

Testing 100 G Transport Networks and Services

Testing 10 G and 100 G services is similar; however, significant differences exist

By Guylain Barlow

The growing need for bandwidth capacity has fueled the development of the 100 G-based network. More specifically, this refers to a network equipped with core interfaces capable of carrying signals at bit rates just above 100 G. The first trials and experiments at 100 G focused on the primary task of transporting information across relatively long distances, which required modulation techniques, with a view to support backbone connectivity. New 100 G network deployments have become commonplace, and are taking place on all continents. Applications are starting to extend beyond backbone transport and are evolving to the wholesale and bundling of 100 G ports. However, the use of 100 G services by enterprise customers, such as banks, remains in its infancy; it will continue to grow over the coming years. In parallel, 100 G is evolving for use in data centers with a roadmap for server connectivity. This paper outlines the fundamentals of 100 G transport network architectures, describes service activation applications applicable to both optics and 100 Gigabit Ethernet (GE) testing, and looks into OTN networking and its evolution.

Network Topology

Loose terminology is often used to classify the main types of 100 G interfaces as “client side” and “line side.” Generically, line side refers to interfaces used for information transmission over relatively long distances such as hundreds or thousands of kilometers, where several signals are transported over a single fiber using wavelength division multiplexing (WDM). This requires the use of advanced modulation techniques that are complex and typically vendor specific. By comparison, client interfaces use simple on-off keying (OOK) modulation and well-defined multi-vendor standard interfaces which call for shorter transmission distances. There are two main types of client interfaces:

- 100 GE at 103.125, defined and standardized in IEEE 802.3ba; this is the most common type
- OTU4 at 111.8, defined and standardized in ITU-T G.709

Line-side interfaces at 100 G do not transmit using any standardized modulation formats; however, the majority of 100 G line interfaces use polarization multiplexed quadrature phase shift keying (PM-QPSK) modulation, also called dual polarization QPSK (DP-QPSK). This modulation technique is typically used in conjunction with coherent detection at the far-end receiver; this is where the receiver is equipped with a frequency-locked local oscillator laser and the signal phase and amplitude is recovered via digital signal processing. An advantage of such coherent receivers is that they can electronically compensate for dispersion in the fiber, which greatly increases the resiliency of the transmitted signals to chromatic dispersion and polarization mode dispersion. Figure 1 shows a high level diagram of DP-QPSK modulation. The transmitter generates two independent optical QPSK signals, each with four different optical phases, which are combined in two orthogonal polarization states. This results in a signal where the baud rate is a quarter of the bit rate.




Figure 1. DP-QPSK signal

Client interfaces are interoperable between vendors and provide connectivity into the network. With 100 GE transport, the most common interface is 100GBase-LR4, which provides fiber connectivity across 4 wavelengths over a maximum distance of 10 km. Each wavelength is equivalent to roughly 25 G of bandwidth providing an aggregate of 103.125 G. There exist alternative 100 GE interfaces as shown in Table 1. Most interfaces are defined in IEEE 802.3ba while the 10x10 MSA outlines a 10 G wavelength interface. Currently, the use of 100GBase-ER4 is in its infancy due to the lack of availability of 40 km optics. The relatively popular 100GBase-SR10 interface is primarily for data center applications and 100GBase-CR10 is an electrical standard.

Figure 2 shows a sample diagram of the relation between client and line 100 G interfaces in a typical 100 G transport network.

