HEAVY READING

White Paper

Deployment & Service Activation at 100G & Beyond

Prepared by

Sterling Perrin
Senior Analyst, Heavy Reading
www.heavyreading.com

on behalf of

JDSU

www.jdsu.com

March 2015

Introduction

The road to 100Gbit/s (100G) transport was a long one that began around 2006, just as the first commercial 40Gbit/s (40G) networks were being deployed. It was also during that year that prominent AT&T Labs Vice President Simon Zelingher (now retired) made a passionate case for the need for 100G transport in his keynote address at Light Reading's Optical Expo conference in Dallas: "We will need 100 Gbit/s by the end of the decade," he stated. AT&T (along with Comcast) pioneered commercial 40G transport in 2006, and his comments were greeted with a great deal of skepticism by many.

However, Zelingher's prediction proved amazingly accurate. The first commercial 100G cards were shipped at the end of 2009 (for Verizon, not AT&T), and the 100G commercial migration was set in motion in 2010. Fast forward to the present, and 100G has risen to dominance in long-haul networks, displacing both legacy 10Gbit/s (10G) and 40G, 100G's former next-generation bit rate rival, which has itself become a legacy technology.

This white paper takes an in-depth look at the market for 100G with a focus on the testing requirements, challenges, and solutions for deploying and activating the technology. In addition, we assess the implications for network test, as the 100G market progresses in two directions:

1. From long-haul applications down to metro and regional networks, and
2. From 100G lines rates to 400G and beyond.

100G Market Overview

We begin with a brief overview of two of the key enabling technologies that made 100G transport a success.

Enabling Technologies

Advanced Modulation Formats

The move away from on-off keying (OOK) modulation, beginning at 40G line rates and continuing at 100G, marked a fundamental change in optical transmission. At 40G, the solution to chromatic dispersion (CD) and polarization mode dispersion (PMD) problems was to break from traditional OOK modulation and introduce new modulation formats that could perform well over existing fiber plants. Phase shift keying (PSK) variants, largely adapted from wireless transmission, became the most popular, including differential PSK (DPSK), differential quadrature PSK (DQPSK) and dual-polarization QPSK (DP-QPSK).

These advanced modulation formats were the trigger that enabled commercially viable long-haul transmission at 40G and higher. While multiple modulation formats were developed and used at 40G, 100G has been developed in a much more cohesive fashion, and suppliers have agreed upon DP-QPSK as the modulation format of choice. (Note that this modulation format is also commonly called polarization multiplexed QPSK, or PM-QPSK.) QPSK is a phase modulation technique used historically in radio communications that transmits two bits per symbol, thus doubling the data rate of a given signal. Dual polarization applies polarization multiplexing to two QPSK signals to double the rate once again (i.e., four bits per symbol in total).

As a result of these modulation advancements, 100G standard transmission at 112G total has become achievable using optics and electronics built for 28G transfer rates.

Coherent Detection With Digital Signal Processing (DSP)

Coherent detection is another key technology enabler for 100G. Like QPSK modulation, coherent detection was also borrowed from radio communications and applied to optical communications. A coherent receiver is able to access the amplitude, phase and polarization of the incoming signal in the electrical domain. DSPs then compensate for CD and PMD impairments – again, all electronically.

The primary benefit of coherent detection with DSPs in 100G is the ability to maintain both high performance, including 2,000+ km transmission distances, and high spectral efficiency – a combination of benefits that is not possible when using direct detection.

In addition, electronic dispersion compensation (EDC) eliminates the need for costly dispersion compensating fiber (DCF) spools, along with their associated amplifiers, that would otherwise be required in long-distance transmission. Eliminating DCFs is also important for reducing latency – an increasingly important consideration for enterprises, particular in finance. Thus, there are applications in metro networks as well.

Global Deployments Overview

In 2010, there were fewer than ten commercial 100G deployments worldwide. By June 2014, Heavy Reading estimates that there were more than 600 commercial 100G wins worldwide (primarily in long-haul networks). **Figure 1** shows Heavy Reading's forecast for long-haul 100G line-side port shipments. We expect 100G shipments to increase at a 50 percent CAGR from 2013-2018.

