

100G / 400G

TESTING GUIDE

About EXFO

EXFO develops smarter network test, monitoring and analytics solutions for the world's leading communications service providers, network equipment manufacturers and webscale companies. Since 1985, we've worked side by side with our customers in the lab, field, data center, boardroom and beyond to pioneer essential technology and methods for each phase of the network lifecycle. Our portfolio of test orchestration and real-time 3D analytics solutions turn complex into simple and deliver business-critical insights from the network, service and subscriber dimensions. Most importantly, we help our customers flourish in a rapidly transforming industry where "good enough" testing, monitoring and analytics just aren't good enough anymore—they never were for us, anyway. For more information, visit [EXFO.com](https://www.exfo.com) and follow us on the [EXFO Blog](#).

Table of Contents

1. High-Speed Technology Overview	2
1.1. Introduction to 100G Technology and Market Drivers	2
1.2. 40G/100G Ethernet (IEEE 802.3ba)	3
1.3. Optical Transport Network/ITU-T G.709	4
1.4. 40G/100G Pluggable Optics	5
1.5. CLR4/CWDM4 100G Transceiver Specifications	7
1.6. 400G Ethernet (IEEE 802.3bs)	8
1.7. 400G Interfaces	11
2. 40G/100G System Testing and Lab Qualifications	12
Line Side	
2.1. Transmitter Compliance Verification	12
2.2. Receiver Performance Verification	14
Client Side	
2.3. Physical Layer Test and Transceiver Testing	16
2.4. OTU Frame Structure, Overhead and Testing	30
2.5. Forward Error Correction (FEC)	34
2.6. Response Test, OTN Equipment Conformance and Interoperability Test	35
2.7. ODU-MUX Testing: Multiplex Mapping Test Method for Constant Bit-Rate Services (e.g., SDH)	37
2.8. Packet/Ethernet over OTN Test: ODU0, GMP ODUflex and Hitless Adjustment of ODUflex (GFP) (HAO)	38
2.9. 40G/100G Ethernet Services over OTN Testing	45
2.10. 40G/100G OTN Service-Disruption-Time Measurement	46
2.11. Delay Measurement	49
2.12. TCM and Performance Monitoring	50
2.13. Advanced Intrusive Through-Mode Analysis	51
2.14. OTN Performance Testing: Long-Term Bit-Error Test	52
2.15. Multichannel OTN	54
2.16. General Communication Channels BERT Test	60
2.17. Ethernet Testing in the Lab: RFC 2544 Test	61
2.18. Ethernet Service Aggregation and Switching Capability of POTN Devices	63
2.19. Multilink	65
3. 40G/100G Commissioning and Turn-Up	67
Line Side	
3.1. Link Characterization and Non-Linear-Effect Validation	67
3.2. OSNR Measurements and Polarization-Dependent OSNR Verification	76
3.3. BER Test Using Receiver-Provided Information	85
Client Side	
3.4. Inspecting Fibers and Connectors	86
3.5. Fault Sectionalization and CFP Health Check	88
3.6. Continuity Testing	91
3.7. OTN Long-Term BERT	91
3.8. ITU-T Y.1564 for Ethernet Services Testing	92
3.9. 40GE/100GE Troubleshooting Using Filtering and Packet Capture	102
3.10. Ethernet Service OAM (SOAM)	105
3.11. POTN and MPLS-TP Test	112
3.12. POTN Packet/Ethernet Service Quality of Service Verification	114
3.13. Clock Performance Test for Packet-Switching-Based POTN Device in Response to ODUk/ODUflex Client Signal	115
3.14. Multiservice Application Characterization in OTN and POTN Networks	117
3.15. Ping and Traceroute from 40GE/100GE Interface	117
4. Acronyms—Transport and Datacom	119

1. High-Speed Technology Overview

1.1. Introduction to 100G Technology and Market Drivers

The move toward 100G technology started a few years ago, when the Optical Internetworking Forum (OIF) published a framework document with guidelines for the industry. Among other things, the OIF wanted to avoid what could be referred to as the “market confusion” surrounding 40 Gbit/s technology due to an absence of standards leading to multiple, divergent technologies coming out to market, including non-return-to-zero (NRZ), return-to-zero (RZ), differential phase-shift keying (DPSK) and differential quadrature phase-shift keying (DQPSK). There was no real volume or payback on the development of any of these strategies, because the industry would have to spread too wide. Knowing from day one that the investments for 100G would be significant, the OIF felt that there would be a need to standardize on different components and elements of 100G transmission.

One aspect that the OIF took into consideration is that these new 100G systems should be able to share the physical medium with existing 10G and 40G signals. In addition, several of the networks already had multiple reconfigurable add-drop multiplexers (ROADMs) and 50 GHz channel spacing and therefore, the new 100G signals would need to be compatible with this architecture. For the most part, these links had been compensated for dispersion, and this factor was also taken into account. All of these considerations contributed to a very precise modulation technique.

All other modulation schemes were quickly eliminated, and the OIF advised the industry to use dual polarization or Pol-Mux QPSK, also known as NRZ-DP-QPSK, DP-QPSK or PM-QPSK, as the modulation of choice for 100G transmission. Consequently, transmission started moving toward a combination of amplitude and phase modulation, contributing to yet another new factor to be taken into account within the context of testing.

One of the challenges that the industry is currently facing in terms of 100G line-side testing is that there are no established standards or recommendations to guide manufacturers and operators on how to test the cards, transponders and systems, or as to how links should be qualified prior to the installation of such systems.

The next section outlines EXFO's recommendations on testing these systems and devices based on experience gained through interaction with customers, vendors and operators since 2008.




Figure 1. Spectral width as a function of the modulation format. The shaded area represents a 50 GHz bandwidth.

1.2. 40G/100G Ethernet (IEEE 802.3ba)

The IEEE 802.3ba standard defines a single architecture capable of supporting both 40G and 100G Ethernet, while producing physical-layer specifications for communication across backplanes, copper cabling, multimode fiber, and singlemode fiber. The new 40G/100G Ethernet architecture is based on the concept of virtual lanes or physical coding sublayer (PCS) lanes that are multiplexed and transported over the four or ten parallel wavelengths on single fiber. For example, 100GBASE-LR4 has four optical wavelengths of 25G on the fiber coming from ten electrical CAUI signals of 10 Gbit/s on the host side.

In the process of transmitting a 40G/100G Ethernet signal (see Figure 2), using a 100GBASE-LR4 CFP, the Ethernet packet is broken into blocks that are mapped into PCS lanes using a round-robin distribution. To ensure a clean realignment of the entire block while managing the skew between blocks at the receiving end, a marker block will be added on each logical lane; this marker is sent at a fixed cycle of 210 μ s. The twenty PCS logical lanes will then be multiplexed down to ten lanes at 10 Gbit/s (also called CAUI lanes) using the physical medium attachment (PMA) located in the C form-factor pluggable (CFP) optical interface. The unique property of the PCS lanes is that regardless of how they are multiplexed, all bits from the same PCS lane follow the same physical path. This enables the receiver to correctly reassemble the aggregate channel. At the same time, the unique lane marker enables the deskew operation in the receiver.

The 40G/100G Ethernet implementation introduces new challenges at both the electrical and optical layers. These challenges are due to the data distribution over multiple channels combined with the skew effect. This triggers the need to test the individual concepts, such as PCS lanes, PCS skew and alignment markers, to ensure that the 40G/100G Ethernet device under test supports the proper realignment capabilities, thus compensating for the physical characteristics of the link. By inserting multiple unit intervals (UIs) of skew on different lanes, engineers can identify receiver buffer issues, which is crucial when validating the skew tolerance of the 100G system.




Figure 2. 100G Ethernet transmission using 100GBASE-LR4 CFP.

1.3. Optical Transport Network/ITU-T G.709

Because many initial deployments of 100G Ethernet are addressing bandwidth needs at the core, the optical transport network (OTN) recommendation for transport applications has become prevalent. The OTU4 (112 Gbit/s) rate in ITU-T G.709 addresses the need to carry 100G Ethernet services over OTN. This is, of course, in addition to the fundamental benefits of OTN in general, such as supporting operations, administration, and maintenance (OAM) procedures, and providing a standardized forward error correction (FEC) mechanism for enhanced network performance and better deployment economics.

Similar to the layers defined by the IEEE P802.3ba Task Force, the ITU-T SG15 defined a new layer on top of the optical channel transport unit (OTU) layer, known as the optical channel transport lane (OTL) protocol. The OTL layer is defined as OTL x.y, where “x” represents the data rate and “y” represents the number of optical lanes. For example, OTL 4.4 represents an OTU4 signal running on four wavelengths at 28G. OTL introduces elements, such as OTL lanes, OTL skew and OTL alignment markers, which are conceptually similar to the PCS lanes, PCS skew and PCS alignment marker described in the IEEE P802.3ba standard. As with Ethernet, the OTL layer uses logical lane markers to realign all sent blocks: the main differences in this case is that the marker is embedded in the frame alignment signal (FAS) and multiframe alignment signal (MFAS) bytes of the OTN frame. Therefore, when validating 100G devices supporting OTN capabilities, it is crucial to test the skew at the OTL layer, in addition to testing all other OTN layers (including OTU4, ODU4, OPU4 and FEC), to ensure proper mapping and demapping of 100G Ethernet client signals and reporting of fault management.




Figure 3. Basic OTN transport structure.

Qualifying the performance of the OTN physical layer by testing the line rate of the signal with a pseudo random bit sequence (PRBS) pattern is a key step. This test is critical for designers and system engineers during the development of the 100G line cards. One of the main tests involves injecting a complex PRBS pattern on each of the physical lanes and analyzing the bit error rate (BER) at the receiver end. The ability to add a different pattern per lane provides a whole overview of the crosstalk between the channels. Moreover, a maximum density pattern can be configured to verify the phase-locked loop (PLL) receiver's response to unbalance bit density.

EXFO's 40G/100G solution offers testing capabilities from the physical layer by way of its signal conditioning interface, which is designed for qualification of CFPs to Internet protocol (IP) testing at 40G/100G line rates and Ethernet mapping into the OTN container. These tests are critical from a NEMs perspective during the development, design and validation stages. This also affects service providers' decisions when it comes time to select a 100G system for deployment. Access to robust and reliable 100G equipment is essential to achieving commercial success and accelerated adoption. Part of the solution lies in the rigorous testing of both 100G Ethernet and OTU4; this will lead to easier and faster deployments, as well as increased confidence in operations.

1.4. 40G/100G Pluggable Optics

The 100G CFP LR4 uses a specific converter module, called the gearbox, to translate the 10 electrical CAUI lanes into 4x25G optical lanes. However, this is not the case for the 40G CFP, whereby four XLAUI lanes are simply and directly translated into 4x10G optical lanes. To properly qualify a CFP, the gearbox, thermal stability, synchronization circuitry, host electrical channel, optical channel laser Tx/Rx power and supported voltage will all need to be tested. This is markedly different from the tests that were performed on lower-rate transceivers, for which a loopback BER test was often sufficient. EXFO's 40G/100G solutions offer a suite of integrated applications as part of its CFP Health Check solution, which supports the above tests and provides the end user with quick and clear test results.

Technology has evolved and allowed the introduction of newer 100G transceivers that consume less power and increase port density at a lower cost, namely CFP2, CFP4 and QSFP28. Unlike the first CFP transceiver, CFP2, CFP4 and QSFP28 are not equipped with the gearbox module, which is now present on the host equipment (i.e., network element equipment or test equipment). Each one of these new transceivers receive 4 CAUI-4 lanes in the electrical interface, which are directly converted into 4x25G optical lanes. These transceivers are all interoperable as long as their optical characteristics are compatible: SMF versus MMF, wavelength, number of lambdas and power level sensitivity. It's worth mentioning that for long reach applications, the most commonly used optical interface is the LC connector; for short reach applications, the most commonly used interface is the MPO (or MPT) connector.




Figure 4. CFP internal structure; 100GBASE-LR4 with LC connector; SMF.




Figure 5. CFP4 internal structure; 100GBASE-LR4 with LC connector; SMF.




Figure 6. QSFP28 internal structure; 100GBASE-SR4 with MPO-12 connector; MMF.

MPO stands for “Multifiber-Push-On” or “Multi-Fiber-Pulloff” and is defined by IEC-61754-7 and TIA-604-5-D. These connectors may have different sizes, with MPO-12 and MPO-24 being the most common ones. The following pictures depict both types. The TX fibers are shown in red, the RX fibers are shown in green and the fibers in white are dark fibers.



Figure 7. 12-fiber MPO/MTP®.



Figure 8. 24-fiber MPO/MTP®.

CFP2, CFP4 and QSFP28 transceivers are also supported by the CFP Health Check application and their qualification and testing remain a very significant task when deploying 100G links.

1.5. CLR4/CWDM4 100G Transceiver Specifications

The growth of the data center market is sparking requirements for new elements to be deployed within data center networks and their interconnections. The amount of information generated, transported and processed per fiber-pair in the upcoming years will demand maximum bandwidth, and transceivers will be key to addressing that demand.

The data center market faces the following specific challenges:

- Lowering power consumption
- Lowering cost of equipment
- Implementing small form factor equipment
- Achieving 2 km reach
- Deploying single-mode fiber
- Achieving high port density

The high volume of transceivers that data centers required to meet the above challenges motivated the development of two new very important transceiver specifications: CLR4 and CWDM4.

CLR4

100G CLR4 is a multivendor specification for 100 Gbit/s (4 x 25.78 GBd) focusing on providing solutions to fill data center requirements. It is a low-power, cost-effective solution for transmission distances up to 2 km using duplexed single mode fiber.

CLR4 transceivers are able to operate with and without forward error correction (FEC). Functioning with FEC ensures interoperation with CWDM4 MSA.

This specification is transceiver form-factor agnostic, however the introductory form factor is a QSFP and is based on a CWDM grid with 20 nm spacing to allow lower temperature operation (as defined by IEEE 802.3 CL87 40GBASE-LR4).

Lane	Center wavelength	Wavelength range
L0	1271 nm	1264.5 to 1277.5 nm
L1	1291 nm	1284.5 to 1297.5 nm
L2	1311 nm	1304.5 to 1317.5 nm
L3	1331 nm	1324.5 to 1337.5 nm

Coarse Wavelength-Division Multiplexing (CWDM)

CWDM4 (4 x 25.78 GBd) is an optical interface specification for Ethernet applications including 100GE that also addresses the data center specifications previously mentioned. Initial form factors included in this specification are CFP4 and QSFP28, where FEC is mandatory and distances reach up to 2 km with single-mode fibers—same wavelength range as CLR4.



1.6. 400G Ethernet (IEEE 802.3bs)

Introduction

Like previous Ethernet speed increases, 400G is designed to satisfy the ever-increasing demand for higher bandwidth. High-speed networks are under intense pressure from consumer markets as data and traffic volumes explode with the proliferation of cloud services, smartphones, social media apps, HUD 4k/8k video streaming and ongoing implementation of the Internet of Things. In response, webscale companies are seeking faster transmission rates that allow them to deliver content in increasingly efficient ways.

More and more data centers are being built to support high-capacity data networks. These centers are interconnected using 100 Gbit/s links but will soon migrate to 400 Gbit/s to increase capacity. According to the high-speed roadmap published by the Ethernet Alliance (Figure 9), 400G Ethernet (400GE) is poised to become the next client rate in the Ethernet ecosystem as the industry ramps up to handle the massive demands of hyperscale data centers, service providers and business users.




Figure 9. Ethernet Alliance roadmap

IEEE is developing a new standard for 400GE (IEEE 802.3bs) offering improved capabilities for key aggregation and high-bandwidth interconnect applications, including cloud-scale data centers, internet exchanges, colocation services, wireless infrastructure, service provider and operator networks and video distribution infrastructure.

This application note will highlight the key objectives of the new 400GE rate and the differences relative to 100GE. We will also take a look at the new supported physical interfaces.

400G Ethernet (400GE)

As was the case previously with the 100GE standard, the 400G task force has established a number of key objectives for the new 400GE rates: support MAC data rates of 200 Gbit/s and 400 Gbit/s aligned with the core router for transport applications; support a BER of 10⁻¹³ or better at the MAC service interface; define new physical layer specifications for 200 Gbit/s and 400 Gbit/s operation (see discussion below); and introduce a mandatory and improved Reed-Solomon FEC. This is a key difference over the 100GE objectives, where the FEC was optional. Another difference is the alignment of the different lanes: M:N gearbox conversion has been eliminated, and PCS, electrical and optical lanes are aligned and are multiples of each other.

The figure below illustrates the TX/RX PCS path structure.




Figure 10. PCS layer

MAC layer traffic is fed directly to the encoders/transcoders and then to the alignment marker and the FEC encoder. Distribution is over 16 PCS lanes to the PMA, with PCS lanes running at a nominal rate of 26.5 Gbit/s.

The CDAUI-n lanes are electrical lanes that feed the optical transceiver, where “n” can be 16 lanes or 8 lanes. The CDAUI-8 will support 25G and 28G for Ethernet and OTN 400G support.

Supported media defined in the IEEE 802.3bs 400G standard that can connect the CDAUI-n are described in the Figure 11 below.

Reach	PCS lanes	Rates/lane	Media	IEEE 802.3bs	Modulation
SR (100 m)	16	26.5625 Gbit/s	Parallel MMF	400GBASE-SR16	16 x 25G-λ NRZ
DR (500 m)	16	26.5625 Gbit/s	Parallel SMF	400GBASE-DR4	4 x 100G-λ PAM4
FR (2 km)	16	53.125 Gbit/s	Duplex SMF	400GBASE-FR8	8 x 50G-λ PAM4
LR (10 km)	16	53.125 Gbit/s	Duplex SMF	400GBASE-LR8	8 x 50G-λ PAM4


Figure 11. Client media interfaces

The 400G PCS lane skew, where we measure the difference between the times of the earliest PCS lane and latest PCS lane for the transition of the alignment marker sync bits.

Skew variation may be introduced due to variations in electrical, thermal or environmental characteristics. The skew points at 400G are 4 times those at 100G: see Figure 12 below for the maximum skew per skew point (SP).

Skew points	Maximum skew [ns] ^a	Maximum skew for 200GBASE-R or 400GBASE-R PCS lane [UI]
SP1	29	≈ 770
SP2	43	≈ 1142
SP3	54	≈ 1434
SP4	134	≈ 3559
SP5	145	≈ 3852
SP6	160	≈ 4250
At PCS receive	180	≈ 4780

Figure 12. PCS layer



The maximum skew and skew variation at physically instantiated interfaces are specified at skew points SP1, SP2 and SP3 for the transmit direction and SP4, SP5 and SP6 for the receive direction. 200GE and 400GE skew monitoring will be more challenging than with 100GE.

Forward error correction (FEC)

In addition to a more stringent skew, 400GE incorporates stronger FEC: RS (544,514) FEC is part of the bj FEC solution, with an estimated gain of about $\approx 8\text{dB}$. Key 400G applications are highly sensitive to latency, a drawback when using an FEC mechanism where coding and decoding increases the transmission delay. Unlike 100GE, 400GE uses a mandatory FEC. This means the link won't be error-free: correctable FEC errors are expected where the FEC will compensate for signal degradation and optimize link quality. Some FEC thresholds (Figure 13) can be set to validate the link in pre-FEC mode, which in some cases can reroute user traffic without dropping it. The signaling used is similar to local/remote fault signaling, allowing for improved interoperability and the ability to show link health prior to packet errors. If the pre-FEC threshold is exceeded, it is treated the same as a sync header error threshold with local degrade (LD) and remote degrade (RD) signaling. It will be critical to validate FEC implementation.


Threshold crossing

FEC excessive threshold

- If the pre-FEC BER exceeds this **mandatory** threshold for longer than the Interval, link fault signaling is generated

FEC degrade threshold

- If the pre-FEC BER is higher than this **optional** threshold for longer than the Interval, a link degrade signal is generated



Link Down	—
L Fault Det	—
L Fault Rcd	—
Remote Fault	—
LOA	—
Hi-SER	—
L Deg SER Det	—
L Deg SER Rcd	—
R Deg SER	—

Figure 13. IEEE P802.3bs 400GE Task Force

1.7. 400G Physical Interfaces

Physical interfaces

The IEEE 802.3bs standard also lists proposed 400G interfaces. Network equipment manufacturers (NEMs) are already providing interface cards with CFP8 support. The CFP8 is based on an 8-lane design (8 x 50 Gbit/s) with a reach of at least 10 km. Two types of PMDs are available: an NRZ modulation for 16 x 25 Gbit/s and a PAM-4 modulation for the 8 x 50 Gbit/s and 4 x 100 Gbit/s versions. The CFP8 supports eight times and four times the bandwidth-density of CFP and CFP2 form factors, respectively.




Figure 14. CFP8 functional block diagram

The following optical 400G interfaces are considered:

- **100 m MMF: 400GBASE-SR16:** 400 Gbit/s transmission over sixteen lanes (i.e., 32 fibers total)
- **500 m SMF: 400GBASE-DR4:** 400 Gbit/s transmission over four lanes (i.e., 8 fibers total)
- **2 km SMF: 400GBASE-FR8:** 400 Gbit/s transmission over an 8 wavelength division multiplexed (WDM) lane (i.e., 2 fibers total)
- **10 km SMF: 400GBASE-LR8:** 400 Gbit/s transmission over an 8 WDM lane (i.e., 2 fibers total)

The following optical 200G interfaces are considered:

- **500 m SMF: 200GBASE-DR4:** 200 Gbit/s transmission over four lanes (i.e., 8 fibers total)
- **2 km SMF: 200GBASE-FR4:** 200 Gbit/s transmission over a 4 WDM lane (i.e., 2 fibers total)
- **10 km SMF: 200GBASE-LR4:** 200 Gbit/s transmission over a 4 WDM lane (i.e., 2 fibers total)



2. 40G/100G System Testing and Lab Qualifications

LINE SIDE

The next section describes the first stage of testing, which primarily involves R&D and manufacturing, and also outlines how manufacturers intend to test both transmitters and receivers once the volume starts to increase.

2.1. Transmitter Compliance Verification

Background

Close examination of a quadrature phase-shift keying (QPSK) modulator reveals that it is relatively simple. However, there are a few considerations that need to be taken into account, because they will influence the testing methods and requirements.

A QPSK modulator is based on two Mach-Zehnder modulators, as illustrated in the middle of the block diagram on the next page. These modulators both use the same carrier-wave laser, and the bottom branch is phase-shifted by 90 degrees to generate the “I” (in-phase) and “Q” (quadrature) branches. If the signal at the output of this transmitter were to be viewed using a traditional oscilloscope, it would theoretically produce a perfect line at level one (assuming everything was normalized to one). But in practice, only small dips would be seen along the way (as shown in the top right-hand graph in Figure 15) due to the response time of the electrical components.

Unfortunately, this information is not adequate. Not only is it essential to know that there is a phase shift, it is also essential to know what phase the transition is in to be able to decode QPSK data. This information is shown on the right-hand side of the following graph.




Figure 15. Block diagram of a DQPSK transmitter; intensity and phase diagrams of DQPSK transmitters.

Because existing testing instruments and methods no longer really apply to dual-polarization QPSK signals, the different approaches used need to be re-evaluated.

Going back a few years, in the context of manufacturing testing, the recommended testing approach by the International Electrotechnical Commission (IEC) for legacy signals such as Gigabit Ethernet and 10G SONET/SDH was using an eye-diagram mask to measure several parameters of the signal, taking timing issues, quality of modulation, adjustments, bias, noise and different levels into account to factor in every possible impairment in a transmitter. As such, one of the challenges with coherent transmission is finding a test method that will take the same types of issues into account.




Figure 16. Constellation diagram.

One fundamental point to keep in mind in terms of phase-modulated signals is that, essentially, a vector is being measured. This is represented in the image on the right, where one of the possible phase states is represented by an orange dot, with the vector pointing to the center of that orange dot. The different phase states have a normalized amplitude of “1”; therefore, the phase states of each data point need to be measured. In the time domain, this is represented by the in-phase (I) part and the quadrature (Q) part.

Looking at the generic block diagram of a DP-QPSK transmitter shown in the figure below, it is easy to see several adjustment points that need to be properly calibrated:

- Four radio frequency (RF) drivers (RF1 to RF4)
- Two delays between the I and Q data branches (D1 and D2)
- Four Mach-Zehnder modulator biases (B1 to B4)
- Two bias levels for the phase shifters used to create the quadrature signals (B5 and B6)




Figure 17. Block diagram of a DP-QPSK transmitter.

This means that there are typically 12 potential sources for problems in a coherent transmitter. And, if any single one is not well adjusted or properly tuned, this could have a fairly severe impact on the quality of the transmitted signal.

Some of the most common impairments are based on the following incorrect adjustments:

- Quadrature errors (the phase error between the “I” and “Q” branches of the polarization)
- Poor modulator bias (will affect the opening of the eye diagram)
- “I” and “Q” gain imbalance (instead of having a square constellation, it should be somewhat rectangular)
- Skew (the timing difference between the “I” and “Q” branches)
- Jitter-related issues (data-dependent or random clock jitter)
- Chirp

2.2. Receiver Performance Verification

Receivers are also built to common standards. Essentially, every manufacturer and test instrument vendor is using a similar approach, which is outlined below.

The first component in a receiver is a polarization-diverse optical hybrid, which is used as a mixing medium between an input signal and a local oscillator. The resulting output is transferred to balance detectors and analog-to-digital converters (ADC), and then to a powerful digital signal processor (DSP). This is in keeping with one of the key elements of coherent transmission; namely, that a large part of the receiver performance is highly dependent on the quality of the signal processing performed at the output of the analogue digital converter.




Figure 18. Components of a real-time electrical-sampling optical modulation analyzer.

As such, the challenge is very different, because most people will have very similar hardware that will pretty much behave in the same way. The challenge therefore lies in determining how to evaluate the performance of a receiver independently from the software, because that is where the real performance resides (e.g., decision algorithms, FEC dispersion compensation algorithms and polarization multiplexing).

It is very difficult to distinguish between impairments generated at the transmitter and impairments generated by the actual receiver itself. As mentioned above, this has a lot to do with the software and the receiver. The data provided by the coherent receiver only provides the effective number of bits and bandwidth needed to recover the signal, but is not enough to provide a sufficient evaluation of the behavior of the receiver itself.

