White paper
Cisco public

The bridge to possible

Growing the Network with 400 Gbps Coherent Pluggable Optics

Interoperability and operational considerations when deploying Digital Coherent Optics

Contents

Executive Summary

The latest generation of Digital Coherent Optics (DCO) pluggable transceivers represents a breakthrough in the optical networking industry. By combining advances in silicon photonics and Digital Signal Processors (DSP) with Quad Small Form-factor Pluggable – Double Density (QSFP-DD) form factor, DCO pluggable transceivers provide a very cost effective, compact, and power efficient option to Dense Wavelength Division Multiplexing (DWDM) transponders at 400Gbps speeds both today and well into the future. Additionally, when combined with other industry innovations in packet processing silicon, networking software, systems, and novel simplified network architectures, DCOs have been shown to provide up to 46 percent savings in an IP transport network's Total Cost of Ownership (TCO) [1]. In other words, DCO technology is redefining the entire networking industry.

Despite the clear benefits of cost reduction, operational simplicity and services agility that are enabled especially by the introduction of 400Gbps DCO over QSFP-DD, service providers and other consumers of this technology have questioned its feasibility for deploying within brownfield DWDM networks.

The arguments against brownfield use cases are based on differences in the technical specifications of the current generation 400Gbps QSFP-DD DCO-based DWDM optics versus those provided by transponders, with the lower optical transmit power (e.g., -10dBm from current QSFP-DD DCO versus 0 to 3dBm from transponders), having the clearest impact in practical designs. Management of the pluggable transceiver has also been considered a major challenge as it's expected that those components will be deployed in host equipment that sits outside of the DWDM system, for instance in routers and switches.

Conversely, there are also many arguments that support the feasibility of 400Gbps QSFP-DD DCO over brownfield DWDM networks. The current generation of 400Gbps QSFP-DD DCO pluggable transceivers is based on widely adopted as well as emerging DWDM standards. Where standards don't yet exist, major industry players have aligned on open systems specifications. Current differences in the technical specifications are within a range that can be compensated by existing DWDM systems by using standard technologies when required. For instance, the difference in the optical transmit power can be simply compensated by standard Erbium Doped Fiber Amplifier (EDFA) optical amplifiers in a variety of configurations. Similarly, management can be addressed by embracing industry-defined, open APIs, data models, and data model-driven software architectures. Hierarchical SDN controllers are also available and can be an integration point for end-to-end network visualization and control.

The conclusion is that while there are differences between current 400Gbps QSFP-DD DCO pluggable transceivers and traditional DWDM transponders' technical specifications, practical, vendor-agnostic solutions exist and interoperable, brownfield, multi-vendor systems are feasible with a variety of brownfield DWDM systems.

This paper compares those technical specifications, analyzes the DWDM photonic layer they connect to, and describes key considerations for interoperability and practical feasibility in existing network designs. It also describes potential solutions for hardware and management interoperability based on existing and emerging standards for DCO deployment over brownfield DWDM networks.

Introduction

400Gbps DCOs are a breakthrough technology that when combined with QSFP-DD transceivers are set to redefine the optical industry. Leveraging state-of-the-art silicon technologies, including silicon photonics, DCO can replace large, costly, and power inefficient DWDM transponders by implementing the same functions in a very compact, low power pluggable form factor that requires no additional space as it can be deployed in a standard 400GE port on host devices that support it, like routers and switches. The result is a solution that's much more cost effective when compared to traditional DWDM transponders. In addition, 400Gbps DCO pluggable transceivers are based on industry standard specifications and are implemented in a variety of form factors beyond QSFP-DD, namely CFP2 and OSFP, allowing equipment vendors to deploy across varying hardware platforms.

400Gbps has become a transition point with DCO in the transport realm. While traditional 400Gbps DWDM transponders are generally closed and proprietary, 400Gbps DCO was standardized by the Open Internetworking Forum (OIF) as 400G ZR while OpenROADM and ITU-T have embraced the same technologies at 400Gbps to define competing standards. Multi-source Agreements (MSAs) are also defined to enable extended reaches of beyond 1100 km, like the OpenZR+ (Open ZR "Plus") MSA specification.

Standardization is also taking place across the optical transport industry with open line systems enabling third-party DWDM optics carried over existing infrastructures to interoperable management and automation interfaces.

From a network architecture perspective, 400Gbps DCO QSFP-DD technology is at the heart of Cisco's Routed Optical Networking solution which combines this technology with innovations in routing silicon and software to provide the simplicity, scale, and cost structure necessary to create a more inclusive internet for decades to come.

Considering all these benefits, the networking industry has embraced 400Gbps DCO technologies, while major packet switching and optical equipment vendors either support or plan to support it in their products. In this context, to drive market adoption and allow end customers to immediately collect benefits, it's important to understand these technologies in more details to develop different insertion strategies, especially for brownfield DWDM networks where careful design and validation work may be required to validate or enable their feasibility.

What challenges need to be addressed?

Despite the clear benefits in network simplification and related reduction in TCO and the standardization work behind 400Gbps DCO pluggable transceivers, service providers and other potential consumers have been concerned with the introduction of this technology in existing, brownfield DWDM networks, potentially in multi-vendor scenarios, given the very large DWDM installed base.

These concerns are based on observations about technical differences, operational aspects, and previous efforts to integrate IP and optical networks, as well as push back from incumbent DWDM vendors who lack the motivation to support third-party interoperability, as this could break their existing commercial models.

On the technology side, 400Gbps QSFP-DD DCO pluggable transceivers are relatively new and while they've been standardized, their technical specifications haven't been considered in the design of most brownfield DWDM networks. In other words, while 400Gbps DCO interoperability with DWDM networks is backed up by standards at the component level, interoperability at the system level requires more engineering work and this is the first challenge we need to address:

Challenge 1: Validate the feasibility of 400Gbps QSFP-DD DCO pluggable transceivers in brownfield DWDM networks and define the additional hardware, if any, required to achieve it.