Figure 1: Worldwide LH DWDM Line-Side Port Shipments

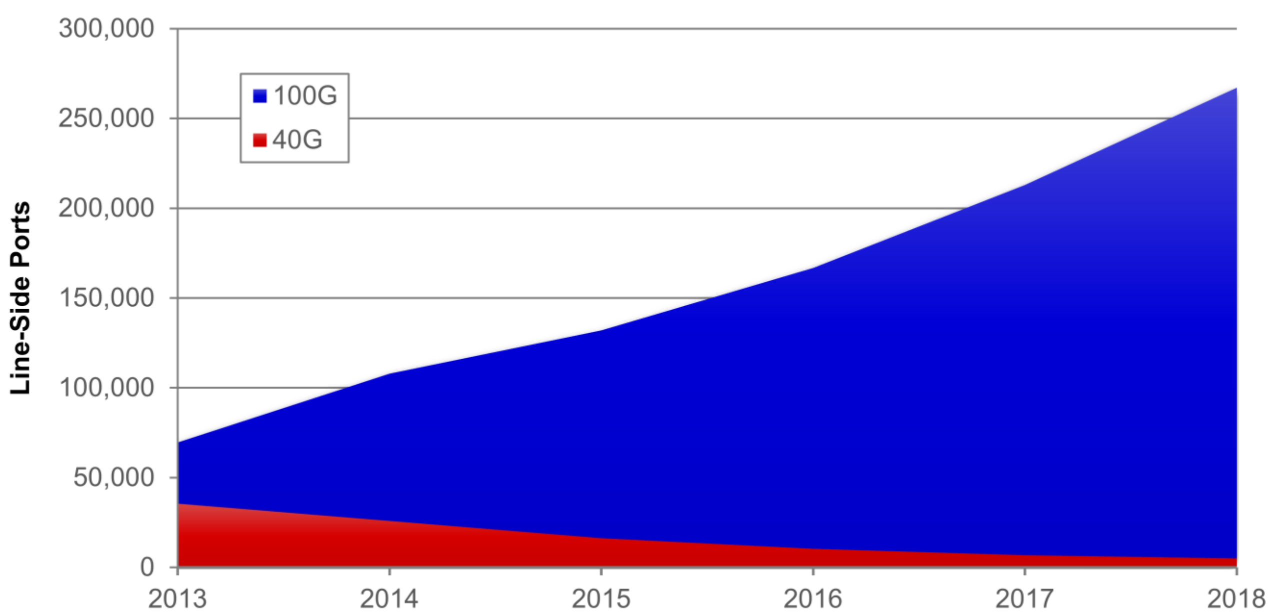

Source: Heavy Reading, 2015

If we look at long-haul DWDM capacity to be shipped during this period (measured by counting Gbit/s of capacity shipped for each line rate), the impact of 100G is even more staggering: By 2018, Heavy Reading forecasts that 100G ports will account for 90 percent of all long-haul capacity shipped.

Most recently, we have seen a migration of 100G from the long haul to the metro. We explore drivers and trends of metro 100G in a later section.

Considerations for Testing 100G

The commercialization of 100G transport heralded new technology innovations, including coherent detection, DSPs and advanced modulation formats.

These innovations overcame CD and PMD sensitivities that imposed crippling distance/reach limitations on early direct detect 100G designs. Coherent 100G with DSPs also changed fiber plant requirements dramatically by eliminating the need for costly dispersion compensating fiber spools and their associated amplifiers, as the high CD and PMD tolerances of coherent detection make this type of dispersion compensation unnecessary. Eliminating DCFs from the line also adds a significant side benefit by lowering latency in transmissions. System vendors estimate that DCFs contribute 8 percent to 10 percent of latency in long-haul networks.

The tradeoff in moving to coherent 100G is that the required technologies added complexity and costs to the transponders.

Key Challenges

In a perfect world, operators would migrate their networks from 10G to coherent 100G simultaneously with new systems and fiber plants, but real-world practice is very different. Although Heavy Reading data shows that 2/3 of *new* long-haul capacity shipped in 2014 was on 100G ports, this does not take into account the massive installed base of 10G systems and transmission cards worldwide. Operators do not rip and replace their 10G systems when they add 100G, so coherent 100G must coexist in a hybrid world.

Networks today are mixed:

- Among new coherent and legacy direct detect-based systems and fibers
- Among 10G, 40G and 100G line rates
- Between fibers in a given cable (if different DWDM systems/vendors are used for legacy 10G/40G and 100G deployments)
- Between different WDM channels with the same fiber (if 10G/40G and 100G cards are housed on the same DWDM system)

Complicating the deployment picture are the following operator network realities:

- Many fibers operating at 10 Gbit/s rates and below do not have documented fiber characterization
- Many metro fibers in operation at 10 Gbit/s rates and below have not been evaluated for PMD or CD, because at these bit rates over metro distances, dispersion was not an issue that merited testing.

