vendors is shown (Figure 11) to demonstrate the diversity that is about to hit the switch market. Each of the companies below have product, feature, and capability differences. The approaches to the designs are different. The switch products that will be produced from these solutions will be different and require characterization (i.e. performance testing). Just the differences between the programmability of these devices and how it relates to the underlying silicon will make for a broad range of switch solutions that must be validated.

In the coming multi-vendor, best-of-breed 400GE marketplace, final end-users of the switches will require independent performance and benchmark testing to validate the equipment before committing CAPEX.

Silicon Chip Vendor	Vendor A	Vendor B	Vendor C	Vendor D
Target Markets	Data Center, Cloud, Hyperscale, Big Data, AI, Financial, Storage	Data Center, Cloud, Hyperscale, HPC, Service Provider Data Center	Data Center, Cloud, Hyperscale	Data Center, Cloud, Hyperscale
Max. Tb/s*	6.4	12.8	6.5	6.4
Silicon Technology	16nm	Public info not available	Public info not available	16nm
SERDES	• 128 x 50Gb/s PAM-4 • 25Gb/s NRZ	• 50Gb/s PAM-4 • 25Gb/s NRZ • 10Gb/s	• 50Gb/s PAM4 – info not available in public • 25Gb/s NRZ • 10Gb/s	• 50Gb/s PAM4 – info not available in public • 256 x 25Gb/s NRZ
Forwarding * (latency)	<300ns Cut-through Store and Forward	<350ns Cut-through Store and Forward	Public info not available	<400ns Cut-through Store and Forward
Packet* Buffer	42MB shared	70MB shared	Public info not available	Public info not available

Silicon Chip Vendor	Vendor A	Vendor B	Vendor C	Vendor D
Port density & Speeds	• 16 x 400 • 32 x 200 • 64 x 100 • 128x50/25/10	• 32 x 400 • 64 x 200 • 128x100 (fan-out) • 50/25 support with 6.4Tb/s option	• 400GE – Public info not available • 200GE – Public info not available • 65 x 100GE • 65 x 40GE • 130 x 50GE • 130 x 40GE (via gearbox) • 260 x 25GE • 260 x 10GE	• 400GE – Public info not available • 200GE – Public info not available • 64x100GE • 128x40/50GE • 128x25/10GE
On Chip CPU	GHz ARM CPU	Public info not available	Public info not available	Public info not available
Flow Management* (mice / elephant)	Line rate capable, Dynamic Flow and queue controls	Line rate	Public info not available	Yes, configurable and dynamic
Congestion Detection/ Management*	Yes	Yes, QoS, spec not available	Public info not available	Yes
Forwarding Table*	512K	Yes, spec not available	0.3-1.2M	Yes, spec not available
ACL	Yes	Public info not available	Yes	Yes, spec not available
Load Balancing*	Dynamic and Static	Yes	Yes	Yes
L2/3 forwarding, tunneling*	VXLAN, NVGRE, Geneve, and MPLS	MPLS, VXLAN, Geneve	Public info not available	VXLAN, NVGRE, MPLS, SPB
Open Source Programmable	Custom protocols, tunnels, packet editor	Yes, details not available	P4, a high degree of programmability	No
Telemetry	Yes	Yes	Yes	Yes
Native O.S.	Yes	Yes	Open source	Yes
OCP support	Yes, SAI	Yes, SAI	Yes, SAI	Not available
NOS	Info not available	Yes	Public info not available	Not available
Linux	Yes	Public info not available	Public info not available	Public info not available

*Note: These features require performance benchmarking. Some are used for the application and RFC benchmark test evaluations shown in the following tables.

Figure 11. Published data from four silicon chip vendors shows a sample of the market diversity.

With 400GE, we are seeing more more silicon makers and more diversity in options and implementations. This begs the question, how do the purchasers of the 400GE equipment know that the networking devices they are buying meet their published specifications? There is only one answer, "They must test the product." How the product is tested is critical. How networking products from different vendors are compared by testing to well-established benchmark and performance test methodologies matters; it matters greatly.

How networking products from different vendors are compared by testing to well-established benchmark and performance test methodologies matters; it matters greatly.

THE NEED FOR PERFORMANCE BENCHMARKING

A LESSON FROM HISTORY

Since the adoption of 100GE in 2014/2015 and the rapid appearance of 25GE and 50GE in 2016/2017, there has been a marked drop in full-port-count performance benchmark testing. Why is that? What happens when this occurs? Let's look at an example.

A study compared two mature 100GE top of tack (ToR) switch products targeted for the data center. Both switches used a different silicon switch chip. One company built their own ASIC for their switch, the other company used merchant silicon for their switch. Both companies are established in the market. These products were performance benchmark tested by a well-known, independent test organization. Both devices were tested with the same test tool and application software. The independent test vendor used standard industry benchmark and performance tests developed by the Internet Engineering Task Force (IETF) in the late 1990's. These tests are known as RFC2544, RFC 2889, and RFC 3819. All of these tests run very specific sets of traffic patterns, packet sizes, rates, and port topologies. The goal is to exercise and stress the device to find its level of performance using a consistent set of test methods so that products will be compared in a fair manner.

A summary of a few of the measurements shows what happened in this test comparison at 100GE speed:

Test Port Configuration	RFC	Desired Result	Switch w/ Silicon Vendor A	Switch w/ Silicon Vendor B	Comment
Many-to-One (3:1)	2544	Even distribution of traffic forwarding on the congested port.	Even forwarding distribution: 33/33/33%. This passed the test.	Unbalanced forwarding distribution: 25/25/50%. This does not pass the test.	What happens in a real network when unbalanced forwarding occurs? End-user performance is negatively affected. An SLA may not be met. This test was run again with more ports and traffic streams and the results were worse. Further tests using a IMIX traffic pattern yielded similar uneven results.
Frame Loss 64B-9216B frame lengths with L2, L3 traffic in full mesh	2544 2889	Switch should forward all traffic at 100% line rate with no frame loss over all frame lengths in the test.	Passed all frames at 100% line rate.	Showed loss between 0.33% and 30.6% for multiple frame lengths.	Frame loss is a direct drop in switch forwarding performance. This affects all end-users.

Test Port Configuration	RFC	Desired Result	Switch w/ Silicon Vendor A	Switch w/ Silicon Vendor B	Comment
Latency at 100GE over 32-ports in full mesh with L2 traffic	2889	Latency should be consistent across all ports on the switch. Some variation acceptable.	Latency was consistent with a range of 50-60ns across all ports.	Latency was dramatically inconsistent across all ports with a very large range reported.	Latency performance is critical in a large data center as it contributes to higher packet forwarding throughput.

Figure 12. Test comparison of two 100GE ToR switches with two different silicon chips.

The switch with silicon vendor A passed the industry specified tests. But the switch with silicon vendor B fell quite short. This was a surprising finding in a market with just two primary chip vendors and a fairly mature technology, after the standards were ratified.