It is therefore necessary to have a golden device somewhere as a point of reference. Accordingly, one highly recommended approach is to use a golden transmitter to transmit an ideal signal (with everything properly balanced, tuned and aligned), making it possible to test corner cases by intentionally adding a certain amount of jitter, offset and skew. In addition to enabling full-scale testing of corner cases, this would also make it possible to determine how robust the receiver is to any impairments that could be generated at the transmitter or along the transmission link.

DSP algorithms could be evaluated by simulating impairments on a link. For example, chromatic dispersion (CD), polarization mode dispersion (PMD) and optical signal-to-noise ratio (OSNR) could be generated using emulators and enabling evaluation of the DSP algorithms.

However, DSP algorithm qualification would not need to be performed in manufacturing, because if qualification passes the first time, it will also pass in all other cases. For instance, it is important to determine the capacity of compensation for CD and PMD, in addition to how it will behave, because it will always behave in the same manner. As such, DSP algorithm qualification is most likely only required at the research stage prior to final system validation.

The receiver can therefore be completely characterized by using a golden transmitter to test its integrity, and then comparing the BER or Q factor before and after equalization.

Unfortunately, many other instruments would require a similar procedure. Like golden transmitters, emulators (CD, PMD and OSNR) are stand-alone equipment; they do not exist in a single, calibrated entity.

2.3. Physical Layer Test and Transceiver Testing

40G/100G Ethernet Interface

System architecture of 40GE/100GE: The IEEE 802.3ba standard includes two Ethernet rates, 40G and 100G, on the same architecture. The 40GE/100GE rates retain the 802.3 MAC frame structure and support only the full duplex mode. The system architecture is shown in Figure 19, consisting of a PCS, PMA and physical-medium-dependent layer, as well as FEC modules and the connection interface bus. The chip buses between MAC and PHY are, respectively, XLAUI (40G) and CAUI (100G), while the on-chip buses are, respectively, XLGMII (40G) and CGMII (100G).




Figure 19. IEEE model with CAUI.

40GE/100GE interface: IEEE 802.3ba defines a number of physical media interface specifications, including 1 m backplate connection 40GBASE-KR4, 7 m copper cable (40GBASE-CR4/100GBASE-CR10), 100 m (OM3) or 150 m (OM4) parallel multimode optical fiber (40GBASE-SR4/100GBASE-SR10), and 10 km WDM-based singlemode optical fiber (40GBASE-LR4/100GBASE-LR4); 100GE defines 40 km WDM-based singlemode optical fiber (100GBASE-ER4) and 2 km serial singlemode optical fiber (40GBASE-FR). The physical interfaces are listed below:

	40G Ethernet	100G Ethernet
40 km singlemode fiber (SMF)		100GBASE-ER4
10 km SMF	40GBASE-LR4	100GBASE-LR4
100 m OM3 multimode fiber (MMF) or 150 m OM4 MMF	40GBASE-SR4	100GBASE-SR10
7 m copper cable	40GBASE-CR4	100GBASE-CR10
2 km SMF	40GBASE-FR	
1 m backplate	40GBASE-KR4	

The 40G/100G architecture is based on the concept of virtual lanes or the PCS, which will get multiplexed and transported over the four parallel wavelengths on a single fiber. For example, 100GBASE-LR4 has four optical wavelengths of 25G on the fiber, coming from 10 electrical CAUI signals of 10 Gbit/s on the host side. The interface wavelength specifications of 40GBASE-LR4 are as follows:

Lane	Center wavelength (nm)	Wavelength range (nm)
L0	1271	1264.5 to 1277.5
L1	1291	1284.5 to 1297.5
L2	1311	1304.5 to 1317.5
L3	1331	1324.5 to 1337.5

The interface wavelength specifications of 100GBASE-LR4/ER4 are as follows:

Lane	Center wavelength (nm)	Wavelength range (nm)
L0	1295.56	1294.53 to 1296.59
L1	1300.05	1299.02 to 1301.09
L2	1304.58	1303.54 to 1305.63
L3	1309.14	1308.09 to 1310.19

The main interface specifications of 100GBASE-SR10/LR4/ER4 are as follows:

Description	100GBASE-SR10	100GBASE-LR4	100GBASE-ER4
Signaling speed per lane and range (Gb/s)	10.3125 ± 100 ppm	25.78125 ± 100 ppm	25.78125 ± 100 ppm
Number of lanes	10	4	4
Working distance (km)	0.1 (OM3) 0.15 (OM4)	10	40
Side-mode suppression ratio (SMSR) (dB)	N. A.	30	30
Total average transmission power (maximum) (dBm)		10.5	8.9
Average transmission power per lane (maximum) (dBm)	2.4	4.5	2.9
Average transmission power per lane (minimum) (dBm)	-7.6	-4.3	-2.9
Maximum difference in power between different lanes (dB)	4	5 (OMA)	3.6 (OMA and average)
Maximum safety power of receiver (dBm)	12.4	5.5	5.5
Maximum receiving power per lane (dBm)	2.4	-10.6	-20.9
Maximum difference in receiving power between different lanes (dB)		5.5 (OMA)	4.5 (OMA and average)
Power budget (dB)	8.3	8.5	21.5
Channel insertion loss (dB)	1.9	6.3	18
Optical return loss (dB)	12	20	20
Encircled flux	≥ 86% at 19 μm ≥ 30% at 4.5 μm	N/A	N/A
Extinction ratio (min.) (dB)	3	4	8
Wavelength (nm)	850	See table above	See table above

In addition to the CFP optical interfaces defined in IEEE 802.3ba, there is another commonly used MSA in the industry, namely 10X10 MSA, with a center wavelength of 1500 nm and 10 lanes. The wavelength parameters for 10x10 MSA are as follows:

Lane	Minimum Wavelength (nm)	Nominal Wavelength (nm)	Maximum Wavelength (nm)
L1	1520	1523	1526
L2	1528	1531	1534
L3	1536	1539	1542
L4	1544	1547	1550
L5	1552	1555	1558
L6	1560	1563	1566
L7	1568	1571	1574
L8	1576	1579	1582
L9	1584	1587	1590
L10	1592	1595	1598

IEEE P802.3bm next-generation 40GE and 100GBASE-SR4: IEEE established the P802.3bm task force in 2012. One of their goals is to define four-channel 100G PMA-to-PMA chip-to-chip and chip-to-module four-channel retiming electrical interfaces. This means CAUI-4 based on 4X25G will pave the way for the production of CFP4. In addition, this project will define 40GE 40 km singlemode fiber 40GEBASE-ER4 and 100GE four-channel 100 m multimode fiber 100GBASE-SR4.

Based on parallel optical transport technology, 40G/100G rate transport is achieved by combining a number of wavelengths (usually four or ten) on the same pair of optical fibers. This is very different from rates lower than 10G Ethernet. The adjacent image shows EXFO's FTB-5240S-P spectrometer for 100GBASE-LR10 testing. From the interface, indicators such as center wavelength, power per lane, and power unevenness can be obtained. The user may also select the relevant standard from the instrument interface. The instrument can generate automatic pass/fail information relating to wavelength and power.



Figure 20. Optical spectral analyzer view.

Interface test using 100G analyzer: Interface metric testing is critical to ensuring the interconnection between devices of the same type, or devices from different manufacturers. The main objective of this type of testing is to verify that the input/output parameters of 40G/100G Ethernet devices under test meet the standard requirements, including the correct wavelength and power, in addition to an Ethernet rate within ± 100 ppm of the specified value. Unlike the original 2.5G and 10G single-wavelength transceiver, power and frequency monitoring are carried out on the parallel channels of the interface. Any indicator that is out of normal range at any wavelength will cause damage to the receiving end of the optical module or timing error. For testing of light wavelength and power indicators, please refer to the sections on line-end spectrometry. Notably, EXFO's 40G/100G solutions integrate a number of physical-layer diagnosis tools, including power and frequency information for each channel, and provides quick and clear test functions for each CFP parallel channel. The figures below show the power and frequency measurement of the received signal reported by the CFP. These interfaces provide CFP status information, which is essential to locating network faults.




Figure 21. The EXFO's 88000 Power Blazer product family integrates a number of physical-layer diagnosis tools. The first picture shows the Rx power level of each optical lane. The second screenshot shows the frequency offset for each physical lane.

CFP Test

I. CFP Health Check

As the deployment of 100G links continues to gather momentum, the demand for increased bandwidth is at an all-time high. Network operators are facing significant challenges. Some of these challenges have been well-documented in the past and are now well understood. However, there is one aspect of 100G technology that is too often overlooked when discussing link deployment: the CFP optical interface.

The optical interface is based on parallel optics transmission, through which a number of wavelengths (usually four or ten) is combined on the same fiber pair to form the 100G rate. This functionality is handled by the CFP optics at both ends of the link. This is very different from 10G transmission, which uses serial transmission with a single wavelength to carry the full line rate. At the moment, the technology behind the pluggable SFP+ and XFP optics used for 10G interfaces is more mature and stable, resulting in a product that is extremely reliable and quite inexpensive. In contrast, 100G CFP technology, which is much more complex than 10G SFP+/XFP, is still evolving and does not yet produce consistently reliable optical interfaces.

Based on our experience with multiple carriers around the world, the likelihood of having a faulty or unreliable CFP is quite considerable. If this is not detected by the carrier during link turn-up, it will result in significant schedule delays, or worse, poor link performance. Furthermore, the cost of CFP interfaces today remains high and their availability is still somewhat limited. Consequently, it is neither practical nor economically feasible for carriers to keep spares of these CFP optics readily available for all 100G links and to replace them when they are suspected of poor performance.

In order to ensure the proper deployment and optimal performance of 100G links, it is imperative that carriers use the right tool to test the stability and reliability of every CFP being used. The following section will present EXFO's 100G test solution and its unique set of functionalities specifically designed for CFP qualification, making it the tool of choice for 100G field deployment. In addition to detailing the benefits of such tests, it will highlight other tests for lab qualification. In cases where certain errors appear intermittently on the network, making the link unstable, CFP Health Check is able to quickly identify these errors thanks to its easy-to-use graphical user interface and suite of applications. Besides the above-mentioned CFP laser, power and frequency test, CFP Health Check includes the processes described below.

a. CFP Identification and MDIO Interface

The CFP information page provides detailed information about the module ID, vendor name and supported rates through the CFP control page; it is no longer necessary to remove the CFP in order to read the module detail. This information is also included in the report to simplify CFP tracking. This information is also needed in the field, because multiple Job IDs are generated throughout the day and, depending on the application, a different type of CFP may be used. The figure below shows one of the CFP4 GUIs available on the FTBx-88200NGE. These interfaces allow the user to verify and manipulate the CFP4 electrical pins, indicating the CFP4 status and any available alarms. On the same line, a complete management data input/output (MDIO) interface allows the user to verify the management interface in the CFP4 through a registered read and write access defined by the CFP multisource agreement (MSA).

The MDIO section can be used to read the CFP4 temperature, enable advanced CFP4 functions and even set the CFP4 to troubleshooting mode.




Figure 22. CFP4 status and MDIO access interface.

b. CFP Stress Testing

As an add-on, all of EXFO's Transport and Datacom 40G/100G solutions include a 100G stress test application that can be run locally or remotely. This tool, which is geared more towards lab qualification, will be useful during transmission tolerance tests such as static skew measurement, crosstalk, electrical amplitude and pattern dependency.

The next section examines some of the tests that can be performed automatically, using the CFP Health Check application—eliminating the need for manual intervention while minimizing the chance of error. The CFP Health Check application menu (shown in Figure 23) supports predefined, but configurable OTN and Ethernet tests; one or several tests can be selected at once, thereby optimizing and reducing test time.

In a typical 100G network, all delays must be minimized, because each component of the network, including the CFP itself, can add a significant level of skew, which is the delay between parallel lanes (PCS or OTL). As such, it is important to qualify the skew that is embedded from the CFP during lab qualification. The IEEE 802.3ba standard defines tolerance skew points with specific values for 40G and 100G. Because the skew between the lanes must be kept within specific limits, these thresholds need to be tested, following which the transmitted information on the lanes can be reassembled by the receiver.






Figure 23. CFP Health Check menu interface.

c. Skew Testing

The CFP stress-testing troubleshooting tool's skew function (shown in Figure 24) can be used to inject different skew levels into one or multiple lanes. With the ramp function, these different tests (whether alternate or lane-per-lane) help verify the far-end receiver's buffer tolerance level to skew, based on the standardized skew points of the 802.3ba standard.



Skew Points	Maximum Skew (ns)	Maximum Skew for 40GBASE-R PCS Lane (UI)	Maximum Skew for 100GBASE-R PCS Lane (UI)
SP1	29	≈ 299	≈ 150
SP2	43	≈ 443	≈ 222
SP3	54	≈ 557	≈ 278
SP4	134	≈ 1382	≈ 691
SP5	145	≈ 1495	≈ 748
SP6	160	≈ 1649	≈ 824
At PCS Receive	180	≈ 1856	≈ 928

Figure 24. Skew point testing.

d. Crosstalk Testing and Unframed Parallel PRBS Testing

The CFP Health Check application also provides crosstalk tests using specific PRBS patterns (PRBS 9 to PRBS 31). This important test will verify any signal integrity issues in the CFP that could cause bit errors. This can be detected when an adjacent pattern is detected on one testing lane. Figures 25 and 26 below outline the PRBS insertion on each channel.

	Test Iteration										
CAUI Lane		0	1	2	3	4	5	6	7	8	9
PRBS Pattern	1	31	23	23	23	23	23	23	23	23	23
PRBS Pattern	2	23	31	23	23	23	23	23	23	23	23
PRBS Pattern	3	23	23	31	23	23	23	23	23	23	23
PRBS Pattern	4	23	23	23	31	23	23	23	23	23	23
PRBS Pattern	5	23	23	23	23	31	23	23	23	23	23
PRBS Pattern	6	23	23	23	23	23	31	23	23	23	23
PRBS Pattern	7	23	23	23	23	23	23	31	23	23	23
PRBS Pattern	8	23	23	23	23	23	23	23	31	23	23
PRBS Pattern	9	23	23	23	23	23	23	23	23	31	23
PRBS Pattern	10	23	23	23	23	23	23	23	23	23	31

Figure 25. PRBS-pattern crosstalk test table.

EXFO's Transport and Datacom 40G/100G solutions now include CFP Health Check capabilities, which can be used to validate the entire network, making it extremely valuable to 100G link deployment. The modular FTB-2 Pro/FTB-4 Pro platforms also include optical physical tools, specifically a visual fault locator and a fiber inspection probe, providing field technicians with everything they need to commission 40G/100G networks in a single portable solution.




Figure 26. Tx and Rx pattern configuration.

II. Support for Multiple Transceivers

The CFP MSA defines the module factors that emphasize flexibility at the expense of the CFP's large size. The first 100G transceivers (CFP) were designed based on 10 lanes of 10 Gbit/s signals. Technology enhancements have led to significant improvements in the design of the transceivers, permitting better power performance and higher port density. As a result, CFP2 and CFP4 specifications have been defined to grant interoperability between all transceivers (CFP, CFP2 and CFP4) based on the compatibility of the optical aspects of each transceiver.

CFP2 transceivers are half the size of original CFPs, providing important space-saving benefits. For the same faceplate surface area, network devices can hold more ports, and therefore pass additional traffic. Moreover, these transceivers consume substantially less power in comparison to CFPs. Unlike CFP transceivers, CFP2s do not hold a gearbox module responsible for the dynamic power and temperature control of the transceiver. Instead, these tasks are carried out by the CFP2 host equipment.




Figure 27. EXFO's CFP-to-CFP2 adapter.

Validation tests of the CFP/CFP2 are still required before each deployment in order to analyze the main status of the CFP/CFP2, and ensure that there are no errors. In many cases, the CFP/CFP2 is considered the weakest link in the network; in such cases, it is critical to ensure that the CFP/CFP2 is completely functional prior to deploying 100G.

EXFO's CFP-to-CFP2 adapter, which can be inserted into a CFP interface, allows field technicians to validate CFP2 transceivers before deployment, thus minimizing risk.

EXFO's FTBx-88200NGE Power Blazer and FTB-890/890NGE NetBlazer solutions now provide support for CFP4 and QSFP+/QSFP28 pluggable interfaces. These interfaces offer higher port density, require less power and are much more cost efficient. For example, a CFP4 transceiver is half the size of a CFP2 transceiver, a major improvement. QSFP interfaces are more suitable for short reach applications, with reaches below 10 km, making them attractive for Web-scale data centers, where cost per bit is an important factor. The FTBx-88200NGE Power Blazer includes CFP Health Check support for CFP4 and QSFP+/QSFP28.



Figure 28. CFP4 transceiver.



Figure 29. QSFP28 transceiver.

III. Validating Optics from 100M to 100G

Our lives are increasingly being affected by the speed, amount and reliability of the information that we receive, transmit and consume on a daily basis. Many factors are contributing to this phenomenon, including the way that people communicate and interact, today's ever-changing work environment and the multitude of collaborative tools that we all use. For example, services provided by multiple industries are being transferred to the cloud, the entertainment industry is heavily advancing in the dematerialization of media, the Internet of Things (IoT) is becoming reality, machine-to-machine (M2M) communications is growing worldwide, and the explosion of mobile devices (e.g., cell phones, smartphones and tablets) is generating constant requests for instantaneous content with high precision and performance.

As a result, data centers and service providers around the globe are experiencing massive growth in data communications through their infrastructures, which are constantly evolving to support current and future traffic needs. The evolution of technology is naturally shaped by the challenges inherent to the industry. Network devices are therefore designed in light of important aspects, such as equipment size, power consumption and cost. One component that has a direct influence on all of these three aspects is the transceivers used in network devices. These transceivers, also known as pluggables, vary based on the data rates that they transmit and receive. Optical transceivers have evolved from SFP to SFP+, XFP, CFP, CFP2, CFP4, QSFP+ and QSFP28. The following table provides a description of each one of these transceivers and the rates at which they operate:

SFP	Small form-factor pluggable	Up to 5 Gbit/s
SFP+	Enhanced small form-factor pluggable	Up to 16 Gbit/s
XFP	10 Gigabit small form-factor pluggable	10 Gbit/s
CFP	C form-factor pluggable	40 Gbit/s and 100 Gbit/s
CFP2	C form-factor pluggable	100 Gbit/s
CFP4	C form-factor pluggable	100 Gbit/s
QSFP+	Quad form-factor pluggable	40 Gbit/s
QSFP28	28 Gbit/s QSFP+ quad form-factor pluggable	100 Gbit/s

The technology used in each one of these pluggables is complex and constantly evolving with the arrival of new manufacturers. Moreover, each type of transceiver is at a different maturity level. These devices are often considered the weakest component of communication links and therefore require special attention. Some, particularly those operating at rates below 16 Gbit/s, have reached low price levels, and for this reason, some service providers prefer to just swap devices whenever something appears to not be working as expected. This is a risky approach, because some devices may only be borderline good, meaning that after deployment and service activation, their performance may fall below the expected thresholds.

The cost of devices operating at 40 Gbit/s and 100 Gbit/s is higher. Web-scale companies and service providers cannot afford to have a significant number of spare units, and the impact of faulty links at higher rates is certainly more important. It is also worth mentioning that these devices (especially the CFP4 and QSFP28) are new to the market, and therefore their reliability level has not yet achieved an acceptable standard. For all these reasons, it is extremely important to acquire and use tools that quickly and accurately assess the quality of these pluggable transceivers in order to ensure that deployment and service activations are done right—in a single test execution and a reasonable amount of time. This will reduce the risk of future failures and the need to perform costly troubleshooting activities later on.

EXFO is once again playing a leading role in understanding the needs of service providers, Web-scale companies and network equipment manufacturers (NEMs) with its launch of the iOptics—Intelligent Pluggable Optics test application designed to be the easiest method on the market to quickly validate the sanity of currently available optical transceivers using minimal configuration. The tested aspects and functionalities are adapted to the type of pluggable transceiver being tested and the optical connector type.

When and Where iOptics Should Be Used

The iOptics—Intelligent Pluggable Optics test application was designed as the first tool to be used to assess any transceiver being installed in a network element deployed in the field or lab environment, and serves to validate proper operation of the optical device. The iOptics test should be run in the following circumstances:

- Prior to installation of a pluggable in a network element
- Before the pluggable is used for a test
- During troubleshooting activities (to confirm failure in suspicious devices)

The test uses a physical fiber loopback to validate the device signal continuity, integrity and excessive skew generation. Appropriate signal attenuation should be used to simulate network condition and protect the optical device.

IV. User Scenarios

Service Providers (FTTA)

Service operators all over the world must incur significant expenses during the deployment of their networks, and have different crews responsible for different elements of their infrastructure. In the case of mobile services companies, the installation of remote radio heads (RRHs) is often cumbersome and requires a very specialized and expensive installation team. Accordingly, it is fundamental to ensure that the SFP or SFP+ installed in the RRH is good prior to deployment of the RRH. Having to send a crew back to a site is costly and can be avoided with the iOptics—Intelligent Pluggable Optics test application, which performs the exact verification required to assure that the pluggable is functional and operating as expected.

Data Centers

Web-scale companies have a myriad of ports in their data-center farms and are gradually deploying hardware equipped with QSFP28 ports that operate at 100G. The quality and maturity of these devices is linked to their development, with issues related to shutoff lanes, power levels and control input/output (I/O) interfaces often arising.

The intelligent iOptics test application enables data centers to quickly validate pluggable optical transceivers and meet the aggressive deployment schedules that are a part of today's business environment.

NEMs

NEMs are continuously introducing new hardware and software features to their designs. As such, their engineers must ensure that all optics are reliable during the design verification phase. NEMs cannot afford to jeopardize the validation of their design through use of a faulty or non-reliable transceiver, and therefore, proper operation of all pluggable transceivers is critical and must be validated.

What does the iOptics application validate?

The iOptics test application performs a sequence of tests that is adjusted according to the port selection and the transceiver under test.

Before the sequence of tests is initiated, the iOptics test application performs a basic module validation by collecting key device information such as vendor, module type, module ID, part number and speed (see Figure 30). This enables the user to visually confirm that the desired device is present in the selected port. The basic module validation also assesses what type of test the optical device under test can support.




Figure 30. The QSFP details dialogue box.

The series of tests is as follows, with the test application fulfilling two tasks: the monitoring and execution of each subtest.

1)Power Consumption Monitoring

The Power Consumption Monitoring test measures the current and the power drawn by the selected optical device. This monitoring is performed for the entire duration of the test, following which, a pass verdict is obtained if the maximum power was equal to or less than the power threshold over the entire duration of the test.

2)Optical Device I/O Interface Quick Check

The Optical Device I/O Interface Quick Check validates the MDIO/I2C and hardware-pin operation using a sample of common commands and controls applied to the selected optical device. The verdict provided by this test is therefore a combination of the MDIO/I2C interface test and the control/status pin check results.

3)Temperature Monitoring

The Temperature Monitoring test determines the internal temperature of an optical device in degrees Celsius until the end of the test. A pass verdict is obtained if the maximum temperature is equal to or less than the temperature threshold over the entire duration of the test.

4)Optical TX Power Level Range Test

This test samples the TX power level, compares it with the device's applicable TX power range and provides a verdict. For devices that operate with multiple lanes, each lane is sampled and compared to the applicable TX power range. Additionally, this test reports both the minimum and maximum TX power levels (dBm) collected during the subtest. A pass verdict is obtained if the measured TX power level is within the TX power range defined by the device manufacturer.

A fail verdict indicates that the pluggable is not operating at the range specified by the manufacturer.

5)Optical RX Signal Presence and Power Level Range Test

A physical loopback is required for this test, which validates the presence of an RX signal and samples the optical RX power level. The test compares the obtained RX level with the device's applicable RX power range and provides a verdict. For devices that operate with multiple lanes, each lane is sampled and compared to the applicable RX power range. This test reports both the minimum and maximum RX power levels (dBm). A pass verdict is obtained if the measured RX power level is within the RX power range defined by the device manufacturer.

6)Bit-Error-Rate Test (BERT)

A physical loopback is required for this test, which validates the bit-error performance of the optical module. The validation is performed at the highest rate/protocol supported by the optical device. Additionally, the validation is performed with and without generation of a frequency offset at the respective protocol boundaries.

The user can configure the duration of the test, and a passing verdict will be assigned to the test if the following conditions are met:

- No loss of signal (LOS)
- No pattern loss
- Bit error count $\leq$ BER threshold

For devices operating at multiple lanes, if any lanes present a problem, this will result in a fail verdict.

7)Excessive Skew Test

This test is only executed for the validation of CFP, CFP2, CFP4, QSFP+ and QSFP28 optical devices. A physical loopback is required for this test, which measures the skew associated to each physical coding sublayer (PCS) lane. The intelligent iOptics test application automatically sets up an OTN BERT or an EtherBERT test based on the highest rate supported by the device.

The skew threshold is also automatically configured by the iOptics test application based on the rate/protocol used. A pass verdict is assigned to the test when the highest measured skew is equal to or less than the skew threshold.

In the event of a fail verdict during the execution of any subtest, the test is aborted and the fault is reported to signal that the faulty transceiver device should not be used.

V. iOptics at a Glance

The iOptics application only requires configuration of a few parameters; the majority of configuration is done for the user by the application. See Figure 31 for an example of the configuration page.




Figure 31. The iOptics test configuration page.

Once the configuration has been completed, the user can start the test, which will bring up the Summary page containing the key results and verdicts associated with each phase in the test sequence, as well as monitoring tasks related to each phase. A progress status is provided for each test. It is worth noting that several tests are executed quite quickly.