Primary Tests for 100G

Today, a myth persists among some network operators that deployment and service activation testing is not needed at 100G, but this is not the case. This section looks at the key deployment and service activation tests that network operators should do when deploying coherent 100G and explains why these tests are necessary.

Fiber Characteristics

Fiber characterization testing is done at the link deployment phase. Key fiber characteristics measured in fiber characterization are connector inspection, insertion loss (IL), optical return loss (ORL), CD, PMD, and spectral attenuation profile (SAP).

Inspection, Attenuation & Reflectance

These conventional fiber link tests are at least as critical at 100G as they are at lower data rates. Coherent 100G has higher optical signal-to-noise ratio (OSNR) requirements than lower data rates, which means that 100G links are more sensitive to power loss (attenuation). Typical contributors to power loss along a fiber are splices, connectors, bends and inherent properties of the fiber itself.

When adding Raman amplification inline, points of excess loss and reflected light can have a significant impact on its efficiency of (e.g., 1dB extra loss = 4dB gain reduction). Therefore, additional care should be taken with the event characterization. Three sets of tools will be required to ensure proper testing:

1. Connector Inspection scope to ensure a pristine optical connector surface
2. IL/ORL meter to provide end-to-end information on the fiber link quality
3. Optical time-domain reflectometer (OTDR) used to detect and locate excessive splice/connector loss, and high reflectance points that could generate multipath interference (MPI) and disturb the Raman amplification

CD & PMD

We know that coherent 100G wavelengths are much more tolerant to CD and PMD than direct detect 100G or even legacy 10G wavelengths, but the tolerances are not unlimited, particularly when it comes to PMD. Maximum PMD tolerance for a coherent receiver is in the range of 25 picoseconds. This is about 2.5x better than 10G. (CD tolerance for coherent 100G, by contrast, is ~20x better than 10G.)

Additionally, CD and PMD tolerances vary from vendor to vendor. These need to be validated up front during the deployment phase in order for both the operator and the equipment suppliers to ensure that the fiber is not at fault if a link fails at a future date. The absence of CD and PMD documentation can lead to much finger pointing (and delays in repair) if there is optical link failure in the future. Another consideration is that PMD tolerance may be acceptable for a given fiber section in isolation, but may have a different reading within the entire fiber cable route that will be considered for the transmission.

While high CD generally does not cause performance issues in coherent 100G networks, it is required to have a minimum amount in order to allow the DSP to properly operate. It also becomes a critical parameter when Raman amplification is inserted into the fiber link. Raman amplification is used to boost ultra-long-haul transmission distances without requiring electrical regeneration. The challenge is that Raman gain varies with fiber type. In a brownfield network, in particular, the right gain applied to one part of the network may be too high or too low elsewhere. Knowing the CD characteristics of each fiber span helps operators apply the right amount of power and adequately select the amplification area.

Optical Signal-to-Noise Ratio (OSNR)

OSNR readings are one of the key parameters that operators have historically used to predict the health of the overall optical signal. Having accurate OSNR data provides operators with the relationship between the loss, noise level, and a good indication of each DWDM wavelength. Two key differences separate OSNR testing at 100G from lower bit rates:

- While coherent 100G is more tolerant to dispersion compared to lower data rates, it is more sensitive to noise.
- While obtaining accurate OSNR readings was relatively straightforward in the past, coherent 100G adds complexity to these measurements and poses challenges.

At the heart of the problem is the fact that the modulated signal bandwidth for 100G is greater than 50 GHz – in the range of 60 GHz to 70 GHz. This results in partially overlapping signals when 100G channels are placed next to each other. With coherent detection and DSPs, this partial signal overlap does not pose limitations on the readability of the signals, but it does create additional crosstalk noise that lowers the OSNR and increases the bit error rate (BER), as shown in **Figure 2**.