But with 400GE, our networking world is taking a major leap that will result not just in more bandwidth, but in challenges to deliver reliable and predictable bandwidth. Consider the diversity of silicon chip vendors, the complexity of the features required to run 400GE/200GE such as FEC and its overhead load, and higher-speed SERDES such as 56Gb/s per lane. And soon 112Gb/s per lane SERDES will be introduced to the market, bringing a whole new set of challenges. What will happen if our industry does not perform benchmark tests for 400GE and 200GE? Will the final end-users of the devices be served well? Undoubtedly, not as well as required.

TESTING IS KEY TO OVERCOME 400GE CHALLENGES

400GE brings forward a whole new set of challenges for Ethernet and they all must be scrutinized before products go to market:

1. With new chip technologies available from multiple vendors, how does the end-user know they are making the right choice in unproven ground.
2. 56Gb/s per lane SERDES with PAM-4 encoding is also new territory. There is not much manufacturing and fabrication history on this technology over large form factor PCBs. Reliability and stress testing will be part of the maturation of this technology.
3. What is the impact of FEC and its overhead on overall switch performance and how does it affect the end-user's applications and network designs?
4. A multitude of vendors have developed FEC engines. How do they interoperate? How many errors does it take to stress these engines? How reliable are these mechanism in stress test?
5. What happens when uncorrectable FEC errors start affecting frame loss and BERs become questionable? How does this affect performance?
6. What is the difference between pre-FEC and post-FEC BER? Why does this matter?

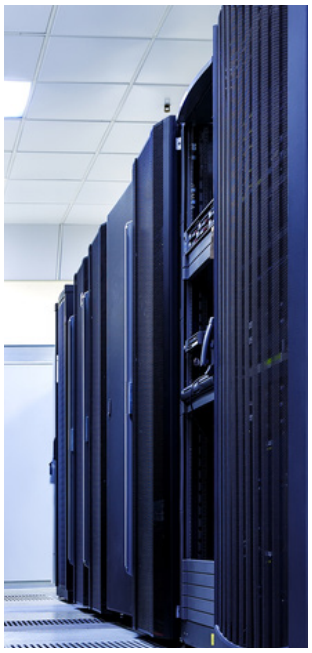

400GE brings forward a whole new set of challenges for Ethernet and they all must be scrutinized before products go to market.

7. There are new elements to the PCS, every new implementation must be evaluated and tested.
8. 400GE brings a broad range of new optical and copper form factors, reach ranges, and varying levels of management, control, and monitoring capabilities. There will be many challenges with integration, and multivendor switch interoperability.

The table below shows which performance tests are most critical for foundational 400GE silicon chip features and capabilities, as they apply to the major applications for 400/200GE switching product implementations. The final end-customer in each of these major application areas will fully expect that the switch they want to deploy will provide full line rate 400Gb/s (about 600M 64B packets per second), non-blocking performance on all ports under all conditions.

Knowing the throughput rating of a switch has major impact. How well a switch performs under full line rate, engaged at full port count under stressed conditions has impact to network designers. They must plan for the total bandwidth requirements in their data center. In leaf and spine and in CLOS network designs, the performance of a switch matters. For every L2- or L3-aware switch that does not perform at full line rate with 100% throughput (i.e., fully non-blocking), additional switches must be added into the network to meet the total bandwidth requirements. And, when the massive scale of today's data center is considered, for every 4 switches that are deployed, each with 80% throughput, an additional switch must be added to meet real 100% bandwidth requirements. Scale that scenario up and it adds CAPEX and OPEX costs to the network provider.

Figure 13 shows the critical switch tests that must be run to benchmark a switch when used in the major application areas that are anticipated for 400GE deployment.

For every L2- or L3-aware switch that does not perform at full line rate with 100% throughput, additional switches must be added into the network to meet the total bandwidth requirements.

Figure 13. The criticality of chip feature performance depends on the targeted end-use-case of each vendor's switch.

400/200GE Application Areas	Port Density/ Speeds (Bandwidth)	Forwarding (Throughput)	Flow Management (Distribution)	Congestion Detection / Management (Oversubscription)	Forwarding Table (Route Capacity)	Load Balancing (Distribution)	L2/3 Forwarding, Tunneling* (Protocol Performance)
Cloud	✓	✓	✓	✓	✓	✓	✓
Hyperscale Data Center	✓	✓	✓	✓	✓	✓	✓
Large Enterprise Data Center	✓	✓		✓	✓	✓	
Big Data		✓		✓	✓	✓	✓
Artificial Intelligence				✓		✓	✓
Financial	✓	✓		✓	✓		✓
Storage	✓	✓		✓	✓	✓	

400GE BENCHMARK AND PERFORMANCE TESTING

Industry-standard RFC benchmark and performance tests provide a repeatable methodology to characterize a switch. The most widely used are RFC 2544 and RFC 2889.

Here is a summary of what Layer 2 and 3 RFC tests measure that will apply to 400GE switches.

400GE Switch Feature / Capability	RFC 2544 Test Required	RFC 2889 Tests Required
Port Density and Speeds	• Throughput • Latency • Frame loss rate • For all ports in a single test	• Fully meshed throughput, frame loss, and forwarding rates • Partially meshed one-to-many/many-to-one • Partially meshed unidirectional traffic • Congestion control • Forward pressure and maximum forwarding rate
Forwarding (Latency)	• Throughput • Latency • Frame loss rate • Back-to-back	• Fully meshed throughput, frame loss, and forwarding rates • Partially meshed one-to-many/many-to-one • Partially meshed unidirectional traffic • Congestion control • Forward pressure and maximum forwarding rate
Packet Buffer	• Throughput • Latency	• Fully meshed throughput, frame loss, and forwarding rates • Forward pressure and maximum forwarding rate
Flow Management (Mice / Elephant)	• Throughput • Latency	• Fully meshed throughput, frame loss, and forwarding rates • Partially meshed one-to-many/many-to-one • Partially meshed unidirectional traffic • Congestion control
Congestion Detection/ Management	N/A	• Congestion control • Forward pressure and maximum forwarding rate
Forwarding Table	• Throughput (Layer 3 test) • Requires large-scale internetworking L3 running stateful protocols to execute the test	N/A
Load Balancing	• Throughput • Latency • Frame loss rate	• Fully meshed throughput, frame loss, and forwarding rates • Partially meshed one-to-many/many-to-one • Partially meshed unidirectional traffic
L2/3 Forwarding, Tunneling	• Throughput (Layer 3 test) • Requires large-scale internetworking L3 running stateful protocols to execute the test	• Fully meshed throughput, frame loss, and forwarding rates

Figure 14. Required industry tests to validate 400GE technologies

RFC PERFORMANCE MEASUREMENT METRICS

RFC 2544	
Throughput	Finds the maximum traffic rate at which no frames are dropped by the switch over a standard set of frame sizes and traffic loads.
Latency	Measures the amount of time for the switch to forward frames: • Store-and-forward - Last bit of input to first bit of output • Cut-through - First bit of input to first bit of output. Typically, based on a threshold of bits received in a packet, or when a certain part of the packet is detected and forwarded.
Frame Loss Rate	Determines the ability to forward frames at different traffic rates. It is the percentage of frames NOT forwarded when the switch is under traffic load.
Back-to-Back	Determines the ability of the switch to forward frames with minimal inter-frame gap size of traffic burst without any frame loss.

