

Testing 40 and 100 Gbps Ethernet

Rev 2

How to deal with challenges at the physical layer as well as of scaling the upper-layer engines when testing 40 and 100 Gbps Ethernet.

APPLICATION NOTE

This application note looks at the challenges of testing 40 and 100 Gbps Ethernet. As with any new technology, the key to success is testing, and 40/100G Ethernet brings a new set of issues to the development cycle.

SCALABLE ETHERNET SERVICE TESTING

This application note looks at the challenges of testing 40 and 100 Gbps Ethernet. As with any new technology, the key to success is testing, and 40/100G Ethernet brings a new set of issues to the development cycle.

Ethernet has migrated from 1G Ethernet to 10G Ethernet to 40G Ethernet and now 100G Ethernet in just a decade, which has also brought an increase in system complexity requiring multiple FPGA and ASSP devices, High-Speed SerDes technologies, and advanced optics.

First, we have the challenges at the physical layer and moving bits at 40 or 100 Gbps speeds. Secondly, we have the challenges of scaling the upper-layer engines to deliver services and Quality of Experience (QoE) at four to ten times the current rate.

There are also challenges associated with equipping the test lab with 40G Ethernet or 100G Ethernet test ports, while maintaining compatibility with network devices capable of converting seamlessly from 10 Gbps to 40 and 100 Gbps speeds.

PHYSICAL LAYER TESTING

Validating 40/100G Ethernet at Layer 1 is about the basics—moving bits error free. At different sub-layers, that means different things.

Physical Medium Dependent (PMD)

Testing the PMD layer is the domain of lower-level testing instruments such as oscilloscopes, jitter stress generators, and CFP to SMA breakout cable adapters.

Physical Medium Attachment (PMA)

Testing the PMA layer involves sending pseudo-random bit sequences (PRBS), through the system and detecting errors, such as bit errors or pattern synchronization issues. The ability to test the physical pathways using unique PRBS BERT patterns reveals physical lane stability and crosstalk issues. Loopback testing is important at this level.

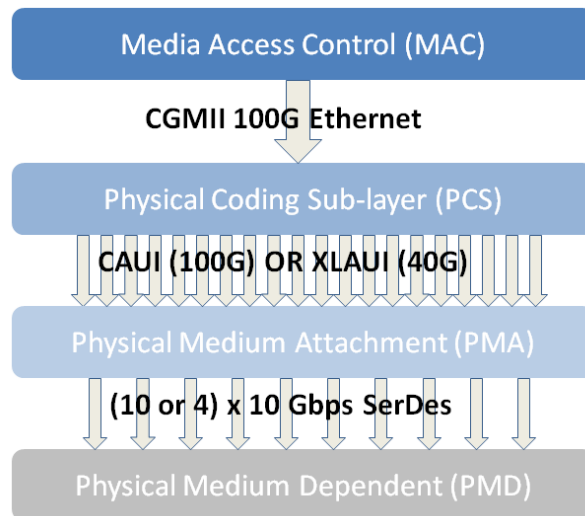


Figure 1 40/100G Physical Layers

The architecture of the 100/40G PHY layers is shown below. The number of lanes is 20 for the 100G CAUI host interface, and 4 lanes for the 40G XLAUI interface, with the specific examples shown illustrating a 100G CAUI host interface with a LR4 100G optical device.