The Summary page includes the following sections:

- The start time of the test
- Subtest sequence
 - General (for each subtest)
 - > Summarized verdict
 - > Progress indication
 - I/O Interface Quick Check
 - > MDIO or I2C pass/fail
 - > Control/status pin pass/fail
 - Optical TX Power Range Test
 - > The minimum/maximum power value with a color status (green or red) indicating whether the value is in or out of range
 - Optical RX Power Range Test
 - > The minimum/maximum power value with color status (green or red) indicating whether the value is in or out of range
 - Bit Error Test
 - > Bit-error count, LOS or pattern loss indication
 - > Green indicator status appears if the bit-error count is equal to or less than the bit-error threshold
 - > Red indicator status appears if the bit-error count is greater than the bit-error threshold

- Excessive Skew Test
 - > Maximum skew (any lane) or LOS indication
 - > Green indicator status appears if the maximum skew is equal to or less than the threshold
 - > Red indicator status appears if the maximum skew is greater than the threshold
- Monitoring
 - Power (W)
 - > Wattmeter representation with a pass/fail verdict region indication (green vs. red region)
 - > Power (current/maximum)
 - > Current (current/maximum)
 - > Pass/fail verdict icon (power only)
 - Temperature (°C)
 - > Thermometer representation with pass/fail verdict region indication (green vs. red region)
 - > Temperature (current/maximum)
 - > Pass/fail verdict icon

Conclusion

In summary, it is paramount that all optics operating from 100M to 100G be validated, regardless of the scenario in which they are used. All players in the industry are seeking the ability to assure sourcing predictability. High-speed optics remain expensive, and companies cannot afford to waste time and resources through service outages and costly visits to faulty sites. The iOptics test application offers a multitude of benefits to NEMs, service providers and Web-scale companies, including:

- Quick sanity check for next-generation 100G network pluggables
- Rapid analysis of the optical interfaces, with quick pass/fail verdicts
- Support of high-speed optics (CFP/CFP2/CFP4 and QSFP+/QSFP28), as well as pluggables operating at lower rates (SFP and SFP+)
- Support for CLR4/CWDM4 100G specifications for data center market

2.4. OTU Frame Structure, Overhead and Testing

The optical channel transport unit (OTU) frame mainly consists of three parts:

- Framing: frame alignment signal (FAS) and multiframe alignment signal (MFAS)
- OTU, optical channel data unit (ODU), optical channel payload unit (OPU) overhead
- OTU FEC

Framing

When transmitting serial blocks of data in an optical transport system, it is essential for the receiving equipment to identify the block boundaries. The ability to identify the starting point in the OTN is accomplished through the use of framing bytes, which are transmitted in every frame. The OTU framing structure is divided into two portions: the FAS and the MFAS.

The **FAS** uses the first six bytes in row one and columns one to six. G.709 uses FAS to provide framing for the entire signal and to identify out-of-frame (OOF) and loss-of-frame (LOF) conditions.

The **MFAS**: G.709 supports the multiframing structure, in which some of the OTUk and ODUk overhead signals could span multiple OTU frames. Examples are the trail trace identifier (TTI) and tandem connection monitoring activation (TCM-ACT) overhead signals. A single MFAS byte is used to extend command and management functions over several frames. The MFAS byte is defined in row one and column seven of the G.709 frame, and is incremented for each OTUk/ODUk frame, providing a 256 multiframe structure.




Figure 32. G.709 frame alignment.

OTU Overhead

The OTU overhead is composed of the section-monitoring (SM), general-communications-channel (GCC0) and reserved (RES) bytes, as shown in the figure on the next page.

- **SM bytes:** These are comprised of the TTI, bit-interleaved parity-8 (BIP-8), backward error indicator (BEI), backward incoming alignment error (BIAE), backward defect indicator (BDI), and incoming alignment error (IAE).
 - **SM-TTI:** This is a one-byte overhead field defined to support 64-byte trace signals. TTI is used to identify a signal from the source to the destination within the network. The TTI includes the so-called access-point-identifiers (API) fields, which are used to specify the source access point identifier (SAPI) and destination access point identifier (DAPI). The APIs include information regarding the country of origin, network operator and administrative details.
 - **SM error BIP-8:** This is a one-byte, error-detection code signal. The OTUk BIP-8 is computed over the OPUk area of a specific frame and inserted in the OTUk BIP-8 overhead two frames later.
 - **SM-BDI:** This is a single-bit signal defined to convey the signal fail status detected in the upstream direction.
 - **SM-BEI and SM-BIAE:** This is a four-bit signal used to convey (in the upstream direction) the number of interleaved-bit blocks detected in error by the section monitoring BIP-8 code. This signal is also used to convey (in the upstream direction) an IAE condition that is detected in the section monitoring IAE overhead.
- **The GCC0 field,** which resembles the data communications channel (DCC) in SONET/SDH, is currently undefined. However, it will likely be used for functions such as network management or control plane signaling for a protocol (e.g., generic multiprotocol label switching [G-MPLS]).
- **The RES fields,** found throughout the overhead, are set aside for future use.

Optical Channel Data Unit Overhead

The optical channel data unit (ODU) overhead is shown in the figure below.

The ODU overhead supports two classes of ODUk maintenance signals, reported using path-monitoring-overhead (PMOH) status (STAT) bits and tandem connection monitoring (TCM) STAT bits. Through either PMOH or TCM STAT bits, the following ODU conditions can be reported: alarm indication signal (ODUk-AIS), open connection indication (ODU-OCI), locked (ODUk-LCK), and generic AIS. In addition, the ODUk overhead supports automatic-protection-switching (APS) functionality. The ODUk overhead is broken into the following fields: RES, PM, TCMi, TCM ACT, FTFL, EXP, GCC1/GCC2 and APS/PCC.




Figure 33. ODU overhead and structure.

- **RES** bytes are undefined and are set aside for future applications.
- **Path monitoring (PM)** enables monitoring of particular sections within the network, as well as fault locations in the network. The PM bytes are configured in row three, columns 10 to 12, and include subfields similar to those in SM, including TTI, BIP-8, BEI, BDI and STAT subfields.
- **The tandem connection monitoring (TCMi)** fields define six ODU TCM sublayers. Each TCM sublayer contains a TTI, BIP-8, BEI/BIAE, BDI and STAT subfield associated with a TCMi level.
- **Tandem connection monitoring activation/deactivation (TCM-ACT)** is a one-byte field located in row two and column four. TCM-ACT is currently undefined in the standard.

- **Fault type and fault location (FTFL)** is a one-byte field located in row two and column 14 of the ODU overhead and is used to transport an FTFL message that is spread over a 256-byte multiframe in order to send forward and backward path-level fault indications. The forward field is allocated to bytes 0 through 127 of the FTFL message. The backward field is allocated to bytes 128 through 255 of the FTFL message.
- **Experimental (EXP)** is a two-byte field located in row three and columns 13 and 14 of the ODU overhead. The EXP field is not subject to standards and is available to network operators to support special applications.
- **General communication channels one and two (GCC1/GCC2):** These are similar to the GCC0 field; GCC1 is located in row four, columns one and two, while GCC2 is located in row four, columns three and four of the ODU overhead.
- **Automatic protection switching and protection communication channel (APS/PCC)** is a four-byte signal defined in row four and columns five to eight of the ODU overhead. The APS/PCC field supports up to eight levels of nested APS/PCC signals.

Optical Channel Payload Unit (OPU) Overhead

The OPU overhead and structure is shown in the figure below.

The OPU overhead is located in rows one to four of columns 15 and 16, and includes the following fields:

- **Payload structure identifier (PSI)** is a one-byte field allocated in the OPU overhead to transport a 256-byte payload structure identifier (PSI) signal. The PSI byte is located in row four, column 15 of the OPU overhead.
- **Payload type (PT)** is a one-byte field defined in the PSI[0] byte and includes the PT identifier that reports the type of payload being carried in the OPU payload to the receiving equipment.
- **The Multiplex structure identifier (MSI)** field is used to encode the ODU multiplex structure in the OPU. Located in the mapping-specific area of the PSI signal, the MSI indicates the content of each tributary slot (TS) of an OPU.
- **Justification-control (JC) overhead** consists of justification control (JC), negative justification opportunity (NJO), and positive justification opportunity (PJO) signals used in the ODU multiplexing process.




Figure 34. OPU overhead and structure.

OTN Overhead (OH) Test

The simplest way to validate an OTN overhead is to connect an OTN analyzer directly to the network element under test. Below is a typical overhead analysis interface generated by EXFO's FTBx-88200NGE OTN analyzer:



Figure 35. OTN overhead configuration.

2.5. Forward Error Correction (FEC)

FEC is a major feature of OTN and uses a 255,239 Reed-Solomon (RS) algorithm code to produce redundant information that gets concatenated with the transmitted signal and used at the receive interface to help identify and correct transmission errors. The FEC algorithm is proven to be effective in systems limited by optical signal-to-noise ratio (OSNR).

The following figure illustrates the process in which the FEC protocol interleaves one overhead byte and 238 data bytes to compute 16 parity bytes, thus forming 255-byte blocks, i.e., the RS (255,239) algorithm. The key advantages of interleaving the information are reducing the encoding rate of each stream relative to the line transmission rate, and reducing the sensitivity to bursts of error. The interleaving, combined with the inherent correction strength of the RS (255,239) algorithm enables the correction of transmission bursts of up to 128 consecutive erroneous bytes.




Figure 36. OTN rows and subrows.




Figure 37. FEC interleaving in OTN.

To determine whether the FEC algorithm is functioning correctly, generate some correctable FEC errors with the test equipment and verify whether the errors can be detected and then correct any bit errors. The signal received by the test instrument should be free of bit errors. You can also use the instrument to generate uncorrectable errors, following which the system under test should be able to properly report the alarm to the network administrator.

Finally, use the instrument to run the FEC pressure test by generating a random bit error in OTU frame. The device under test (DUT) should be able to correctly identify all the bit errors and return an error-free OTN signal.

Test Connection and Method

The diagram below illustrates the EXFO FTB-2 Pro host connection for the OTN FEC test.




Figure 38. OTN FEC test configuration.

2.6. Response Test, OTN Equipment Conformance and Interoperability Test

Verification of correct alarm and error responses is an important aspect of verifying the performance of OTN network element devices. This test is performed by using the test instrument to generate different alarms and errors, and then checking whether the DUT is responding properly. One single OTN stimulus signal should be able to trigger the DUT into generating a number of upstream and downstream responses simultaneously. For instance, if an instrument generates a path alarm indication signal (AIS-P) alarm, the DUT should send an OTU-BDI alarm upstream, and an OTU-AIS alarm downstream.

Device Connection Diagram:




Figure 39. Connection diagram of response test using EXFO's FTB-2 Pro.

The figure below shows common stimulus signals, along with corresponding upstream and downstream responses of the DUT:

Stimulus	Upstream Alarm/Error	Downstream Alarm/Error
LOS-P, LOF, AIS-P, LOM	OTU BDI	ODU AIS
OTU BIP-8	OTU BEI	—
OTU TIM	OTU BDI	ODU AIS
OTU IAE	OTU BIAE	—
ODU AIS	OTU BDI	ODU AIS
ODU BIP-8	OTU BEI	—
ODU TIM	OTU BDI	—
ODU OCI	OTU BDI	ODU AIS
ODU PLM	—	ODU AIS

Figure 40. Upstream and downstream DUT response.

Conformance and Interoperability Test

Unlike response tests, as part of the conformance test, the instrument must generate a series of stimulus signals while the DUT detects different events under the conditions, specified in the standard. The standard generally stipulates conditions for generating and eliminating an alarm, e.g., a certain number of frames or milliseconds. Therefore, as part of the conformance test, the stimulus signal must generate various numbers of frames in order to verify that DUT meets the standard requirements.

2.7. ODU-MUX Testing: Multiplex Mapping Test Method for Constant Bit-Rate Services (e.g., SDH)

The OTN standard defines various synchronous and asynchronous payload-bearing methods. The client signals include a synchronous digital hierarchy (SDH) service, OTUk service, Ethernet service, Fibre Channel (FC) service, common public radio interface (CPRI) service, Gigabit-capable passive optical network (GPON) service, etc. Service types mainly consist of constant-bit-rate (CBR) signals and packet signals. The adaptation of packet service to ODUk service is the key to supporting testing of the data-service passive optical transport network (POTN), mainly including adaptation and mapping of the Ethernet/MPLS-TP signal to the ODUk channel.

ODU multiplex functionality testing is also a key parameter that needs to be validated as part of G.709. In order to determine the appropriate multiplexing capability of the network element under test, the test equipment is used in OTN-decoupled mode to generate a lower-rate signal on the transmit side. The transmitted signal then gets multiplexed within a higher-order ODU signal on the G.709 network element, along with the proper overhead and FEC bandwidth to compose the final OTU signal. Finally, the received OTU signal is checked at the test equipment in order to verify the proper multiplexing with the proper frequency justification and synchronization, as shown in Figure 41 below.

ODU multiplexing structures are identified by the OPU multiplex-structure-identifier (MSI) bytes.

The multiplex-structure-identifier-mismatch (MSIM) alarm is used to identify a mismatch between the accepted and expected MSI.




Figure 41. ODU-multiplexing test configuration.

In the OTN standard, a method is designed to transport synchronous or asynchronous payload. An instrument can be operated in decoupling mode to verify whether the OTN device has properly mapped/demapped the loaded CBR service signal to OPU, where the difference between the SDH/SONET signal and OPU rate is adjusted by the justification byte. A typical instrument connection is shown in Figure 42 below:




Figure 42. Client signal mapping test configuration: both the transmitting and receiving ports require the decoupling mode.

In the preceding figure, the meter sends SDH/SONET signals and receives OTN signals containing SDH/SONET payload. Of course, testing can also be carried out reversely, i.e., by sending the OTN signal and receiving the SDH/SONET signal in order to verify that the mapping action of the OTN device is correct.

2.8. Packet/Ethernet over OTN Test: ODU0, GMP ODUflex and Hitless Adjustment of ODUflex (GFP) (HAO)

The Ethernet, which was responsible for the first ever successful deployment of OTN devices for SDH/SONET signal transport, continues to break the application limit of the enterprise network realm and extend towards the public network realm. In response, the International Telecommunication Union (ITU) has done a lot to support Ethernet by defining and standardizing the transport of multirate Ethernet signals over the OTN network, including Gigabit Ethernet, 10G Ethernet, 40G Ethernet and 100G Ethernet.

ODU0

One important concept for Ethernet support is to define a transmission container with a suitable size for Gigabit Ethernet. In the initial definition of OTN, the optical channel (OCh) data unit ODU1 is the smallest transmission container, exclusively used to transport single STM-16 signals with a payload capacity of 2 488 320 kbit/s. This also signifies that if the ODU1 is merely transporting single Gigabit Ethernet signals, a lot of bandwidth will be wasted. Hence, the ITU defines the ODU0 as half the payload bit rate of the OCh payload unit OPU1, i.e., 1 238 954.310 kbit/s.

The PCS of Gigabit Ethernet adopts the 8B/10B line-coding method and generates a bit rate 25% higher than the 1 Gbit/s information rate. In order to implement bit pass-through and clock pass-through of Gigabit Ethernet, the PCS signal must be kept and transported. However, the OPU0 payload bit rate is not sufficient in carrying 1.25 Gbit/s PCS signals. Therefore, it is necessary to perform timed transparent transcoding (TTT) on the PCS signals at the OTN network entry to reduce the bit rate of the transported signals while retaining the information required to restore the PCS signals at the OTN network exit. TTT treatment employs the transparent mapping mechanism in the ITU G.7041 specification for 8B/10B payload coding, as defined in the general framing specification (GFP-T). The GFP frame header is added, but there is no GFP-based rate adjustment or GFP payload frame check sequence (FCS). An EXFO OTN analyzer is designed to test ODU0 and GFP-T; an example of the test interfaces is shown below.





Figure 43. GFP-T overhead-byte analysis interface on EXFO's FTBx-88200NGE OTN analyzer.

ODUflex

With the application of multiservice interfaces (e.g., Ethernet, FC, CPRI and GPON), the existing OTN container and conventional mapping procedures, such as the asynchronous mapping procedure (AMP) and the bit synchronous mapping procedure (BMP), have been unable to meet the requirements needed to carry OTN full service. ODUflex offers a flexible rate adaptation mechanism that enables OTN to efficiently carry full services (including IP) and maximize the utilization of line bandwidth. Currently, ITU-T defines two types of ODUflex, one is ODUflex based on CBR service, where the rate is arbitrary and CBR service is packaged to this kind of ODUflex through synchronous mapping; the second type is ODUflex based on packet service, where the rate is N times that of HO ODU slot, and the packet service is packaged to the ODUflex via generic framing procedure (GFP). ODUflex gives OTN the capacity to carry different services in future, especially to improve the efficiency of packet-service support.

ODUflex is able to provide flexible containers of variable sizes and client signals of variable sizes, using a principle similar to virtual concatenation (VCAT) in multiservice-transport-platform (MSTP) technology. ODUflex offers an efficient yet simple way to map variable-rate packet services or fiber-channel-like constant rate services to ODU. ODUflex employs a 1.25G tributary slot (ODTUGk) to create a variable container, thus ensuring that the client signal is mapped into this container and then transported via the ODUk signal. For non-CBR service signals, ODUflex employs the frame-mapped generic framing procedure (GFP-F) to map signals. The main advantage of ODUflex is that it enables unoccupied slots to be reused and provides similar bandwidth-variable block-oriented-device (BOD) services.




Figure 44. Reusable capacity via ODUflex.

Hitless adjustment of ODUflex (HAO) is a scheme for hitless adjustment of the bandwidth of an ODUflex (GFP) connection in an OTN network, and is based on the G.709 recommendation. Through the HAO protocol, the system is able to increase or decrease the client signal data rate carried by the ODUflex (GFP) over the end-to-end path. The process closely resembles the dynamic link capacity adjustment scheme (LCAS) for adjustment of VCAT in the MSTP system. However, it should be noted that the HAO requires the involvement of all the nodes on the entire path, as opposed to just two ends (which is the case for VCAT, a merely logical concatenation). In general, every VCAT tributary can have different paths, with variable time delay differences being off set on the receiving end with cache. However, HAO requires all ODUflex tributaries to go along one path, thus addressing the significant need for cache that affected LCAS in the past.

HAO overhead is mainly implemented by resize control overhead (RCOH). RCOH is located in the HO OPUk tributary slot overhead (TSOH) and the OPUflex overhead. RCOH mainly consists of two parts, i.e., link connection resize (LCR) control protocol and bandwidth resize (BWR) control protocol, respectively.

Test example: mapping paths of EXFO FTBx-88200NGE tester Pattern-ODUFlex-ODU4.




Figure 45. Mapping of Pattern-ODUFlex-ODU4.

The ODUflex-related overhead analysis interface is as follows:



Figure 46. ODUflex overhead byte analysis interface on EXFO's OTN analyzer.

In an OTN network, the packet Ethernet signal will be the main service type. The GE/10GE/40GE/100GE Ethernet signal can be mapped to OPU0/1/2/3/4/flex and OPU1/2/3-X by means of GFP-F. In addition, the 40GE Ethernet signal can be mapped to OPU0/3 by the generic mapping procedure (GMP). The 10GE signal can also be mapped to OPU2e by the BMP. The 100GE signal is restored as 64B/66B codes by the GMP and is mapped to OPU4.

GMP

The OPU payload rate (initially defined for OTN) matches the STM-n (n=16, 64, 256) client signal rate very well. Using the simple AMP method, the STM-n signal can be mapped into OPU, and lower-order ODU can be multiplexed to the tributary slot of the higher-order OPU. With the advent of the new client signal based on definition of OPU4 for 100G Ethernet, in many cases, the adjustment scope of AMP mapping fails to cover the rate difference between the client signal and the server signal (note that a lower-order ODU can be considered as the client signal of a higher-order OPU service layer).

A more flexible or more generic method exists, namely, the GMP. In all cases (e.g., the maximum frequency offset of the client signal in ppm and the minimum frequency offset of the server signal in ppm), this method is able to map any client signal rate into a corresponding payload container as soon as the server signal rate is determined to be higher than the client signal rate.

In the process of mapping the 40G/100G Ethernet client signal into OTU3/OTU4, where the service frame (or multiframe) can contain data or stuffing information, the Sigma-Delta data/stuffing algorithm is used to distribute the data as evenly as possible within the service frame. During this process, to guarantee that the client signal is restored completely without any demapping errors, the receiving end must be aware of the data stuffing method for every frame/multiframe at the transmitting end.

In the GMP, the OPUk overhead contains the PSI, PT, and six justification bytes (JC1, JC2, JC3, J4, J5, and J6) as the justification-control bytes of the GMP, as shown in the figure below. Bytes JC1, JC2 and JC3 contain 14-bit Cm identifiers (C1, C2...C14), an increment indicator (II), decrease indicator (DI), and an 8-bit error-check CRC-8 that is used to check bytes JC1, JC2 and JC3 for errors.

The GMP uses these justification-control bytes to notify the receiving end of the number of bytes in the OPUk payload, and this process will be carried out in every frame. The size of the payload bytes is represented by the Cm value, whereas the variation in the byte size of every OPU frame is represented by two specific indicators, the II and the DI. Bytes JC4, JC5 and JC6 contain the 10 bit CnD byte (D1, D2 ...D10), bytes four through eight are 5-bit error-check CRC-5, and byte nine is reserved (RES) for later use. The CnD parameter is used in the GMP, providing a scheme that can be used to carry some client signals with more stringent restrictions on jitter, e.g., SONET/SDH signals.




Figure 47. GMP justification-control bytes.

The GMP solves mapping of the client signal to LOODU, and LOODU to HOODU, e.g., signals such as STM-1/STM-4 and GE are mapped into LOODU0 via the GMP, and ODUflex is mapped into HOO-DUK ($k=2,3,4$), etc. Transport of the 40G/100G Ethernet signal in the OTN network is implemented by GMP mapping, followed by the corresponding transcoding.

Ethernet Signal Mapping and Multiplexing Test

In a POTN network, the packet Ethernet signal will be the main service type. The GE/10GE/40GE/100GE Ethernet signal can be mapped to OPU0/1/2/3/4/flex and OPU1/2/3-X by means of GFP-F. The 40GE signal can also be mapped to OPU0/3 by the GMP. The 10GE signal can also be mapped to OPU2e by BMP, whereas the 100GE signal is restored as 64B/66B codes by the GMP and mapped to OPU4.

Mapping and Multiplexing of a GE Service

There are two major mapping methods for a GE service. The signal can be mapped into OPU0 via the GMP, and into ODTU01 via the AMP, and then multiplexed into ODU1 and OTU1; or it can be mapped into OPU0 via the GMP, and into ODTU2.1, ODTU3.1 and ODTU4.1 via the GMP, and then multiplexed into ODU2/3/4 and OTU2/3/4. Therefore, multiplexing of the GE service into 100G OTN can be done through multiple paths.

The adjacent figures show two typical GE signal paths multiplexed into 100G OTN, as verified using EXFO's 100G analyzer:




Figure 48. GE-ODU0-ODU1-ODU2-ODU3-ODU4 signal structure interface.




Figure 49. GE-ODU0-ODU1-ODU4 signal structure interface.

10GE Signal Mapping Test

The 10GE signal can be divided into the wide-area network (WAN) and local-area network (LAN), in which the WAN signal can be directly mapped into the OPU2 via the AMP or BMP, such as the 10G SDH signal; whereas the 10GE LAN signal can be mapped into the OPU2 via GFP-F, or into the ODUflex via GFP-F. The other method is to carry the signals via OTN overlocking, i.e., transporting the 10G Ethernet LAN service on the OTN at line rates of 11.0491 Gbit/s (OTU1e) and 11.0957 Gbit/s (OTU2e).

Because the 10 GigE rate is higher than the OPU2 payload rate, the overclocked OTN brings the data rate higher than 10.709 Gbit/s, which is standard for OTU2, in order to carry 10 GigE LAN client signals. Because the OTN line rate has changed, there is a compatibility and interconnection problem between the overclocking approach and the standard OTN rate. The other problem is that the OTU2 rate fails to be multiplexed into OTU3 after being overclocked. Therefore, ODU3e is defined for multiplexing of the overclocked signal.

Overclocked OTN compensates for the rate mismatch between 10 GigE LAN and the OPU2 payload by raising the overall OTU2 data rate from the standard 10.709 Gbit/s to fit the 10 GigE LAN client signal. Obviously, this modification of the standard OTN line rate will bring about interoperability issues, and the option for aggregating OTU2 signals into OTU3 will be lost; however, ODU3e does allow for multiplexing. On the positive side, overclocked OTN offers real bit transparency of 10 GigE LAN signals—a necessity for the mass deployment of 10G services.

Overclocked OTN supports OTU1e, OTU2e, OTU3e1 and OTU3e2 optical line rates for mapping of 10 GigE LAN signals. Furthermore, OTU1f and OTU2f line rates are used for mapping Fibre Channel signals. The following figure illustrates an overclocked OTN application scenario.




Figure 50. Application scenario of overclocked OTN.

The transparent transportation of 10 GigE LAN signals means that the full 10 GigE data rate (10.3125 Gbit/s) is transported over OTN, including the PCS 64B/66B coded information, inter-packet gap (IPG), MAC FCS, preamble, start of frame delimiter (SFD), etc. The OTN clocking is derived from the Ethernet customer signal (± 100 ppm) rather than that of a standard OTU2 signal (± 20 ppm). Therefore, overclocked OTN is mainly used for point-to-point data paths.

The OTN overclocked interface rates and corresponding client rates are listed in the table below:

G.709 Interface	OTN Line Rate	Corresponding Client Rate
OTU-1e	11.0491 Gbit/s (without stuffing bits)	10 GigE LAN (direct mapping over OTN)
OTU-2e	11.0957 Gbit/s (with stuffing bits)	10 GigE LAN (direct mapping over OTN)
OTU-1f	11.27 Gbit/s	10G Fibre Channel
OTU-2f	11.3 Gbit/s	10G Fibre Channel
OTU-3e1	44.57 Gbit/s	4 x ODU2e (uses 2.5G TS; total of 16)
OTU-3e2	44.58 Gbit/s	4 x ODU2e (uses 1.25G (ODU0) TS; total of 32)

Figure 51. OTN overclocked interface rates and corresponding client rates.

10GBASE-R Signal Mapping into OPU2e/OPU1e

The OTU2e is a mapping mechanism that uses the mapping scheme of CBR 10G signals into OPU2, defined in G.709 subclause 17.1.2. The 10GBASE-R client signal with fixed stuff bytes is accommodated into an OPU2-like signal, and then further into an ODU-like signal, and then even further into an OTU-like signal. These signals are denoted as OPU2e, ODU2e and OTU2e, respectively.

The signal rate is different from that of the previously defined OTU2, and the PT value is “0x80” in the OTN frame. OTU1e is similar to OTU2e. The only difference is that it does not contain any fixed stuff bytes, and therefore, the rate is slightly lower. The diagrams are as follows:




Figure 52. OTU2e with fixed stuff bytes.