Figure 15. RFC 2544 Benchmarking Methodology for Network Interconnect Devices.

RFC 2889	
Fully Meshed Throughput, Frame Loss, and Forwarding Rates	Measures the Layer 2 throughput, frame loss, and forwarding rates when all ports in a switch are forwarding to all the other ports. Measures the maximum capacity over standard frame sizes the switch can handle at full 100% line rate.
Partially Meshed One-to-Many/ Many-to-One	Determines the throughput when transmitting from/to multiple ports and to/from one port. Measures the capability of the switch to switch frames without frame loss at full line rate. It measures when a single port becomes congested and oversubscribed.
Partially Meshed Unidirectional Traffic	Determines the throughput of the DUT/SUT when presented multiple streams of unidirectional traffic with half of the ports on the DUT/SUT transmitting frames destined to the other half of the ports. Measures frame loss and flood count. Looks for uneven distribution of traffic, throughput, and forwarding rate under full line rate load.
Congestion Control	Measures how a switch handles congestion. Determines if the device implements congestion control and if congestion on one port affects an uncongested port. This procedure determines if head-of-line-blocking and/or backpressure are present. Useful to know how even the distribution of traffic is across all ports of the switch and how much frame loss is occurring and flooding. It also measures flow control. A non-blocking switch must be able to pass this test at full line rate for all packet sizes.

Figure 16. RFC 2889 Benchmarking Methodology for LAN Switching Devices.

RFC 2889	
Forward Pressure and Maximum Forwarding Rate	Measures the peak value of the forwarding rate right at the point where the switch can no longer forward packets at full line rate.
Address Caching Capacity	Measures the point where the offered addresses per port were successfully forwarded without flooding.
Address Learning Rate	Determines the rate of MAC address learning at L2. Measures whether the offered addresses per port were successfully forwarded without flooding at the offered learning rate.
Errored Frames Filtering	Determines the behavior of the switch under error or abnormal frame conditions. Indicates if the switch filters the errors, or simply propagates the errored frames along to the destination. Frames that are undersize, oversize, and that have CRC errors must be handled correctly by the switch.
Broadcast Frame Forwarding and Latency	Determines the throughput and latency of the DUT when forwarding broadcast traffic. Measures the frame loss rate and the forwarding rate.

SUMMARY

Today's race to higher bandwidth is no longer about the core network. It is about data center and server I/O computing that has a 4x larger growth rate than the previous traditional non-data-center IP traffic. With 200/400GE technology the best option to meet the demand, the once limited-supplier market is now flooding with new communications chip manufacturers. At the same time, 400GE by its very nature, introduces new challenges—increased bit error rate, mandatory FEC, and the performance of the PCS in the presence of errors that 56Gb/s SERDES and PAM-4 encoding inherently provide.

At Layer 2 and 3, the industry is going to see a variety of 400/200/100/50GE solutions facilitated by many different merchant silicon providers—all with different designs and implementations. The switches that will be built with merchant silicon will have to interoperate with the ASICs that the larger NEMs will design and develop in-house. If the switch industry is to address the explosive bandwidth demand, then high-density 400GE switches will have to perform well to meet these requirements.

400GE is destined to be adopted by the industry at least 2x faster than was 100GE. With all the new optics and cables, switches, and servers driven by 400/200/100/50/25GE speeds, the need for testing overall has never been greater. Testing not just by chip and equipment makers, but by independent benchmark and performance test organizations. The end users—data centers, service providers, and network operators—need empirical data on which equipment will satisfy their particular needs.

Fortunately, industry-standard benchmarks are available for the high-speed solutions about to come to market. Early 400GE interoperability testing in 2016-2017 has also relied on a few L2/3 test vendors that have products capable of such testing. So, look to those that have proven success in validating 400GE.

This is a critical time to renew efforts to compare and benchmark the performance of these new network products and devices. Welcome to the new era of industry benchmarking and performance testing!

This is a critical time to renew efforts to compare and benchmark the performance of these new network products and devices.

IXIA WORLDWIDE

26601 W. Agoura Road
Calabasas, CA 91302

(Toll Free North America)
1.877.367.4942

(Outside North America)
+1.818.871.1800

(Fax) 1.818.871.1805

www.ixiacom.com

IXIA EUROPE

Clarion House, Norreys Drive
Maidenhead SL64FL
United Kingdom

Sales +44.1628.408750
(Fax) +44.1628.639916

IXIA ASIA PACIFIC

101 Thomson Road,
#29-04/05 United Square,
Singapore 307591

Sales +65.6332.0125
(Fax) +65.6332.0127

K400 QSFP-DD 400GE LOAD MODULES

DATA SHEET

PUT YOUR 400GE PRODUCT TO WORK NOW!

400GE is undergoing rapid development to meet the growing bandwidth requirements of ever-evolving data networks. Core networking applications are increasing by a factor of 10 every 5 years. 400GE addresses the broad range of bandwidth requirements for key application areas such as cloud-scale data centers, Internet exchanges, co-location services, wireless infrastructure, service provider and operator networks, and video distribution infrastructure.

TEST TOMORROW'S TECHNOLOGY TODAY

Ixia's K400 QSFP-DD-400GE load module will accelerate your development of IEEE 802.3bs-compliant 400GE networking systems. Giving development teams the options they need to create the networking technologies of the future, these load modules are available in two different models:

- Full feature and scale QSFP-DD-400GE
- Reduced feature and scale QSFP-DD-R400GE

Both models enable full line-rate traffic generation transmit, receive, and capture functionality from 64B to 16,000B frame lengths. This facilitates stress testing, hardware/ASIC bring-up, optics and cable qualification, interoperability, and functional test and Layer 2/3 routing protocol emulation as required.

K400 QSFP-DD-400GE: A full-featured model with 1-port, 1-slot, native QSFP-DD 400GE load module designed for enterprise and data center switch and router testing.

K400 QSFP-DD-R400GE: A reduced-featured model with 1-port, 1-slot, native QSFP-DD 400GE load module designed for hardware, ASIC, cable/optics qualification, and interoperability testing. The QSFP-DD-R400GE scales down the L2/3 feature set and L2/3 networking protocol scaling, while increasing affordability.

HIGHLIGHTS

- World's first 400GE native QSFP-DD test system from the leader in Ethernet test
- World's first 400GE test system running IEEE 802.3bs RS-544 Forward Error Correction (FEC) – Special Award Winner for Testing in 2016 at Interop Tokyo, Japan
- World's first live 400GE line rate interoperability demonstration with FEC with two independent MAC/PCS/FEC implementations, in 2016 at ECOC in Dusseldorf, Germany
- Ixia developed the intellectual property for the critical elements of 400GE: MAC, PCS, FEC, and FEC error injection and statistics — resulting in faster response times as industry standards evolve
- FEC error injection, PCS lanes Tx/Rx test capability, and classical L1 BERT support provide essential tools for hardware development and bring-up
- The hardware is capable of 200GE, 100GE, and 50GE fan-outs as lower-speed testing evolves

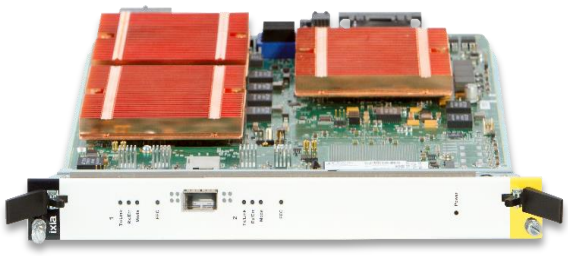

944-1152 QSFP-DD-400GE, 1-port, 1-slot, load module

The QSFP-DD-R400GE load module may be upgraded in the field to have a higher-capacity L2/3 feature set and L2/3 networking protocol scaling performance through a purchasable upgrade option. This extends the reuse of the load module and improves the ROI.