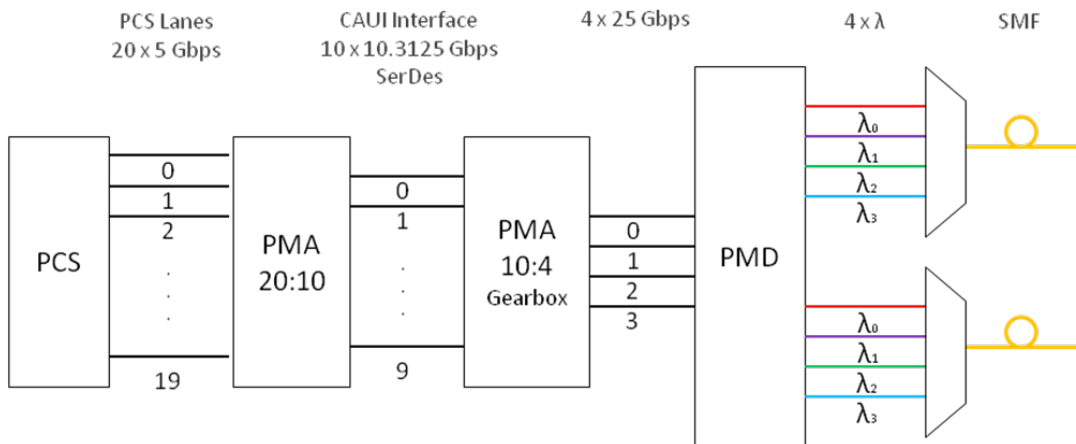


Figure 2 PHY (PCS and PMA) for 100G CAUI with LR4 optical transceiver

Physical Layer testing (PMA)

The increasing complexity of the Phy layer requires screening of optical devices during development and production. Use the PMA PRBS test function to determine if your 100/40 Gbps optical transceivers function correctly. Insert the transceiver optics into the Xena tester, and insert an optical loopback cable in the CFP module. Then enable PRBS for all lanes, and let the test run for at least 24 hours before verifying that there are no PRBS errors detected. The loopback cable must be clean and able to pass error-free traffic on any other known good transceiver.

PRBS configuration:		
	PRBS on	Error inject
All lanes:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 0:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 1:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 2:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 3:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 4:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 5:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 6:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 7:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 8:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 9:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 10:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 11:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 12:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 13:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 14:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 15:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 16:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 17:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 18:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Physical lane 19:	<input checked="" type="checkbox"/>	<input type="checkbox"/>

Figur 3 Enable PMA PRBS per lane Tx

PRBS status:			
	PRBS lock	PRBS errors	Error rate
Physical lane 0:	LOCK	0	
Physical lane 1:	LOCK	0	
Physical lane 2:	LOCK	0	
Physical lane 3:	LOCK	0	
Physical lane 4:	LOCK	0	
Physical lane 5:	LOCK	0	
Physical lane 6:	LOCK	0	
Physical lane 7:	LOCK	0	
Physical lane 8:	LOCK	0	
Physical lane 9:	LOCK	0	
Physical lane 10:	LOCK	0	
Physical lane 11:	LOCK	0	
Physical lane 12:	LOCK	0	
Physical lane 13:	LOCK	0	
Physical lane 14:	LOCK	0	
Physical lane 15:	LOCK	0	
Physical lane 16:	LOCK	0	
Physical lane 17:	LOCK	0	
Physical lane 18:	LOCK	0	
Physical lane 19:	LOCK	0	

Figure 4 Verify PMA PRBS per lane error statistics

Repeat this test for all optics. By measuring potential errors over a long duration, this test determines whether the optics function correctly.

To determine correctly whether the PMA layer on the DUT functions without bit errors, put the DUT in a PMA loopback (PCS loopback does not work) and then enable PRBS for all lanes on the Xena tester, and let the test run for at least 24 hours before verifying that there are no PRBS errors detected. The primary function of the PMA is to multiplex input lanes to output lanes, but this test also verifies clocking and clock recovery mechanisms in the DUT.

Physical Coding Sub-layer (PCS)

Testing the PCS layer involves introducing lane swapping and lane skew. Measuring the ability of the DUT to manage the virtual lane to physical SerDes lane translation is critical because lane swapping errors can lead to interface link-down problems that are difficult to debug in deployed systems.

First, configure the Xena tester to transmit L2/3 traffic from the tester to be sent across the DUT and back to the same tester port. Then, verify that header lock and alignment lock works on all Rx lanes, to verify that all lanes are recognized and are aligned. Then swap lanes randomly and verify the new Rx lane mapping. Any alignment lock issues indicate that the DUT PCS layer is not inserting the lane markers properly, and incorrect RX lane mapping indicates an issue with how the DUT transmits frames. Note that L2/3 traffic loss can be expected when the lane mapping is changed, and also when the Tx skew is changed.