Figure 53. OTU1e without fixed stuff bytes.

ODU2e Signal Mapping into ODU3e

The OTU3e is a mechanism that allows 10 GigE LAN signals to be carried directly over 40G OTN networks. Multiplexing ODU2e in OTU3e provides granularity of the 40G circuits, optimizing the payload and simplifying network provisioning and maintenance. The frame structures of the OTU3e2, ODU3e2 and OPU3e2 are the same as those for the OTUk, ODUk and OPUk specified in ITU-T G.709. The OPU3e2 carries one or more ODUj (j=2e) signals.

GFP-F Mapping of 10GigE Signals

The generic framing procedure (GFP) is another mechanism for transporting 10 GigE LAN or WAN client signals over OTN. GFP-F encapsulates 10 GigE frames into GFP-F frames first and then into OTU2. GFP-F maps the Ethernet MAC traffic, eliminating the 64B/66B PCS sublayer. In addition, mapping the GFP-F frames into OPU2 uses the entire OPU2 payload area, meaning that the fixed stuff bytes of the CBR 10G mapping are not present. Finally, the key benefit of using GFP-F over OTN is to support various data-packet services on the same network.

EXFO's Solution

With the EXFO FTBx-8880/8870 Power Blazer, the user can support different 10G Ethernet mapping and carrying approaches, including GFP and overclocking. The instrument is able to generate overclocked OTN signals with mapped 10G Ethernet LAN client-end signals, which include an intact Ethernet frame (including the frame size, transport rate, source/destination MAC addresses and VLAN ID). This test gives network operators an adequate understanding of the OTN transport layer and its alarms, errors, trace message and overhead bytes. In addition, the test provides complete statistics of the 10G Ethernet information flow, throughput (bandwidth) and line-rate utilization. Once the transport layer test is performed, the user can then perform the RFC 2544 test to provide measurements of available bandwidth, transport delay, link burstability and service integrity. The operator can use the RFC 2544 test result to ensure that the provided working parameters for the 10G Ethernet LAN service comply with the service-level agreement (SLA).




Figure 54. Mapping 10GE to ODU2e, and to OTU4.

2.9. 40G/100G Ethernet Services over OTN Testing

The 100G Ethernet client signal rate is 103.125 Gbit/s, and can be directly mapped via the GMP into OPU4 104.35 Gbit/s. However, the case is more complicated for mapping the 40G Ethernet client signal into OPU3, primarily because the rate of 40GE is greater than that of OPU3, and has to be restored by 64B/66B coding followed by 1024B/1027B transcoding to decrease the rate to 40.117 Gbit/s and then be mapped via the GMP into OPU3. 40G/100G Ethernet mapping approaches are shown in the figure below:




Figure 55. 40G/100G Ethernet mapped into OPU3/OPU4.

EXFO's 40G/100G portable solution, composed of the portable FTB-2/FTB-4 Pro platforms and the FTBx-88200NGE Power Blazer multiservice test module, offers a complete suite of testing capabilities for qualification of 40G and 100G transponders in carrier labs.




Figure 56. EXFO's 40G/100G portable solution.

Using EXFO's FTBx-88200NGE Power Blazer, service providers can map 40G/100G Ethernet services over OTN with different traffic characteristics, run bit-error-rate (BER) tests across the OTN, and measure the ratio of the error bits compared to the number of sent bits. In this testing configuration, the FTBx-88200NGE Power Blazer module provides complete analysis of the OTU3/OTU4 (including the OPU, ODU, OTU and OTL layers), related alarms, errors and skew measurement. The 40G/100G module also provides GMP-related measurements, including Cm and CnD statistics, to ensure proper recovery of the client signal at the receive end. Once the OTN with the embedded 40G/100G Ethernet client service is qualified from end to end, service providers can take their service validation a step further by qualifying the actual 40G/100G Ethernet service from client to client. This includes validating the 40G/100G Ethernet IP traffic transmission with 100-percent throughput, and ensuring that latency measurement is not impacting customers' SLAs.

2.10. 40G/100G OTN Service-Disruption-Time Measurement

The service disruption time (SDT) test is usually used to validate a system's automatic protection switching (APS) and measure the disruption time of the signal received at the transport layer to ensure that the operating circuit is able to switch to the protection circuit within the 50 ms window specified in the ITU G.841 standard. A common method is to set the detection layer as a bit error pattern and the fault selection mode as bit error before SDT measurement in order to secure the integrity of the transported payload (i.e., the client information flow). The cause of a long service-disruption time cannot be analyzed by simply using the pattern bit error test. Another common way in the POTN system-device test is to set fault triggering methods in different OTN layers, i.e., set different fault types, such as OOF, LOF and AIS, as the triggering condition for calculation of disruption time in order to conduct detailed analysis of the protection-switching process. The figure below lists common fault trigger types that are configurable at various layers by test instruments in an OTN protection-switching test.

Layer	Defect
OTL	LOF, OOF, LOL, LOR, OOR, Inv. Marker, FAS
FEC	FEC-CORR, FEC-UNCORR
OTU	AIS, LOF, OOF, LOM, OOM, BDI, IAE, BIAE, BIP-8, BEI, FAS, MFAS
ODU	AIS, OCI, LCK, BDI, BIP-8, BEI, FSF, BSF, FSD, BSD
OTU	AIS, CSF
Pattern	Pattern Loss, BER

Figure 57. Common defect triggers.

The function monitors the presence of a specific defect on an OTN layer to ensure that network survivability mechanisms managed by the control plane can recover the affected optical path and its traffic within the 50-millisecond industry standard. When the SDT function is activated, the test instrument scans for defects. The measurement is triggered when a defect is detected (as shown in the figure below). The test period is initiated, and the SDT begins to measure the time spent in the test-period window for the duration of the disruption. When no defects are detected over a period exceeding the user-configurable no-defect-time parameter, the SDT stops measuring. The SDT measurement is then calculated as the time between the detection of the first defect and the end of the last defect.




Figure 58. SDT measurement.

ROADM-Based Network SDT Measurement

The flexibility of the POTN service is reflected by the application of ROADM in the optical layer. When network operators are initially deploying ROADM, they face serious challenges in terms of network performance maintenance. Undoubtedly, client SLA compliance suffers. Because ROADM is able to flexibly add and drop any wavelength at any port, this technical trend requires that the test instrument be configured as OTN-intrusive Through mode (as shown in the next figure) in order to monitor any selected optical channel transparently. In this process, after the optical channel is bypassed, the tester will check the OTU, ODU and OPU alarms and errors of the signal, and then resend the channel as a plugged channel. For instance, by monitoring error-detection-code information, i.e., the bit interleaved parity-8 (BIP-8) from the OTU, ODU and TCM overhead bytes, any uncorrectable error existing on the post-FEC optical link can be detected and the network performance can be assessed as per specified threshold values. These modules test functions by using the supported insertable intrusive Through modes (with this test functionality; the user can insert OTN alarms and errors into the normally transported OTN signal). The insertable intrusive Through mode is usually used to check OTN-element fault detection, generate reports in order to deploy appropriate response measures, and check the APS functionality of the network element. In addition, the insertable intrusive Through mode can be used during field trial run and activation in order to verify the interoperability of OTN elements.




Figure 59. ROADMs-based network SDT measurement.

Multichannel Service-Disruption Test in the POTN Network

Mesh topological structures will be extensively employed in POTNs. In core networks, OTN cross-connect devices will give rise to delay, causing the switching time of each channel to vary. Therefore, in contrast to the traditional point-to-point networks, it is necessary to monitor every channel simultaneously during POTN protection-switching tests: the test instrument is required to have the ability to support the multichannel service disruption test, and test all channels simultaneously.

Characteristics of EXFO'S Tester for SDT Implementation

The main advantage of EXFO's SDT feature is its precision and ability to identify multiple disruptions on a network, even when they occur within a short span of time. This is controlled via the user-configurable no-defect-time parameter, which is automatically set to 300 ms and supports values between 5 μ s and 2 s, as shown in the figure below. The no-defect-time parameter defines the threshold that is used to determine when it is safe to report that a defect is no longer present. When properly configured, it allows the FTBx-88200NGE Power Blazer to correctly terminate the test-period window and get ready for the next defect. In a situation where multiple disruptions are present and the time between defects is less than the no-defect-time, a single disruption is reported.

Figure 60. SDT trigger (defect) and no-defect-time configurations.

The SDT feature of EXFO's OTN test instrument also supports a parent defect approach, where the SDT measurement is triggered when other specific defects on that layer or higher are detected; this eliminates the need for the user to manually select different defects at their corresponding layers. The SDT function supports multiple statistics, as shown in the figure below. These statistics include the shortest, longest, last, average and total duration of all the disruptions. All SDT measurements are provided with a resolution of 1 μ s for all supported OTN layers.

Service Disruption						
Verdict Threshold						50.0
Longest (ms)	Shortest (ms)	Last (ms)	Average (ms)	Total Duration (ms)	Count	Verdict
15.462	0.419	1.219	4.432	17.730	4	PASS

Figure 61. SDT statistics.

To meet mesh testing requirements, EXFO's OTN tester is able to run the service disruption test for all lanes, with granularity as small as ODU0.

2.11. Delay Measurement

Delay is an important indicator of network performance, especially in a 100G OTN system that adopts the coherence technique and requires dispersion-compensating fiber, thereby greatly improving time-delay performance. The delay can be verified via the SONET/SDH interface or the OTN interface. The typical connection diagram is as follows:




Figure 62. Delay measurement test configuration.

The test interface and results are shown in the following figure:

Round Trip Delay

Mode Single Measure Delay Status Ready Reset

	Last	Minimum	Maximum	Average	Unit
Delay (ms)	< 0.001	< 0.001	< 0.001	< 0.001	ms

	Successful	Failed
Count	1	0

Figure 63. Round-trip-delay test configuration and results.

2.12.TCM and Performance Monitoring

OTN supports multiple traces, including ODU, path monitoring (PM), TTI, OTU, segment monitoring (SM) TTI, and six tandem-connection-monitoring i (TCMi) TTI traces. During commissioning of a new OTN link, trace messages are usually used to monitor the integrity of the connection route between terminals during the establishment of the connection. In addition, when a connection is activated, the TTI message is needed to ensure the connectivity. TTI messages for OTN contain network-related information in the form of the SAPI and DAPI. Usually, when it is detected that the received SAPI and/or DAPI in the TTI fail to match the expected preset value(s), a trace identifier mismatch (TIM) alarm will be generated. With the TTI test functionality of the test module, the user can provide SAPI, DAPI and operator-specific information fields in SM, PM or TCM TTI, and verify whether these fields have been properly sent over the whole network. These modules can also be used to monitor TTI messages offered by the network management system.




Figure 64. Testing of TTI SAPI/DAPI/operator information using the FTBx-88200NGE Power Blazer.

The TCM function has begun to be adopted in SONET/SDH networks and is becoming more prevalent in OTN networks. The main concept of TCM is to split the channel trail into a series of tandem-connected channel intervals, each of which is managed by a different network operator. With tandem connection monitoring (TCM), in the event of any error or defect, the operator can locate the interval of the fault very quickly and conveniently. The OTU layer of the OTN supports six TCM fault-monitoring levels. Through the TCM, the operator can monitor the quality of the information flow being transported over different parts of the network, and track errors and faults on a specific part of the trail. This functionality is particularly important when an optical signal trail passes multiple networks of one or more operators. Figure 70 shows a scenario in which a client is monitoring the end-to-end path connection with TCM1, where network operators A, B and C are monitoring the connection passing their respective subnets with TCM2. The TCM monitoring functionalities of the FTBx-88200NGE Power Blazer include: TCM bit-interleaved parity-8 (TCM-PIB-8) and TCM backward error indication (TCM-BEI) error monitoring, TCM alarm indication signal (TCM-AIS), TCM loss of tandem connection (TCM-LTC), TCM open connection indication (TCM-OCI), TCM locking (TCM-LCK), TCM trace identifier mismatch (TCM-TIM), TCM backward defect indication (TCM-BDI), TCM incoming alarm error (TCM-IAE) and TCM backward incoming alignment error (TCM-BIAE) alarms. With TCM errors and alarms, operators can monitor their network performance from various different intervals or as a whole.




Figure 65. TCM allocation.

2.13. Advanced Intrusive Through-Mode Analysis

OTN and Ethernet Through mode and OTN Intrusive Through mode allow the end user to monitor the network and verify network issues by introducing stress errors and alarms that help validate upstream and downstream effects.





Figure 66. Through mode test configuration.

2.14. OTN Performance Testing: Long-Term Bit-Error Test

The most significant component of the OTN performance test is the long-term BER test, which has also become the most fundamental element of the transport network test. The BER test is conducted by placing a pseudo-random bit sequence (PRBS) in the OTN frame and measuring the ratio of the number of incorrect bits after passing the network over the total number of sent bits. Before the advent of the OTN system, there were several common test standards, as follows:

- **G.821:** Basically an error evaluation based on the N x 64 kbit/s digital circuit-switched system, where test parameters, objectives and methods are detailed.
- **G.826:** Basically an assessment of end-to-end bit error performance of an international digital path at 2M and above, where test parameters, objectives and methods are detailed. Unlike G.821, the block-based test method is used, which can be carried out online for services above 2M. This method is primarily used to conduct tests on plesiochronous-digital-hierarchy (PDH) links.
- **G.828:** Basically indicators regarding the SDH international digital path; however, this allocation principle can also apply to a domestic or specific design of a synchronous digital-path error. A block-based measurement, it utilizes an inherent error-monitoring code inside the path under test in the SDH system; the block repeat rate complies with the spherical harmonics (SH) technique. This method simplifies non-stop service measurement. Corresponding stipulation is made for events, parameters and indicators. It also includes path performance estimation, tandem connection monitoring, etc.
- **G.829:** Although G.828 defines bit error performance, it is mainly based on the hypothetical reference path (HRP) of the international digital path with a length of 27 500 km. G.828 does not include index requirements for multiplexing segments and regeneration segments. The G.829 standard is a complement to G.828, and ensures that the performance allocated to the multiplexing and regeneration segments is consistent with the previous G.828 objective.
- **G.8201:** Basically defines the bit error performance of the international ODUk channel in the OTN network. Although G.8201 is related to the indices of the international ODUk channel, its allocation principle applies to the domestic or network-specific ODUk channel as well. As per G.828/829, the inherent detection code in the channel under test of the OTN system is used for block-based measurements. And, as per G.828/829, G.8201 makes corresponding provisions for events, parameters and indicators, enhancing content such as tandem connection monitoring to enable measurement of non-interrupted service. The hypothetical reference path of bit error performance is 27 500 km, covering both the starting point and the end point of the channel, as shown in the following schematic.



Basic concepts of the OTN-system bit error test:

- **Errored block (EB):** Blocks with one or more bit errors
- **Severely errored second (SES):** $\geq 15\%$ errored blocks or at least one defect in one second. In G.8201, defects are described in detail. The AIS, TIM and OCI alarms of the channel layer will all be treated as defects.
- **Background block error (BBE):** Errored blocks are not ascribed to the SES.
- **Background block error ratio (BBER):** The ratio of BBE within the available time over the total number of blocks within the available time during a fixed time period. To calculate the total number of blocks, all blocks in the SES have to be deducted.

The threshold value of the number of severely errored blocks in one second varies with ODUk. Its main threshold values are listed below:

Bit rate (kbit/s)	Channel type	Threshold values of the errored block number for determination of the SES (number of errored blocks in 1 s)
1 244 160	ODU0	1 526
2 498 775	ODU1	3 064
10 037 273	ODU2	12 304
10 399 525	ODU2e	12 748
40 319 218	ODU3	49 424
104 794 445	ODU4	128 459
Variable rate $X \geq 1\,244\,160$	ODUflex	Max. $\{[150 \times X]/122\,368\}$

Figure 67. Threshold values of the errored block number in one second for determination of the SES.

Content of the G.8201 Bit Error Test

The G.8201 bit error test can be divided into in-service monitoring and out-of-service measurement. In principle, the test is achieved through the bit-interleaved parity (BIP) check.

In-Service Monitoring

In-service monitoring is used to monitor each block by way of the internal BIP-8 error detection code (EDC). Each EDC bit is physically isolated from the block it monitors. It is usually impossible to determine whether the error is in the block or its controlling EDC bit. If there is any discrepancy between the EDC and its controlled block, the controlled block is always deemed erroneous. With a device tester or OTN tester, the operator can monitor the channel performance and status.

Out-of-Service Measurement

The block-based bit error test can also be out of service, which is also the basis for a bit error test with an OTN tester. As a basic requirement of the engineering test, the capability of the out-of-service bit error test must be stronger than it is online.

With EXFO's tester, the user can also set ten user-defined patterns. In addition to the PRBS, different kinds of user information signals (such as GE, 10GE, FC or SDH signals) can be mapped into OUT for testing. The bit error test normally requires long-term, error-free verification, e.g., a 24-hour or 72-hour error-free test.




Figure 68. Long-term OTN BER test.

2.15. Multichannel OTN

OTN testing

As mentioned earlier, OTN boasts unique characteristics that make it a very resilient technology. The next challenge for users is to validate the logical configuration of circuits on network elements during development, installation and troubleshooting. In response, the global network testing and data experts at EXFO have introduced the FTBx-88200NGE Power Blazer along with the multichannel OTN test application enhanced with mixed mapping capabilities. It's a smart and easy way to extensively test customer OTN implementations.



Let's have a look at points to keep in mind while testing OTN and also the key points addressed by enhanced multichannel OTN with mixed mapping.


Connectivity validation

Different methodologies are available to validate connectivity. From a physical perspective, the goal is to generate a test signal to the device under test (DUT) and analyze the response:


Scenario 1: The FTBx-88200NGE (test equipment) is directly connected to the terminating network element (DUT).



Scenario 2: Both the FTBx-88200NGE and the DUT are connected to the network



Scenario 3: Multiple DUTs are connected to the network with a single test apparatus that generates the stimulus. In the following example, the FTBx-88200NGE is connected to the network via an OTU4 link. The interoperability monitoring of multiple DUTs is done via OTU2 links.



Multiplexing

Multiplexing is a very important OTN function that allows different data containers (ODUs) to be multiplexed into higher-rate containers. The multichannel OTN test application supports single-stage multiplexing for fitting low-rate containers into an ODU4.

The following table shows the maximum number of channels for carrying the OTU4 signal:

Multiplexing	Number of channels
OTU4/ODU4/ODU0	80
OTU4/ODU4/ODU1	40
OTU4/ODU4/ODU2	10
OTU4/ODU4/ODU3	2

OTN's multiplexing capabilities allow it to carry more data containers inside an ODU4. Users wishing to test network element multiplexing functionality more exhaustively can use multichannel OTN enhanced with mixed mapping to test a basic scenario by saturating the pipe with the same type of data units. This is equivalent to testing the capacity of a highway with the same-sized cars (testing scenario 1 in Figure 69). An example of such a test would be to configure 80 ODU0 channels on an ODU4 to be transported on an OTU4, as in Figure 70.




Figure 69. Multichannel OTN enhanced with basic mixed mapping scenario




Figure 70. Configuration of 80 ODU0 channels on an ODU4

Mixed mapping also enables users to select different types of data units to be inserted into an ODU4 container. Looking back at the example from highway scenario 2 in Figure 69, we could test highway capacity with different sizes and types of vehicles. If we apply this analogy to data units, we can picture the car as an ODU0 rate container, the ambulance as an ODU2 and the truck as an ODU3, all of them traveling on a highway equivalent to an OTU4 transport container. This mixing of ODU container types is closer to what users would transport on a real OTN link, enabling NEMs to test the response of their network elements, and allowing service providers to verify link operation during turn-up. A sample configuration using this test application is shown in Figure 71, where we can see that the ODU channels include a different quantity for each container.




Figure 71. Sample test configuration

Configurable channels/tributary slots

Another feature available with the multichannel OTN test application enhanced with mixed mapping is the possibility of coupling and decoupling Tx tributary slot/channel assignment from the Rx tributary slot/channel assignment. This provides users with the flexibility to select either a different channel configuration or a tributary slot to validate, including the possibility of choosing different data unit containers (ODU0, ODU1, ODU2 or ODU3). Selection can be done manually or by the application, which can be instructed to automatically identify all available tributary slots and assign them to a channel. Once the channel configuration is complete, it can be easily edited or deleted using the elements on the graphical interface as shown in Figure 72 below.



Figure 72. Configurable channels/tributary slots

During testing, users have the option of assigning channels manually using the graphical interface to configure traffic so that it is as close as possible to what will be transported. The test application can help users save time by automatically identifying and configuring the frames used for the test based on the ones received. This simplifies the user configuration process and reduces the complexity of detailed configuration (see Figure 73).




Figure 73. Automatic identification and configuration of tested frames

Once the channel configuration is complete, the test page displays a summary of the ODU container types configured as well as the quantity of channels used and any unused capacity, helping the user to graphically confirm the structure to be tested (Figure 74).




Figure 74. Summary of ODU container types

Concurrent BER analysis

Once the channel configuration between the DUT and the test equipment (FTBx-88200NGE) is ready, the next test configuration step is the bit error rate (BER) test. This allows the user to validate whether the overall link meets the performance specifications of the OTN standard over a specific period of time. The test application provides the necessary tools to support concurrent BER testing. The user can define the pattern to be used and the bit error rate considered (Figure 75), and start the BER evaluation for the required time period. There is an easy-to-configure interface as well as an optimum view of the pass/fail status of the test. The application constantly monitors traffic and, in case of errors, compares results against the threshold. Performance is assessed on all composing channels during the test period.




Figure 75. BER testing

Service disruption time

As previously mentioned, one of OTN's biggest strengths is fast and seamless recovery in the event of failure. In addition to BER testing, the multichannel OTN test application enhanced with mixed mapping includes special features that allow users to measure service disruption time (SDT) and verify how quickly network elements switch to the protection path.

To calculate disruption time, the application considers a specific defect, a no-defect time of 300 ms, and an SDT threshold value provided by the user (Figure 76). SDT duration is calculated as the time between the detection of the first defect and the end of the last defect, excluding the no-defect time.




Figure 76. User provided SDT threshold

Let us suppose that the defect type is set to AIS and a no-defect time configured to 300 ms. The first AIS event lasts 20 ms, followed by an absence of the alarm for 150 ms and another AIS event lasting 45 ms. No other alarm is observed. Since the no-defect time is set to 300 ms, the 150 ms interval is insufficient for the system to clear the service disruption. Therefore, the total SDT is 215 ms (20 + 150 + 45), with one single disruption event counted (Figure 77).




Figure 77. Example SDT calculation

If the user were to select the pass/fail verdict for the SDT measurement and define the SDT threshold as 50 ms, the first scenario would result in a failing verdict with an SDT measurement of 215 ms.

Per-channel testing and validation with this test application makes it possible to easily measure SDT on every channel in order to detect defective channels during turn-up or advanced troubleshooting. Overall test time is reduced because each channel is tested individually, providing its own statistics.

Monitoring, alarms and errors

Once all parameters have been configured and the test initiated, the user-friendly multichannel OTN application graphical interface allows users to navigate through the OTN alarms and errors on the optimized summary page. This feature provides a single-page overview of all channels and layers and displays the responses from the DUT, as shown in Figure 78. The user can easily validate connectivity between the DUT and the test equipment.




Figure 78. OTN summary page

The interface offers a simple and unique monitoring system that detects channel alarms and errors in real time, indicating with a specific color code the status of every channel in the test configuration (Figure 79). Green indicates no channel errors, orange shows the channel is selected, yellow indicates some alarms or errors have been recorded since the test started, and red signals that there is currently an active error or alarm on that specific channel.





Figure 79. Status of every channel in the test configuration is color coded


Lastly, to make absolutely sure that the DUT is detecting and reacting correctly to the alarms and errors (e.g., BIP-8, AIS) that have been injected, users can inject those alarms and errors in a very intuitive way—using only a single button, either on a specific channel or on all channels included in the configuration.

2.16. General Communication Channels BERT Test

The General Communication Channel 0 (GCC 0) is a two-byte field that is part of the OTU overhead, as defined by G.709, which resembles the data communication channel (DCC) in SONET/SDH, and it is currently undefined. However, it will likely be used for certain functions, such as network management or control plane signaling for G-MPLS protocols.




General Communication Channel 1 and 2 are two-byte long fields inside the ODUk overhead and are similar to the GCC0 field; GCC1 is located in row four, columns one and two; while GCC2 is located in row four, columns three and four, at the ODU overhead. They are used to carry end-to-end client information.



Recently, GCC channels have been used for different purposes, including synchronization and encryption, and therefore, their integrity is prime for testing.

The GCC BERT test provides a tool for measuring the bit error rate on these fields, providing statistics that include a count of any detected errors and pattern loss on each individual GCC channel.



2.17. Ethernet Testing in the Lab: RFC 2544 Test

Service providers worldwide are actively turning up new services based on Carrier Ethernet technology in a fierce competition to attract premium subscribers. The need for quality services has never been more important, making comprehensive Ethernet testing immediately at service turn-up vital to ensuring service quality and increasing customer satisfaction.

Customer SLAs dictate certain performance criteria that must be met, with the majority documenting network availability and mean-time-to-repair values that are easily verified. However, Ethernet performance criteria are more difficult to prove, and demonstrating performance availability, transmission delay, link burstability and service integrity cannot be accomplished accurately by a mere PING command alone.

Portable RFC 2544 test equipment enables field technicians, installers and contractors to immediately capture test results and demonstrate that the Ethernet service meets the customer SLA. These tests also serve as a performance baseline for future reference.

What is RFC 2544?

The RFC 2544 standard established by the Internet Engineering Task Force (IETF) standards body is the de facto methodology outlining the tests required to measure and prove performance criteria for Carrier Ethernet networks. The standard provides an out-of-service benchmarking methodology for evaluation of network-device performance using throughput, back-to-back, frame loss and latency tests, with each test validating a specific part of an SLA. The methodology defines the frame size, test duration and number of test iterations. Once completed, these tests will provide performance metrics for the Ethernet network under test.