K400 KEY FEATURES

- Line-rate 400Gbps packet generation, capture, and analysis of received traffic to detect and de-bug data transmission errors
- Line-rate per-port and per-flow statistics
- High latency measurement resolution at 0.625ns
- RS-544 (KP4) Forward Error Correction (FEC) support
- Native 56Gb/s SERDES with PAM4 encoding support that is IEEE P802.3bs and IEEE P802.3cd compliant
- FEC error injection with a comprehensive set of FEC corrected and uncorrected statistics, including Bit Error Rate statistics for pre- and post-FEC operations
- Inject packet errors: CRCs, runts, giants, alignments, checksum errors, and out of sequence
- Standard Ixia instrumentation including timestamp, sequence number and flow identification, and data integrity
- Layer 1 BERT: 8 independent lanes of 56Gb/s PRBS pattern generation, error checking, and statistics
- 400G PCS lanes Transmit, error injection testing and receive measurement:
    - Per-lane controls and status, FEC and error monitoring, error insertion, lane mapping and skew insertion
- +/-100 PPM line frequency adjustment
- Mid-to-high-range L2/3 networking protocol emulation to validate performance and scalability of L2/3 routing/switching and data center test cases using the Ixia's IxNetwork application
- An excellent test platform for full line-rate 400Gb/s to evaluate the new 400GE ASIC designs, FPGAs, and hardware switch fabrics that use the new 8x56Gb/s electrical interface with PAM4 encoding
- Supports benchmarking of networking devices and equipment using industry-standard RFC benchmark tests at line-rate 400GE
- Supported by the XGS12 and XGS2 High Performance (HSL) and Standard Performance (SDL chassis models with the Native IxOS
- Application support including: IxExplorer, IxNetwork, and related Tcl and automation APIs

SPECIFICATIONS

PRODUCT DESCRIPTION	QSFP-DD-400GE FULL FEATURE	QSFP-DD-R400GE REDUCED FEATURE
Part Number	944-1152	944-1153
Hardware Load Module Specifications		
Slot / Number of Ports	1-slot, 1-port 400GE	
Physical Interfaces	Native QSFP-DD physical port	
Supported Port Speeds	400GE/port: 400GE-capable fiber and passive copper cable media 200/100/50GE port capable with a purchasable option	
CPU and Memory	Multicore processor with 4GB of CPU memory per port	
IEEE Interface Protocols for 400GE	IEEE P802.3bs 200GE & 400GE, 400GBASE-R IEEE P802.3cd 50 Gb/s, 100 Gb/s, and 200 Gb/s Ethernet	
Advanced Layer 1 support	400GE: • KP4 (RS-544) Ethernet Forward Error Correction, Clause 119 • Pre- and post-FEC statistics: Comprehensive per-port and per-lane statistics • FEC error injection • PCS lanes Tx and Rx test and statistics • Classical Layer 1 BERT test and statistics	
Transceiver Support	• Capable of support for 400GBASE-DR8, 400GBASE-FR8, and 400GBASE-LR8 when they are available in the QSFP-DD MSA compliant form factor. Other configurations may be supported—consult factory. • QSFP-DD MSA compliant optical transceivers are supported when the power consumption of the optical transceiver does not exceed 7 watts.	
Cable Media	400GBASE-CR8, passive, copper Direct Attached Cable (DAC) up to 2 meters in length	

PRODUCT DESCRIPTION	QSFP-DD-400GE FULL FEATURE	QSFP-DD-R400GE REDUCED FEATURE
Load Module Dimensions	• 16.4” (L) x 1.3” (W) x 12.0” (H) • 417mm (L) x 33mm (W) x 305mm (H)	
Load Module Weights	• Module only: 8.9 lbs. (4.04 kg) • Shipping: 15.7 lbs. (7.12 kg)	
Temperature	• Operating: 41°F to 95°F (5°C to 35°C) • Storage: 41°F to 122°F (5°C to 50°C)	
Humidity	• Operating: 0% to 85%, non-condensing • Storage: 0% to 85%, non-condensing	
Chassis Capacity: Maximum Number of Cards and Ports per Chassis Model		
XGS12-HSL Chassis (940-0016)	12 load modules: • Rackmount chassis • 12-ports of 400GE	
XGS2-HSL Chassis (940-0014)	2 load modules: • Rackmount chassis • 2-ports of 400GE	
XGS12-SDL Chassis (940-0015)	12 load modules: • 12-slot rackmount chassis • 12-ports of 400GE	
XGS2-SDL Chassis (940-0013)	12 load modules: • 2-slot rackmount chassis • 12-ports of 400GE	
Transmit Feature Specifications		
Transmit Engine	Wire-speed packet generation with timestamps, sequence numbers, data integrity, and packet group signatures	
Max. Streams per Port	400GE: 128	400GE: 32

PRODUCT DESCRIPTION	QSFP-DD-400GE FULL FEATURE	QSFP-DD-R400GE REDUCED FEATURE
Max. Streams per Port in Data Center Ethernet	400GE: 128	400GE: 32
Stream Controls	• Rate and frame size change on the fly • Advanced stream scheduler	
Minimum Frame Size	400GE: • 60 bytes at full line rate • 49 bytes at less than full line rate	
Maximum Frame Size	16,000 bytes	
Maximum Fame Size in Data Center Ethernet	9,216 bytes	
Priority Flow Control	• 4 line-rate-capable queues, each supporting up to 9,216-byte frame lengths • 1 line-rate-capable queue, non-blocking supporting up to 9,216-byte frame lengths	
Frame Length Controls	Fixed, increment by user-defined step, weighted pairs (up to 16K), uniform, repeatable random, IMIX, and Quad Gaussian	
User Defined Fields (UDF):	Fixed, increment or decrement by user-defined step, sequence, value list, and random configurations; up to 10, 32-bit-wide UDFs are available	
Value Lists (Max.)	400GE: 1M / UDF	
Sequence (Max.)	400GE: 32K / UDF	
Error Generation	• FEC symbol error-injection allows the user to inject FEC symbol errors using various weighted methods to achieve specific error rates • Generate good CRC or force bad CRC, undersize and oversize standard Ethernet frame lengths, and bad checksum	