In addition to the interoperability at the hardware level, operational aspects must also be addressed. Optical transport networks in general are known for their dependence on graphical software tools for Fault, Configuration, Accounting, Performance, and Security (FCAPS) functions. Optical networks are also known for using closed, proprietary software architectures for those functions, with each transport vendor providing management tools specific to their technology.

The networking industry has worked hard to address the proprietary nature of optical transport management. In today's context, the focus is on automation systems by developing open software architectures supported by industry standards which provide APIs and data models for building multi-vendor systems. Once again, these standards are relatively new and while the industry adopts them, flexibility is key to co-exist with brownfield systems. For that, different solutions exist to aggregate management information from disparate systems and allow for end-to-end network visualization and services configuration.

In this context, as 400Gbps QSFP-DD DCO pluggable transceivers are deployed as external components, potentially from a different vendor to the DWDM network, the second challenge we need to address is:

Challenge 2: Managing a heterogeneous, end-to-end network consisting of 400Gbps QSFP-DD DCO pluggable transceivers hosted in external equipment, DWDM transport network, and existing management (potentially automation) tools.

In a multi-vendor environment, the end-to-end network can take multiple forms and combinations. A non-exhaustive list is described below and captured in Figure 1, where each color represents a hypothetical technology vendor:

Scenario A:

- Source and terminating host networking equipment and DCO pluggables from vendor X

- DWDM network infrastructure or line system from vendor Y

Scenario B:

- Source host networking equipment and DCO pluggables from vendor X

- Terminating host networking equipment and DCO pluggables from vendor Y

- DWDM network infrastructure or line system from vendor Z

Scenario C:

- Source and terminating host networking equipment from vendor Y

- DCO pluggables from vendor X

- DWDM network infrastructure or line system from vendor Z

Scenario D:

- Source and terminating host networking equipment from vendors X and Y

- DCO pluggables from vendor Z

- DWDM network infrastructure or line system from a vendor X

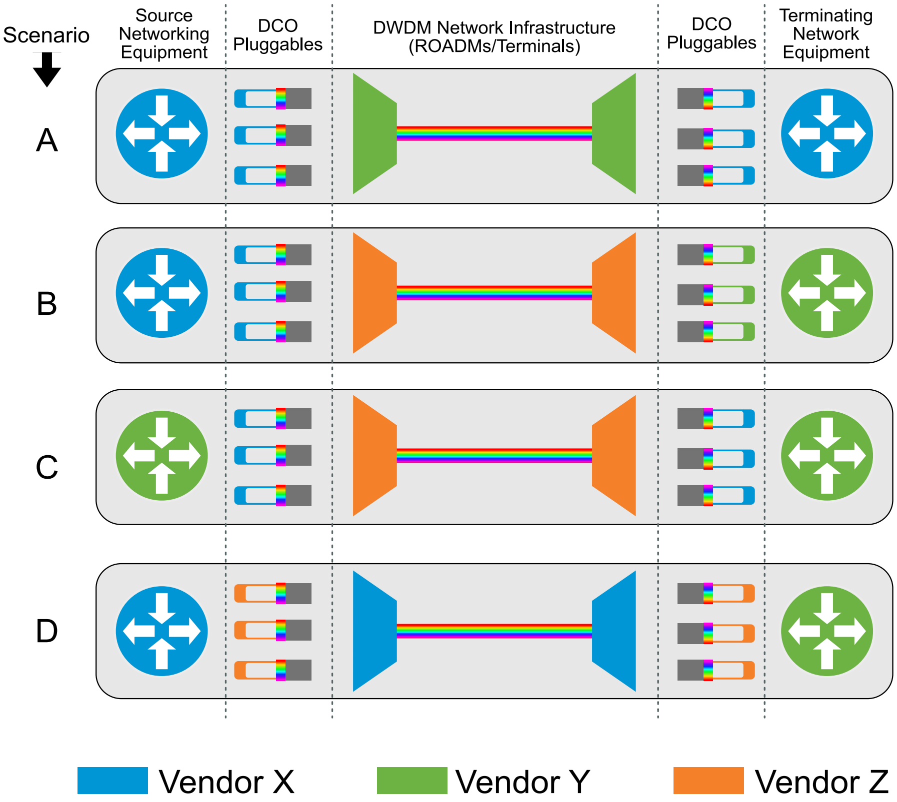

Figure 1. Host equipment, DCO and DWDM systems - interoperability scenarios (not exhaustive)

The next challenge we need to address is not related to technology. It results from the fact that over the years, DWDM technology vendors have built business models that resemble the traditional “razor and razor blade” model, where the common components of the DWDM network, known as the photonic layer, or in some cases the Reconfigurable Add/Drop Multiplexer (ROADM) layer, represents the razor and the DWDM transponders represent the razor blades.

In this model, DWDM vendors count on the transponder revenues to close their business cases and drive profit. As DCO pluggable transceivers represent a compelling, lower cost alternative to transponders that can be sourced in the market, it breaks this vendor lock-in model. Traditional DWDM vendors, on the other side, will resist to support DCO technology even though technically they may be entirely feasible in their networks.

Another non-technology related challenge is how to couple this open and more integrated network architecture principle supported by DCO to a traditionally siloed organizational model that’s pervasive in large service providers where optical and packet transport teams are separated. In this respect, the industry needs to provide practical solutions that respect operational boundaries while companies evolve to be able to realize the financial benefits of the integrated model.

Challenge 3: Counter non-technical challenges including organizational boundaries incumbent to DWDM vendor push backs by understanding the technology details behind 400Gbps QSFP-DD DCO transceivers and vendor-neutral solutions to achieve full interoperability.