Physical lane	Header lock	Align lock	Virtual lane	Skew bits
Physical lane 0:	LOCK	LOCK	1	0
Physical lane 1:	LOCK	LOCK	0	0
Physical lane 2:	LOCK	LOCK	3	0
Physical lane 3:	LOCK	LOCK	2	0
Physical lane 4:	LOCK	LOCK	5	0
Physical lane 5:	LOCK	LOCK	4	0
Physical lane 6:	LOCK	LOCK	7	0
Physical lane 7:	LOCK	LOCK	6	0
Physical lane 8:	LOCK	LOCK	9	0
Physical lane 9:	LOCK	LOCK	8	0
Physical lane 10:	LOCK	LOCK	11	0
Physical lane 11:	LOCK	LOCK	10	0
Physical lane 12:	LOCK	LOCK	13	0
Physical lane 13:	LOCK	LOCK	12	0
Physical lane 14:	LOCK	LOCK	15	0
Physical lane 15:	LOCK	LOCK	14	0
Physical lane 16:	LOCK	LOCK	16	0
Physical lane 17:	LOCK	LOCK	17	0
Physical lane 18:	LOCK	LOCK	18	0
Physical lane 19:	LOCK	LOCK	19	0

Figure 5 Verify alignment lock and errors on all Rx lanes

Physical lane	Virtual lane	Skew bits
Physical lane 0:	0	0
Physical lane 1:	1	0
Physical lane 2:	2	0
Physical lane 3:	3	0
Physical lane 4:	4	0
Physical lane 5:	5	0
Physical lane 6:	6	0
Physical lane 7:	7	0
Physical lane 8:	8	0
Physical lane 9:	9	0
Physical lane 10:	10	0
Physical lane 11:	11	0
Physical lane 12:	12	0
Physical lane 13:	13	0
Physical lane 14:	14	0
Physical lane 15:	15	0
Physical lane 16:	16	0
Physical lane 17:	17	0
Physical lane 18:	18	0
Physical lane 19:	19	0

Figure 6 Swap physical to virtual lane mapping randomly

To validate skew tolerance and compensation algorithms, insert a skew of 10 ns on multiple lanes and verify that the PCS markers remain aligned. Then, increase the skew until lane alignment lock is lost to find the maximum skew for the DUT. On the receive port, the amount of skew present after compensation by the system under test is reported. The Rx skew is measured with a resolution of 66 bits, which is equivalent to the size of one PCS codeword. Skew testing compares the ability of the system under test to support the standard, or report the degree to which it doesn't match the specification.

Lane configuration:		
	Virtual lane	Skew bits
Physical lane 0:	0	240
Physical lane 1:	1	0
Physical lane 2:	2	0
Physical lane 3:	3	500
Physical lane 4:	4	0
Physical lane 5:	5	0
Physical lane 6:	6	0
Physical lane 7:	7	0
Physical lane 8:	8	1000
Physical lane 9:	9	0
Physical lane 10:	10	0
Physical lane 11:	11	0
Physical lane 12:	12	400
Physical lane 13:	13	0
Physical lane 14:	14	0
Physical lane 15:	15	0
Physical lane 16:	16	700
Physical lane 17:	17	0
Physical lane 18:	18	0
Physical lane 19:	19	0