Figure 2: Inter-Channel Cross Talk Caused by Overlapping 100G Signals

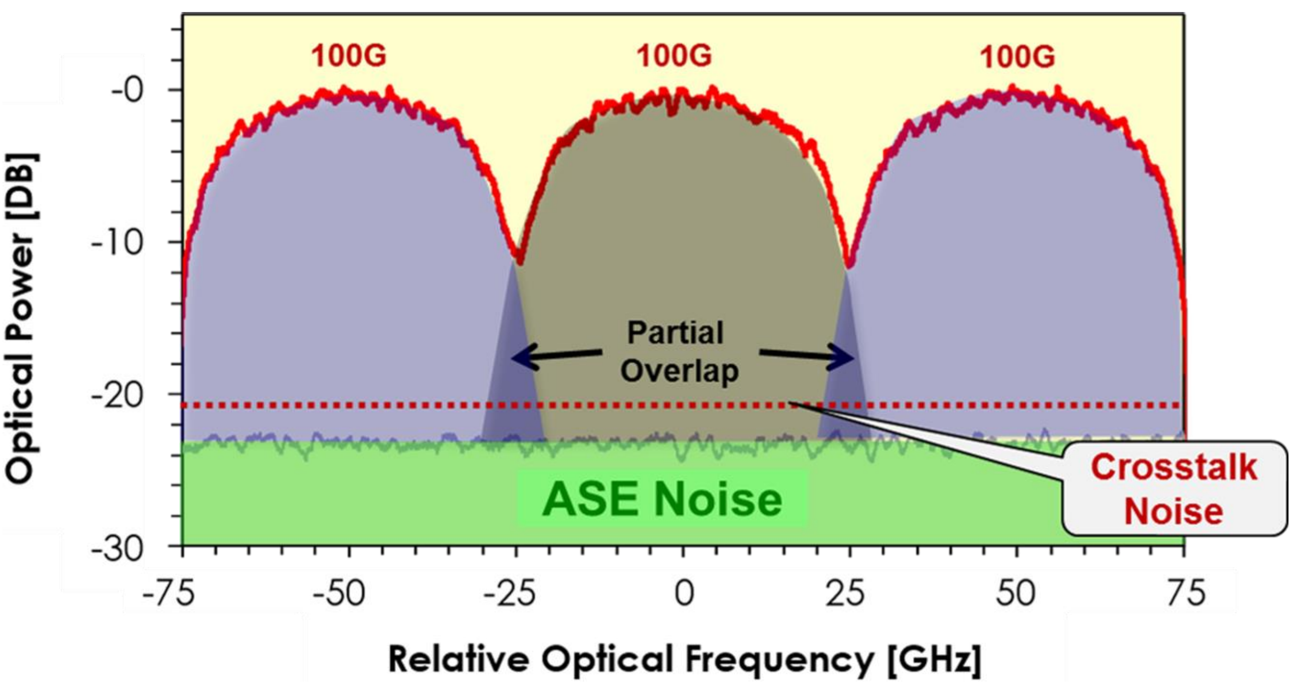

Source: JDSU, 2015

Note that inter-channel crosstalk is not an issue at 10G rates, because modulated 10G signals fit easily into the allotted 50 GHz of spectrum – occupying just 25 GHz to 30 GHz of the 50 GHz total.

The partial overlap in 100G signals also creates problems for obtaining accurate OSNR measurements, as standard out-of-band OSNR measurements will yield inaccurate results due to the lack of space between adjacent 100G channels, as illustrated in **Figure 3**. The current practice is to shut down a channel to obtain in-band OSNR measurements, but this will be an out-of-service method. In the future, operators would like to be able to conduct in-service, in-band OSNR measurements at 100G without switching off optical channels. Commercial in-service, in-band OSNR analyzers for coherent 100G do not exist, but test equipment manufacturers are working on commercializing this important function.