This paper provides information necessary to address these problems from a technology and network engineering perspective and provide enough details to remove some of the roadblocks that can appear in discussions between vendors and network operators.

At the end of the day, the optical transport industry is set to be redefined by the introduction of 400Gbps QSFP-DD DCO technology with no way back and everyone will benefit from finding practical solutions to these problems sooner rather than later.

Anatomy of 400ZR/OpenZR+ Coherent Pluggable Optics

Optical interfaces for DWDM applications are defined based on some key attributes that equally apply to transponders and 400Gbps QSFP-DD DCO pluggable interfaces. For the latter, those attributes are specified by standards bodies and industry organizations including the OIF 400ZR specification and OpenZR+ Multi-Source Agreement (MSA). These are the main attributes of optical interfaces in DWDM systems:

- Optical transmit power: Defined in dBm (decibels in relation to 1mW), with minimum, maximum, and typical values. When engineering networks, we use the minimum value as a worst-case condition for performance while the maximum value is used to evaluate possible signal interference when these signals are multiplexed and amplified. Industry standard software-configurable Variable Optical Attenuators (VOAs) or fixed optical attenuators are commonly used before DWDM channels are multiplexed and amplified to avoid major discrepancies between them and to compensate for the non-linear gain curve of EDFAs, a process called “pre-emphasis”.

- Optical receive power: Also defined in dBm, with minimum and maximum values. When engineering networks, we use the minimum value and add system penalties to define the worst-case condition and the minimum amount of power required to maintain “error free” transmission, or in other words, to maintain the specified system performance in Bit Error Rate (BER).

- Optical channel spectral width: Defined in GHz, the optical channel spectral width is related to the interface’s baud (e.g., the symbol rate used to carry information into the physical medium). Using different modulation schemes, a 400 Gbps interface can have different baud rates. The OpenZR+ MSA specification version 1.0 defines the spectral width @400Gbps as 60.14 GHz while the OIF 400ZR defines a signal with 59.84 GHz width. The actual spectrum window required by the optical signal will be higher than those values as it must account for the signal shape and a guard band (e.g., a space between the optical channels to minimize interference between them). Both specifications define 75 GHz as requirement for 400Gbps channels.

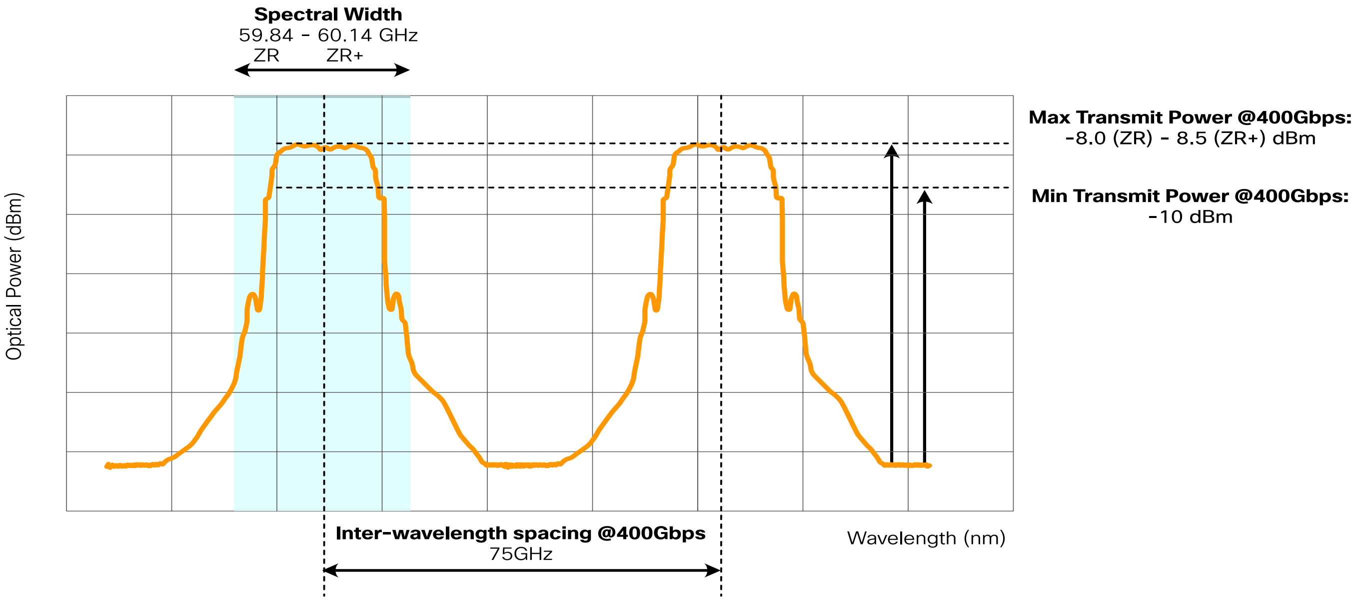

Figure 2. Anatomy of 400ZR and OpenZR+ optical signal

- Modulation scheme: Different modulations can be used to produce an optical signal and modern systems allow operators to change the modulation through software configuration. The modulation scheme will be responsible for defining the symbol rate or baud rate and the number of bits carried per symbol, which in turn defines the transmission rate and the interface performance. Examples of optical modulation schemes are DP-QPSK and DP-8QAM used by the 400Gbps OpenZR+ MSA specification and DP-16QAM used by both OIF 400ZR and OpenZR+ MSA.

- Data encapsulation: At the electrical domain, optical interfaces use different encapsulations or digital wrappers to add OAM tools on top of the payload traffic. For high performance optical interfaces, the encapsulation also includes tail-end Forward Error Correction (FEC) information that allows the receiver to identify and recover transmission errors. As a result, the optical signal is more resilient to noise produced by the optical line system amplifiers (e.g., it's optical signal to noise tolerance is improved and the system performance and reach are boosted). This additional budget for optical noise is called Net Coding Gain (NCG). Examples of digital wrappers include ITU-T G.709 for traditional transponders and FlexO for DCO pluggable, while FEC methods include staircase "hard decision" FEC, "soft decision" FEC, and concatenated FEC (C-FEC) that mixes both soft and hard decision FEC. OIF 400ZR uses C-FEC while OpenZR+ uses Open FEC (O-FEC), which is a block code-based encoder and iterative soft decision-based decoder introduced by the OpenROADM initiative.

