

100G High-Speed Testing: A Primer

Jean-Marie Vilain, Product Line Manager, Transport and Datacom, EXFO

INTRODUCTION

Given that many carriers have begun mass deployment of their 100G networks, a couple of questions arise: first, is 100 Gigabit Ethernet (GigE) testing different than 10 GigE testing? Second, what needs to be tested during the deployment of a new 100 GigE network? In the following application note, we will address the key tests that need to be performed at 100 GigE during commissioning, turn-up, and even troubleshooting.




Figure 1. Overview of 100G Network

Figure 1 illustrates some of the network complexity in today's high-speed networks, with rates from 10M all the way up to 100G. Multiple tasks must be performed to qualify the network, and to this end we will be covering the key testing points from the interface all the way to services testing.

a.) CFP/CFP2 Validation

CFP quality is improving, and as a result we are now seeing fewer and fewer issues due to incompatibility or electrical crosstalk. However, a quick validation test is still required before each deployment to analyze the main status of the CFP and ensure that there are no errors. The main tests involve analyzing the laser power to ensure that the laser's transmitted and received optical power is within the specific range; as part of the same validation process, the received per-lane frequency will need to be tested.

In some cases, the CFP is considered the weakest link in the network; in such cases, it will be critical to ensure that the CFP is completely functional before conducting any CFP deployment.

Unlike the single-wavelength transceivers that were used for legacy 2.5G and 10G, each CFP parallel optical channel must be monitored for transmitted and received power levels. Even if the overall power across all channels (4 or 10) is within the acceptable range, this might be the result of averaging a channel with very low power, which could in turn impact the transmission and performance on the optical channel and a high-power channel, which could damage the optical receiver at the other end. The add-on of a simple attenuator on the optical link can help reduce the RX power. Figure 2 below shows a quick CFP power validation.

In some cases, a per-lane view may be required to identify any suspicious lanes that could eventually cause bit errors. For instance, lane 9 as shown in Figure 3 below is slightly out of range.

100GE (10 Lanes)			
LINK	Lasers 10 of 10 ON		
	LO	HI	
RX Power (dBm)	1.1	2.4	
RX Offset (ppm)	0.0	0.0	
TX Offset (ppm)	0.0		
IP	10.10.172.37		

P2 100GE (10 Lanes) LINK ↑ Power ⚠

Figure 2. Quick CFP Power Validation

Physical Interface						
Optical Lane	Laser	TX Power (dBm)	Wavelength (nm)	RX Power (dBm)	Min RX Power (dBm)	Max RX Power (dBm)
0	ON	2.05	0.00	-7.55	-8.51	-7.54
1	ON	1.66	0.00	-7.11	-8.06	-7.11
2	ON	2.25	0.00	-6.28	-7.23	-6.28
3	ON	1.94	0.00	-6.61	-7.61	-6.60
4	ON	2.03	0.00	-6.97	-7.98	-6.96
5	ON	2.32	0.00	-6.98	-8.09	-6.94
6	ON	2.31	0.00	-6.89	-7.96	-6.89
7	ON	2.03	0.00	-6.81	-7.73	-6.81
8	ON	2.04	0.00	-6.88	-7.96	-6.88
9	ON	1.64	0.00	-8.00	-8.80	-7.98

Laser ON/OFF Laser OFF at Start-Up Power Range (dBm) -8.1 To 5.2

Figure 3. Per-lane power and frequency measurement

The above tests are also applicable to the upcoming CFP2 interface, which is backward-compatible with CFP technology. There is no gearbox in the CFP2—it was moved to the host, reducing the physical size and power consumption. However, this also introduces some new challenges, and at the same time new issues. EXFO's CFP2 adapter, which can be inserted into a CFP interface, allows field technicians to validate CFP2 transceivers before deployment, thus minimizing risk. Given that more than 60% of deployment issues are related to CFP or CFP2, it is now critical to test each CFP. EXFO's CFP health check supports both interfaces with a basic and advanced mode for quick pass/fail validation.

b) Ethernet BER Test (EtherBERT)

Once the physical-layer test complete, the next step is to ensure that an Ethernet packet can be transmitted in an error-free manner over the entire network. This could be considered a basic test, but is helpful in validating that all of the traffic is able to transition to the far end and loopback without any errors. Failing this would result in bit errors. The same test can also provide complete Ethernet statistics and protection switching time data from both paths (working and protection). This data helps qualify Ethernet switches or routers that identify end-to-end connectivity issues from the MAC layer and up.