Interface	Reach	Medium	No. of WL/Fibers	Benefit
100GBase-LR4	10 km	SMF	4 λ / dir	IEEE 802.3ba
100GBase-ER4	40 km	SMF	4 λ / dir	IEEE 802.3ba
100GBase-SR10	100 m 125 m	OM3 MMF OM4 MMF	10 fibers / dir	IEEE 802.3ba
100GBase-CR10	7 m	Twin-axial electrical	10 cables / dir	IEEE 802.3ba
10x10 MSA	2 km 10 km 40 km	SMF	10 λ / dir	10x10 MSA Tech Spec

Table 1. 100 GE interfaces




Figure 2. Line and client interfaces

The application illustrated in Figure 2, where the transport equipment is equipped with a 100 G CFP (C-form factor pluggable) on the client interface is generally referred to as a transponder application. In such a case, a service such as 100 GE is transported across the network using WDM core links. This is in contrast to a muxponder application where multiple lower-speed links, typically multiple 10 G links such as 10 GE interfaces, connect into a core backbone to form an aggregate of 100 G consolidated onto a wavelength. Basically, the client traffic goes through a multiplexing stage to take advantage of the efficiencies provided by higher-rate core backbone links. In terms of application testing, activating such a muxponder service implies generating and analyzing traffic into the network from one or multiple 10 G sources rather than from a 100 G source. There exist additional 100 G transport applications; the transponder and muxponder variants are the two most common at this point in time. Figure 3 illustrates sample connectivity in a muxponder application.




Figure 3. Muxponder application

Line-Side Testing

Testing online interfaces involves analyzing the optical parameters related to the signal and medium, which in this case is fiber optics. Concerning dispersion, a higher line rate such as 100 G implies a higher level of sensitivity at the receiver. For chromatic dispersion, the sensitivity grows proportionally to the square of the line rate, which is significant, while the polarization mode dispersion grows linearly with the line rate. However, because 100 G systems use coherent detection coupled with high-speed electronic signal processing in the receiver, the signal distortions caused by chromatic and polarization mode dispersion are well mitigated. The tolerance of these signals to chromatic and polarization mode dispersion is substantially higher than that of conventionally detected 10 G or 40 G signals (typically more than 30,000 ps/nm of accumulated chromatic dispersion and more than 30 ps of polarization-mode dispersion). Nevertheless, characterization and documentation of the fiber infrastructure remains a highly recommended step in system turn up, especially in WDM system with mixed signals, where 10 G and 40 G signals may co-propagate with 100 G signals through the same optical fiber.

An important parameter to consider when commissioning 100 G line-side links is the optical signal to noise ratio (OSNR) of the signal along the fiber link and in particular at the end of the link. Ensuring that the OSNR of a 100 G signal falls within the tolerance limits is critical in ensuring network operation with a minimal information error rate. However, it is difficult to measure OSNR directly on the transmitted 100 G line signals, because the signal spectrum extends over the entire width of the WDM channel. Furthermore, polarization discrimination techniques (like polarization nulling) used to measure OSNR on single-polarization signals (like 10 or 40 G OOK-NRZ) cannot be applied in this case, because a PM-QPSK signal does not exhibit a well-defined polarization state. Therefore, the OSNR of 100 G PM-QPSK signals cannot be measured when the network is in service. In some cases, it may be possible to measure the optical noise level in a selected WDM channel by turning off the transmitted 100 G signal. This requires the network to be out of service and is not practical in ROADM networks which employ optical channel monitors (OCMs) that block channel transmission when a signal goes missing. Research is being conducted at the industry level in order to identify an in-band OSNR measurement method for PM-QPSK signals.

100 GE Client Testing

The basis of 100 GE client interface activation testing involves the generation and analysis of traffic while monitoring alarms and errors. To run network tests, the typical approach is to use a test suite which simplifies and automates execution. The common traffic service activation tests suites are defined in IETF RFC 2544 and ITU-T Y.1564. Most often, these test suites are run with a loopback at the far end of the network as illustrated in Figure 1. The far-end loop may take the form of a hard loop, typically a fiber optics patch cord, which only applies to Layer 1 transmission networks. When Layer 2 or 3 functions are present, meaning Ethernet switching or routing, the far-end loopback must be a logical function. Such functionality requires an active device, such as test equipment, where Layer 2 and as required Layer 3 source and destination addresses are swapped.