PRODUCT DESCRIPTION	QSFP-DD-400GE FULL FEATURE	QSFP-DD-R400GE REDUCED FEATURE
Physical Coding Sublayer	• PCS lane skew injection • PCS lane re-mapping • PCS lane marker error injection • PCS bit error generation	
L1 BERT	Classical, line rate, un-encapsulated transmit and receive of various PRBS patterns, controls over the patterns that help to produce BER statistics	
Hardware Checksum Generation	Checksum generation for IPv4, IP over IP, ICMP/GRE/TCP/UDP, L2TP, GTP, and multilayer checksum. Support for protocol verification for control plane traffic	
Link Fault Signaling	Reports, no fault, remote fault, and local fault port statistics; generate local and remote faults with controls for the number of faults and order of faults, plus the ability to select the option to have the transmit port ignore link faults from a remote link partner	
Latency Measurement Resolution	400GE: 0.625 nanoseconds	
Intrinsic Latency Compensation	Removes inherent latency error from the 400GE port electronics	
Transmit Line Clock Adjustment	Ability to adjust the parts-per-million-line frequency over a range of -100 ppm to +100 ppm on the 400GE port	
Transmit/Receive loopback	Internal and line loopback support	
Receive Feature Specifications		
Receive Engine	Wire-speed packet filtering, capturing, real-time latency, and inter-arrival time for each packet group, with data integrity, sequence and advanced sequence checking capability	
Trackable Receive Flows per Port	400GE: • 32K with the full statistics	

PRODUCT DESCRIPTION	QSFP-DD-400GE FULL FEATURE	QSFP-DD-R400GE REDUCED FEATURE
Minimum Frame Size	400GE: • 64 bytes at full line rate • 49 bytes at less than full line rate	
Filters (User-Defined Statistics, UDS)	2 SA/DA pattern matchers, 2x16-byte user-definable patterns. 6 UDS counters are available with offsets for start of frame.	
Hardware Capture Buffer	256KB	
Standard Statistics and Rates	Link state, line speed, frames sent, valid frames received, bytes sent/received, fragments, undersize, oversize, CRC errors, 6 user-defined stats, capture trigger (UDS 3), capture filter (UDS 4), data integrity frames, data integrity errors, sequence and advanced sequence checking frames, sequence checking errors, ARP, and PING requests and replies	
FEC Statistics	• FEC port statistics: Total Bit Errors, Max Symbol Errors, Corrected Codewords, Total Codewords, Uncorrectable Codewords, Frame Loss Ratio, Pre-FEC Bit Error Rate, and Codeword error distribution analysis • FEC per lane Rx statistics: FEC Symbol Error Count, Corrected Bits Count, Symbol Error Rate, Corrected Bit Rate	
Latency / Jitter Measurements	Cut-through, store & forward, forwarding delay, up to 16 time bins latency/jitter, MEF jitter, and inter-arrival time	
Receive-side PCS Lanes Port Statistics Counters	PCS: Sync Errors, Illegal Codes, Remote Faults, Local Faults, Illegal Ordered Set, Illegal Idle, and Illegal SOF	
400GE Physical Coding Sublayer (PCS) Receive-Side Statistics and Indicators	IEEE 802.3bs-compliant per-lane PCS receive capabilities include: • Receive – per-lane PCS receive statistics; Physical Lane assignments, Lane Marker Lock, Lane Market Map, Relative Lane Skew, Lane Marker Error Count • Receive – per-lane FEC receive statistics; FEC Symbol Error Count, FEC Corrected Bits Count, FEC Symbol Error Rate, FEC Corrected Bit Rate	

PRODUCT DESCRIPTION	QSFP-DD-400GE FULL FEATURE	QSFP-DD-R400GE REDUCED FEATURE
Layer 2-3 Protocol Support		
Routing and Switching	BGP-4, BGP+, OSPFv2/v3, ISISv4/v6, EIGRP, EIGRPv6, RIP, RIPng, BFD, IGMPv1/v2/v3, MLDv1/v2, PIM-SM/SSM, PIM-BSR, STP/RSTP, MSTP, PVST+/RPVST+, Link Aggregation (LACP), LLDP	Complete protocol coverage with reduced session scale: • 100 routing & switching sessions • 2000 host/access sessions
Software Defined Network	VXLAN, EVPN VXLAN, OpenFlow, ISIS Segment Routing, OSPF Segment Routing, BGP Segment Routing, BGP Link State (BGP-LS), PCEP, OVSDB	Complete protocol coverage with reduced session scale: • 100 routing & switching sessions • 2000 host/access sessions
Basic	IPv4/IPv6, DHCP, PPPoE, L2TP, IGMP, MLD, 802.1x	
MPLS	RSVP-TE, RSVP-TE P2MP, LDP/LDPv6, mLDP, PWE, VPLS-LDP, VPLS-BGP, BGP auto-discovery with LDP FEC 129 support, L3 MPLS VPN/6VPE, 6PE, BGP RT-Constraint, BGP Labeled unicast, L3 Inter-AS VPN Options (A, B, C), MPLS-TP, MPLS OAM, Multicast VPN (GRE, mLDP, RSVP-TE P2MP), EVPN, PBB-EVPN	Complete protocol coverage with reduced session scale: • 100 routing & switching sessions • 2000 host/access sessions
Broadband and Authentication	PPPoX, DHCPv4, DHCPv6, L2TPv2, Radius attributes for L2TP, ANCP, IPv6 Autoconfiguration (SLAAC), IGMPv1/v2/v3, MLDv1/v2, 802.1x	Complete protocol coverage with reduced session scale: • 100 routing & switching sessions • 2000 host/access sessions
Industrial Ethernet	Link OAM IEEE 802.3ah, CFM IEEE 802.1ag, Service OAM ITUT-Y.1731, PBT/PBB-TE, Sync-E ESMC, PTP IEEE 1588 with G.8265.1 Telecom Profile, ELMI	Complete protocol coverage with reduced session scale: • 100 routing & switching sessions • 2000 host/access sessions

PRODUCT DESCRIPTION	QSFP-DD-400GE FULL FEATURE	QSFP-DD-R400GE REDUCED FEATURE
Data Center Ethernet	FCoE/FIP, Priority Flow Control IEEE 802.1Qbb (PFC), LLDP/DCBX	Complete protocol coverage with reduced session scale: • 100 routing & switching sessions • 2000 host/access sessions

APPLICATION SUPPORT

QSFP-DD-400GE / QSFP-DD-R400GE
• **IxExplorer:** Layer 1-3 wire-speed traffic generation, capture, and analysis with Forward Error Correction and error injection with statistics, PCS Lanes Tx/Rx with statistics, and Layer 1 BERT test and reporting capability. • **IxNetwork:** Wire-rate traffic generation with service modeling that builds realistic, dynamically controllable data-plane traffic. IxNetwork offers the industry's best test solution for functional and performance testing by using comprehensive emulation for routing, switching, MPLS, IP multicast, broadband, authentication, Carrier Ethernet, and data center Ethernet protocols. • **Tcl API:** Custom user script development for Layer 1-3 testing

ORDERING INFORMATION

LOAD MODULE

944-1152

K400 QSFP-DD-400GE 1-port, 1-slot QSFP-DD 400GE load module with the native QSFP-DD physical interface, L2-3 support (944-1152). Compatible with the XGS12-HSL rack mount chassis (940-0016), the 2-slot high performance XGS2-HSL chassis (940-0014), the XGS2-SDL, 2-slot Standard Performance chassis (940-0013), and the XGS12-SDL, 12-slot chassis (940-0015).