Figure 4. Typical EtherBERT 100G Network Test


c) Testing Multiple Types of Services

The BER test is primarily based on pseudo-random traffic, whereas in reality, live traffic tends to be based more on different types of services, thereby driving the need for more multiservice tests in which the user can simulate multiple streams with predefined data, voice and video traffic profiles, or user-configurable profiles. These tests make it possible to simulate all of the types of services that will be running on the network while simultaneously qualifying all the key service-level agreement (SLA) parameters for each of these services. Moreover, these tests validate the quality of service (QoS) mechanisms provisioned in the network for prioritization of the different service types, which results in better troubleshooting, more accurate validation and much faster deployment. This capability includes the generation, shaping and monitoring of Ethernet and IP traffic with throughput, frame loss, sequencing, packet jitter, latency, frame size, traffic type and flow control.

d) Network Performance Validation: RFC 2544 and Y.1564 Test

To ensure that a 100 Gigabit Ethernet network is capable of supporting a variety of services (such as VoIP and video), an RFC 2544 test can be executed, whereby the 100G pipe will be tested and confirmed for 100% throughput and low-latency frame loss by testing at different frame sizes.

The predefined frame sizes will simulate various traffic conditions, because small frame sizes increase the number of frames transmitted, thereby stressing the network device. Because carriers are currently faced with stressful constraints in terms of time to deployment, the industry can now initiate a shift toward ITU-T Y.1564 (EtherSAM) for turn-up testing. This test offers simultaneous tests (as opposed to RFC 2544's sequential tests), and validates multiple services at the same time, thereby qualifying the services as defined in the SLA. ITU-T Y.1564 is a crucial test that must be executed, and is based on two phases. The first phase, which is called the network configuration test, consists of validating the QoS mechanisms and limits based on the key performance indicator (KPI) threshold: the committed information rate (CIR), the excess information rate and the discard traffic conditions. Figure 5 below shows the KPIs that will be measured during the network configuration test.



Bandwidth: The maximum amount of data that can be forwarded in one second. Generating traffic that exceeds the bandwidth limit leads to frame loss or service outages.

Packet Jitter: Refers to the variability in arrival time between deliveries; this translates into stress on the receiving buffers of the end nodes, and impacts real-time applications and QoE.

Latency: A measurement of the time delay between a packet's transmission and its reception. This is typically a round-trip measurement, and is critical for voice applications.

Figure 5. Key Performance Indicators

The second phase, called the service performance test (see Figure 6), runs all services simultaneously and validates each service individually based on the specified CIR and SLA. This means that each selected service will be run simultaneously in order to measure the throughput, jitter and latency of each service. A pass/fail result will then be generated based on the service-specific thresholds provided.

As part of the process of minimizing test times and reducing the amount of truck rolls, a **dual test set** (DTS) configuration can measure the performance of the system under test (SUT) by measuring both transmission directions. Additionally, this test (which is also known as a head-to-toe test) offers a lot of value when the networks feature multiple routers.



Figure 6. Service Performance Test

In a DTS configuration, the user can have one test set controlling the other by simply dedicating one as the “local” test set and the other as the “remote” test set. This makes it easy to determine in which direction the traffic is flowing. It also provides asymmetric results that can easily identify network limitations. For example, Figure 7 below shows a local-to-remote direction in red, and a remote-to-local direction in green. This bidirectional flow of traffic enables the test sets to transmit and receive at the same time.




Figure 7. Dual Test Set Configuration with Bidirectional Measurements

100G Network Troubleshooting Tips

In some cases, the above tests may fail to qualify the network performance. In such cases, troubleshooting tools may be needed to help identify the root cause. The next section will provide an overview of the various solutions available on the EXFO FTB/IQS-881000NGE and FTB/IQS/IQS-85100G (including basic IP test tools such as ping, traceroute and Ethernet packet capture) to assist with troubleshooting of particular problems

Packet Filtering and Capture

In most troubleshooting cases, only a particular type of traffic is of interest, whereas any other traffic will take up memory without providing any useful information. EXFO's FTB-88100NGE has the capability needed to filter Ethernet traffic in such a way that only traffic fitting a specific profile is captured, thus maximizing efficient use of available memory. Furthermore, the FTB-88100NGE provides an innovative capture capability limiting capture to a specific number of bytes starting from the first bit of the Ethernet frame. This enables network engineers to limit capture to the first few bytes of the header, or add more bytes to include higher-layer information.

Per-Lane Alarm and Error Monitoring

For advanced troubleshooting needs, the FTB-88100NGE offers per-lane alarm and error monitoring capabilities that enable tier-2 engineers to quickly pinpoint and resolve issues. These capabilities include interface-type alarms such as loss of signal (LOS), and loss of continuity (LOC) and frequency, in addition to PCS-related alarms and errors.

CONCLUSION

In summary, although 100G networks are more complex, the required tests are fairly similar to those performed at lower Ethernet rates, such as 1 GigE and 10 GigE. EXFO's FTB-88100NGE, which offers advanced Ethernet and OTN testing from 10M all the way to 100G, provides network engineers with a quick and accurate way to pinpoint issues in the field and speed up the troubleshooting process for quick service recovery.