The choice in using RFC 2544 or Y.1564 can depend on the nature of the activation test or simply on the existing procedures for the commissioning or troubleshooting team. RFC 2544 is a well-established method which includes tests for throughput, latency, frame loss rate, and bursting. It is generally accepted that although outside of the recommendation, a packet jitter test is added. The main role fulfilled by RFC 2544 is to activate a new link as a single service. The execution time of RFC 2544 test suites actually depends on specific implementations; optimizations are possible to increase the speed and efficiency of this test.

The focus of ITU-T Y.1564 primarily is on service activation; resulting test suite implementations have the main benefit of providing automated test execution for multiple services simultaneously. In a test suite, such services are simulated as traffic streams that are used, for instance, to discriminate between service traffic assigned to different classes of service or even between regular service traffic and mission-critical control plane traffic. Y.1564 focuses on 3 key performance indicators (KPIs) which provide metrics in the form of frame loss rate (FLR), frame delay (FD, latency), and frame delay variation (FDV, packet jitter). The service performance tests from Y.1564 are well-suited to service level agreement (SLA) testing by verifying each service against its bandwidth profile. The traffic parameters associated with these profiles always include committed information rate (CIR) and there is growing interest in committed burst size (CBS); excess information rate (EIR) testing is part of Y.1564 although it is not always exercised in practice while the use of excess burst size (EBS) is less common. The JDSU Y.1564 test suite is called SAMComplete. Like RFC 2544, Y.1564 can be used to test bi-directional traffic (with loopback) and even uni-directional or asymmetric traffic. Table 2 provides a high-level summary comparison of RFC 2544 and Y.1564. In the case of errors detected while running network tests, one of the troubleshooting steps may involve testing the client optics. Such CFP tests can be run on test equipment and should include low-level pattern testing across the supported clock offset range which for Ethernet is ± 100 ppm.

	RFC 2544	Y.1564
Main 100 GE Application	Turn up new backbone link or single end-to-end connection	Turn up Ethernet connection(s) or service(s) in the core
Parameters Measured / Validated	Throughput Latency Frame loss rate Burstability (Extra: CBS) Extra: packet jitter	FLR Latency: FD Packet jitter: FDV CIR EIR Traffic policing CBS
Key Focus	Test one stream (address pair); identify maximum performance	Test one or multiple service, validate frame transfer performance including against SLA parameters

Table 2. RFC 2544 and Y.1564 Applications

An important item to consider as part of RFC 2544 and Y.1564 is the accuracy and resolution of latency measurements. Accuracy refers to how close a measurement is to the true value, within a margin of error, and resolution refers to the level of information provided, for instance 100 ns or 0.01 μ s resolution. To be able to accurately correlate fiber length to latency requires both high resolution and accuracy. This correlation is important, especially when supporting end customers involved in applications such as high-frequency trading in the financial sector or time-sensitive wireless services. At 103.125 G, the actual 100 GE rate, it takes 100 ns to transmit a 1289 byte frame, providing frame-level latency accuracy is within the reach of current technology.

An additional recommended procedure for service activation is Layer 2 transparency testing. In Ethernet networks, where there are elements of switching or even routing, this test is an important consideration which can eliminate long and arduous troubleshooting. Switches and routers, including Ethernet/MAC modules found in transport equipment, generate a relatively small amount of vital traffic in the form of control-plane information. These control-plane frames are mainly destined for other switches and routers. When providing bandwidth services in a network, the idea is to ensure that all control-plane protocols will be treated appropriately: implying their being tunneled, peered, or discarded as per MEF 6.1.1. Layer 2 control-plane protocols are identified by their protocol type from the Ethertype and subtype fields or logical link control (LLC) code. Examples of such protocols include spanning tree protocol (STP) and link layer discovery protocol (LLDP). An effective method to run this type of test is to emulate these control plane protocols in the network while making sure that corresponding frames do not get improperly intercepted or modified. An additional application example for Layer 2 transparency tests is when an operator provides bandwidth services which traverse a third-party network.

In summary, testing a network using test equipment connected to client interfaces has the advantage of testing the full and complete data path in the network. Using such equipment provides an independent assessment of performance in addition to complete reporting capabilities.