Figure 3: In-Band & Out-of-Band OSNR Measurements at 100G

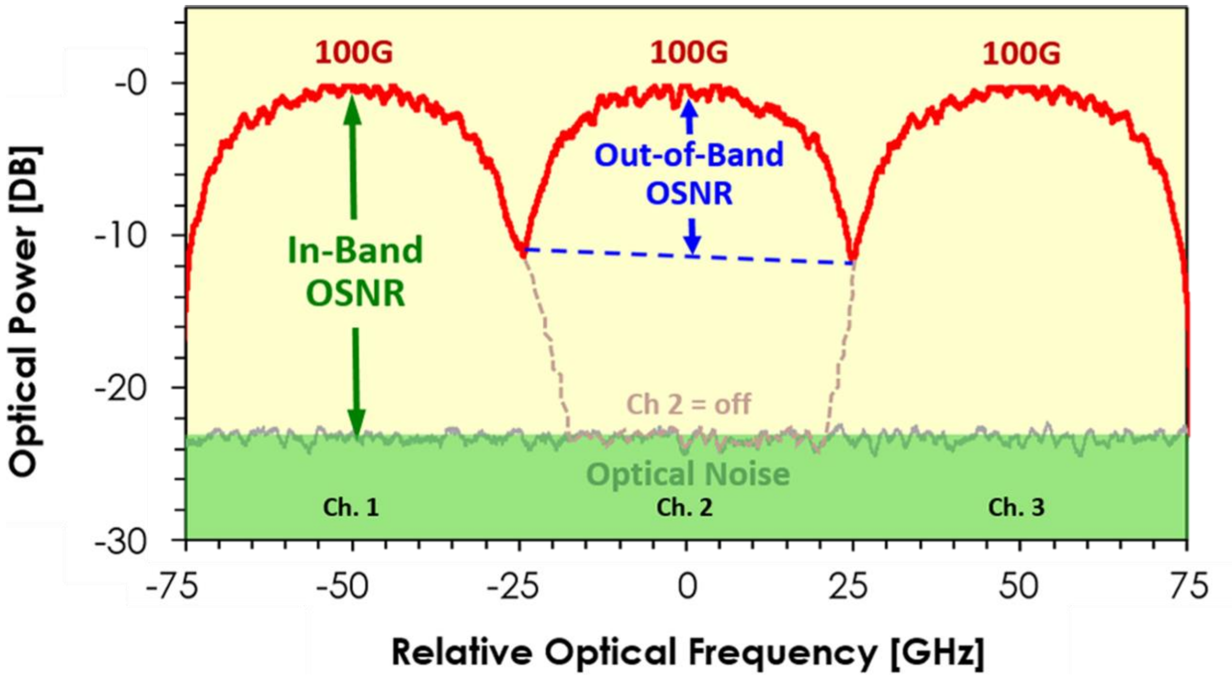

Source: JDSU, 2015

Metro 100G

While 100G growth in long-haul networks continues (as shown in **Figure 1**), the industry is also at the early stages of a metro 100G boom. An Alcatel-Lucent Bell Labs traffic study published in December 2013 yields the following conclusions, which underscore what is driving the new metro network boom:

- Total metro network traffic will increase 560 percent by 2017
- Video traffic (including fixed and mobile) will increase 720 percent by 2017
- Cloud data center traffic will increase 440 percent by 2017

- Total metro traffic will grow approximately twice as fast as traffic going into the backbone network by 2017

The vast majority of metro networks are based on 10G and lower rates today, and 100G is widely viewed as the best means of addressing the explosive growth occurring in the metro. **Figure 4** shows Heavy Reading's global forecast for metro 100G line-side port shipments through 2018.