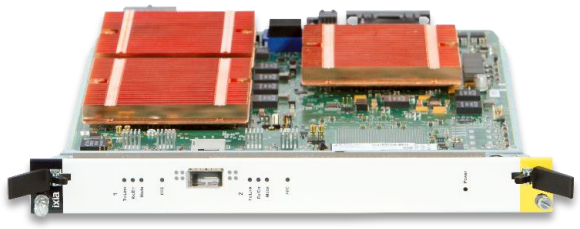

944-1153

K400 QSFP-DD-R400GE 1-port, 1-slot QSFP-DD 400GE reduced load module with the native QSFP-DD physical interface, L2-3 support and reduced protocol emulation scale (944-1153). Compatible with the XGS12-HSL rack mount chassis (940-0016), the 2-slot high performance XGS2-HSL chassis (940-0014), the XGS2-SDL, 2-slot Standard Performance chassis (940-0013), and the XGS12-SDL, 12-slot chassis (940-0015).

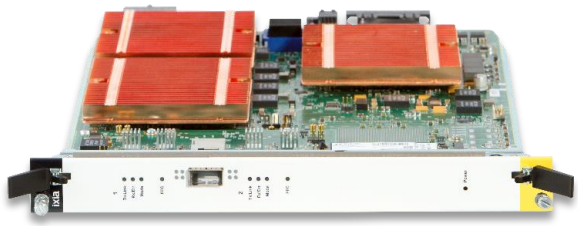

LOAD MODULE UPGRADE OPTION

905-1032

UPG-QSFP-DD-R400GE is a FIELD UPGRADE that upgrades the reduced performance K400 QSFP-DD-R400GE 1-port, 1-slot, L2-3 load module (944-1153) to have an increased number of transmit streams and higher L23 IxNetwork protocol emulation of the QSFP-DD-400GE full feature, L2-3 load module (944-1152) model. Note: At the time of order placement of the purchase of the upgrade, please provide the serial number of the desired QSFP-DD-R400GE load module to install the upgrade on.

FAN-OUT OPTIONS

905-1023

200GE/100GE/50GE FAN-OUT FACTORY INSTALLED Option for the K400 QSFP-DD-400GE (944-1152) or the QSFP-DD-R400GE (944-1153), 400GE QSFP-DD 1-port load modules. Note: This option is REQUIRED ON NEW PURCHASES to enable the 200GE/100GE/50GE speeds on the QSFP-DD-400GE or the QSFP-DD-R400GE 400GE QSFP-DD 1-port load modules.

905-1024

200GE/100GE/50GE FAN-OUT FIELD UPGRADE Option for the K400 QSFP-DD-400GE (944-1152) or the QSFP-DD-R400GE (944-1153), 400GE QSFP-DD 1-port load modules. Note: This option is REQUIRED ON FIELD UPGRADE PURCHASES to enable the 200GE/100GE/50GE speeds on the QSFP-DD-400GE or the QSFP-DD-R400GE, 400GE QSFP-DD 1-port load modules. Note: For the 200GE/100GE/50GE upgrade purchase please provide the serial number of the desired load module to install the option on at the time of order placement.

CABLES & TRANSCEIVERS

QSFP-DD-1M-CBL

QSFP-DD-to-QSFP-DD 400GE 400GBASE-R passive copper, Direct Attach Cable (DAC), point-to-point cable, 1-meter length. This cable is compatible with these load modules: QSFP-DD-400GE 1-port, full performance 400GE (944-1152) and QSFP-DD-R400GE 1-port, reduced 400GE (944-1153) load modules.

IXIA WORLDWIDE
26601 W. AGOURA ROAD
CALABASAS, CA 91302

(TOLL FREE NORTH AMERICA)
1.877.367.4942
(OUTSIDE NORTH AMERICA)
+1.818.871.1800
(FAX) 818.871.1805
www.ixiacom.com

IXIA EUROPE
CLARION HOUSE, NORREYS DRIVE
MAIDENHEAD SL6 4FL
UNITED KINGDOM

SALES +44.1628.408750
(FAX) +44.1628.639916

IXIA ASIA PACIFIC
101 THOMSON ROAD,
#29-04/05 UNITED SQUARE,
SINGAPORE 307591

SALES +65.6332.0125
(FAX) +65.6332.0127

K400 CFP8 400GE LOAD MODULES

DATA SHEET

PUT YOUR 400GE PRODUCT TO WORK NOW!

400GE is undergoing rapid development to meet the growing bandwidth requirements of ever-evolving data networks. Core networking applications are increasing by a factor of 10 every 5 years. 400GE addresses the broad range of bandwidth requirements to support these key application areas such as cloud-scale data centers, Internet exchanges, co-location services, wireless infrastructure, service provider and operator networks, and video distribution infrastructure.

TEST TOMORROW'S TECHNOLOGY TODAY

Ixia's K400 CFP8-400GE load module will accelerate your development of IEEE 802.3bs-compliant 400GE networking systems. Giving development teams the options they need to create the networking technologies of the future, these load modules are available in two different models:

- Full feature and scale CFP8-400GE
- Reduced feature and scale CFP8-R400GE

Both models enable full line-rate traffic generation transmit, receive, and capture functionality from 64B to 16,000B frame lengths. This facilitates stress testing, hardware/ASIC bring-up, optics and cable qualification, interoperability, and functional test and Layer 2/3 routing protocol emulation as required.

K400 CFP8-400GE: A full-featured model with 1-port, 2-slot, native CFP8 400GE load module designed for enterprise and data center switch and router testing.

HIGHLIGHTS

- Industry-leading 400GE test system from the leader in Ethernet test
- World's first 400GE test system running IEEE 802.3bs RS-544 Forward Error Correction (FEC) – Special Award Winner for Testing in 2016 at Interop Tokyo, Japan
- World's first live 400GE line rate interoperability demonstration with FEC with two independent MAC/PCS/FEC implementations, in 2016 at ECOC in Dusseldorf, Germany
- Ixia developed the intellectual property for the critical elements of 400GE: MAC, PCS, FEC, and FEC error injection and statistics — resulting in faster response times as industry standards evolve
- FEC error injection, PCS lanes Tx/Rx test capability and classical L1 BERT support provide essential tools for hardware development and bring-up

CFP8-400GE 1-port, 2-slot load module

K400 CFP8-R400GE: A reduced-featured model with 1-port, 2-slot, native CFP8 400GE load module designed for hardware, ASIC, cable/optics qualification, and interoperability testing. The CFP8-R400GE scales down the L2/3 feature set and L2/3 networking protocol scaling, while increasing affordability.

The CFP8-R400GE load module may be upgraded in the field to have a higher capacity L2/3 feature set and L2/3 networking protocol scaling performance through a purchasable upgrade option. This extends the reuse of the load module and improves the ROI.