- ITU-T frequency or wavelength: Also referred to as the optical signal "color" or "lambda", this is the central frequency in THz or the wavelength in nanometers used by the DWDM channel. In coherent optical systems it also defines the frequency at which the interface processes the received signal. The frequencies are defined as per ITU-T G.694.1 Standard, which specifies the central frequencies for optical signals for different channel spacing or grids (e.g., 50GHz, 100GHz, or 200GHz). The initial definition was created when DWDM systems supported only optical channels with fixed spectral width. Modern DWDM systems based on ROADM with gridless technology, also known as Flex-Spectrum, are no longer restricted to fixed spectral widths. Gridless systems allow DWDM channels with variable spectral width by assigning them slices of spectrum with configurable sizes based on their requirements. Understanding the channel center frequency, its spectral width, and the channel spacing requirements is very important to define the

Compliant Standard		**OIF 400G ZR [2]**	**400G OpenZR+ MSA [3]**
Parameter	**Unit**	**Value**	**Value**
Possible client rates	-	*Nx100GE, 400GE*	*nx100GE, 400GE*
Rate per wavelength	Gbps	*400*	*100/200/300/400*
Client mapping	-	*OIF ZR Framing*	*OIF ZR Framing*
Adaptive line rate modulation	Gbps	*NA*	*100G/200G/300G/400G*
Number of Carriers	-	*Single*	*Single*
Modulation of wavelength	-	*DP-16QAM*	*DP-16QAM/DP-8QAM/DP-QPSK*
Baud Rate	Gbaud	*59.84*	*60.14 for 200G/300G/400G and 30.07 for 100G*
Count of wavelength per fiber pair in the C-Band @75GHz spacing	-	*64*	*64*
Spectral excursion	GHz	*-32 / + 32*	*-32 / +32*
Inter-wavelength spacing @400Gbps	GHz	*75*	*75*
Spectral efficiency	bits/s/Hz	*5,33*	*5,33*
Capacity per fiber pair	Tbps	*25.6 Tbps*	*25.6 Tbps*
Guard band	GHz	*NA, Supports 75GHz grid*	*12,5*
Central wavelength accuracy	GHz	*+/-1.8*	*+/-1.8*
FEC coding gain / Overhead	dB / %	*10.80 / approx. 14.8%*	*11.60 / approx. 15.3%*
Transmit Power	dBm	*-6 to -10 dBm*	*-4 to -8 dBm for 100G* *-5 to -9 dBm for 200G* *-6 to - 10 dBm for 300* *-6 to - 10 dBm for 400*
Input power sensitivity	dBm	*-12 dBm*	100G: *-18 dBm* | 200G: *-18 dBm* 300G: *-15 dBm* | 400G: *-12 dBm*
Receiver OSNR tolerance *At FEC threshold. Referenced to an optical bandwidth of 0.1 nm at 193.7 THz or 12.5 GHz.*	dB	*26*	100G: *12.5* | 200G: *16* 300G: *21* | 400G: *24*
Default CD Tolerance	ps/nm	*2.400*	100G: *100000* | 200G: *50000* 300G: *40000* | 400G: *20000*
DGD Max	ps	*33*	100G: *83* | 200G: *66* 300G: *66* | 400G: *50*
PMD Tolerance - Mean	ps	*10*	100G: *30* | 200G: *25* 300G: *25* | 400G: *20*
Pre-FEC BER Threshold	Rate	*1.25E-2*	*2,00E-02*
Post-FEC BER	Rate	*1,00E-15*	*1,00E-15*

Table 1. 400 Gbps DCO initial industry specifications

characteristics for the DWDM line system. Both OIF 400ZR and OpenZR+ MSA specifications define 75 GHz of spectrum for 400Gbps channels, which can be achieved using fixed filters or gridless ROADMs.

- Optical Signal to Noise Ratio (OSNR): Defined as a target in dB units, OSNR figures represent how much noise the receiver tolerates before it can no longer maintain the desired performance or the required signal power for the receiver to separate it from the noise. The sources for optical signal noise are many but the main one is from the EDFA-induced noise, commonly known as Amplifier Spontaneous Emission (ASE). Some DWDM systems are capable of filtering some of the “side” noise (e.g., the noise adjacent to the signal itself as the optical signal crosses filter components, especially in colored DWDM systems).

Certainly, there are other specifications and parameters associated with optical interfaces. However, for the sake of understanding what’s generally required to evaluate interoperability between a given DCO pluggable technology with DWDM network infrastructures, the list above is sufficient.

Table 1 provides a summary of the specifications for OIF 400ZR and OpenZR+ MSA interfaces.

To validate interoperability at the physical layer (e.g., hardware interoperability between DCO and the DWDM system), the network engineering goal is to validate that these optical interface specifications are within the range delivered or supported by the DWDM network infrastructure at the photonic layer (the pure optical layer) composed of ROADMs or other DWDM line systems, for instance an open line system. The engineering challenge is to account for differences in these specifications when they do exist and find solutions to compensate for them when necessary to allow for the interoperability.