Figur 7 Insert skew on random selected lanes on the Tx PCS

Lane status:				
	Header lock	Align lock	Virtual lane	Skew bits
Physical lane 0:	LOCK	LOCK	1	66
Physical lane 1:	LOCK	LOCK	0	264
Physical lane 2:	LOCK	LOCK	3	594
Physical lane 3:	LOCK	LOCK	2	66
Physical lane 4:	LOCK	LOCK	5	66
Physical lane 5:	LOCK	LOCK	4	0
Physical lane 6:	LOCK	LOCK	7	66
Physical lane 7:	LOCK	LOCK	6	66
Physical lane 8:	LOCK	LOCK	9	0
Physical lane 9:	LOCK	LOCK	8	990
Physical lane 10:	LOCK	LOCK	11	0
Physical lane 11:	LOCK	LOCK	10	0
Physical lane 12:	LOCK	LOCK	13	0
Physical lane 13:	LOCK	LOCK	12	396
Physical lane 14:	LOCK	LOCK	15	0
Physical lane 15:	LOCK	LOCK	14	0
Physical lane 16:	LOCK	LOCK	16	726
Physical lane 17:	LOCK	LOCK	17	66
Physical lane 18:	LOCK	LOCK	18	66
Physical lane 19:	LOCK	LOCK	19	66

Figur 8 Verify the skew insertion on Rx PCS

The IEEE 802.3ba specification suggests a max PCS skew tolerance of 180 ns. The IEEE defined skew tolerances are shown below.

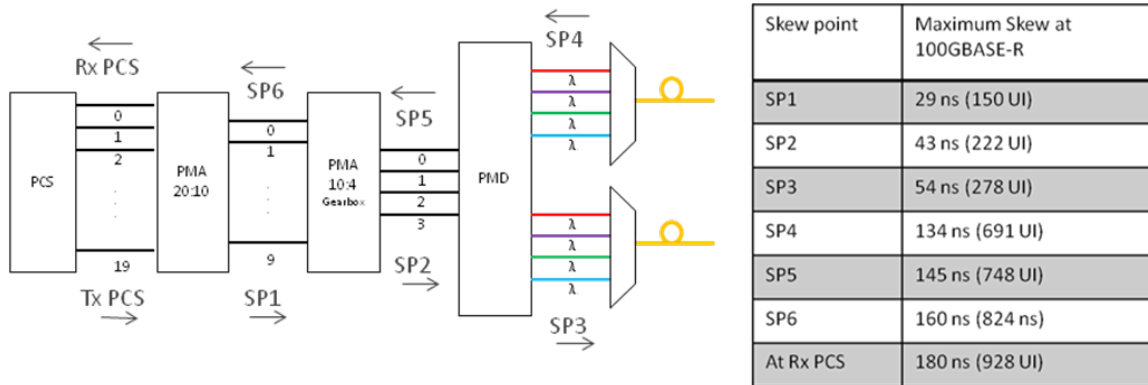


Figure 9 IEEE 802.3ba defined skew tolerances

Note that for 40GBASE-R, 1 UI is equal to 97 ps at PCS lane signaling rate of 10.3125 Gbps

Note that for 100GBASE-R, 1 UI is equal to 194 ps at PCS lane signaling rate of 5.15625 Gbps

Media Access Control (MAC)

The Layer 2 MAC layer has similar testing issues to Layer 1. Testing the MAC layer involves measuring the ability to transmit and receive frames and handle frame check sequence (FCS) errors at all possible rates, frame lengths, and frame inter-frame gaps. The test system also introduces FCS errors, as well as counting FCS errors received.

BENCHMARK TESTING WITH RFC 2544 AND RFC 2889

To provide a standard way of measuring and evaluating a solution, the industry has created a series of benchmarks. These benchmarks establish a uniform procedure for the generation of traffic to and from the system under test with a normalized procedure of analysis and reporting.

The ubiquitous RFC 2544 is used to verify forwarding performance and includes throughput (performance availability), back-to-back (link burstability), latency (transmission delay) and frame loss (service integrity) measurements. The Xena testers can perform the RFC 2544 test suite for 100G and 40G Ethernet interfaces at all frame sizes and at full-line rate, allowing service providers to certify that the circuit is efficient and error-free at 100% utilization at Layer 2 (VLAN) or layer 3 (IP).