RFC 2544 Tests

In order to ensure that an Ethernet network is capable of supporting a variety of services (such as VoIP, video, etc.), the RFC 2544 test suite supports seven predefined frame sizes (64, 128, 256, 512, 1024, 1280 and 1518 bytes) to simulate various traffic conditions. Small frame sizes increase the number of frames transmitted, thereby stressing the network device, which must switch a large number of frames.

Throughput Test

The throughput test defines the maximum number of frames per second that can be transmitted without any errors. This test is done to measure the rate-limiting capability of an Ethernet switch as found in Carrier Ethernet services. The methodology involves starting at a maximum frame rate, and then comparing the number of transmitted and received frames. If frame loss occurs, the transmission rate is divided by two and the test is relaunched. If there is no frame loss during this trial, the transmission rate is increased by half of the difference from the previous trial. Known as the half/doubling method, this trial-and-error methodology is repeated until the rate at which there is no frame loss is found.

The throughput test must be performed for each frame size. Although the test time during which frames are transmitted can be short, the duration of the final validation must be at least 60 seconds. Each throughput test result must then be recorded in a report, using frames per second (f/s or FPS) or bits per second (bit/s or BPS) as the measurement unit.

Back-to-Back Test

The back-to-back test (also known as burstability or burst test) assesses the buffering capability of a switch and measures the maximum number of frames received at full line rate before a frame is lost. In Carrier Ethernet networks, this measurement is quite useful, because it validates the excess information rate (EIR) as defined in many SLAs. A burst of back-to-back frames is transmitted across the network with minimum interframe gap. If a frame is dropped, the burst length is shortened. If it is received without any errors, the burst length will be increased. The trial length must be at least two seconds long, and the measurement should be repeated at least 50 times, with the average of the recorded values being reported for each frame size. The average measurement should be logged in the report.

Frame Loss Test

The frame loss test measures the network's response in overload conditions—a critical indicator of the network's ability to support real-time applications, in which a large amount of frame loss will rapidly degrade service quality. Because there is no retransmission in real-time applications, these services might become rapidly unusable if frame loss is not controlled.

The test instrument sends traffic at maximum line rate and then measures whether or not the network dropped any frames. If so, the values are recorded, and the test will restart at a slower rate (the rate steps can be as coarse as 10%, although a finer percentage is recommended). This test is repeated until there is no frame loss for three consecutive iterations, at which time a results graph is created for reporting. The results are presented as a percentage of the frames that were dropped, i.e., the percentage indicates the variable between the offered load (transmitted frames) vs. the actual load (received frames). Again, this test must be performed for all frame sizes.

Latency Test

The latency test measures the time required for a frame to travel from the originating device through the network to the destination device (also known as end-to-end testing). This test can also be configured to measure the round-trip time, i.e., the time required for a frame to travel from the originating device to the destination device, and then back to the originating device.

When the latency time varies from frame to frame, it causes issues with real-time services. For example, latency variation in VoIP applications would degrade voice quality and create pops or clicks on the line. Long latency can also degrade Ethernet service quality. In client-server applications, the server might time out, or poor application performance could occur. For VoIP, this would translate into long delays in the conversation, producing a “satellite call feeling”.

The test procedure begins by measuring and benchmarking the throughput for each frame size in order to ensure that the frames are transmitted without being discarded (i.e., the throughput test). This fills all device buffers, therefore measuring latency in the worst conditions. The second step is for the test instrument to send traffic for 120 seconds. At the midpoint of the transmission, a frame must be tagged with a time-stamp and when it is received back at the test instrument, the latency is measured. The transmission should continue for the rest of the time period. This measurement must be taken 20 times for each frame size, and the results should be reported as an average.

Complementary Testing

Although not included in the defined RFC 2544 standard, another crucial measurement in Ethernet networking is packet jitter. This measurement is critical, because excessive jitter can cause major failures in real-time applications such as VoIP or streaming video. In VoIP applications, excessive jitter will cause dropout effects. In video applications, the image will become choppy and pixelization could occur.

As covered in the latency test section above, frame latency will vary over time. It is this variation in inter-packet arrival time that is referred to as packet jitter. At the start of the transmission, the inter-frame gap between all frames was identical. As the frames navigate through the network, being buffered or routed through different network devices, the inter-frame gap begins to vary. The packet jitter measurement should be performed at maximum frame rate, as this is where the most variation will occur.

Implementing RFC 2544 Testing: Important Considerations

Although RFC 2544 is highly effective for assessing the performance of interconnected network devices, it was originally developed for performance benchmarking of individual network elements, which didn't demand time-sensitive testing scenarios. As an example, the latency test alone calls for testing 20 iterations of 7 different frame sizes: that's 140 tests running 2 minutes each for a minimum test time of 4.6 hours just for the latency test! Because manually completing the entire series of RFC 2544 tests would be quite time-consuming, a certain level of automation is required from today's test solutions in order to thoroughly test and turn up Ethernet services as quickly and efficiently as possible.

Service providers will usually benchmark their network and network elements in the lab and then, from this knowledge, customize their RFC 2544 test routine for field turn-up and installation. For some, the number of frame sizes to test is too large, so they streamline their turn-up testing to include only the extremes, i.e., 64 and 1518 byte frames (could be longer if VLANs are used). Some may keep all frame sizes, but decide to conduct only two tests, such as throughput and latency. Others will decide to limit the number of tests and frame sizes. The critical consideration is for service providers to comprehensively test before turning the circuit over to the customer, because this will be the last time that they can fully access it without affecting subscriber services.

2.18. Ethernet Service Aggregation and Switching Capability of POTN Devices

Among POTN-provided services, one very important aspect consists of aggregation-type Ethernet services, such as the special group client line, mobile backhaul service, and OLT uplink carrier in PON networks. In terms of test interfaces, GE/10GE to 40GE/100GE aggregation test functionalities are included, and POTN should be provided with Ethernet and MPLS-TP switching functions to ensure compatibility with existing PTN devices.

Because central-switching POTN devices have the statistical characteristic of packet switching, packet services need to be subjected to quality of service (QoS) classification. In other words, higher-demand services are classified as guaranteed service traffic (GST) in order to guarantee that such services have no packet loss. Packet delay and packet delay variation (PDV) must be provided with circuit-switching characteristics for use in transporting higher-priority services such as the synchronous signal, real-time service, VIP client special line, mobile backhaul, etc. Low-priority services are classified as the statistical multiplexed (SM) type, for which a best-effort strategy is adopted.

The diagram in the following figure represents a structural schematic of an SM-based device. The GSTs on the left are guaranteed (5 GE) signals. The one at the bottom right is an SM signal (one 10G signal), and the one at the top right is of 10G output with an SM signal added.

2.19. Multilink

At EXFO, we are the global network test, data and analytics experts. For the last 30 years, we have been working alongside our customers and have come up with many innovative solutions. Multilink is one of our most recent tools that helps our clients address the following lab and testing environment-related challenges.

Noisy Environments

Use of multiple lab equipment is only generating increased amounts of heat. For example, high-speed optical transceivers, such as CFP, CFP2, CFP4, need a very stable temperature to operate and more powerful fans are required to reduce lab temperatures. This also applies to many other network elements. Unfortunately, all the cooling fans being used tend to generate noise that may disturb users located in the same area, making this environment a hard one to work in.

Space


Having adequate space in lab environments is a great challenge due to the different types of technologies and pieces of equipment that must be extensively tested in the same facilities. As a result, there's not much leeway for users to move around or to be comfortable.

Remote Locations

Time to market is more crucial than ever and network equipment manufacturers are required to have different teams working on the same project from various remote locations. Fast developing strategies should not jeopardize security and teams from different regions of the world should be able to have secure and efficient access to testing environments. There should not be any risk incurred from sharing proprietary and confidential information when units are accessed from remote locations.

Lack of Flexibility


Hardware equipment must be extensively tested in lab environments during long periods of time; additionally, many users and testers must have access to different platforms and modules at the same time; securing the utilization of specific cards or test sets becomes a complex, yet fundamental, administration task.



What is Multilink?

EXFO MultiLink is a unique lab test management system with a multi-user interface that offers remote access to multi-modules and multi-platforms (LTB-8/FTB-2 Pro/FTB-4 Pro) across numerous locations. It is the only true Web-based interface that enables this capability.

- EXFO is the only company in the industry offering this capability.
- EXFO Multilink simplifies testing and offers very important features that enable it to be a very powerful tool.
- Multi-User/Chassis/Module – Using the same tool, multiple users within the organization can securely and easily access several chassis hosting many modules, even from remote locations.
- Web-Based Interface – Users need only access a Web page without having to install any additional software to use all the chassis and modules managed by Multilink.
- Inventory/Module Management – Directly from the same Web page, a user can access a quick inventory of the hardware Multilink is managing at that time, including detailed information about such hardware.
- Software/Firmware – This tool allows users to perform software upgrades directly from the Web page to all the chassis and modules managed by Multilink.



3. 40G/100G Commissioning and Turn-Up

LINE SIDE

The following section details different deployment strategies. At this stage, the cards and systems will be ready and fully tested during manufacturing. The next step is deployment, with the main objective being to ensure that everything functions properly.

The system itself contains three main components: the transmitter, the link and the receiver.



Figure 82. Main block components of a network.

The system's performance is highly dependent on the DSP. Every vendor will have its own unique proprietary test methodology, so it is important to define the actual parameters that will be needed, and whether or not testing should be performed before or after deployment. The same applies with regard to identifying and troubleshooting potential issues.

3.1. Link Characterization and Non-Linear-Effect Validation

In addition to standard link characterization, tests such as bidirectional optical time-domain reflectometry (OTDR), optical return loss (ORL), loss and connector-checking validation will not be discussed in this document, because these tests are irrelevant to transmission speed and modulation format. However, the following two tests are of paramount importance to error-free 100G deployment: CD and PMD.

Chromatic Dispersion (CD)

Good digital signal processing can compensate for a very large amount of CD. As such, the OIF recommends (in accordance with the application and distance) target CD tolerance values ranging from 2000 ps/nm for short links to above 40 000 ps/nm for ultra-long haul.

The range of offerings is therefore quite wide, depending on the quality of the DSP, the optics, the local oscillator and other variables. While the OIF recommends a certain amount of compensation per application, many carriers will not willingly deploy the most expensive DSP if they can make do with a cheaper one.

So while in theory, an unlimited amount of CD can be compensated for, often a simple test for CD will determine the amount that is residual on the link, which will in turn optimize the purchase of compensation electronics.

Test Conditions

a. Greenfield

In greenfield deployments, which involve brand-new fiber, typically only 100 Gbit/s coherent signals will be carried. As such, the CD will be known, albeit not with high accuracy. However, knowledge of the fiber type and length will provide an accurate assessment, making it easy to get a good idea of the total CD per span and to select the proper DSP technology. Accordingly, CD testing is not required in most greenfield applications.

b. Brownfield

Brownfield installation involves deploying 100G coherent signals on fiber that has already been deployed, but that is most likely carrying non-coherent traffic. The distance for this non-coherent traffic, most likely 10G on-off keying (OOK), is typically greater than 80 km to 100 km, requiring that dispersion compensation modules (DCMs) be deployed. Although required for legacy OOK, this low residual CD can be a source of cross-phase modulation (XPM) when 100G is deployed.

XPM is a non-linear effect (NLE) in which one carrier (wavelength) negatively influences a neighboring wavelength. In OOK, only the amplitude of the transmitter is modulated.



Figure 83. Amplitude modulation in OOK.

Due to the power density, as this propagates (locally when a “1” is propagated), the medium (glass) becomes locally heated. And although this local heating is very minute, it changes the index of refraction (IOR) of the fiber very locally and propagates along the fiber, making it a medium modulated in IOR.

Because OOK only uses amplitude information to detect the signal, this medium modulation has no real impact. But, the IOR has a direct impact on the speed of propagation and, therefore, on the phase of the modulated signal (currently most high-end CD analyzers use phase-shift detection to measure the difference in group delay).

In coherent-based transmission, i.e., dual polarization state phase modulation (DP-QPSK), it is not the amplitude, but the phase that is modulated, encoded and used to retrieve the information (among other variables, such as polarization):




Figure 84. Phase modulation in QPSK.

Accordingly, the OOK causes phase modification effects on a local scale, and the adjacent transmission is phase-encoded. As such, this phase modulation will have an impact on the phase encoding, inducing phase noise known as XPM.

This would not be an issue in the presence of strong CD, because every wavelength travels at different speeds and the medium modulation created by the OOK would not be time-synchronized (in-phase) with the DP-QPSK, so it would not create any XPM. But, as mentioned above, most brownfield applications are optimized to result in very low residual CD, making these systems prone to XPM.

In addition, the longer the fiber, the longer the duration of this interaction, resulting in more significant XPM.

A practical solution for bypassing this problem is to use guard bands, e.g., freeing a few channels between OOK and DP-QPSK.

As such, in the transmission scheme illustrated below, adding a channel of a different modulation (represented by the dark blue channel)...




Figure 85. Adding a channel of a different modulation scheme may lead to NLEs.

... would require removal of a few neighboring channels, as shown in the figure below:




Figure 86. Guard bands on either side of the new channel.

Although this does work, it is not extremely efficient. For example, adding a 100 Gbit/s channel would require removal of two 10 Gbit/s channels, which would result in a bandwidth gain of only 80 Gbit/s.

The noise created by XPM is not amplified spontaneous emission (ASE) noise and is therefore not detectable by standard optical spectrum analyzers (OSAs). The OSA tests might indicate sufficient OSNR, even if high BER is present. The only way to detect the presence of XPM is to use EXFO's WDM Investigator, which would indicate the presence of any XPM on the *Non-linear Depolarization* line, as shown in the figure below:




Figure 87. An example of non-linear depolarization in WDM Investigator.

Other papers and lab experiments have also documented this issue. Indeed, ClariPhy published a paper on the topic to demonstrate the relationship between OSNR and launch power at different channel spacings and with different amounts of FEC.




Figure 88. Relationship between OSNR and launch power.

As power increases and channel spacing decreases, OSNR requirement increases. This is a direct cause of XPM.

As discussed earlier, XPM occurs due to local heating of the transmission medium. It is therefore normal that higher power creates more XPM, as shown in the graphs above.

However, power is not the actual limiting factor. The real issue is power density; the same power on a smaller core will create more XPM than on a larger core. In commercial fibers, G.652 dispersion-shifted fiber has a larger core and is less prone to high power density, but is also more frequently compensated for. Therefore, G.652-based networks typically have lower residual CD. Non-zero dispersion-shifted fiber (specified in ITU-T G.655) has a much smaller core, but because it naturally has lower CD, it is frequently not compensated for. For this reason, the residual CD can be slightly higher, which is offset by the higher power density. Some G.655 fiber has a larger effective area than other G.655 fiber and, as a result, could be less prone to this issue.

Accordingly, fiber type must be taken into consideration prior to creating a channel allocation plan. In greenfield, this is not an issue. In the brownfield environment, the fiber type is often unknown.

One of the best ways to identify the fiber type in question is to test for CD, because each fiber type has a CD, a CD coefficient, and a CD slope signature:

Fiber Type	Abbreviation	Lambda Zero	Dispersion at 1550 ps/(nm*km)	Slope at 1550 nm (ps/(nm*nm)*km)
Standard Singlemode	SM	1300-1324	16 to 18 (17 Typical)	≈0.056
LS	LS	≈1570	-3.5 to 0.1 (-1.4 Typical)	≈0.07
Dispersion Shifted	DS	≈1550	≈0	≈0.07
True-Wave Classic	TW C	≈1500	0.8 to 4.6 (2 Typical)	≈0.06
True-Wave Plus	TW Plus	≈1530	1.3 to 5.8	
True-Wave Reduced Slope	TW RS	≈1460	2.6 to 6 (4 Typical)	<0.05 (0.045 Typical)
E-LEAF	E-LEAF	≈1500	2 to 6 (4 Typical)	≈0.08
Teralight	Teralight	≈1440	5.5 to 9.5 (8 Typical)	≈0.058
True-Wave Reach	TW Reach	≈1405	5.5 to 8.9 (7-8 Typical)	<0.045

Figure 89. Chromatic dispersion signature of various commercially available fibers.

Tools such as EXFO’s FTB-5700 indicate the CD, in addition to measuring the length and coefficient. Accordingly, they are ideal for the fiber-type identification step required prior to wavelength allocation on brownfield.

Once the allocation is complete, the WDM Investigator (which is pictured above and available as an optional feature on EXFO’s FTB-5240S-P OSA) will show any residual XPM.

c. Mesh Network

Metro mesh networks, which include ROADMs, add flexibility, optimize bandwidth, allow for faster disaster recovery, and allow wavelengths from different transmitters and different systems to be routed differently as time goes by.

When these networks were deployed (with OOK in mind), it was important for each drop port of the ROADM to be dropping a well-compensated wavelength. At the network design stage, each ROADM node was equipped with internal DCMs, each compensating sufficiently for the length and type of fiber between itself and the previous ROADM.




Figure 90. Mesh network, short individual spans, and long total transmission distance.

It was mentioned earlier that greater distances make for longer paths on which OOK and DP-QPSK interaction occur, and longer length creates more potential for XPM. Therefore, it is easy to conclude that the individual spans in a metro environment, such as the mesh network shown above, have little chance of creating XPM, because they are short and fully compensated. This is true from a span-to-span perspective; however, data travels through many spans from the Tx to the Rx, thus extending the length of the link.

As such, qualification of residual dispersion on a per-span basis and fiber-type information are two critical pieces of information that the FTB-5700, for example, can retrieve from this type of network.

d. Raman Networks

Considering that the main predictable limitation of DP-QPSK and coherent detection is OSNR, the extra distance and amplification needs to be obtained with low-noise amplification piggybacked with the current amplifiers. This also needs to distribute the gain to lower the risk of NLEs. Raman amplification is therefore very popular in most coherent systems, whether they are greenfield or brownfield.

While erbium-doped fiber amplifiers (EDFAs) use a specially-doped fiber as an amplification medium, Raman gain is distributed throughout the transport fiber itself. Since the most gain is needed at the location where the signal is weakening (i.e., not right after the EDFA), Raman pumps are launched backwards with their light propagating in the opposite direction. Gain is strong at the Raman pump site, whereas the Raman pump light becomes weaker and provides very little amplification as it approaches the EDFA site.




Figure 91. Counter-propagating Raman amplification.

Now of course, like all amplifiers, the power of the pump can be adjusted. In counter-propagation schemes like the one highlighted in Figure 80, gain (in dB) is proportional to the Raman pump power (in mW), provided there is no saturation.

But, the gain is also a function of the effective area of the fiber. Smaller fiber cores will translate to higher power density, and therefore, to higher gain.

It is therefore a common procedure to identify the fiber type and determine its effective area when installing the Raman pump, and to adjust it accordingly, as shown in the adjacent figure.




Figure 92. Raman gain distribution for various commercially available fibers.

As previously demonstrated, the key for proper Raman gain adjustment, and therefore proper gain, depends widely on the fiber type, which is a requirement at the design level (Raman gain). Also mentioned previously, this is usually not an issue in greenfield, because the technician is aware of what he or she is installing. However, brownfield systems contain several old fibers that have either changed hands many times through various mergers and acquisitions, been rerouted without proper documentation, or whose records and documentation have been lost or misplaced.

Today's procedure might involve tweaking the Raman gain on-site and monitoring the gain at the receiving end; however, this requires time, effort and team synchronizations involving many individuals. More often than not, the Raman scripts are based on the expected or assumed fiber type. If any surprises do occur, like higher-than-expected loss or misidentification of the fiber type, the entire turn-up will shut down and be reconfigured. Raman amplification is for long-haul, multispan systems. The technician starts with the first span, and then moves on to the next until he or she reaches the end, at which point the return signal is initiated. If any one span was incorrect in the Raman gain adjustment (or the fiber-type assumption), the entire link will be faulty. And, finding the wrong span can be a huge challenge that is likely to delay installation and commissioning of the link by several days, in addition to requiring several truck rolls in an attempt to isolate the issue.

As mentioned earlier, CD testing provides good insight into the fiber type at hand. Accordingly, CD testing for this aspect of 100 Gbit/s coherent transmission is recommended.

Potential test tools for CD issues in 100 Gbit/s systems are as follows:

- FTB-5800 CD Analyzer for ultra-long haul
- FTB-5700 CD Analyzer for metro links, Raman amplification and fiber-type information
- FTB-5240S-P OSA with WDM Investigator for XPM and guard-band optimization

Polarization Mode Dispersion (PMD)

PMD values are often much smaller than CD values. This could lead to a perception that a powerful DSP compensating for huge amounts of CD would be able to take care of PMD issues, which most often is the case.

However, “most often” does not mean “always” due to the nature of PMD: a stable average, yet very high instantaneous maximums.

So, while most coherent DSPs will compensate for a large amount of average PMD (as per CD, the values of tolerance differ depending on the DSP and optics quality, etc.), some high values will nonetheless cause network failure.

PMD is the average (or RMS) of differential group delay (DGD). This variation over time (wavelengths) typically corresponds to a Maxwellian distribution, meaning that there is always a possibility that at any given time and wavelength, an extreme, highly instantaneous DGD can occur, or that the change in DGD could be very abrupt and rapid.




Figure 93. Relationship between PMD and instantaneous DGD.

Because overcompensation can create a delay as significant as no compensation at all, the real-time impact of DGD per wavelength is monitored and compensated for at the receiver site. This implies a certain limitation in tracking, both in terms of the range and transition/reaction speed.

This basically means that a high local extreme or rapid transition in DGD could imply that the DSP is failing to track and compensate. The system (DSP algorithms) will attempt to find and retrieve DGD information in order to resume compensation, but it will not compensate while searching, even if the instant DGD falls back within the compensation range.

PMD compensation in coherent systems might also fail for reasons related to state of polarization. Polarization, which is a property of light, is defined in terms of the pattern traced out in the transverse plane by the electric field vector as a function of time. The light is 100% polarized if it has a defined, repeatable trace that represents its state of polarization (SOP). Polarized light can be classified into the following three groups: linearly polarized, circularly polarized and elliptically polarized light. Poincaré sphere representation is used to describe the polarization and changes in polarization of a propagating wave. Any given polarization state corresponds to a unique point on the sphere, as shown in Figure 95. Like DGD, SOP can change very fast in a link, which implies the SOP tracking must be quick and in real time.




Figure 94. DGD vs. PMD Maxwellian distribution.




Figure 95. Poincaré sphere showing different SOPs.

In summary, PMD compensation can fail due the following reasons:

- Fast SOP changes
- Abrupt SOP jump
- Loss of SOP orthogonality
- Fast PMD changes
- Sudden PMD jump
- Large PMD values
- The presence of PDL

Any of these reasons could lead to one of the following impacts:

- Burst of error, increasing the BER
- Loss of tracking and long recovery time (at durations of up to 30 seconds)

The best way to reduce the probability that any of these issues occur is to use fibers with low PMD. PMD testing is therefore paramount for successful operation of 100G systems.

Exception Case

An exception case occurs when a section offers high-polarization scrambling, and occurs right before a section of slightly higher PMD. High-polarization scrambling can occur due to any of the following:

- Aerial fiber
- Fiber under bridges
- Underground fiber alongside highways, railways, etc.
- Any situation that can cause vibration in the fiber

When this scenario occurs, the scrambling caused by the vibrating fiber can greatly influence the state of the polarization conditions entering the high PMD section. This could cause abrupt changes in PMD, and much higher and more frequent DGD Max than expected at random times and wavelengths. The Maxwellian distribution referenced above does not hold true in those scenarios. Any link that exhibits fairly high PMD (even if it is below the specified 100 Gbit/s tolerance) and contains vibrating sections should be inspected closely using, for example, a distributed PMD analyzer at the sections bordering the vibration.

Therefore, despite the fact that vendors claim extreme tolerance to PMD, testing for PMD remains a required part of 100 Gbit/s testing. For obvious reasons, a recognized method for buried and aerial fibers that is traceable to standards is preferred.

In addition, the following two tools may prove useful:

- 1- EXFO's FTB-5600: If indeed high PMD is detected, and the probability of instantaneous PMD is too high for the agreed QoS and SLAs, EXFO's unique polarization OTDR, the FTB-5600, can be used to find and replace the highest PMD section on a given route. This method is the simplest and least expensive way to mitigate PMD issues.
- 2- EXFO's WDM Investigator: In cases where the PMD has been tested along with the OSNR and CD, and is within perceived tolerance, but some high BER still occasionally appears on one or many channels, local DGD spikes might be the culprit. The WDM Investigator, which is available as an option on EXFO's OSAs, is able to track this, making it an indispensable tool for 100 Gbit/s troubleshooting. Results appear on a per-channel basis as PMD pulse spreading, as shown in the adjacent figure.




Figure 96. Example of PMD pulse spreading in WDM Investigator.

Potential test tools for PMD issues in 100 Gbit/s systems are as follows:

- FTB-5500B PMD Analyzer for ultra-long-haul testing
- FTB-5700 PMD Analyzer for metro links
- FTB-5600 P-OTDR for PMD mitigation
- FTB-5240S-P OSA with WMD Investigator for troubleshooting

3.2. OSNR Measurements and Polarization-Dependent OSNR Verification

Signal-to-noise ratio (SNR) is the ratio between a signal (meaningful information) and the background noise. OSNR provides indirect information about the BER, making it the most useful parameter available from the measured spectrum, which is why it is listed as an interface parameter in ITU-T Recommendations G.692 and G.959.1.

The importance of OSNR is well-known. OSNR assesses the signal quality, providing signal-quality indication at all wavelengths simultaneously (whereas BER only assesses one wavelength at a time). In most cases, it can predict whether the BER is likely to pass or fail, and it is therefore of enormous importance during commissioning for proper system turn-up, optimization and time savings. In troubleshooting cases, it is the perfect tool for determining the presence and cause of high-noise sources.

In a nutshell, better OSNR leads to lower BER, which in turn leads to higher QoS. As such, in the very worst cases, low OSNR can be the root cause of lack of transmission, but more commonly leads to errors in transmission, increased downtime, extra truck rolls, lower QoS and SLA penalties.

Today's 100G coherent detection systems are, as previously discussed, much more robust in withstanding physical impairments. Provided that the transmitter and receiver synchronize, transmission should be error-free. However, it is an accepted fact that the main cause for Tx-Rx losing synchronization is OSNR, and therefore OSNR should be the de facto test conducted prior to any coherent deployment.