Interoperability considerations

Before we talk about DCO interoperability considerations for existing DWDM systems, let’s first compare the optical signal generated by a traditional transponder, 400ZR and OpenZR+. This will help us understand the

Parameter	**Unit**	**400G Transponder [4]**	**OIF [2]** **400ZR**	**MSA [3]** **400G OpenZR+**
Optical Modulation		PM-16QAM	PM-16QAM	PM-16QAM
Baud rate	GBaud	46.3 to 72	59.84	60.14
Transmit Power	dBm	+3 to -10 dBm	-6 to -10 dBm	-6 to -10 dBm
Receiver sensitivity	dBm	-17 dBm	-12 dBm	-12 dBm
Receiver OSNR sensitivity (0.1nm resolution)	dB	22	26	24
Max CD Tolerance @400Gbps	ps/nm	150,000	2,400	20,000
DGD Max	ps	64	33	50

Table 2. Typical optical specifications for 400Gbps DCOs and transponders

main differences that we'll need to address in our search for solutions.

A key difference between transponders and the current generation of 400Gbps DCOs is in the optical transmit power. While transponders have a typical 0 to 3dBm optical transmit power, OIF's 400ZR and the first generation of 400Gbps OpenZR+ MSA compliant transceivers have a -6 to -10dBm as the minimum value for optical launch power. This typical 10dB difference will be the first point for attention while engineering the interoperability with brownfield DWDM networks that were designed for transponders. We'll get back to this item later as we discuss different strategies to compensate for differences in optical launch power.

Another difference is on the target OSNR (e.g., the optical signal-to-noise ratio that we should obtain at the receiver side to guarantee the desired system performance). This difference is around 2.0 dB, a value that can be addressed by adopting novel network architectures with reduced the optical distances that DCO interfaces will have to go through. One example of such novel network architectures is Routed Optical Networking. In this architecture, shorter optical distances are achieved by promoting a shorter, direct connectivity between routers (e.g., router-to-router connectivity, hence relaxing the performance requirements that otherwise would be required when the signal must traverse multiple ROADM nodes to reach thousands of kilometers in the so called "optical bypass" architecture).

For longer distances in ultra-long haul and submarine DWDM networks, this difference can be a deal breaker given the fact that such systems are designed to operate at the edge of the system specifications. For such applications, high performance transponders will continue to be used.

If the DWDM system can't provide the 75GHz or wider spectrum for the 400Gbps channels, the interface speed will need to be throttled down to fit in the available spectrum. The OpenZR+ DCO specification allows that configuration supporting from 100Gbps to 400Gbps in 100Gbps steps. For example, some older DWDM networks have been designed for 50GHz spaced channels with a fixed grid, which was fine for the 100Gbps and below speeds available at that time. This applied to both DWDM terminals as well as ROADMs. For these networks, the network engineer should pay special attention to the 400Gbps DCO 75GHz spectrum requirement.

Now, let's focus the discussion on how to compensate for the lower optical transmit power from DCO transceivers in applications that require the optical channels to be balanced like in long distance DWDM networks. Multiple strategies can be employed to address this engineering requirement and the right choice will depend on the ROADM or terminal DWDM filters used and how the network operator plans to address the multi-vendor scenario.

Let's review five solutions or methods to deploy DCOs in brownfield networks:

1. Using external DWDM filters and EDFAs
2. Using internal DWDM filters and EDFAs
3. Using ROADM built-in channel equalization capabilities
4. Dedicating a ROADM degree to DCO channels
5. Using 400G DCOs with 0dBm output power

1. Using external DWDM filters and EDFAs

In this solution, a set of DCO generated channels is first aggregated through a multiplexer or other type of optical filter and the resulting aggregated signal is amplified by an EDFA. Both the multiplexer and the EDFA are provided by the DCO, router, or switch vendor or from a third party instead of the existing DWDM vendor, which could also be an option. Given the fact that the amplified signal is an aggregate with multiple wavelengths, this solution is feasible over multi-degree ROADM systems where one degree can be dedicated to receiving these channels, or it's also feasible over systems with optical channels coming from multiple multiplexers, when interleavers/de-interleavers are used to mix odd and even DWDM channels or when bandpass filters are used in combination with multiple multiplexers or add/drop units.