Maximum port-pair throughput with no loss				
Frame size	64	128	256	512
Maximum physical port speed (Mbps)	100000	100000	100000	100000
Configurable maximum speed (Mbps)	n/a	n/a	n/a	n/a
100000 Mbps maximum Rate (fps)	148809524	84459459	45289855	23496241
Passed rate (%)	100.0	100.0	100.0	100.0
Acceptable loss percent	0.00	0.00	0.00	0.00
<0-to-0> (fps)	148809523	84459459	45289855	23496240
<0-to-0> (%)	100.00	100.00	100.00	100.00
<0-to-0> (Mbps)	100000.0	100000.0	100000.0	100000.0
Throughput SUMMARY: Total Port-Pairs				
Frame size	64	128	256	512
Passed rate (fps)	148809523	84459459	45289855	23496240
Passed rate (%)	100.00	100.00	100.00	100.00
Passed rate (Mbps)	100000.0	100000.0	100000.0	100000.0

▪ Figure 10 Sample RFC 2544 test report for 100 Gbps DUT

Several vendors are now offering 40 Gbps enabled layer 2 switches. The RFC 2889 test suite is ideally suited for validating the performance of such switching devices, by testing at full line-rate, large-scale full-mesh test beds, and characterizing the switch's ability to meet the demand for the low-latency and throughput performance for lossless data center Ethernet transmission.

With multicast traffic growing in importance in cloud networking applications, the RFC 2889 benchmark tests can also be used to characterize multicast-only throughput test cases where latency and throughput are key performance parameters.

TESTING MULTI-RATE NETWORK PORTS

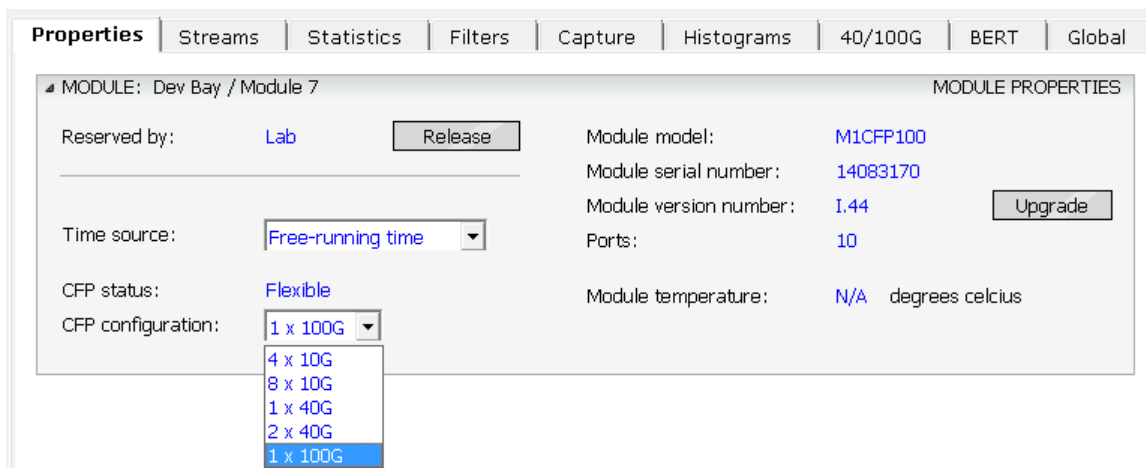
New commercially available Layer 2 and layer 3 network devices offer a smooth network upgrade path, by providing full flexibility to configure network ports as either 4 x 10 Gbps, or 1 x 40 Gbps. This powerful capability allows the customers to deploy 10 Gigabit Ethernet to meet current needs and be ready to transition to 40 Gigabit Ethernet without any disruption. Such devices include the Cisco Catalyst 6500 series, Juniper's Nexus 7000 switches, as well as several others commercially available solutions.

The Xena 40 Gbps and 100 Gbps testers are the only commercially available Ethernet testers offering tri-speed test ports. Xena's CFP based testers can support 100 Gbps, 40 Gbps, and 10 Gbps port speeds. A summary of the port speeds supported by the 100/40G Xena testers is shown in Table 1.