K400 KEY FEATURES

- Line-rate 400Gbps packet generation, capture, and analysis of received traffic to detect and de-bug data transmission errors
- Line-rate per-port and per-flow statistics
- High latency measurement resolution at 0.625ns
- RS-544 (KP4) Forward Error Correction (FEC) support
- Native 28Gb/s SERDES with NRZ encoding support that is IEEE 802.3bs compliant
- FEC error injection with a comprehensive set of FEC corrected and uncorrected statistics, including Bit Error Rate statistics for pre- and post-FEC operations
- Inject packet errors: CRCs, runts, giants, alignments, checksum errors, and out of sequence
- Standard Ixia instrumentation including timestamp, sequence number and flow identification, and data integrity
- Layer 1 BERT: 16 independent lanes of 28Gb/s PRBS pattern generation, error checking, and statistics
- 400G PCS lanes Transmit, error injection testing and receive measurement:
    - Per-lane controls and status, FEC and error monitoring, error insertion, lane mapping and skew insertion
- +/-100 PPM line frequency adjustment
- Mid-to-high-range L2/3 networking protocol emulation to validate performance and scalability of L2/3 routing/switching and data center test cases using the Ixia's IxNetwork application
- An excellent test platform for full line-rate 400Gb/s to evaluate the new 400GE ASIC designs, FPGAs, and hardware switch fabrics that use the new 16x28Gb/s electrical interface with NRZ encoding
- Supports benchmarking of networking devices and equipment using industry-standard RFC benchmark tests at line-rate 400GE
- Supported by the XGS12 and XGS2 High Performance (HSL) and Standard Performance (SDL chassis model with the Native IxOS
- Application support including: IxExplorer, IxNetwork, and related Tcl and automation APIs

SPECIFICATIONS

PRODUCT DESCRIPTION	CFP8-400GE FULL FEATURE	CFP8-R400GE REDUCED FEATURE
Part Number	944-1150	944-1151
Hardware Load Module Specifications		
Slot / Number of Ports	2-slot, 1-port 400GE	
Physical Interfaces	Native CFP8 physical port	
Supported Port Speeds	400GE over 400GE-capable fiber media	
CPU and Memory	Multicore processor with 4GB of CPU memory per port	
IEEE Interface Protocols for 400GE	IEEE P802.3bs 400GbE, 400GBASE-R	
Advanced Layer 1 support	400GE: • KP4 (RS-544) Ethernet Forward Error Correction, Clause 119 • Pre- and post-FEC statistics: Comprehensive per-port and per-lane statistics • FEC error injection • PCS lanes Tx and Rx test and statistics • Classical Layer 1 BERT test and statistics	
Transceiver Support	Capable of support for 400GBASE-SR16, 400GBASE-LR8 10km and 400GBASE-FR8 2km optical transceivers	
Cable Media	400GBASE capable multimode fiber and single mode fiber cables that are compatible with the optical receptacle on the CFP8 transceiver	
Load Module Dimensions	• 17.3” (L) x 1.3” (W) x 12.0” (H) • 440mm (L) x 33mm (W) x 305mm (H)	
Load Module Weights	• Module only: 11.8 lbs. (5.35 kg) • Shipping: 18.6 lbs. (8.44 kg)	

PRODUCT DESCRIPTION	CFP8-400GE FULL FEATURE	CFP8-R400GE REDUCED FEATURE
Temperature	• Operating: 41°F to 95°F (5°C to 35°C) • Storage: 41°F to 122°F (5°C to 50°C)	
Humidity	• Operating: 0% to 85%, non-condensing • Storage: 0% to 85%, non-condensing	
Chassis Capacity: Maximum Number of Cards and Ports per Chassis Model		
XGS12-HSL Chassis (940-0016)	6 load modules: • 12-slot rackmount chassis • 6-ports of 400GE	
XGS2-HSL Chassis (940-0014)	1 load module: • 2-slot rackmount chassis • 1-port of 400GE	
XGS12-SDL Chassis (940-0015)	6 load modules: • 12-slot rackmount chassis • 6-ports of 400GE	
XGS2-SDL Chassis (940-0013)	1 load module: • 2-slot rackmount chassis • 1-port of 400GE	
Transmit Feature Specifications		
Transmit Engine	Wire-speed packet generation with timestamps, sequence numbers, data integrity, and packet group signatures	
Max. Streams per Port	400GE: 128	400GE: 32
Max. Streams per Port in Data Center Ethernet	400GE: 128	400GE: 32
Stream Controls	• Rate and frame size change on the fly • Advanced stream scheduler	

PRODUCT DESCRIPTION	CFP8-400GE FULL FEATURE	CFP8-R400GE REDUCED FEATURE
Minimum Frame Size	400GE: • 60 bytes at full line rate • 49 bytes at less than full line rate	
Maximum Frame Size	16,000 bytes	
Maximum Fame Size in Data Center Ethernet	9,216 bytes	
Priority Flow Control	• 4 line-rate-capable queues, each supporting up to 9,216-byte frame lengths • 1 line-rate capable queue, non-blocking supporting up to 9,216-byte frame lengths	
Frame Length Controls	Fixed, increment by user-defined step, weighted pairs (up to 16K), uniform, repeatable random, IMIX, and Quad Gaussian	
User Defined Fields (UDF):	Fixed, increment or decrement by user-defined step, sequence, value list, and random configurations; up to 10, 32-bit-wide UDFs are available	
Value Lists (Max.)	400GE: 1M / UDF	
Sequence (Max.)	400GE: 32K / UDF	
Error Generation	• FEC symbol error-injection allows the user to inject FEC symbol errors using various weighted methods to achieve specific error rates • Generate good CRC or force bad CRC, undersize and oversize standard Ethernet frame lengths, and bad checksum	
Physical Coding Sublayer	• PCS lane skew injection • PCS lane re-mapping • PCS lane marker error injection • PCS bit error generation	
L1 BERT	Classical, line rate, un-encapsulated transmit and receive of various PRBS patterns, controls over the patterns that help to produce BER statistics	

PRODUCT DESCRIPTION	CFP8-400GE FULL FEATURE	CFP8-R400GE REDUCED FEATURE
Hardware Checksum Generation	Checksum generation for IPv4, IP over IP, ICMP/GRE/TCP/UDP, L2TP, GTP, and multilayer checksum. Support for protocol verification for control plane traffic	
Link Fault Signaling	Reports, no fault, remote fault, and local fault port statistics; generate local and remote faults with controls for the number of faults and order of faults, plus the ability to select the option to have the transmit port ignore link faults from a remote link partner	
Latency Measurement Resolution	400GE: 0.625 nanoseconds	
Intrinsic Latency Compensation	Removes inherent latency error from the 400GE port electronics	
Transmit Line Clock Adjustment	Ability to adjust the parts-per-million-line frequency over a range of -100 ppm to +100 ppm on the 400GE port	
Transmit/Receive loopback	Internal and line loopback support	
Receive Feature Specifications		
Receive Engine	Wire-speed packet filtering, capturing, real-time latency, and inter-arrival time for each packet group, with data integrity, sequence and advanced sequence checking capability	
Trackable Receive Flows per Port	400GE: • 32K with the full statistics	
Minimum Frame Size	400GE: • 64 bytes at full line rate • 49 bytes at less than full line rate	
Filters (User-Defined Statistics, UDS)	2 SA/DA pattern matchers, 2x16-byte user-definable patterns. 6 UDS counters are available with offsets for start of frame.	
Hardware Capture Buffer	256KB	