Troubleshooting Network Issues

Beyond service activation testing is the need for troubleshooting. At the Ethernet level, implying Layer 2 and Layer 3, the test methodologies and procedures are fairly similar to what is done at 10 GE (like RFC 2544, Y.1564, and Layer 2 transparency). However, the Layer 1 implementation on 100 G modules and systems is fundamentally different from lower line-rate systems, and these differences extend to client optics across multiple wavelengths. As a result, field troubleshooting incorporates the requirement to view information at the electrical lane level for reporting. These electrical lanes are found in the back of optical modules like CFPs. Such information, such as per lane errors and alarms (for example, loss of synchronization) can, for instance, help point to a hardware issue isolated to a specific physical lane. In any case, the need to perform deep testing for items such as skew testing are beyond field requirements and are instead performed in laboratory environments. Should serious physical layer issues be present in the field, they will be detected via traffic test procedures.

A troubleshooting capability which is not yet required in the field is the need for field-level packet capture. Such capabilities do require the presence of mirror ports for capturing, a capability not readily available at 100 GE. In addition, this level of troubleshooting is performed on an end customer circuit, such as an enterprise customer, where very few 100 GE circuits are currently available.

The Optical Transport Network (OTN) Future is Upon Us

The evolution to 100 G networks is promoting the deployment of a greater number of OTN ports. The OTN was created more than a decade ago with the primary purpose of carrying SONET/SDH clients over long distances. It features a forward error correction (FEC) algorithm associated with each OTN frame that helps extend the distance between terminating nodes. The main standard, ITU-T G.709, has evolved to accommodate diverse payload clients such as Ethernet, Fibre Channel, and common public radio interface (CPRI); this is where the 100 G OTU4 interface is defined.

Going forward, the main change from a test perspective is that more OTN ports are being used as client interfaces. This in turns enables test application as client ports are fully standardized and interoperate with test equipment. Most legacy OTN applications, in particular digital wrappers where a client signal is wrapped in OTN, applied to line side ports where test traffic generation is often not achievable due to compatibility issues. These arise on line-side ports because of two main reasons: the use of vendor-specific advanced modulation techniques and of proprietary FEC algorithms. All OTU4 client ports support the generic G.709 RS(255,239) FEC; this is mandated in G.709, thereby providing common interfaces usable as test access points.

Emerging OTN client applications include OTN handoffs directly from a router port, where Ethernet is wrapped in OTN, or simply between two operators. Among reasons for the use of OTN are the advanced operations, administration, and maintenance (OAM) capabilities that provide robust alarm and error monitoring. In addition, of note is the fact that there is no SONET/SDH standard for 100 G. In the near future, we will see emerging OTN field applications such as Layer 1 OTN switching based on ODU multiplexing. This is an upcoming wave of deployment where lower speed signals can be switched and aggregated onto high-speed OTN links up to OTU4. The key engine behind all this development remains the transport of high-speed Ethernet signals.

Conclusion

In summary, 100 G deployments are accelerating in all regions of the world. Compared to 10 and 40 G transport, there are a number of advances including modulation schemes, and differences such as new physical layers. The main test parameters to consider, such as dispersion and OSNR on line interfaces, and frame loss rate, latency, packet jitter, protocol transparency on client interfaces, remain similar. However, the sheer quantity of information and the criticality of 100 G services make testing a key deployment factor. In addition, the global shift toward the use of OTN client interfaces makes the adoption of stringent test procedures an important consideration.

These elements contribute to the evolution of the telecommunications industry and move us toward the next innovation cycle based on even higher-capacity backbones.



Test & Measurement Regional Sales

NORTH AMERICA	LATIN AMERICA	ASIA PACIFIC	EMEA	www.jdsu.com/test
TOLL FREE: 1 855 ASK-JDSU 1 855 275-5378	TEL: +1 954 688 5660 FAX: +1 954 345 4668	TEL: +852 2892 0990 FAX: +852 2892 0770	TEL: +49 7121 86 2222 FAX: +49 7121 86 1222	