In addition to all of the above-mentioned virtues, there is a direct relationship between OSNR and BER. However, a BER test takes much longer to perform (several hours or even days) than an OSA test (often less than a few minutes).

The following two charts, which are sourced from Fabio Moliner García's "Analysis of the OSNR of the Links of Optical Fibre," *Universidad Autónoma de Madrid, 2009*, clearly show the OSNR and BER relationship (each curve is system-specific, depending on dispersion, modulation format, rate, channel count, etc.; however, all other things being equal, the OSNR-BER relationship is clear).




Figure 97. Relationship between OSNR and BER: a higher OSNR means a lower BER.

The importance of OSNR in coherent systems is further reinforced by the fact that these systems are not planned for metro or long-haul applications, but rather for ultra-long-haul applications spanning over several hundred or even thousands of kilometers. Each amplifier, whether it is EDFA or Raman, increases the noise. This, combined with the fact that for a similar Tx power, the power is divided into two phases and two polarizations that cause each symbol to have 1/4 of the total power, or 6 dB less than non-DP-QPSK systems, meaning that the pressure on the OSNR makes it a prime candidate for network failures.

Traditional OSNR Definition

The conventional method of determining the OSNR, called the interpolation method, requires that the noise level be measured at the midpoint between two peaks, and that a linear interpolation be performed. The noise under the peak can then be estimated and the OSNR calculated.




Figure 98. IEC 61280-2-9 method for measuring OSNR.

The value measured at the output of the first multiplexer (transmit side) should be greater than 40 dB for all channels. This value will be greatly affected by the optical amplifiers in the link, and will drop to about 15 dB to 20 dB at the end of the link (receiver side), depending on the link length, the number of cascaded EDFAs, and the bit rate.

In non-coherent systems, and in the absence of ROADMs, the interpolation method is currently used and generally accepted.

The IEC method described above assumes that the noise is visible between channels, and that it is flat throughout the spectrum. These two conditions are true in most legacy networks, but false in next-generation networks.

In the first case, when there are filters such as ROADMs in the optical path, they will not only filter the signal, but some of the noise as well. As illustrated by the third channel in the chart to the right, the shoulders or humps on both sides of the peak represent filtered noise. The IEC OSNR method would measure the noise at the red arrow location, while in actual fact it is found at the green arrow level.

Accordingly, in such cases the IEC method underestimates noise, and users are led to believe they have better OSNR than is actually the case, thereby creating a dangerously false sense of security.




Figure 99. Noise level of a filtered channel. Red arrow: noise level using the IEC method. Green arrow: real noise level.

The second case arises when 40G/100G is present. The spectral shape of these wavelengths is much larger than that of a 10G wavelength (typically two and a half to four times wider). This prevents the spectrum analyzer from seeing the noise level between two large signals. Again, in the example below, the red arrow represents the point at which the IEC method would measure the noise, whereas the green arrow represents a more probable noise level.




Figure 100. Noise level of 40G channels. Red arrow: noise level using the IEC method. Green arrow: real noise level.

In contrast to the previously shown filter case, in this example, the IEC method overestimates the noise, leading users to believe that they have worse OSNR than is actually the case, thereby creating a false sense of a problem, which can be quite costly.

The following chart represents a true operating transmission line, and is divided into three sections to highlight the previously explained details.




Figure 101. Spectrum with 10G channels in a ROADM network and 40G channels.

The first section to the left in the chart above (i.e., from 1527 nm to 1545 nm) represents ASE noise. Normally, ASE noise is somewhat flat and uniform. The example in the chart has gone through a certain amount of ROADMs, and therefore the shapes of the different residual filters can be observed. This proves that in such systems, one of the assumptions of the IEC OSNR measurement method, i.e., that noise is flat and uniform, is false.

The second section (i.e., from 1545 nm to 1553 nm) displays narrow 10G transmission in a ROADM mesh network, which will create a false sense of security. This example clearly indicates the level of noise that the IEC method would measure versus the true noise experienced by the channel. In this case, the error is about 3 dB, which in many cases is a toggle between system acceptance and failure. This system could experience BER despite the fact that all the parameters seem OK.

Finally, according to the information displayed in the third section (i.e., from 1553 nm to 1562 nm), the user would get an inaccurate representation of the problem if using a traditional IEC-based OSA, as demonstrated by the different noise levels highlighted. The network manager could spend hours trying to understand, from a network point of view, why the OSNR is so low, when in fact it could be just fine. So, a single link can contain wavelengths that are over or underestimated based on the testing method.

It is worthy to note that most vendor management systems (VMS) use this OSNR definition to compute the OSNR, and to balance and attempt to optimize the system, which could lead to significant and costly mistakes.

Next-Generation OSNR Definition

In cases such as those previously described, a new test method must be developed and a new definition standard defined. IEC/TR 61282-12 is a technical report titled "*Fibre optic communication system design guides – Part 12: In-band optical signal-to-noise ratio (OSNR)*," which discusses the definition of in-band noise.

The purpose of this document is to provide a definition of OSNR in next-generation networks. This standard suggests integration of a local OSNR over a certain bandwidth as defined below:

• OSNR = 10log(R)

$$R = \frac{1}{B_r} \int_{\lambda_1}^{\lambda_2} \frac{s(\lambda)}{\rho(\lambda)} d\lambda$$

- s(λ): time-averaged signal spectral power density, not including ASE, expressed in W/nm
- ρ(λ): ASE spectral power density, independent of polarization, expressed in W/nm
- B_r: reference bandwidth expressed in nm (usually 0.1 nm)
- λ₁ to λ₂: signal spectral range

Figure 91 depicts how this calculation is done.





Figure 102. OSNR definition from IEC 61282-12 standard.

As previously mentioned, historically, OSNR has been based only on noise coming from the ASE. But, only measuring OSNR in this manner could lead to serious errors. For example, boosting the power to increase the ASE-based OSNR may actually lead to worse BER.

The following figure shows that once a certain level of OSNR is reached (21 dB to 22 dB in this example), the spectral efficiency starts to decrease again, creating additional BER.



WDM non-linearities responsible for limiting capacity

[R.J.Essiambre et al, JLT 2010]

Figure 103. Impact of non-linear effects on spectral efficiency as a function of power.

The pink line in Figure 103 (above) represents the pure ASE non-linearity trace. The blue line represents WDM non-linearities, and the green line represents single ASE plus optical filtering. The red line represents the resulting trace from a real signal combining the different effects of ASE noise and non-linearity impairments.

As indicated, at a certain point the transmission quality will be limited by non-linearity. As such, if the BER is insufficient, it will be difficult to determine whether the cause of the impairment is ASE noise or non-linearity. Most optical spectrum analyzers deliver OSNR based on ASE noise only, which is part of the actual definition, but do not account for non-linear noise and OSNR contribution.

As discussed in the section devoted to chromatic dispersion, NLEs can now be monitored using the WDM Investigator dashboard while the OSA delivers a direct ASE noise measurement independent of other noise sources. In addition, noise is defined two-fold (whenever distinguishable and independently measurable), as follows: ASE noise and extended noise (taking into account ASE and non-linear noise). The BER estimate should be based on this extended noise and extended OSNR.

That said, a final receiver BER test is always recommended.



Figure 104. WDM Investigator reading showing a non-linear depolarization contributing to extended noise (OSNR_e).

Pol-Mux OSNR Definition

For some 40G systems and for the vast majority of 100G+ systems, polarization multiplexing is used. The next step is to determine whether the IEC method or the in-band methods work for these two Pol-Mux signals. As shown in figures 93 and 94, the IEC method won't work for Pol-Mux signals, because it does not take into account the impact of ROADMs (i.e., filtered noise), and because the large spectral width of Pol-Mux signals prevents measurement of the noise level at the midpoint between them.

The next step involves in-band OSNR applied to Pol-Mux signals. The basic setup of an in-band OSA, as shown in Figure 105, consists of using a polarization controller followed by a polarization beam splitter. In the case of non Pol-Mux signals, given that the signal is polarized and that the noise is unpolarized, adjustments to the polarization controller will change the proportion of signal split into each of the two branches, SOP-1 and SOP-2. To compute the in-band OSNR, complex algorithms can then be applied to the OSA traces found in branches one and two.




Figure 105. In-band OSNR setup.

For Pol-Mux signals, in-band approaches such as polarization nulling or EXFO's WDM-Aware will not work, because the signal appears unpolarized due to the fact that it consists of two orthogonal polarizations, a fact acknowledged by ITU-T Recommendation G.697. So, both IEC and in-band methods fail with OSNR measurements of 40G/100G Pol-Mux signals, thereby requiring a third method, called Pol-Mux OSNR.

There are currently two standards applicable for Pol-Mux OSNR: the IEC 61282-12 mentioned previously, as well as the China Communications Standards Association (CCSA) YD/T 2147-2010.

The Chinese CCSA YD/T 2147-2010 standard recommends calculating Pol-Mux OSNR as follows:

$$Pol\ Mux\ OSNR = 10 \log_{10} \left(\frac{P-N}{n/2} \right)$$

where, for a 50 GHz channel:

P = integrated power (signal + noise) over the 0.4 nm channel bandwidth

N = integrated power (noise) over 0.4 nm bandwidth

n = integrated power (noise) inside 0.2 nm, then normalized to 0.1 nm

Without going into too much detail, both of these standards are based on the channel turn-off method, which consists of measuring the signal plus noise power with the channel turned on, and then measuring the noise power with the channel turned off.

EXFO's Pol-Mux OSNR measurement tool, called the commissioning assistant, is a software option available in EXFO's FTBx-5245/5255 OSA that complies with both the IEC 61282-12 standard and the CCSA standard.

To utilize the commissioning assistant, the user must first take a measurement at the receiver with all of the channels turned on (as shown in Figure 106), and then acquire a series of traces that have each been taken with one channel turned off (as shown in Figure 107). The commissioning assistant therefore requires a total of m+1 traces, where m represents the number of channels. The commissioning assistant then performs the Pol-Mux OSNR calculations via a user-friendly wizard.




Figure 106. Trace with all channels on, ready for Pol-Mux OSNR calculation with the commissioning assistant.




Figure 107. Trace with channel number seven turned off, ready for Pol-Mux OSNR calculation with the commissioning assistant.

Although, in theory, a user who has all of the $m+1$ traces can perform these measurements manually with the OSA markers and compute the Pol-Mux OSNR value, the procedure is long and tedious. In fact, the commissioning assistant enables time savings and reduces the risk of human error associated with the manual method.

In-service Pol-Mux OSNR

Although the commissioning assistant is perfect for OSNR measurements of 40G/100G+ signals at turn-up, it is not suitable for measurements on live networks, because it requires turning off channels. Therefore, in-service Pol-Mux OSNR is required. EXFO has developed a non-intrusive, reference-based method in that it relies on a detailed spectral comparison of a reference measurement trace against an active measurement trace (the signal of interest). The reference measurement trace may have been acquired some time ago and/or at another location via the available taps (monitor ports) in the transmission system. A reference contains two elements: a spectrum, and a knowledge of the Pol-Mux OSNR values of the channels under analysis. These Pol-Mux OSNR values can either be contained in the OSA files (e.g., obtained with the commissioning assistant), or from another source and manually entered into the OSA file.

Figures 108 to 110 show the most common combinations of reference and active traces. Figure 108 describes a case where a reference has been acquired at time t_0 , commonly at turn-up, and an active trace is taken at a later time (t), say 6 months later. On the other hand, figure 109 displays a case where both the reference trace and active trace are taken at roughly the same time (t), in which the reference Pol-Mux OSNR values are known. Figure 110 shows a case that can be very useful if no prior information about the network is available. In this case, the reference acquired at the transmitter Tx can be assumed to be noise free, because almost all the noise comes from the optical amplifiers, which the signal has not yet travelled through. The OSNR value can then be manually entered as 35 or 40 dB in the reference trace.




Figure 108.




Figure 109.




Figure 110.

The basic measurement principle consists of subtracting a noise-free reference from the active trace, which contains signal + noise, in order to obtain the noise level (figure 111).




Figure 111. Reference-based Noise Measurement

In general, a normalization step is needed when the power of the active and the reference traces are different.

A noise-free reference is only readily available at the transmitter location. However, one can compute a noise-free reference from any location provided that the OSNR or noise level at this location is known. To do so:

- The first step consists of computing the noise level of the reference from the OSNR value
- The second step consists of subtracting the reference noise level from the reference spectrum
- A noise-free reference is then obtained (figure 112)




Figure 112. Computing a noise-free reference

All these steps are carried out using a user-friendly wizard in EXFO OSA— the in-service Pol-Mux assistant.

3.3. BER Test Using Receiver-Provided Information

Measuring EVM_{rms} at the receiving point would be very similar to measuring Q on legacy systems. EVM_{rms} can then be directly related to BER before the FEC at the receiver point. For more information about EVM_{rms} and receiver validation, please read section 2.1.3 and 2.2.




Figure 113. BER vs. EVM relationship for different modulation formats.

3.4. Inspecting Fibers and Connectors

Inspection and cleaning of optical fiber connector endfaces.

According to one NTT Worldwide Telecommunications survey, more than 70% of optical network faults arise from fiber endface inspection. A dirty or damaged endface of an optical fiber connector leads to increased insertion loss, return loss and even connector burnout in extreme cases.

High Insertion Loss (IL) Impact

Compared with 1G/10G Ethernet, 40G/100G Ethernet total channel IL for 50/125 micron multimode fiber was reduced to 1.9 dB (100 m OM3) or 1.5 dB (150 m OM4) for 40GBASE-SR4 or 100GBASE-SR10. Consequently, a maximum connector loss of 1.0 dB is required for a 150-meter channel containing multiple connector interfaces and high-bandwidth OM4 fiber. Therefore, in data centers, upgrades to higher data rates such as 40G/100G may fail, because the tolerance to IL becomes much tighter.

	IEEE	Designation	Mbit/s	Fiber Type	Number of Fibers	Maximum Link Length (m)	Maximum Channel Insertion Loss [dB]
10 Gbit Ethernet	802.3ae	10GBASE-SR4	10 000	OM3	2	300	2.6
40 Gbit Ethernet	P802.3ba	40GBASE-SR4	40 000	OM3	8	100	1.9
				OM4		150	1.5
100 Gbit Ethernet	P802.3ba	100GBASE-SR10	10 000	OM3	20	100	1.9
				OM4		150	1.5

High ORL impact

Every system has a maximum ORL, to which clean connectors are vital. One area where ORL can be extremely detrimental is in high-speed coherent transmission (40G and 100G transport). In most of these deployments, whether they are greenfield or brownfield, low-loss amplification is required in order to optimize distance. This means deploying a mix of EDFAs and more recently, Raman amplifiers. Raman is a low-noise amplification that uses the fiber itself as the amplifying media. Raman can easily be added to any existing infrastructure with little engineering, but because the fiber is the amplifier, every light travelling within it will be amplified (i.e., the signal and unwanted reflections). Reflections must therefore be kept to a minimum in every single Raman-amplified system.

In the end, all of these elements are reflected in network quality, i.e., increased bit error or even service disruption. Therefore, inspection and cleaning of the optical fiber endface is vital to the optical network. For optical fiber endface inspection, common instruments for optical fiber endface inspection include optical fiber microscopes and optical fiber endface detectors. However, the optical fiber microscope may be eliminated due to its inconvenient operation and inherent safety risks (danger of eye damage due to possibility of direct eye exposure to lit fiber). As such, the fiber endface detector is the mainstream product in the current market.



Figure 114. Probe for optical fiber endface inspection.

Two problems might be encountered when inspecting the optical fiber endface: endface damage or contamination.





Figure 115. Endface scratches (left) and clad breakage (right).

Endface scratches and breakage cannot be repaired by cleaning. If serious, it is necessary to replace the connector. Contamination can usually be removed by cleaning, which consists of either dry rubbing or wet rubbing.

The hybrid method combines dry and wet cleaning. First, the connector endfaces need to be cleaned with solvent, after which any residual solvent must be wiped off with a wipe or cotton swab (depending on the type of connector being cleaned). The inspection clean inspection clean (ICIC) method is recommended in practical operation. The diagram on the next page compares optical fiber endfaces before and after cleaning. In terms of optical fiber endface inspection, related international standards (such as those of the IEC and IPC) have very detailed requirements; meanwhile, operators can also define their own standard for optical fiber endface inspection according to their respective situations.



It should be noted that the current 40GBASE-SR4 for the CFP module adopts the multimode fiber MPO/MTP interface. The MPO interface has twelve cores, among which eight cores (four cores on the outer sides each) are used to send and receive Ethernet signals. The MPO/MTP connector is pictured in the first image on the right.




100GBASE-SR10 adopts the 24-core multimode MPO/MTP interface, among which the 20-core fiber is used to send, receive and transport Ethernet signals. The 24-core MPO/MTP connector is pictured in the second image at the right.



EXFO's Fiber Inspection Solution

EXFO is the first company in the industry to supply fully automatic optical-fiber-endface inspection and test instruments. Our fiber inspection probe (FIP) product range offers automatic fiber centering, automatic focusing, automatic capture of endface images and automatic determination of pass/fail information as per the relevant standard. In particular, its multicore optical fiber functionality can be used to inspect 40G/100G multicore MPO/MTP connectors. The test interface is as follows:



3.5. Fault Sectionalization and CFP Health Check

Internal Loopback

The FTB-88100NGE field tester supports an internal loopback capability that makes it possible to isolate the test set at any time in order to ensure error-free operation before the next component on the optical link is tested. Once the FTB-88100NGE is confirmed to be running error-free, the 40G/100G CFP can be looped back using an optical fiber to identify faulty CFPs, which can then be replaced if necessary. In this case, confirmation can be obtained to ensure that the CFP loopback configuration is also running error-free. The next step consists of looking into the 40G/100G NE or router itself, as well as its pluggable CFP, and so on.

Per-Optical-Channel Laser Control and Tx/Rx Power Measurements

Unlike the single wavelength transceiver that was used for legacy 2.5G and 10G, each CFP parallel optical channel must be monitored for transmitted and received power levels. Although the overall power across all channels (4 or 10) might be within the acceptable range, this might be the result of averaging a channel with very low power, possibly affecting the transmission and the performance on the optical channel and a high power channel, which might damage the optical receiver at the other end. Making it simpler, the FTB-88100NGE, FTBx-88200NGE and FTB-890/890NGE support per-optical-channel color-coded Tx/Rx power measurements.

Per-Lane Frequency and Frequency Offset Measurements

Field testers, such as the FTB-88100NGE, FTBx-88200NGE and FTB-890/890NGE, support per-lane frequency, frequency offset, maximum positive and maximum negative offset in ppm measurements to ensure proper network clock and timing recovery.

The CFP/CFP2/CFP4/QSFP Information Page provides detailed information on the plugged-in CFP/CFP2/CFP4/QSFP: module ID, vendor name, supported rates, firmware and hardware revision of the transceiver, among other valuable information. As a result, removal of the transceiver is no longer required in order to read the CFP/CFP2/CFP4/QSFP module details. This information helps to identify interoperability issues between two transceivers from the same or different suppliers. Furthermore, all this information is captured in the test report generated by the FTB-88100NGE, FTBx-88200NGE or FTB-890/890NGE. Different types of transceivers can be used while multiple job IDs are being performed throughout the day, in accordance with the service and its rate.

User-Configurable Skew Threshold and Excessive Skew Monitoring

The 40G/100G implementation introduces new challenges due to data distribution over multiple lanes with potential differential delay in arrival, which is referred to as skew. Although this effect triggers the need to test individual concepts such as PCS lanes, PCS skew and the alignment markers in the lab environment before rolling the equipment in the field (since this is definitely not standard commissioning and turn-up practice), EXFO's FTB-88100NGE, FTBx-88200NGE and FTB-890/890NGE offer the flexibility to provision the skew alarm threshold value and monitors per-lane excessive skew as part of its standard BERT and RFC 2544 test applications. The advanced test results view is categorized under advanced functions, distinct from the default simplified test results pages.

CFP MDIO Read/Write and Advanced CFP Troubleshooting

The user interface of EXFO's FTB-88100NGE, FTBx-88200NGE and FTB-890/890NGE field testers provides the full access needed to verify and manipulate the CFP electrical pins, as well as full management data input/output (MDIO) read/write access, allowing the user to read the CFP temperature, enable advanced CFP functions, and even set the CFP to troubleshooting mode.

Furthermore, and consistent with EXFO's simplified testing approach, the advanced CFP/CFP2/CFP4/QSFP testing capabilities (including MDIO/i2C read/write, among others) are distinctly categorized under the Functions category in the GUI, separate from the default setup and results pages. The Functions category offers various 40G/100G testing capabilities required by tier-2 engineers for advanced field troubleshooting, eliminating the need for a second test instrument.

Per-Lane BER Testing and Advanced CFP Troubleshooting

EXFO's FTB-88100NGE, FTBx-88200NGE and FTB-890/890NGE field testers featuring advanced per-lane BER capability provide the mechanism for testing crosstalk using specific PRBS patterns (PRBS 9 – PRBS 31). This important test verifies any signal integrity issues in the CFP that could cause bit errors.

Transceiver testing

The Intelligent Pluggable Optics test application—iOptics—was designed as the first tool to be used to assess any transceiver being installed in a network element deployed in the field or lab environment, and it serves to validate the proper functioning of the optical device.


Using minimal configuration, it quickly validates the soundness of an optical device and it includes a sequence of basic sub-tests, where parameters and thresholds are automatically set or are user configurable.

It supports a wide range of optical modules: QSFP28, QSFP+, CFP, CFP2, CFP4, XFP, SFP+, SFP.

The automated test sequence is conducted as follows:

- 1. Power Monitoring
- 2. Quick Check (MDIO/I2C, Common Status/Control Pins)
- 3. Temperature Monitoring
- 4. Optical TX Power Level Range Test (each lane is individually tested for multi-lane devices)
- 5. Optical RX Signal Presence & Level Range Test (each lane is individually tested for multi-lane devices)
- 6. Stress BERT
- 7. Framed Excessive Skew Test (40G/100G devices only)

For more details about this tool, please refer to Section 2.3.



3.6. Continuity Testing

OTN supports multiple traces, including OCh data unit (ODU), port monitoring (PM), trail trace identifier (TTI), OCh transport unit (OTU), segment monitoring (SM) TTI, and six tandem connection monitoring (TCMi) TTI traces. When commissioning a new OTN link, trace messages are usually used to monitor the integrity of the connection route between terminals during establishment of the connection. In addition, when a connection is activated, a TTI message is needed to ensure the connectivity. TTI messages for the OTN contain network-related information in the form of the source access point identifier (SAPI) and destination access point identifier (DAPI). Usually, once it has been confirmed that the received SAPI and/or DAPI in the TTI fails to match the preset expected value(s), a trace identifier mismatch (TIM) alarm is generated. The test module's TTI test functionality enables the user to provide SAPI, DAPI and operator-specific information fields in SM, PM or TCM TTI, and verify whether these fields have been properly sent over the whole network. These modules can also be used to monitor TTI messages initiated by the network management system.




Figure 116. TTI SAPI/DAPI/operator information test using the FTB-88100NGE Power Blazer.

3.7. OTN Long-Term BERT

The typical performance test for OTN projects is the 24-hour long-term error-free test. This bit-error test can be carried out with the OTU4 interface by carrying a 100 GigE client signal. Test configuration:




Figure 117. OTN long-term BER testing configuration

Test procedure:

- Connect test configuration as per the above diagram. All OTU4 services must be concatenated respectively. The services should be activated normally.
- Configure the client signal as 100 Gbit/s, then start the long-term BER test.
- Record the test results after 24 hours.

3.8. ITU-T Y.1564 for Ethernet Services Testing

With Ethernet continuing to evolve as the transport technology of choice, networks have shifted their focus from purely moving data to providing entertainment and new applications in the interconnected world. Ethernet-based services such as mobile backhaul, business and wholesale services need to carry a variety of applications, namely voice, video, e-mail, online trading and others. These latest applications impose new requirements on network performance, and on the methodologies used to validate the performance of these Ethernet services.

The following sections cover EtherSAM or ITU-T Y.1564, the ITU-T standard for turning up, installing and troubleshooting Ethernet-based services. EtherSAM is the only standard test methodology that allows for complete validation of Ethernet service-level agreements (SLAs) in a single, significantly faster test with the highest level of accuracy.

The Reality of Today's Networks

Ethernet networks are now servicing real-time and sensitive services. This reference to services refers to the various types of traffic that the network can carry. Generally, all network traffic can be classified under three traffic types: best-effort, real-time and high-priority. Each traffic type is affected differently by the network characteristics, and must be groomed and shaped to meet their minimum performance objectives.

Traffic Type	Main Application	Examples of Service
Best-Effort Data	Non real-time or data transport	• Data • Internet access • FTP download/upload • Server, storage applications
Real-Time Data	Real-time broadcast that cannot be recreated once lost	• VoIP • IPTV, video on demand • Internet radio, TV • Internet gaming • Videoconference
High-Priority Data	Mandatory traffic used to maintain stability in the network	• OAM frames • Switching/routing control frames • Network synchronization such as SyncE, 1588v2

Figure 118. Network traffic types.

To assure QoS, providers need to properly configure their networks to define how the traffic inside will be prioritized. This is accomplished by assigning different levels of priority to each type of service and accurately configuring network prioritization algorithms. QoS enforcement refers to the method used to differentiate the traffic of various services via specific fields in the frames, thus prioritizing frames for certain services over other frames. These fields make it possible for a network element to service and discriminate between high- and low-priority traffic.

The Importance of SLAs

An SLA is a binding contract between a service provider and a customer that guarantees the minimum performance that will be assured for the services provided. These SLAs specify the key forwarding characteristics and the minimum performance guaranteed for each characteristic.

Key Performance Indicators	Best-Effort Data (Internet Access)	Real-Time Data	High-Priority Data
CIR [Mbit/s] (green traffic)	2.5	5	10
EIR [Mbit/s] (yellow traffic)	5	0	5
Frame delay [ms]	<30	<5	5-15
Frame delay variation [ms]	n/a	<1	n/a
Frame loss [%]	<0.05	<0.001	<0.05
VLAN	300	100	200

Figure 119. Key performance indicators (KPIs) for various traffic types.