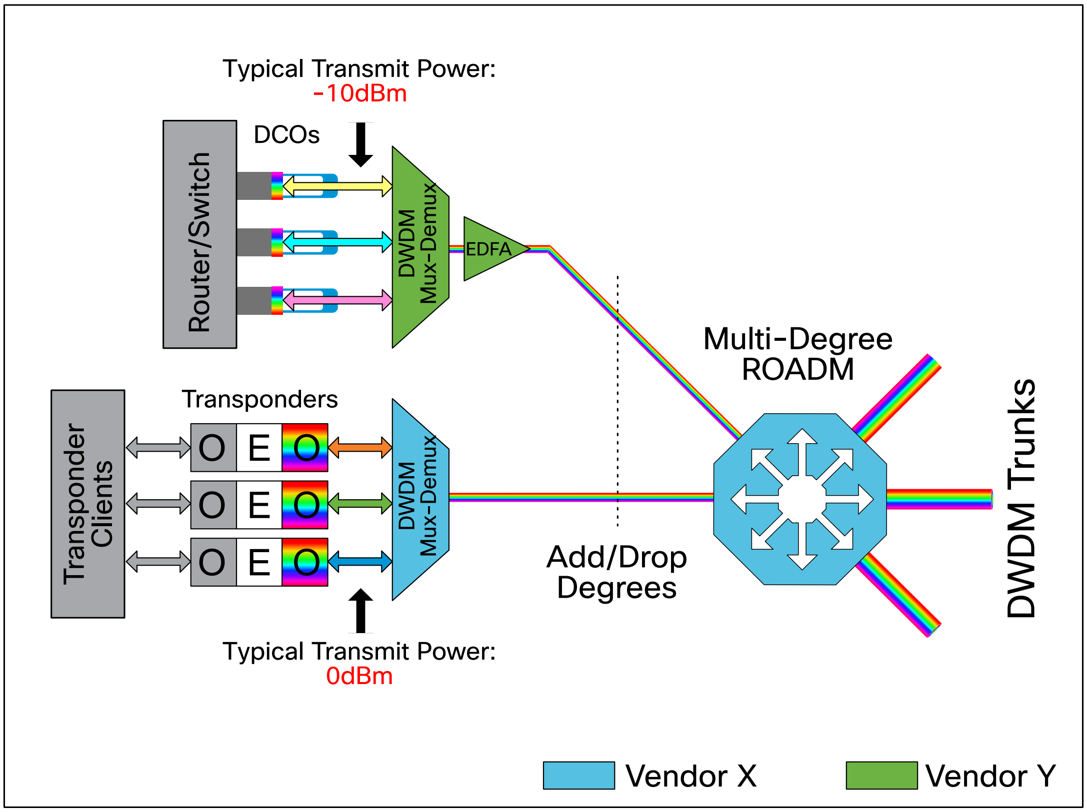

Figure 3. External Mux and EDFA connected to a ROADM degree

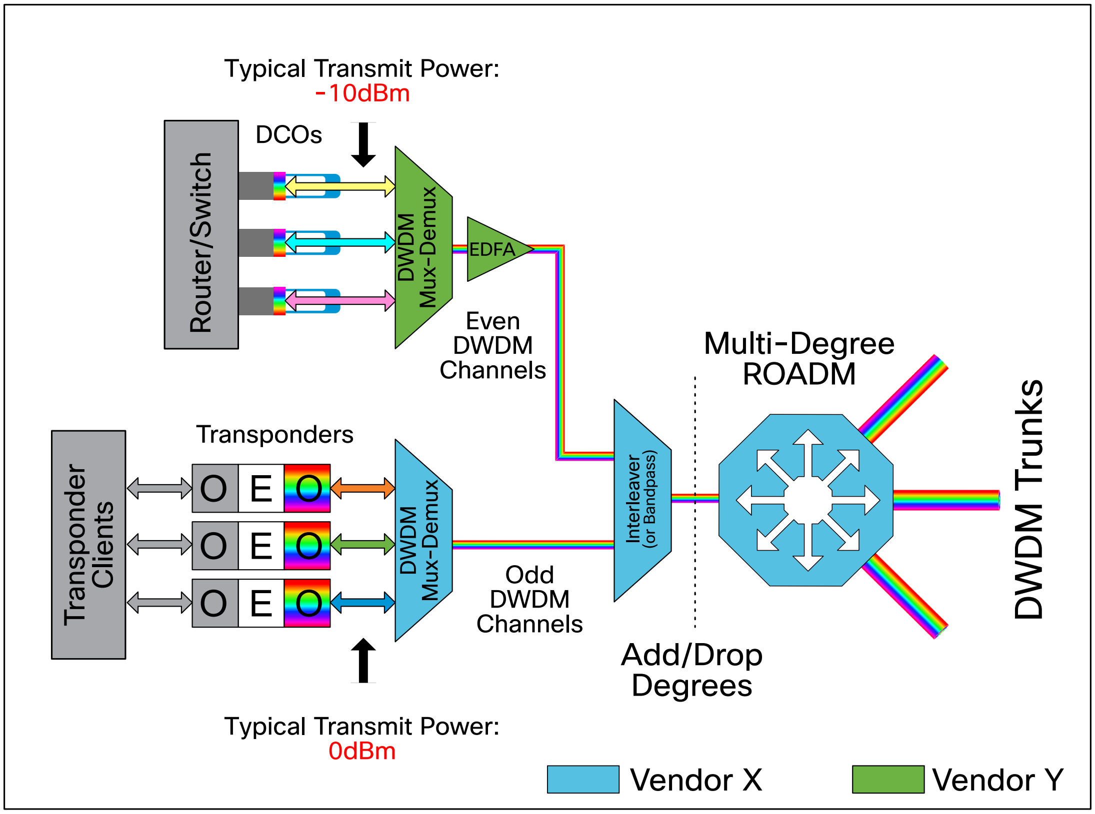

Figure 4. External Mux and EDFA connected via DWDM interleaver or bandpass filter

The advantage of this solution is that, from an optical channel perspective, from the point of view of the photonic layer, it's much simpler as no compensation must be done by the existing ROADM or DWDM line system. All compensation is done externally. In addition to that, the components used are relatively simple and inexpensive.

