

Validating optics from 100M to 400G:

Quickly and accurately evaluate
the health of optical devices

app
note

EXFO

Validating optics from 100M to 400G:

Quickly and accurately evaluate the health of optical devices



app
note

By [Marcos Vasconcelos](#),
Product Line Manager
T&D, EXFO

Table of contents

Introduction	2
When and where iOptics should be used	3
User scenarios	3
Service providers (FTTA)	3
Data centers	4
NEMs	4
What does iOptics validate?	4
iOptics at a glance	6
Conclusion	7

Introduction

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 optical interfaces used in network devices. These pluggables vary based on the data rates that they transmit and receive. Optical transceivers have evolved from SFP to QSFP-DD. The following table provides a description for 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
SFP28	Enhanced small form-factor pluggable	Up to 28 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
CFP8	C form-factor pluggable	400 Gbit/s
QSFP56-DD	Quad small form-factor pluggable - Double density	400 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. For additional flexibility, Data centers often use 100G AOC cables to connect interfaces on the same rack or between different racks. On the other hand, due to the emerging 400G technology, NEMs and carrier labs are using 400G DAC cables to interconnect different network elements.

The cost of devices operating at 100 Gbit/s and 400 Gbit/s is higher. 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 CFP8 and QSFP56-DD, also known as QSFP-DD) 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, webscale companies and network equipment manufacturers (NEMs) with its launch of the iOptics - intelligent pluggable optics test application designed to be the easiest method in the industry to quickly validate optical transceivers using minimal configuration. The tested aspects and functionalities are adapted to the type of pluggable 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 that should 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.

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, SFP+ or SFP28 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

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

iOptics 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. As an example, in the 400G ecosystem, DAC cables are often used because of their lower price and easier availability. NEMs must be assured that these cables operate as expected. Validating them becomes a paramount step in the development of new network elements.

What does iOptics validate?

iOptics 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, iOptics performs a basic module validation by collecting key device information such as vendor, module type, module ID, part number and speed (see Figure 1). 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.

QSFP Parameters	
Module ID	QSFP28
Vendor Name	FINISAR CORP
Part Number	FTLC9551SEPM
Serial Number	XV20RCL
Hardware Revision	A0
Revision Compliance	7
Connector Type	MPO
Speed	103.125 Gbit/s, 111.8 Gbit/s
Type/Compliance Code	100GE-SR4-100m
Mode	MMF (OM3 + OM4)
Power Class	Power Class 4 Module (<= 3.5W)
CLEI Code	

Figure 1. 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

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

Temperature monitoring 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. Stress test

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, CFP8, QSFP+, QSFP28 and QSFP-DD optical devices. A physical loopback is required for this test, which measures the skew associated to each physical coding sublayer (PCS) lane. iOptics 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 iOptics 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.

iOptics at a glance

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

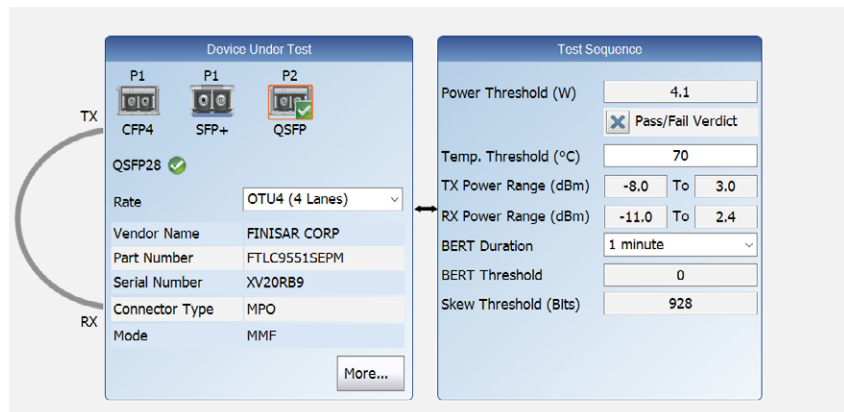


Figure 2. The iOptics test configuration page

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.

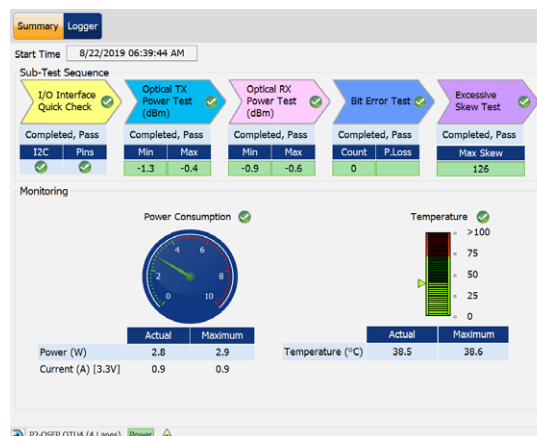


Figure 3. Results Summary page

Conclusion

In summary, it is paramount that all optics and cables operating from 100M to 400G be validated, regardless of the scenario in which they are used. High speed optics remain expensive, and companies cannot afford to waste time and resources through service outages and costly visits to faulty sites. iOptics offers a multitude of benefits to NEMs, service providers and web-scale companies, including:

- Accurate evaluation of the health of optical devices from 100M to 400G
- Rapid analysis of the optical interfaces, with quick pass/ fail verdicts
- Support of high-speed optics (CFP, CFP2, CFP4, CFP8, QSFP+, QSFP28, QSFP-DD), as well as pluggables operating at lower rates (SFP, SFP+, XFP and SFP28)
- Support of 100G AOC and 400G DAC