Customer traffic is classified into three traffic classes, each of which is assigned a specific color in Figure 120 (below): green for committed traffic, yellow for excess traffic and red for discarded traffic.

- **Committed information rate (CIR), or green traffic:** Refers to bandwidth that is guaranteed available at all times for a specific service; for green traffic, minimum performance objectives (i.e., key performance indicators or KPIs) are guaranteed to be met.
- **Excess information rate (EIR), or yellow traffic:** Refers to excess bandwidth above CIR that may be available depending on network loading and usage; for yellow traffic, minimum performance objectives are not guaranteed to be met.
- **Discarded or red traffic:** Refers to traffic that is above the CIR or the CIR/EIR rate, and that cannot be forwarded without disrupting other services; red traffic is therefore discarded.

Traffic Class	Bandwidth	Performance Objective	KPI
Green traffic	0 to CIR	Guaranteed forwarding	KPIs are guaranteed
Yellow traffic	CIR to EIR	Best effort	KPIs are not guaranteed
Red traffic	> EIR or CIR	Discarded traffic	Not applicable

Figure 120. Traffic classes.

Key Performance Indicators

KPIs are specific traffic characteristics that indicate the minimum performance of a particular traffic profile. Under green traffic condition, the network must guarantee that these minimum performance requirements are met for all forwarded traffic. Typical KPIs are as follows:

- **Bandwidth:** Bandwidth refers to the maximum amount of data that can be forwarded. This measurement is a ratio of the total amount of traffic forwarded during a measurement window of one second. Bandwidth can either be “committed” or “excess,” with different performance guarantees. Bandwidth must be controlled, because multiple services typically share a link. Therefore, each service must be limited to avoid affecting another service. Generating traffic over the bandwidth limit usually leads to frame buffering, congestion, and frame loss or service outages.
- **Frame Delay (Latency):** Frame delay, or latency, is a measurement of the time delay between the transmission and the reception of a packet. Typically, this is a round-trip measurement, meaning that it simultaneously calculates the near-end to far-end directions, as well as the far-end to near-end directions. This measurement is critical for voice applications, because too much latency can affect call quality, leading to the perception of echoes, incoherent conversations, or even dropped calls.
- **Frame Loss:** Frame loss can occur for numerous reasons, including transmission errors and network congestion. Errors due to a physical phenomenon can occur during frame transmission, resulting in frames being discarded by network devices such as switches and routers based on the frame-check-sequence field comparison. Network congestion also causes frames to be discarded, because network devices must drop frames to avoid saturating a link in congestion conditions.
- **Frame Delay Variation (Packet Jitter):** Frame delay variation, or packet jitter, refers to the variability in arrival time between packet deliveries. As packets travel through a network, they are often queued and sent in bursts to the next hop. Random prioritization may occur, resulting in packet transmission at random rates. Packets are therefore received at irregular intervals. This jitter translates into stress on the receiving buffers of the end nodes, which may result in buffers becoming overused, or underused in the event of large swings of jitter.

Real-time applications such as voice and video are especially sensitive to packet jitter. Buffers are designed to store a certain quantity of video or voice packets, which are then processed at regular intervals to provide a smooth and error-free transmission to the end user. Too much jitter will affect the quality of experience (QoE), because packets arriving at a fast rate will cause the buffers to overfill, resulting in packet loss. Packets arriving at a slow rate will cause buffers to empty, leading to still images or sound.

Testing Methodology: EtherSAM (ITU-T Y.1564)

To resolve issues with existing methodologies, the ITU-T has introduced a test methodology standard: ITU-T Y.1564, which is aligned with the requirements of today's Ethernet services. EXFO was the first to implement EtherSAM—the Ethernet service testing methodology based on this standard—into its Ethernet testing products.

EtherSAM enables complete validation of all SLA parameters within a single test, thus ensuring optimized QoS. Contrary to other methodologies, it supports new multiservice offerings. In fact, EtherSAM can simulate all the types of services that will run on the network, and simultaneously qualify all key SLA parameters for each of these services. In addition, it validates the QoS mechanisms provisioned in the network to prioritize the different service types, resulting in more accurate validation, and much faster deployment and troubleshooting. Moreover, EtherSAM offers additional capabilities, such as bidirectional measurements.

EtherSAM (ITU-T Y.1564) is based on the principle that the majority of service issues are found in two distinct categories: a) in the configuration of the network elements that carry the service, or b) in the performance of the network during high-load conditions, when multiple services cause congestion.

Service Configuration

Forwarding devices, such as switches, routers, bridges and network interface units, are the foundation of any network, because they interconnect segments. These forwarding devices must be properly configured to ensure that traffic is adequately groomed and forwarded in accordance with its SLA.

If a service is not correctly configured on a single device within the end-to-end path, network performance can be greatly affected. This may lead to service outages and network-wide issues, such as congestion and link failures. Therefore, a very important part of the testing effort is ensuring that devices are properly configured and able to handle the network traffic as intended.

Service Performance

Service performance refers to the ability of the network to carry multiple services at their maximum guaranteed rate without any degradation in performance; i.e., KPIs must remain within an acceptable range.

As network devices come under load, they must make quality decisions. They must prioritize one traffic flow over another in order to meet the KPIs of each traffic class. This is necessary, because as the amount of traffic flow increases, so does the likelihood of performance failures.

Service performance assessment must be conducted over a medium- to long-term period, because problems typically occur in the long term and will probably not be seen with short-term testing.

The focus of EtherSAM (ITU-T Y.1564) is therefore threefold:

- First, the methodology serves as a validation tool, ensuring that the network complies with the SLA by ensuring that a service meets its KPI performance objectives at different rates, and within the committed range.
- Second, the methodology ensures that all services carried by the network meet their KPI performance objectives at their maximum committed rate. This proves that under maximum load, the network devices and paths are able to service all the traffic as designed.
- Third, service performance testing can be performed over medium- to long-term test periods to confirm that network elements are able to properly carry all services while under stress during a soaking period.

EtherSAM: Tests and Subtests

EtherSAM is comprised of two tests: the service configuration test and the service performance test. The frame size used for the service configuration test and the service performance test can be constant, or a distribution of multiple frame sizes. The ITU-T Y.1564 has defined a variable frame-size sequence format named EMIX, or Ethernet Mix. The EMIX frame sequence format can be configured from two to eight frames, with configurable frame sizes ranging from 64 bytes to 16000 bytes. EMIX's main purpose is to emulate real-life network traffic and uncover potential issues that may not arise when testing with a constant frame size.

Service Configuration Test

The service configuration test is a per-service test that verifies the bandwidth and performance requirements of a specific service, as defined by the user. The process follows three key phases and monitors all performance indicators during these steps to ensure that they are all met at the same time.

Phase 1: Minimum Data Rate to CIR

In this phase, bandwidth for a specific service is ramped up from a minimum data rate to the committed information rate (CIR). This ensures that the network is able to support this specific service at different data rates while maintaining the performance levels. In addition, it provides a safe and effective way to ramp up utilization without overloading a network, in the event that the service is not configured correctly. As the service is gradually ramping up to the CIR, the system automatically measures the KPIs at each step to ensure that the minimum performance objectives are always met. For this phase to pass, all performance objectives must be met at each step all the way up to the CIR.

Phase 2: CIR to EIR

In this phase, the service is ramped up from the CIR to the excess information rate (EIR). This ensures that the service's EIR is correctly configured, and that the rate can be attained. However, as per accepted principles, performance is not guaranteed in the EIR rates; therefore, no KPI assessment is performed. At this stage, the system only monitors the received throughput. Since EIR is not guaranteed, bandwidth may not be available for all traffic above the CIR. A pass condition corresponds to the CIR as the minimum received rate, and the EIR as a possible maximum. Any measured rate below the CIR is considered as having failed.

Phase 3: Traffic Policing Test

One of the attributes of packet transport is the capability to handle bursty traffic. EIR can occur in conditions of burst, or conditions that surpass the committed bandwidth, and this usually leads to discarded traffic. In this step, traffic is sent above the EIR, and the received rate is monitored. At the very least, the CIR must be forwarded. The EIR traffic should be forwarded depending on the availability of resources. Any traffic above this maximum should be discarded in order to avoid overloading the network. If the received traffic exceeds the EIR, this means that a device is not properly configured and a fail condition will be signaled.

These three phases are performed per service; therefore, if multiple services exist on the network, each service should be tested sequentially. This ensures that there is no interference from other streams, and that the bandwidth and performance of the service alone are measured specifically. At the end of the Ethernet service configuration test, the user has a clear assessment of whether or not the network elements and path have been properly configured to forward the services while meeting minimum KPI performance objectives.

Phase 4: Burst Testing

The burst test is a subtest within the service configuration test. In the context of SLA assessments, the objective of burst testing is to verify that the expected burst size can be transmitted through the network equipment with minimal loss. The bandwidth profile of the network equipment contains attributes of committed burst size (CBS) and excess burst size (EBS) that service providers should test at the time of service activation to verify proper attribute configuration. The most common protocol used today to transport data in IP-based networks is transmission control protocol (TCP), which by nature is a bursty protocol. Therefore, it is very useful for service providers to perform burst testing for TCP-based applications (e.g., FTP, HTTP and e-mail services) during their service turn-up and troubleshooting phases. The burst test phase is composed of two parts: the CBS and EBS test. The CBS is the number of allocated bytes available for bursts transmitted at rates above the CIR while meeting the SLA requirements. The EBS is the number of allocated bytes available for bursts transmitted at rates above the CIR+EIR while remaining EIR-conformant. The following graph shows an example of the CBS and EBS burst test sequence.




Figure 121. CBS and EBS burst test sequence.

Because the CBS and EBS attributes on the network equipment may be configured differently for each service direction, testing CBS and EBS in a round-trip configuration (one end in loopback) has little to no value. It is essential that these parameters be tested independently for each service direction. Leveraging EXFO's simultaneous bidirectional testing, network operators can stress and emulate real-life network traffic during their service turn-up and troubleshooting phases. This is truly the only way to accurately test and validate proper network configuration and operation—especially when testing with a bursty traffic type such as the previously mentioned TCP. Along with the CBS and EBS burst size, the burst sequence parameters are fully configurable on EXFO's PowerBlazer series.

Figure 122. Burst testing—SLA parameters configuration.

Service Performance Test

While the service configuration test concentrates on the proper configuration of each service in the network elements, the service performance test focuses on the enforcement of the QoS parameters under committed conditions, replicating real-life services. In this test, all configured services are generated at the same time and at the same CIR for a soaking period that can range from a few seconds to a maximum of 30 days. During this period, the performance of each service is individually monitored. If any service fails to meet its performance parameters, a fail condition is signaled.

The combination of these two tests provides all the critical results in a simple and complete test methodology. The service configuration test quickly identifies configuration faults by focusing on each service and how it is handled by the network elements along the paths. The service performance test focuses on the network's capacity to handle and guarantee all services simultaneously. Once both phases have been successfully validated, the circuit is ready to be activated and placed into service.

EtherSAM Test Topologies: Loopback and Bidirectional (Dual Test Set)

EtherSAM can also perform round-trip measurements with a loopback device. In this case, the measured value reflects the average of both test directions, from the test set to the loopback point and back to the test set. In this scenario, the loopback functionality can be performed by another test instrument in Loopback mode, or by a network interface device (NID) in Loopback mode. The same test can also be launched in Dual Test Set mode. In this case, two test sets, one designated as local and the other as remote, are used to communicate and independently run tests simultaneously for each direction. This provides much more precise test results, such as independent assessment per direction and the ability to quickly determine which direction of the link is experiencing failure. It is important to point out that EXFO's EtherSAM test application performs a simultaneous bidirectional test, which means that traffic is active in both directions simultaneously. Testing today's advanced network paths simultaneously in both directions is crucial.

This emulates real-life network traffic, and can uncover network-equipment configuration issues that could go undetected with non-simultaneous bidirectional testing. Furthermore, performing simultaneous bidirectional testing significantly reduces costs by decreasing test time by 50%.

Carriers and service providers face the constant challenge of ensuring the proper delivery of services to customers. Ethernet services need to be delivered to customers in compressed time frames, while proving to be more reliable than ever. EtherSAM bridges the gap between service validation and performance assessment by providing an intuitive and easy approach for confident control and management of networks while reducing OPEX and growing revenues. EtherSAM is the only standard test methodology that allows for complete validation of SLAs in a single, significantly faster test, while offering the highest level of accuracy.

iSAM: Simplifying Ethernet Service Activation

How many times have field technicians started a test application only to realize that someone with a high level of technical expertise is required to configure and perform the test? The goal of iSAM is clear-cut: to simplify Ethernet service-activation testing for field technicians of all skill levels.

iSAM does this by making service turn-ups what they should be: simple and quick. Inspired by the feedback and field-testing knowledge of customers around the world, iSAM was designed to simplify field technicians' daily tasks by focusing on the most common Ethernet service turn-up configurations used in the industry today, from 10M all the way up to 100G.

Imagine that field technicians could set up an Ethernet service test within a single page, automatically connect to their remote unit, complete testing, and then have the test report automatically uploaded to a cloud-based storage location. iSAM makes this possible.

In today's test and measurement (T&M) market, field technicians often waste valuable time going through a multitude of test setup pages. Thanks to its innovative new approach, iSAM can be configured entirely within a single page for maximum efficiency.

The screenshot displays the iSAM Test Configurator interface, which is organized into three main sections: Local, Network, and Remote.

- Local Section:** Contains a 'LINK' status indicator, a 'Port 1' dropdown menu set to '10GE LAN', and an 'IP' field set to '10.10.106.160'. A 'More...' button is located at the bottom right.
- Network Section:** Features a central cloud icon with three numbered circles (1, 2, 3) and a summary box showing 'Total: 875.0 Mbit/s' and 'Duration: 12m35s'. Below this, a 'Layer' dropdown is set to 'L3/L4', and a 'Services' dropdown is set to '3'. A table lists the services and their CIR (Mbit/s):

Service	CIR (Mbit/s)
1 Best Effort	125.0
2 Priority	275.0
3 User Profile	475.0

At the bottom, there are checkboxes for 'Configuration Test', 'Performance Test', and 'RFC 6349 Test on Service 1', along with a 'More...' button.
- Remote Section:** Displays the 'Status' as 'Available', the 'Name' as 'FTB-880Q81', and the 'IP' as '10.10.200.129'. A 'Connect' button with a blue arrow icon is present. Below this, it shows 'Remote Found 1' and a 'More...' button.

Arrows indicate the flow of data and configuration between the Local, Network, and Remote sections.

Figure 123. iSAM-simple one-page test setup.

Tackling the Transmission-Control Protocol (TCP) Layer

Given the intense competition in the telecom industry today, service providers need to ensure the best possible quality of experience for their customers in order to prevent customer churn and increase customer loyalty. As such, complete validation of Ethernet and Internet protocol (IP) layers (layers 2 and 3), as well as TCP layer 4, is a must for service providers who need to understand their network and the quality of services being delivered to their customers.

Due to the shift in customer usage patterns and the demand for bandwidth-intensive applications, service providers must also adapt their testing requirements to meet customer expectations. Fifteen years ago, service providers rarely talked about testing or validating the TCP layer. However, today they need to go beyond traditional RFC 2544 testing, and even ITU-T Y.1564 service activation testing.

With the simplicity that iSAM brings to the table, service providers also have the ability to easily validate their services right up to TCP layer 4. Simply enabling the RFC 6349 test in the iSAM setup page will allow the TCP Throughput validation test to be performed automatically after layers 2 and 3 have been tested. This will ensure validation of the base layers (Ethernet and IP) along with the TCP layer, thereby delivering a true end-to-end service activation test.

The RFC 6349 test will automatically determine the path maximum transmission unit (MTU) to avoid packet fragmentation. The test will also measure the baseline round-trip time (RTT), which constitutes the minimum round-trip time required for a TCP packet to be sent out and acknowledged by the receiver. From the baseline RTT and the user-entered committed information rate (CIR), the bandwidth delay product (BDP), or optimum window size, will be calculated. The BDP (optimum window size) will be used to perform the TCP throughput test. Field technicians will then be able to quickly determine whether the circuit under test is operating at the expected level.




Figure 124. Layer-2/3 and Layer-4 (TCP) service validation

Performance Criteria

In addition to the simplicity it brings to the Ethernet testing world, iSAM also includes the latest cutting-edge Metro Ethernet Forum (MEF) standard. Using predefined performance metric profiles from MEF 23.1, field technicians can quickly select the MEF service type they need to test, and validate Ethernet services in a consistent, repeatable and automated manner. In this respect, iSAM uses a simple test methodology specifically adapted to today's next-generation Ethernet networks in order to properly qualify the complete service-level agreement (SLA).

Designed according to the same principles as EtherSAM, iSAM also integrates the concept of two subtests: the configuration test and the performance test. The configuration test verifies the performance according to per-service-basis criteria. This test verifies that the performance metrics defined by the MEF, such as frame loss ratio (FLR), frame delay (FD) and inter-frame delay variation (IFDV), are met at the CIR for each individual Ethernet service. The performance test, on the other hand, verifies that the performance metrics are met while all services are operating simultaneously at the combined CIR.

Performance-Metric Naming Convention	
Traditional	MEF
Frame loss	Frame loss ratio (FLR)
Latency	Frame delay (FD)
Jitter	Inter-frame delay variation (IFDV)

Over the years, many service providers have searched for guidance regarding the proper performance metric settings needed to test Ethernet services. For instance, should the FLR be below 0.5% for a short-span service such as the metro service? Should the FD be less than 10 ms, 20 ms, or 30 ms? There has been quite a bit of confusion concerning the right metrics needed based on the type of Ethernet service provided. iSAM clears up this confusion by providing MEF performance profiles to field technicians. With iSAM, service providers and field technicians can now easily test according to these guidelines based on the Ethernet service being delivered.

MEF 23.1 specifies three performance objectives and four performance tiers (PTs). The performance objective relates to the class of service (CoS) being delivered, whereas the performance tier relates to the span of the Ethernet service.

Performance Objectives	
H (High)	Stringent performance, such as timing distribution
M (Medium)	Real-time applications, e.g., voice-over-Internet protocol
L (Low)	Best-effort applications (web traffic)

Performance Tier (PT)—Span (km)	
Metro	< 250 km
Regional	< 1200 km
Continental	< 7000 km
Global	< 27 500 km

The combination of these performance objectives and performance tiers defines 12 different classes of Carrier Ethernet services. The table below outlines examples of the performance criteria for some of these services.

	FLR (%)	FD (ms)	IFDV (ms)
Metro L	≤ 0.1%	≤ 37 ms	Not specified
Continental M	≤ 0.025%	≤ 115 ms	≤ 40 ms
Global H	≤ 0.05%	≤ 230 ms	≤ 32 ms

Because iSAM includes an extensive list of Ethernet services, service providers no longer have to play a guessing game when turning up and validating Ethernet services. Field technicians can easily select the desired performance criteria based on the Ethernet service being delivered to the end customer.

Performance-Criteria Customization

What happens when the MEF predefined performance criteria do not meet the user's needs? With iSAM, users also have the ability to modify the predefined MEF performance criteria or create their own user-defined performance-criteria profiles. Some service providers may even have their own Ethernet service offering (e.g., Gold, Silver and Bronze) with associated performance criteria for these services. With iSAM, custom performance-criteria profiles can easily be created.

These custom performance-criteria profiles are created using master templates that are accessible on the platform (not within the test application), and can be restricted to managers and/or supervisors. Also, if a user has modified an MEF performance-criteria template, an asterisk will appear near the performance-criteria name to indicate that some modifications to the metrics have been carried out.

3.9. 40GE/100GE Troubleshooting Using Filtering and Packet Capture

With 100G set for mass deployment in order to address the growing demand for bandwidth, both service providers and network operators are facing a number of challenges. Some of these challenges are related to the multiservice offering of 100G networks, the nature of each and every one of its supported services, and its implications on SLAs. And, as the service offering scales to the 100G rate and becomes more and more complex, network engineers are being called on to perform even more troubleshooting and service calls.

100G network troubleshooting involves performing a number of complex procedures to identify where and why network failures are occurring. New networks are increasingly complex, usually offering very little information about the event to network engineers, who need to investigate various probable causes of failure.

The difficulty of the associated tasks is compounded by the pressure of limited investigation time and the risk of customers being affected.

One of the options available to network engineers is to capture and decode traffic on the affected circuit. Decoding traffic usually involves analyzing the header traffic to identify any issues, such as modifications or incorrect content.

Filtering

The first step in capturing traffic involves the filter. The key aspect is to focus on one particular traffic stream, since there is a lot of traffic in a normal transmission. Filtering will isolate the significant traffic while removing any unnecessary traffic, thus optimizing the search engine.

The filter engine is based on basic and advanced filtering capabilities. The filter can be configured based on different logical operands (AND, OR, NOT). The figure below shows the filtering flow, where the Ethernet packet can be passed through and/or selected for capture and analysis.




Figure 125. Traffic filtering mechanism.

In most captures, the payload information is typically proprietary information that cannot be decoded by the test equipment. As such, network engineers usually focus on the header information, which is decoded and used for more in-depth troubleshooting. Again, this eliminates the capture of unnecessary data that would otherwise consume memory without providing any additional useful information.

The FTB-88100NGE, FTBx-88200NGE and FTB-890/890NGE field testers provide an innovative truncation capability that limits capture to a specific number of bytes, starting from the first bit of the Ethernet frame. Network engineers can therefore limit capture to the first few bytes of the header, or add more bytes to include higher-layer information.


Capture Trigger

A very common issue with typical capture tools is that capture starts as soon as the capture tool is enabled. However, the event of interest may occur later, allowing previously captured traffic to fill the memory buffer without providing any useful information. In some cases, the testing opportunity can be completely missed due to the high amount of captured data and the short event window.

EXFO's Ethernet capture tool solves this issue by including a set of triggering capabilities that allow customers to fine-tune and specify when the capture process should start. This powerful capability simplifies the troubleshooting process by filling the memory only when the event of interest is detected. The memory and the troubleshooting time are therefore efficiently used, resulting in meaningful capture data that yields more important information.


Users can capture traffic based on the following three types of triggers:

- **Manual trigger** is the simplest type of trigger, and basically starts the capture as soon as it is enabled. This is the default mode of operation that mimics traditional capture tools.



- **On-error trigger** is a trigger that starts to capture the operation when a specific event is detected. These events are typically Ethernet errors such as frame check sequence (FCS) errors. This mode enables on-event capture, a scenario where a capture device monitors the circuit until the specific event is detected and the capture is triggered.

- **Field-match trigger** launches the capture when a frame with a specific filtered condition is detected. This condition uses a system similar to the traffic filter system and enables the user to monitor the circuit and start the capture as soon as a specific frame condition is detected.



Triggering Position

The triggering position is used to determine the position of the triggered frame within the captured data, solving a common problem with traditional capture tools, whereby the event of interest is often located within the capture data.

A typical use for the triggering position is to perform pre- and post-analysis. In network troubleshooting, it is very important to understand the events that led to the failure and to view the events that followed the failure. These two critical phases provide a wealth of information about the failure, what caused it, and how the network reacted to it.

The triggering-position capabilities allow the user to specify where the trigger event will be located in the capture, making it possible to select the frames to be captured depending on their position relative to the trigger event. Traditional capture tools do not offer the ability to perform a mid-trigger or pre-trigger; instead, they only provide post-trigger capabilities. As such, users are left to manually search within the captured sequence to identify the event and perform the analysis. This, in combination with the lack of a trigger mechanism, means that it is possible to completely miss the event of interest when using traditional capture tools, resulting in an inefficient capture process.

EXFO's Ethernet capture tool provides three triggering positions, as follows:


- **Post-Trigger mode:** the first frame of the capture is always the trigger, and the remaining frames are the frames that follow the trigger event. This mode is typically used to analyze content after the event.



- **Pre-Trigger mode:** the last frame of the capture is the trigger event; therefore, the captured output contains all the frames leading up to the event. This mode can be used to determine what led to the specific event.



- **Mid-Trigger mode:** a very powerful application that provides a snapshot of the traffic before and after the trigger event. In this mode, the trigger event is usually in the middle of the captured traffic.



Exporting Capture and Analysis

Once a capture is completed, the captured data can be exported to the platform's internal memory for decoding. The exporting process generates an industry-standard packet capture (PCAP) file that can be used by a variety of open-source decoding tools. Decoding and post-analysis is performed using the Wireshark application.

Conclusion

With 100G moving toward mass deployment, the Ethernet capture and decode capability offered by the FTB-88100NGE, FTBx-88200NGE and FTB-890/890NGE field testers allow network engineers to quickly pinpoint issues in the field and speed up the troubleshooting process to ensure quick service recovery.

Other tools for troubleshooting: IP ping

3.10. Ethernet Service OAM (SOAM)

Service providers and carriers around the world spend billions of dollars to build their networks. As networks become increasingly sophisticated, effective management capabilities become a necessity. Reliability, availability, fast fail over and recovery are all very important parameters that service providers and carriers rely on to ensure SLAs. One of the biggest challenges facing service providers as they deploy Ethernet-based solutions lies in achieving the same level of operations, administration and maintenance (OAM) support required in traditional carrier networks. Essentially, SLAs for specific services need to be met, regardless of the underlying technologies used to provide them. Ethernet OAM is one part of the capability needed to meet these SLAs.

What is Ethernet OAM?

OAM is a set of functions that provides system or network fault indication, performance monitoring, security management, diagnostic functions, configuration and user-provisioning. The purpose of these management tools or capabilities is to enable the monitoring and quick restoration of a network in case of failure. Given that a network is typically comprised of equipment owned by different operators and built by many different manufacturers, OAM has to be standardized to ensure consistency and interoperability. OAM entities are network-aware in that they use information from and provide information to other network entities. In addition, they cooperate to provide the consistency and conformity that are critical to an entity's OAM functions. Ethernet OAM injects OAM packets into the normal stream of data packets at layer 2, and uses endpoints to process those packets to determine performance using parameters such as improperly configured nodes, unidentified and out-of-place nodes, disconnected or failed nodes, frame loss, frame delay, end-to-end path identification and bit error rates.