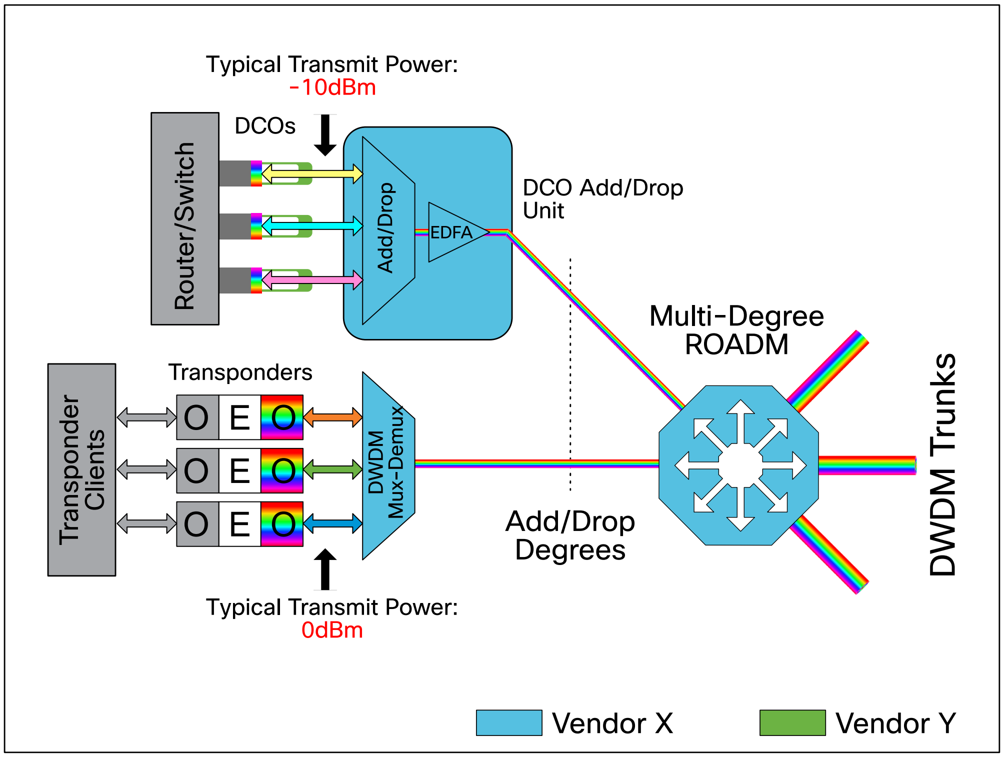

Figure 5. Existing ROADM is augmented to support DCO channels with lower transmit power

The caveats of this approach are related to the management of the solution, which will be the topic of the next section of this document. For now, it's sufficient to understand that the optical components (e.g., the additional EDFAs and the mux/demux/filters will be potentially managed independently of the ROADM or DWDM line system). Some operators may consider that an operational hurdle.

However, given the simple nature of this solution and the static nature of the configuration when deployed with constant gain optical amplifiers that maintain the per channel gain as more DWDM channels are added, the operational overhead is minimum and it's mainly towards fault management of the additional devices.

2. Using internal DWDM filters and EDFAs

This solution is like the previous solution, but instead of using external components such as from a third party, the additional optical components required for aggregating the DCO channels and increasing their optical power are provided by the existing ROADM or DWDM line system as part of the system itself.

This approach has the advantage of being supported directly by the DWDM system vendor and managed by the DWDM EMS/NMS system without the need for additional tooling or systems integration.

From an implementation perspective, there are two options:

a) The solution is implemented using discrete components for the add/drop and amplifier functions: In this case the DWDM vendor adds multiplexers/demultiplexers or optical add/drop and EDFAs on top of an existing ROADM or DWDM line system to cater to the specific amplification requirements of the 400ZR/OpenZR+ transceivers. The DWDM vendor may decide to support this configuration as part of an optical node itself, resulting is simpler management, or as a custom configuration that may not be supported as a single DWDM node. In both cases the network operator gets the benefit of being able to manage the additional hardware using an existing EMS/NMS from the DWDM system vendor.

b) The solution is implemented using integrated components optimized for ZR/ZR+ applications: In this case the vendor provides an optimized solution that combines the multiplexer/demultiplexer or add/drop and EDFAs into one unit that is supported as part of the optical node configuration. The management of the entire solution is simplified as there's no need to add new optical nodes to the network and the hardware can be optimized for lower power and space requirements. As an example, Cisco has developed an add/drop unit that provides built-in amplification for up to six QSFP-DD 400ZR/OpenZR+ channels per unit, as part of the NCS 2000 ROADM platform. The solution requires only one slot and it's optimized for this application.

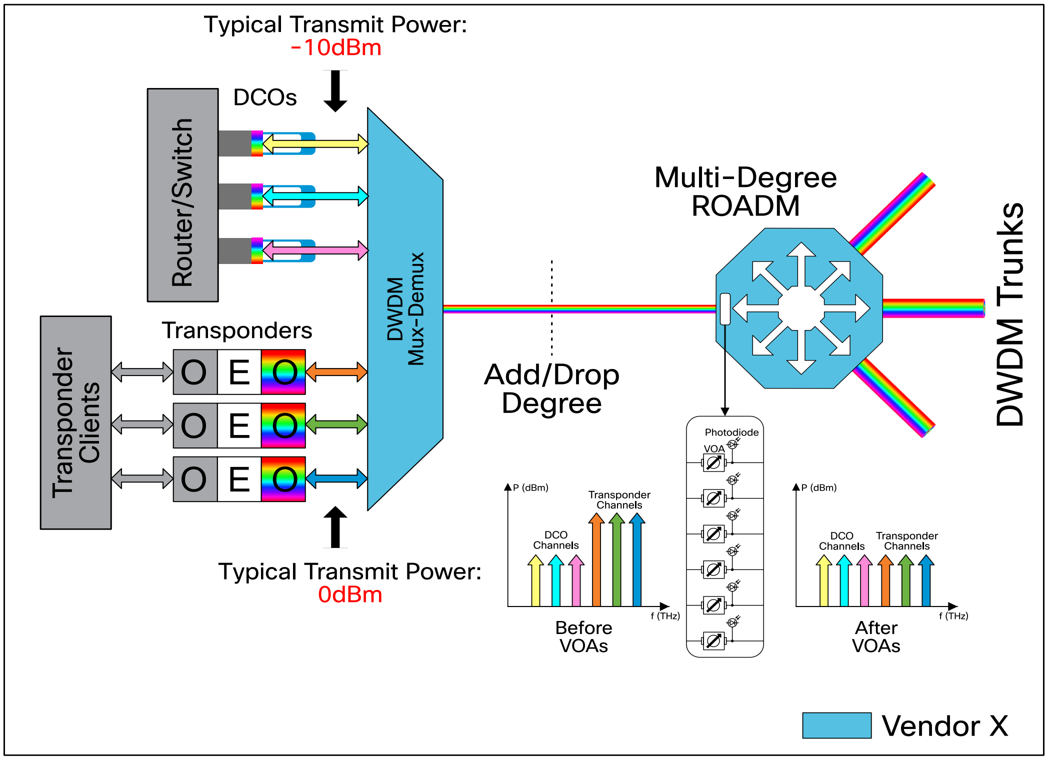

Figure 6. Attenuation of existing channels to match DCO with -10dBm using built-in ROADM VOAs

3. Using ROADM built-it channel equalization

In this approach, instead of using amplifiers to boost the -10dBm optical signal from a DCO interface to levels closer to the ones generated by traditional transponders that will co-exist in the DWDM system, the solution takes the opposite approach. Using Wavelength Selective Switches (WSS), which provide built-in configurable Variable Optical Attenuators (VOAs) and photodiodes for per-wavelength power level adjustments and monitoring, ROADM systems can attenuate the existing channels to lower power levels closer to the DCO provided -10dBm.