Figure 4: Worldwide Metro DWDM Line-Side Port Shipments

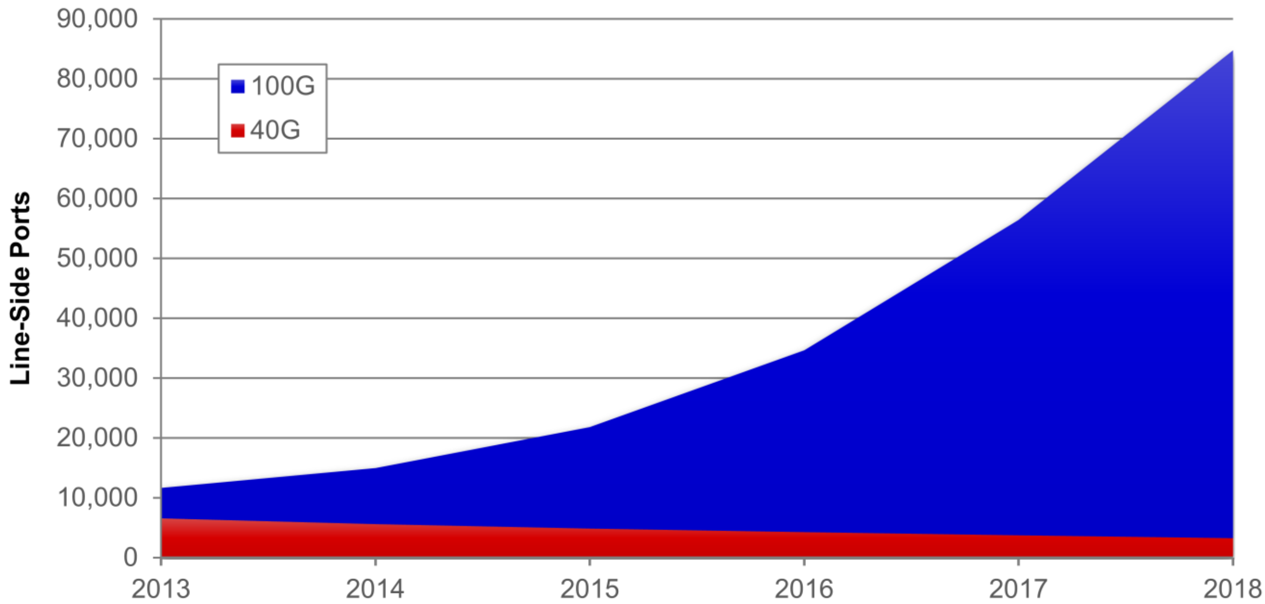

Source: Heavy Reading, 2015

Testing Metro 100G

From a deployment and service activation testing standpoint, we make the following observations:

- The long-haul 100G market is about three years ahead of the metro in its deployment cycle. This means that demand for deployment and activation testing will be high, as service providers with established 100G long-haul networks extend into the metro, and as new service providers focus on the metro deployment of 100G for the first time.
- For the most part, testing requirements for coherent 100G in the long haul apply equally to coherent 100G in the metro – including the fiber characterization and OSNR tests highlighted in this paper. One difference is reach-extending Raman amplification, which is applicable to some ultra-long-haul networks, but will not be needed for metro 100G distances.
- Both metro and long-haul networks contain ROADMs, but they are much more prevalent in the metro. ROADMs are active elements in metro networks, with associated amplifiers and filters, and they contribute to node loss. Loss created by multiple ROADM nodes in metro networks must be understood.

Beyond 100G

As the industry looks beyond 100G – at bit rates including 400G, 1Tbit/s, and higher – there is work to be done. A key technology enabler for bit rates beyond 100G is the super channel – an evolution in DWDM in which several optical carriers are combined to create a composite signal of the desired capacity. Techniques such as orthogonal frequency division multiplexing (OFDM) can be applied to pack super channels closer together. In addition, a new filtering technique called Nyquist filtering reduces signal width compared to conventionally filtered signals. The combination of super channels and Nyquist filtering yields high-capacity channels that also improve the spectral efficiency of the fiber – meaning not just more bits per WDM channel, but also more bits per fiber in aggregate, which is critical for operators.

New advanced modulation formats also come into use at rates beyond 100G. There is particular supplier activity currently using 16QAM modulation, which doubles data rates from 100G to 200G per channel, but results in reduced transmission distances. Although 16QAM can be used in sub-carriers that each form part of a greater super channel, Heavy Reading sees far more near-term interest in applying 16QAM modulation to single-carrier wavelengths to get a quick doubling of capacity from 100G to 200G in metro/regional links (less than 1,000 km), where 16QAM's inherent distance limitations are not a problem.

Beyond 100G rates – particularly when tightly spaced super carriers and Nyquist filtering are used – will require new testing techniques and new innovations. At this point, equipment suppliers have not standardized what exactly needs to be tested at beyond 100G rates.

But the good news is that the industry has some time to settle these requirements. Heavy Reading does not see significant demand for 400G and above channel rates for at least the next five years, as 100G will provide the most efficient and economical means of boosting capacity during this period. (However, the use of advanced modulation to increase channel rates to 200G in metro networks could become prevalent during this time.)

Conclusions

After a long road to commercialization, 100G transport is here now in both long-haul and metro networks. While coherent 100G removed a host of technical barriers that had plagued the commercialization of direct-detect 100G, operators are mistaken to believe that deployment and service activation testing are not needed for coherent 100G. Critical fiber characterization tests at the deployment phase of 100G are: connector inspection, insertion loss, optical return loss, CD, PMD and spectral attenuation profile. In addition to fiber characterization tests, accurate OSNR readings are required to predict the overall health of the optical signal.

With 100G entrenched in long-haul networks, we see two evolutionary paths moving forward. First, we are in the midst of a massive migration to 100G in the metro, a migration that is just beginning. Here, the testing requirements that apply to long-haul networks will continue to apply in the metro. Second, we see an evolution of 100G to higher bit rates in both networks. There is still work to be done here from a testing perspective, particularly when Nyquist filtering is used. The good news, however, is that the industry has some time to settle these requirements, as 100G will have a long lifespan in operator networks.