Ethernet service OAM (SOAM), on the other hand, provides management solutions to guarantee and measure end-to-end performance. SOAM enables end-to-end SLAs for standardized Ethernet services, in-service SLA verification, as well as network monitoring and troubleshooting from the central office. Also, SOAM protocols support two sets of functions, the first of which is connectivity fault management (CFM), which detects, verifies and isolates connectivity failure as defined in ITU Y.1731, IEEE 802.1ag and MEF 30.1. This is performed end-to-end, although some functions can isolate faults in segments. The second set of supported functions is performance monitoring (PM), which provides the required capability for performance monitoring as defined in ITU Y.1731 and MEF 35. This is performed from end to end. Other OAM protocols include IEEE 802.1ab for link-layer discovery, IEEE 802.3ah for Ethernet in the first mile, and ITU-T G.8113.1 for multiprotocol label switching—transport profile (MPLS-TP) OAM. The Metro Ethernet Forum has defined additional service OAM requirements, namely MEF 17.

Definitions

• **Maintenance domain (MD):** The portion of a network typically owned and operated by a single entity, over which connectivity faults can be managed. MDs are configured with names and eight levels ranging from 0 to 7. A hierarchical relationship based on these levels exists between domains.




Figure 126. Customer maintenance domain.

- **Maintenance entity group end point (MEP):** The boundary points of a maintenance domain, they can initiate and terminate OAM frames. End-to-end tests are initiated and terminated by MEPs.
- **Maintenance entity group intermediate point (MIP):** The intermediate points in a maintenance domain, they do not initiate OAM frames, but can respond to some OAM frames (loopback and link trace) in order to isolate faults.
- **Maintenance entity (ME):** This entity requires management and defines a relationship between two maintenance entity group endpoints.
- **Maintenance entity group (MEG):** A group of MEs that are in the same administrative boundary, with the same MEG level, and belonging to the same point-to-point or multipoint Ethernet connection.
- **Maintenance association (MA):** A set of MEPs established to verify the integrity of a single service instance.




Figure 127. Hierarchical maintenance domain.

SOAM functions are performed from end to end (i.e., from customer to customer). Therefore, when a connectivity fault occurs, it is possible to locate it. In Figure 127 on the preceding page, a fault can be located in any one of three possible segments. In the service provider's maintenance domain, SOAM functions are performed from end to end within the MEG, which comprises two MEPs and two MIPs. In this location as well, a fault can be located in any one of the three possible segments. The other two segments belong to operators A and B, respectively. Therefore, operators A and B can use SOAM functions, but only within their respective MEGs.

Ethernet OAM Standards

This section covers the standards that are available for Ethernet OAM, including their different functionalities and use cases.

- **IEEE 802.1ag:** 802.1ag focuses on the end-to-end connectivity and continuity of nodes in an Ethernet network. This is why it is referred to as connectivity fault management, or CFM. Because it applies to bridges and bridge applications, it specifies a lot of multicast packets in addition to unicast packets, and also handles both multipoint and point-to-point connections.

802.1ag has three main functions: continuity check messages (CCM), loopback messages (LBMs) and loopback responses (LBRs), and link trace messages (LTMs) and link trace responses (LTRs).

- **CCM:** A continuity check message is an OAM protocol data unit (OAMPDU) that provides service monitoring from one endpoint to another (MEP to MEP). The CCMs exchanged between MEPs are analogous to a "heartbeat" message, and can be configured to be sent at one of seven standard intervals: 3.3 ms, 10 ms, 100 ms, 1 s, 10 s, 1 min, and 10 min. CCMs can be either multicast or unicast, although the use of multicast packets, which run continuously until turned off, is preferred. When the reception of CCM messages at a node is lost, this constitutes a loss of connectivity to that node. The reception of CCM messages from an unknown node represents a possible misconfiguration of nodes.





Figure 128. Continuity check message.

- **LBM/LBR:** LBMs and LBRs are used to determine the integrity of the data path. Once a network fault has been established, the service provider can verify the loss of service by initiating an 802.1ag loopback test. A loopback is similar to a layer-3 ping request/reply. In a loopback test, an MEP sends messages to another MEP or MIP to verify connectivity across a given MA. These messages can be unicast or multicast, although the use of unicast is preferred.
- **LTM/LTR:** LTMs and LTRs are used by an MEP to verify the complete link trace to a peer MEP. Each MIP and the peer MEP will respond to an LTM. Once a network fault has been confirmed by the service provider, the link trace feature can be used to isolate its specific location. This feature traces the service from one MEP to another MEP or MIP using its MAC address, and to all the MIPs along the MA.
- **ITU-T Y.1731:** The ITU-T Y.1731 protocol is used mainly for fault management and performance monitoring, defining performance monitoring measurements (i.e., frame loss ratio, frame delay, frame delay variation, service availability) in order to assist with SLA assurance and capacity planning. This protocol applies to both multipoint and point-to-point connections, and relies on the 802.1ag protocol for transport, making it a type of extension to 802.1ag. The ITU-T Y.1731 protocol measures the following performance parameters: frame loss, delay, delay variation and service availability. Its main features are alarm indication signals (ETH-AIS), remote defect indication (ETH-RDI), locked signal (ETH-LCK), test signal (ETH-Test), performance monitoring (ETH-PM), frame loss measurement (ETH-LM), frame delay measurement (ETH-DM) and client signal fail (ETH-CSF).



- **ETH-AIS:** This message is sent to the far end when the near end detects an alarm via the CCM packets. The AIS frame is transmitted periodically until the fault condition is resolved.
- **ETH-RDI:** This indication is used for fault notification. If an MEP is defective, it will transmit an RDI to its peer MEPs to inform them that a defect condition has been encountered.
- **ETH-LCK:** This signal is used to communicate an administrative lock to the far end, resulting in an interruption in data traffic. The signal tells the far end that the near end is present, but unavailable for use. LCK frames are also transmitted periodically until the administrator clears the lock.

- **ETH-Test:** This test signal, which is generated by an MEP, is sent to a peer MEP for verification of the integrity of the received test signal from the peer MEP. In addition, it is used for bit-error-rate (BER) and throughput measurements.
- **ETH-PM:** This functionality is used to monitor the performance of traffic from point to point or from end to end on a given domain.
- **ETH-LM:** This feature is used by an MEP to measure the frame loss with the peer MEP in both directions from a single endpoint.
- **ETH-DM:** This functionality is used by an MEP to measure the roundtrip delay with the peer MEP. This is also used to measure the delay as well as the delay variation.
- **ETH-CSF:** This feature is used for fault notification. It is used by an MEP to propagate to a peer MEP the detection of a failure or event in an Ethernet client signal. It is used when the client does not support fault detection mechanisms.
- **IEEE 802.3ah:** This standard is also referred to as Ethernet for the First Mile (EFM) OAM. While 802.1ag specifies verification of end-to-end connectivity across multiple domains and hops, 802.3ah specifies verification of point-to-point connectivity across only one hop. It uses the following mechanisms: discovery, loopback, polling of management information base (MIB) variables, remote failing indication and link monitoring.
- **Discovery:** This is the first phase. This is where the devices in the network are identified along with their OAM capabilities. If the discovery fails for any reason, all other OAM activity is aborted until the discovery can be re-established.
- **Loopback:** An OAM entity can put its remote peer into loopback mode using loopback control payload. This helps the administrator ensure link quality during installation or troubleshooting.
- **Polling of MIB Variables:** IEEE 802.3ah provides read-only access to remote MIB, but this is limited to specific MIB branches and leaves. This feature is based on the premise that the administrator can also retrieve/reset MIB variables at the far end. These variables are one of the primary sources of OAM information available to the system administrator.
- **Remote Failure Indication:** This mechanism enables the OAM entity to convey the degradation of an Ethernet link to its peers via specific tags in OAM payload.
- **Ethernet link trace:** This functionality, which is used for fault isolation, allows an MEP to verify the complete link trace to a peer MEP. Each MIP and its peer will respond to the link trace message.

The Importance of Testing Ethernet OAM

Ethernet OAM impacts every aspect of Carrier Ethernet, and is essential to the network, because it enables the automated provisioning, monitoring and fault isolation that makes Carrier Ethernet a truly integrated, scalable and interconnected service.

Network equipment supporting OAM features is being massively deployed in networks carrying Ethernet services due to the essential functions it provides in those networks. As such, it is imperative to ensure that this equipment is properly configured and delivering on the promised features and functionality. Testing OAM services prior to deployment will therefore help service providers and carriers save time and money.

How to Test Ethernet OAM

The test methodology or function used will depend on the OAM feature to be validated and on the network architecture. Depending on the OAM standard used, the feature set and functionalities will vary between few and many. This section highlights the main OAM functionalities that must be tested. The different tests can be grouped into two main categories: fault management and performance monitoring. The service lifecycle has three parts: Provisioning and turn-up (tested and validated using service activation tests, such as ITU-T Y.1564), performance monitoring (validated using an OAM test), fault management (tested using an OAM test). In order to test and validate services and networks carrying these services, it is necessary to not only test the OAM functionalities of the network elements, but also the network and services themselves during provisioning and service turn-up. Therefore, only a test combining ITU-T Y.1564 and SOAM can achieve these objectives.

- **Fault Management:** Depending on the OAM standard being used, several types of tests can be performed to allow the detection, verification, location and notification of various defect conditions.
- **Continuity Check:** This function is used to confirm the presence of an MEP or test equipment on a network, and to verify the presence of peer MEPs. During this test, transmitted frames are received either by a peer MEP using a unicast destination address, or by all the MEPs in the MEG using a multicast destination address.
- **Loopback Test:** During this test, the tester transmits an LBM payload with a specific sequence number. The peer MEP responds to the LBM with an LBR payload. The test validates each received LBR and reports any invalid LBR, invalid payload and LBR timeout.
- **Link Trace Test:** This test verifies that the complete link trace reaches the peer MEP. The tester will send LTM messages and receive LTR messages.
- **Test Function:** During this phase, the tester generates frames with a specific test pattern and sequence number in order to verify the integrity of the signal received by the peer MEP. This test requires two testers, one at each end.
- **RDI Test:** During this phase, the tester generates a remote defect indicator (RDI) to simulate a defect and validate the reaction and behavior of the peer MEP.

- **Lock Signal Test:** This test is used to generate and detect locked signals. When a lock frame is received, the tester sounds an alarm.
- **CSF Test:** During this phase, the tester generates and detects a client signal fail (CSF) in order to validate the reception and behavior of the peer MEP.




Figure 130. Test configuration for client-signal fail phase.

- **Performance Monitoring:** Performance monitoring is used to measure parameters such as frame delay, frame loss and synthetic loss.
- **Frame Delay Test:** In this test, which is a critical OAM metric, the tester measures the roundtrip delay to the peer MEP. In order to simulate real-world conditions, different frame sizes should be used to validate the frame delay.
- **Frame Loss:** This test checks the bidirectional frame loss occurring with a peer MEP from a single endpoint, and should be done over as long a period of time as possible in order to obtain a good idea of how the network will behave. To simulate real-world conditions, different frame sizes should be used to validate the frame loss.
- **Synthetic Loss:** This test uses synthetic frames to check the bidirectional frame loss occurring with a peer MEP from a single endpoint. This test should also be done using different frame sizes.

Conclusion

As multiservice networks get more complex, new technologies will necessarily emerge. Even though the concept of operating, administrating and maintaining networks has existed since the early days of synchronous optical networks and synchronous digital hierarchy (SONET/SDH), Ethernet OAM is continually evolving, with new standards being developed as stakeholders are just beginning to understand existing ones.

OAM plays an important role in networks, and as such, these networks must be properly configured to support all the features it offers. However, this can only be achieved by thoroughly testing all the applicable parameters mentioned in this document. From a practical standpoint, because the latest technology is not always fully understood by all users, the most efficient approach available to network operators is to use testers equipped with all the required OAM features and metrics.

3.11. POTN and MPLS-TP Test

The OTN transport network is experiencing unprecedented change. In terms of service requirements, it is necessary to inherit traditional superior performance based on circuit switching technology while ensuring that the packet service remains efficient and flexible. In terms of functional features, it will have a fusion of packet and light, i.e., layer-2 or even layer-3 aggregation and switching features on the OTN network. Currently, OTNs supporting packet service are named differently, e.g., POTS, EOTN, MS-OTN, PEOTN or POTN. Although these systems have different technologies and support different packet services, the current general consensus in the industry is to synchronize the layer-2 feature to OTN devices, among which the POTN network technically featured in PWE3 and MPLS-TP will be favored by the industry.

POTN is built to the OTN G.709 standard based on ITU-T and is also featured by the packet network. The conventional G.709 is based on circuit-switching technology, which is advantageous in its ability to guarantee zero packet loss, low delay, small packet delay variation and superior synchronous clock performance. The conventional packet network improves transport efficiency through statistical multiplexing and features a flexible service, but it cannot offer high clock performance as per the OTN network, e.g., high-quality transport. This is where the POTN network comes into play: a single device integrates circuit-switching characteristics inherent to the OTN network while incorporating the highly-efficient statistical multiplexing feature of packet switching.

POTN networks adopt MPLS-TP technology, which is a simplified version of MPLS for transport networks from which some of the MPLS functions have been removed, including penultimate hop popping (PHP), label-switched path (LSP) merge, and equal cost multipath (ECMP). MPLS-TP does not require MPLS control-plane capabilities, and enables the management plane to set up LSPs manually.

The functions of OAM for MPLS-TP networks are intended to reduce the operational complexity associated with network performance monitoring and management, fault management and protection switching. One of the goals of MPLS-TP OAM is to provide the tools needed to monitor and manage the network with the same attributes offered by legacy transport technologies.

Two important components of the OAM mechanisms are generic associated channel (G-Ach) and generic alert label (GAL), which allow an operator to send any type of control traffic into a pseudowire (PW) or label-switched path (LSP). G-Ach is used in both PWs and MPLS-TP LSPs, whereas GAL is used in MPLS-TP LSPs to flag G-Ach.

The OAM functionality of MPLS-TP used in fault management consists primarily of two kinds of management, namely proactive FM OAM and on-demand FM OAM functions.

Proactive FM OAM includes the following functionalities:

- **Continuity check (CC) and connectivity verification (CV):** In active operating mode, this function sends CC/CV messages at the source-end maintenance association endpoint periodically, and checks the loss-of-clock (LOC) fault at the host-end maintenance association endpoint and along with connectivity faults, such as mismerge and misconnection. The minimum device-supported transmission period is 3.3 ms in order to guarantee that fault detection can be done in as little as 10 ms.
- **Remote defect indication (RDI):** Used to notify the opposite end of the fault information detected locally.
- **Alarm indication signal (AIS):** After detecting any fault, the local service sublayer inserts this alarm into the client layer upstream and sends it to the downstream maintenance end. This alarm is mainly used to suppress the secondary alarms of the client layer in order to avoid a lot of redundant alarms.
- **Link down indication (LDI):** An indication information after an AIS alarm occurs.
- **Lock report (LKR):** The message the source end sends to the host end after service disruption. It is primarily used to suppress unnecessary redundant alarms.

On-demand FM OAM consists of the following functionalities:

- **Connectivity verification (CV), route trace (RT), transport plane loopback, lock indication (LI)**
- **RT** and transport-plane loopback messages allow all MIPs and MEPs to return corresponding OAM messages after receiving the message, and can be used to maintain bidirectional connectivity between the MEP and MIP so as to detect faults between and within nodes and locate the faults.

OAM functions for performance management basically include the following:

Packet loss measurement (PLM), packet delay measurement (PDM), throughput measurement and delay variation measurement. Among them, the packet loss measurement can be implemented with continuity check and connectivity verification messages; the throughput measurement is carried out with packet loss measurement and can be conducted in both in-service and out-of-service modes; the packet delay measurement is implemented with exclusive messages by sending and receiving timestamps and is able to measure unidirectional or bidirectional time domains; the delay variation measurement can be obtained by performing a packet delay test.

Typical test method: EXFO's FTB-890/890NGE, FTB-88100NGE or FTBx-88200NGE can be used for OAM verification and testing of the MPLS-TP network. The instrument can be used separately for testing of the MEP or MIP point, or used for the end-to-end cross test, as shown in the figure below:




Figure 131. Test setup for quality-of-service verification.

3.12. POTN Packet/Ethernet Service Quality of Service Verification

The reality of today's network is the presence of different types of traffic, each with different performance requirements requiring QoS mechanisms in order to provide segregation and prioritization of the most sensitive traffic class.

Traditional TDM networks were devoid of such parameters and only provided guaranteed throughput. Inefficient and complex, they focused only on plain transmission of data. There was no recognition of traffic type, and all traffic was serviced with the same level of service. With the deployment of packet-based networks, new applications must now be serviced by the networks; however, these different applications have different performance requirements. Voice over TDM makes very different assumptions than over-the-top video, broadcast TV, the Web or e-mail. To allow different applications to share a single backhaul circuit, the circuit must support more than one class of service (CoS), each with its own QoS requirements. Packet-based technologies, on the other hand, provide differentiation mechanisms that allow the network to prioritize certain transmissions over others and provide some transmission guarantees.

POTN retains the IP/MPLS QoS systems such as IP differentiated services (DiffServ) and MPLS labeling, allowing the network to carry voice, video and data packets. The network is then able to classify packets according to their priority, and service them to ensure that they meet their performance requirements.

The Ethernet QoS test in the POTN network includes the following items: service flow classification and priority mapping ability test; flow-classification-based access control list (ACL) capability test, access rate control test; access rate control granularity, color sensitivity functional verification, connection admission control (CAC) scheme, and client service priority remapping ability.




Figure 132. Test setup for quality-of-service verification.

Test method: A typical network configuration is shown above. Create an LSP and a PW between NE1 and NE3. Multiplex PW into LSP. Create flow classification criteria, map them into PW and then verify the following flow classification criteria:

- | | |
|--------------------------------------|----------------------------------|
| a. Ports | f. DiffServ code point (DSCP) |
| b. VLAN | g. IP TOS |
| c. VLAN PRI | h. TCP/IP port number (optional) |
| d. Source or destination MAC address | i. Any combination of the above |
| e. Source or destination IP address | |

The flow creation interface of EXFO's FTB-8880/88100NGE and FTBx-88200NGE is as follows:




Figure 133. EXFO's flow creation interface.

3.13. Clock Performance Test for Packet-Switching-Based POTN Device in Response to ODUK/ODUflex Client Signal

In the POTN network, both frequency and time synchronization must be supported. The frequency synchronization should support the frequency pass-through mode of the client signal and network clock synchronization; the time synchronization shall support optical supervisory channel (OSC) out-of-band mode or ODUK in-band mode. For network elements based on oldest-packet-first (OPF) packet switching, OIF-OPF-01-0 documentation requires that the packet-switching-based device be capable of keeping the clock information of the ODUK/ODUflex signal, as follows: the maximum delay for packet switching must be less than 100 μ s; the delay variation must be less than 50 μ s; a scheme shall be used to compensate for the delay arising from packet switching; this preset value and the maximum delay follow a relationship that the maximum switching delay $\leq$ preset delay value $\leq$ 100 μ s.

EXFO's SychWatch-110 is a tester for frequency and time synchronization. The SyncWatch-110 test unit has a precise reference based on its high-precision rubidium standard clock or built-in GPS receiver. The accuracy of its main standard clock exceeds the primary-reference-clock (PRC) performance specified in G.811. Each standard clock can provide precise retention functionality when the GPS signal is interrupted, as well as measurement stability in the field test.



Figure 134. POTN equipment and network clock performance test.

EXFO's SyncWatch-110 device can also employ any external standard provided by the client, thereby making full use of the existing synchronizing infrastructure and ensuring that the measurement results adopt the same source as the local standard. The SyncWatch-110 unit provides real-time time-interval-error (TIE) measurement, and enables calculation and display of maximum-time-interval-error (MTIE) and time deviation (TDEV) measurements at any point.

Test method and connection diagram: The test mainly includes parameters relating to the synchronization network, SyncE and IEEE 1588v2, among which the most important parameters are MTIE, TIE and TDEV for SyncE, and unidirectional packet loss, delay and delay variation performance for 1588v2 clock distribution.

EXFO's SyncWatch-110 can be connected to a single high-precision reference clock (e.g., rubidium interface), following which one or two input ports of the unit can be used to run typical device benchmark tests. In addition, the SyncWatch-110 unit is able to test performance at both input ports simultaneously, and run accurate benchmark tests on two devices according to the benchmark. The test connection is shown in the adjacent diagram on the right-hand side:

Test result: Figure 135 below shows the TIE test result, and Figure 136 shows the MTIE test result.





Figure 135. TIE test result.




Figure 136. MTIE test result.

For the IEEE 1588v2 PTP test, the main concerns are delay asymmetry and delay variation. The images below represent typical examples of delay symmetry. In the graph on the left-hand side, the master-to-slave delay is much higher than the slave-to-master delay, and therefore the result is a fail. In the middle graph, one-way delay is seriously degraded, which also leads to a fail result. Ideal results are obtained when delays for each direction are symmetric.



The graph on the far left is an example of one-way delay variation; the delay variation in this example is too high, leading to a fail result.

3.14. Multiservice Application Characterization in OTN and POTN Networks

To be viable and long-term investments, OTNs or POTNs must be able to accommodate different types of traffic. TDM circuits are still widely used as legacy connections. These circuits continue to provide effective bandwidth and guaranteed services in many applications, and their widespread deployment makes it impossible to replace them on a short-term basis. Ethernet and IP-based packet networks now constitute the technology of choice for next-generation networks due to their very low cost per bit, ease of deployment, maintenance and wide acceptance.




Figure 137. OTN and POTN environment.

The POTN provides the capability to handle both types of traffic through the implementation of pseudo wire edge-to-edge emulation (PWE3) forwarding. This encapsulation method has been specifically designed to segment and encapsulate TDM and packetized transmission onto packet networks, allowing both types of traffic to be transparently transported on the POTN network, while ensuring independence between the transport encapsulation and the encapsulated data.

The FTBx-88200NGE's multiport testing function is able to meet network maintenance requirements, as shown below:




Figure 138. FTB-2 Pro platform and FTBx-88200NGE module.

3.15. Ping and Traceroute from 40GE/100GE Interface

Ping and traceroute can be used to send packets of information to remote units for the purpose of retrieving information. Ping can test the speed of the connection, the “distance” to target, and whether or not the connection is even up and running. Traceroute tracks the path that a packet takes from the source unit to a destination address. The examples below feature EXFO’s GUI.

☐ Automatic IP (DHCP)

Source IP Address

10.10.122.11

Destination IP Address

10.10.122.11

Quick Ping

☒ Resolve MAC Address

Resolved

Subnet Mask

255.255.0.0

☐ Default Gateway

TTL

128

IP TOS/DS

0x00

TOS/DS Config

Ping & Trace Route

40/100G Advanced

Source IP Address

10.10.0.0

Destination IP Address

0.0.0.0

☐ Use Stream

Ping

Data Size (Bytes)

32

Timeout (ms)

4000

Ping

TTL

128

Delay (ms)

1000

IP TOS/DS

0x00

Attempts

n-Attempt

4

Trace Route

Max Hop Count

128

Timeout (ms)

4000

Trace Route

Results

No.	Status	Reply Details

Statistics

Packets Transmitted

--

Packets Received

--

Percentage Lost (%)

--

Min Round Trip Time (ms)

--

Max Round Trip Time (ms)

--

Avg. Round Trip Time (ms)

--

Figure 139. Ping and traceroute.

4. Acronyms—Transport and Datacom

ADC	Analog-to-digital converters
APS	Automatic protection switching
BDI	Backward defect indicator
BEI	Backward error indicator
BER	Bit error rate
BIAE	Backward incoming alignment error
BIP-8	Bit-interleaved parity-8
CAUI	100 Gbit/s attachment unit interface
CD	Chromatic dispersion
CFP	C form-factor pluggable
CGMII	100 Gbit/s media independent interface
CPRI	Common public radio interface
DAPI	Destination access point identifier
DPSK	Differential phase-shift keying
DQPSK	Differential quadrature phase-shift keying
DSP	Digital signal processor
EVM	Error vector magnitude
EXP	Experimental
FAS	Frame alignment signal
FC	Fibre Channel
FCS	Frame check sequence
FEC	Forward error correction
FTFL	Fault type and fault location
GMP	Generic mapping procedure
GPON	Gigabit-capable passive optical network
IAE	Incoming alignment error
IEC	Electrotechnical Commission
IP	Internet protocol
JC	Justification control
LAN	Local-area network
LLC	Logical link control
LWDM	LAN wave division multiplexing
MAC	Media access control
MDI	Media-dependent interface
MDIO	Management data input/output
MFAS	Multiframe alignment signal
MSI	Multiple structure identifier
MSIM	Multiple structure identifier mismatch
NRZ	Non-return-to-zero
ODTU	Optical channel data tributary unit
ODTUG	Optical channel data tributary unit group
ODU	Optical channel data unit

OH	Overhead
OIF	Optical Interworking Forum
OPU	Optical channel payload unit
OSNR	Optical signal-to-noise ratio
OTL	Optical channel transport lane
OTN	Optical transport network
OTU	Optical channel transport unit
PCC	Protection communication channel
PCS	Physical coding sublayer
PM	Path monitoring
PMA	Physical media attachment
PMD	Polarization mode dispersion
PMOH	Performance monitoring overhead
POTN	Passive optical transport network
PRBS	Pseudo random bit sequence
PSI	Payload structure identifier
PT	Payload type
QPSK	Quadrature phase-shift keying
RES	Reserved
RF	Radio frequency
RMS	Root mean square
ROADM	Reconfigurable add-drop multiplexer
RS	Reconciliation sublayer
RZ	Return-to-zero
SAPI	Source access point identifier
SDH	Synchronous digital hierarchy
SFD	Start of frame delimiter
SNC	Subnetwork connection
SNR	Signal-to-noise ratio
SOAM	Service operations, administration, and maintenance
SOP	State of polarization
SP	Skew point
STAT	Status
TC	Tandem connection
TCM	Tandem connection monitoring
TCM ACT	Tandem connection monitoring activation
TS	Tributary slot
TTI	Trail trace identifier
UI	Unit interval
VLAN	Virtual local-area network
GCC	General communication channel

No part of this guide may be reproduced in any form
or by any means without the prior written permission of EXFO.

Printed and bound in Canada

ISBN 978-1-55342-111-5

Legal Deposit—National Library of Canada 2017

Legal Deposit—National Library of Quebec 2017

For details about our products and services, or to download technical and application notes, visit our website at www.EXFO.com.