Xena tester P/N	100 Gbps Interfaces	40 Gbps Interfaces	10 Gbps Interfaces
Odin-100G-3S-1P* (1 CFP cage)	1 x 100GBASE-LR4/ER4 1 x 100GBASE-SR10	1 x 40GBASE-LR4 2 x 40GBASE-SR4	4x10GBASE-LR 8x10GBASE-SR
Odin-40G-2S-2P-B* (1 CFP cage)	-	1 x 40GBASE-LR4 2 x 40GBASE-SR4	4x10GBASE-LR 8x10GBASE-SR

▪ Table 1 Port speeds supported by Xena’s CFP based testers

To change the port speed, the user must simply configure the port speed and change the physical configuration of the CFP transceiver and optical cables.

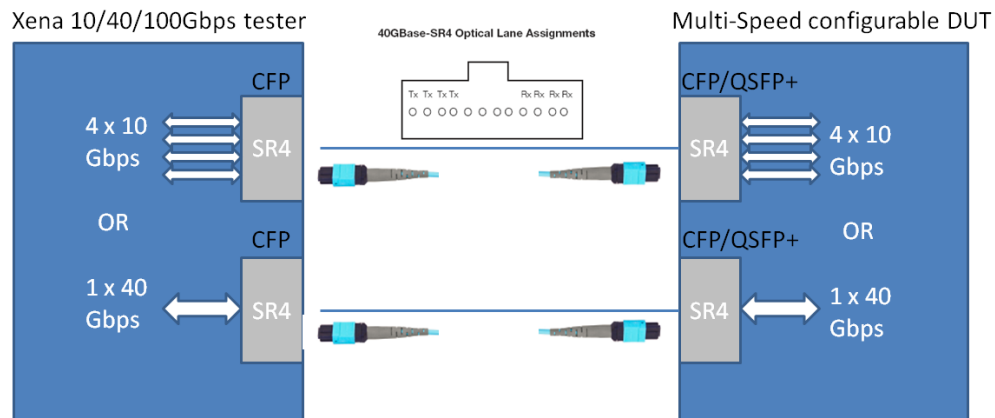


▪ Figure 11 Select 1 x 100 Gbps CFP test module configuration in ValkyrieManager

The unique flexibility of the Xena testers is not only useful for scaling your test lab test infrastructure from 10 Gbps, to 40 Gbps, and 100 Gbps as your requirements change, but also allows you to verify multi-speed network devices using a single lab testbed.

It is possible to setup a single lab test-bed for verifying network devices with configurable port interface speeds of 4 x 10 Gbps or 1 x 40 Gbps, by using SR4 parallel optics based transceivers and MPO/MTP cabling.

Existing product offerings from Cisco, Juniper, and several other vendors provide multi-speed configurable network devices based on supporting SR4 parallel optics in either CFP or QSFP+ form factor. A unified 10 Gbps and 40 Gbps testbed is shown in Figure 12.



▪ Figure 12 unified 10 and 40 Gbps testbed using SR4 optics in CFP or QSFP form factor

This unified testbed can be used for verifying both 10 Gbps and 40 Gbps modes of operation without having to change the physical testbed setup, since both the DUT and the Xena tester are configurable for either 4 x 10 Gbps or 1 x 40 Gbps modes of operation per optical transceiver cage (CFP or QSFP+ form factor).

Next generation switch silicon and network devices entering the market in 2012 will support multi-speed port configurations, where one CXP interface can be configured as either 1 x 100 Gbps, 3 x 40 Gbps, or 12 x 10 Gbps. For these 100G enabled devices, it will also be possible to setup a testbed which can support any mix of speeds using a single unified testbed setup.

In addition to the advantage of using a single physical test bed to test both 10 Gbps, 40 Gbps, and 100 Gbps modes of operation of the DUT, the unique 10/40/100G tri-speed capability of the Xena testers will also reduce the need for extra investments in test equipment to cover the different speed modes, and result in a more compact and less space demanding lab testbed installations.

LIVE DEMO

It's easy to try our solutions live right now - all you need is a PC running Windows XP or later! Here's what you do, go to http://xenanetworks.com/html/live_demo.html

**End of slide modules*