PRODUCT DESCRIPTION	CFP8-400GE FULL FEATURE	CFP8-R400GE REDUCED FEATURE
Standard Statistics and Rates	Link state, line speed, frames sent, valid frames received, bytes sent/received, fragments, undersize, oversize, CRC errors, 6 user-defined stats, capture trigger (UDS 3), capture filter (UDS 4), data integrity frames, data integrity errors, sequence and advanced sequence checking frames, sequence checking errors, ARP, and PING requests and replies	
FEC Statistics	• FEC port statistics: Total Bit Errors, Max Symbol Errors, Corrected Codewords, Total Codewords, Uncorrectable Codewords, Frame Loss Ratio, Pre-FEC Bit Error Rate, and Codeword error distribution analysis • FEC per lane Rx statistics: FEC Symbol Error Count, Corrected Bits Count, Symbol Error Rate, Corrected Bit Rate	
Latency / Jitter Measurements	Cut-through, store & forward, forwarding delay, up to 16 time bins latency/jitter, MEF jitter, and inter-arrival time	
Receive-side PCS Lanes Port Statistics Counters	PCS: Sync Errors, Illegal Codes, Remote Faults, Local Faults, Illegal Ordered Set, Illegal Idle, and Illegal SOF	
400GE Physical Coding Sublayer (PCS) Receive-Side Statistics and Indicators	IEEE 802.3bs-compliant per-lane PCS receive capabilities include: • Receive – per-lane PCS receive statistics; Physical Lane Assignments, Lane Marker Lock, Lane Market Map, Relative Lane Skew, Lane Marker Error Count • Receive – per-lane FEC receive statistics; FEC Symbol Error Count, FEC Corrected Bits Count, FEC Symbol Error Rate, FEC Corrected Bit Rate	
Layer 2-3 Protocol Support		
Routing and Switching	BGP4/BGP4+, OSPFv2/v3, ISISv4/v6, EIGRP/EIGRPv6, RIP/RIPng, BFD, IGMP/MLD, PIM-SM/SSM, STP/RSTP/MSTP, PVST+/RPVST+, Link Aggregation (LACP), LISP	Complete protocol coverage with reduced session scale: • 100 routing & switching sessions • 2000 host/access sessions

PRODUCT DESCRIPTION	CFP8-400GE FULL FEATURE	CFP8-R400GE REDUCED FEATURE
Software Defined Network	OpenFlow, Segment Routing, BGP Link State (BGP-LS), PCEP, VXLAN, EVPN VXLAN, OVSDB, GENEVE, BGP FlowSpec, BGP SR TE Policy	Complete protocol coverage with reduced session scale: • 100 routing & switching sessions • 2000 host/access sessions
MPLS	RSVP-TE, RSVP-TE P2MP, LDP/LDPv6, mLDP, PWE, VPLS-LDP, VPLS-BGP, BGP auto-discovery with LDP FEC 129 support, L3 MPLS VPN/6VPE, 6PE, BGP RT-Constraint, BGP Labeled unicast, L3 Inter-AS VPN Options (A, B, C), MPLS-TP, MPLS OAM, Multicast VPN (GRE, mLDP, RSVP-TE P2MP), EVPN, PBB-EVPN	Complete protocol coverage with reduced session scale: • 100 routing & switching sessions • 2000 host/access sessions
Broadband and Authentication	PPPoX/L2TPv2, DHCPv4/DHCPv6, ANCP, IPv6 Autoconfiguration (SLAAC), IGMP/MLD, 802.1x	Complete protocol coverage with reduced session scale: • 100 routing & switching sessions • 2000 host/access sessions
Industrial Ethernet	Link OAM (IEEE 802.3ah), CFM/Y.1731, PBB/PBB-TE, ELMI, Sync-E ESMC, IEEE 1588v2 (PTP)	Complete protocol coverage with reduced session scale: • 100 routing & switching sessions • 2000 host/access sessions
Data Center Ethernet	DCBX/LLDP, FCoE/FIP, PFC (IEEE 802.1Qbb), TRILL, Cisco FabricPath, SPBM, VEPA	Complete protocol coverage with reduced session scale: • 100 routing & switching sessions • 2000 host/access sessions

APPLICATION SUPPORT

CFP8-400GE / CFP8-R400GE
• **IxExplorer:** Layer 1-3 wire-speed traffic generation, capture, and analysis with Forward Error Correction and error injection with statistics, PCS Lanes Tx/Rx testing with statistics, and Layer 1 BERT test and reporting capability. • **IxNetwork:** Wire-rate traffic generation with service modeling that builds realistic, dynamically controllable data-plane traffic. IxNetwork offers the industry's best test solution for functional and performance testing by using comprehensive emulation for routing, switching, MPLS, IP multicast, broadband, authentication, Carrier Ethernet, and data center Ethernet protocols. • **Tcl API:** Custom user script development for Layer 1-3 testing

ORDERING INFORMATION

LOAD MODULE

944-1150

K400 CFP8-400GE 1-port, 2-slot CFP8 400GE load module with the native CFP8 physical interface, L2-3 support (944-1150). Compatible with the XGS12-HSL rack mount chassis (940-0016), the 2-slot high performance XGS2-HSL chassis (940-0014), the XGS2-SDL, 2-slot Standard Performance chassis (940-0013), and the XGS12-SDL, 12-slot chassis (940-0015).

944-1151

K400 CFP8-R400GE 1-port, 2-slot CFP8 400GE reduced load module, with the native CFP8 physical interface, L2-3 support and reduced protocol emulation scale (944-1151). Compatible with the XGS12-HSL rack mount chassis (940-0016), the 2-slot high performance XGS2-HSL chassis (940-0014), the XGS2-SDL, 2-slot Standard Performance chassis (940-0013), and the XGS12-SDL, 12-slot chassis (940-0015).

LOAD MODULE UPGRADE OPTION

905-1033

UPG-CFP8-R400GE is a FIELD UPGRADE that upgrades the reduced performance K400 CFP8-R400GE 1-port, 1-slot, L2-3 load module (944-1151) to an increased number of streams and have the higher L23 IxNetwork protocol emulation of the CFP8-400GE full feature, L2-3 load module (944-1150) model. Note: At the time of order placement of the purchase of the upgrade, please provide the serial number of the desired CFP8-R400GE load module to install the upgrade on.

IXIA WORLDWIDE
26601 W. AGOURA ROAD
CALABASAS, CA 91302

(TOLL FREE NORTH AMERICA)
1.877.367.4942
(OUTSIDE NORTH AMERICA)
+1.818.871.1800
(FAX) 818.871.1805
www.ixiacom.com

IXIA EUROPE
CLARION HOUSE, NORREYS DRIVE
MAIDENHEAD SL6 4FL
UNITED KINGDOM

SALES +44.1628.408750
(FAX) +44.1628.639916

IXIA ASIA PACIFIC
101 THOMSON ROAD,
#29-04/05 UNITED SQUARE,
SINGAPORE 307591

SALES +65.6332.0125
(FAX) +65.6332.0127