For new deployments, a solution can be designed for the worst-case power levels from day one hence network feasibility can be validated during the design phase. On the other side, for brownfield networks, changing power levels of existing DWDM channels may generate undesired consequences due the extra attenuation impact on OSNR due to the and hence it must be done with extreme caution and just after proper validation. It also possible that such a solution is not feasible at all for a brownfield network that has been designed to stretch system performance.

This is a relatively simple solution to implement when it's feasible to attenuate the -0dBm or higher power channels closer to -10dBm. In any case, it's also an option for lab environments, proof of concept, qualification, and interoperability tests.

4. Dedicating a ROADM degree to DCO channels

Modern ROADM systems provide many degrees that can be used as DWDM trunk ports to carry wavelengths to adjacent DWDM nodes or as add/drop ports. For instance, Colorless, Directionless, Contentionless, Flex-Spectrum/Gridless (a.k.a. CDC-FS or CDC-G) ROADMs leverage ROADM degree ports to provide those capabilities and, in some configurations, each degree port can be dedicated to a DWDM channel. Given the ROADM insertion losses, for example the attenuation imposed by the ROADM components to the wavelengths that go through the optical components (typically in the nine to 16 dBs range), a common solution is to use built-in EDFAs and optical amplifiers to compensate for them.

Considering those capabilities, another solution that can be considered to facilitate interoperability is to dedicate one or more of the ROADM degrees for DCOs with -10dBm and using the built-in amplifiers to boost the signal to the power levels required to achieve the required equalization with existing wavelengths. If the desired ROADM solution requires CDC ports, then each ROADM port will be dedicated to a given DCO. For other deployments, a DWDM multiplexer or an add/drop filter can be used to aggregate various DCO generated wavelengths before they reach the ROADM port, exactly in the same way it's done for transponders in "colored" ROADM implementations, providing a more scalable solution.

On the management side, this solution should see the DCO wavelengths as "alien", a concept that's widely known by the optical industry and supported by standards recommendations.

All these solutions can be applied to address eventual optical power differences between the DCO transceiver and existing transponders to resolve interoperability at the hardware level. So far, we've discussed interoperability of 400Gbps DCO pluggable transceivers and brownfield DWDM at the component level. Network-wide feasibility of this solution is the main goal, and it requires a broader discussion about network modeling, which will be the topic of the next section.

5. Using 400G DCOs with 0dBm output power

The interoperability solutions discussed so far aim at compensating for the lower DCO output power, i.e. -10dBm, by adding EDFAs and DWDM filters, another solution is to leverage DCO pluggable transceivers that can generate 0dBm output power.

Currently, Open ROADM has defined a 400Gbps DWDM interface with 0dBm output power. Initially, this interface is supported by CFP2 DCO pluggable transceivers that, when compared to QSFP-DD, require more power and space in the host equipment.

Moving forward the industry is extending the support for 0dBm also to QSFP-DD pluggable transceivers. In this case, Open ROADM can be used as the standard specification to facilitate interoperability especially with Open ROADM compliant DWDM line systems. Given the higher output power which is comparable to the ones generated by 400Gbps transponders, this solution will also facilitate multivendor deployments and open network architectures in greenfield and brownfield networks.

Network feasibility considerations

Network-wide feasibility of 400Gbps DCO pluggable transceivers over an existing network is the ultimate objective of any interoperability assessment. To validate network feasibility, multiple strategies are available:

- Running a theoretical exercise where the existing network is modeled, and the network performance is estimated based on industry-proven mathematical models and algorithms
- Simulating a real network in a controlled lab environment as a Proof of Concept (PoC) exercise
- Running a field trial or first office application, where an available wavelength can be used to validate the 400Gbps DCO interface performance, after careful analysis is done beforehand following the principles and characteristics discussed earlier in this document.

As part of the feasibility analysis for critical, large-scale networks, network engineers may use all these strategies, starting with a network modeling exercise.

Modeling optical networks can be very complex and requires a software tool that can take topology information (e.g., sites/fiber spans and traffic demands as main inputs and then route traffic across the available links and assess optical feasibility over the defined paths). The Dijkstra's algorithm is commonly leveraged to calculate shortest path routing where needed. Traffic engineered paths can also be leveraged while optical feasibility can be computed by taking advantage of the industry agreed upon Gaussian Noise (GN) Model as an alternative to proprietary, vendor specific models.

The GN Model computes linear and nonlinear performance of each Optical Multiplexing Section (OMS) in the DWDM network and can be calculated using open-source tools. For example, the GNPy open-source project provides a GN modeling Python library and has been used by major industry players including large service providers. Using such models, from a given DWDM network topology and the associated back-to-back interface performance, the proper line side data rate is chosen by the modeling algorithm by taking the worst-case computed end-to-end linear and nonlinear performance.

Computing the linear and nonlinear performance requires the definition of interface baud rate, channel spacing, number of channels, and fiber characteristics including length, type, and loss. With that information, a modeling tool can compute the optical launch power into the fiber, followed by the linear OSNR per OMS and with the GN Model the nonlinear contribution, plus any additional penalties. The tool can then account for the multiplexer and demultiplexer or ROADM structure and any additional penalty the network designer would like accounted for. The results can then be compared to the back-to-back interface performance at a defined rate to associate the proper data rate.

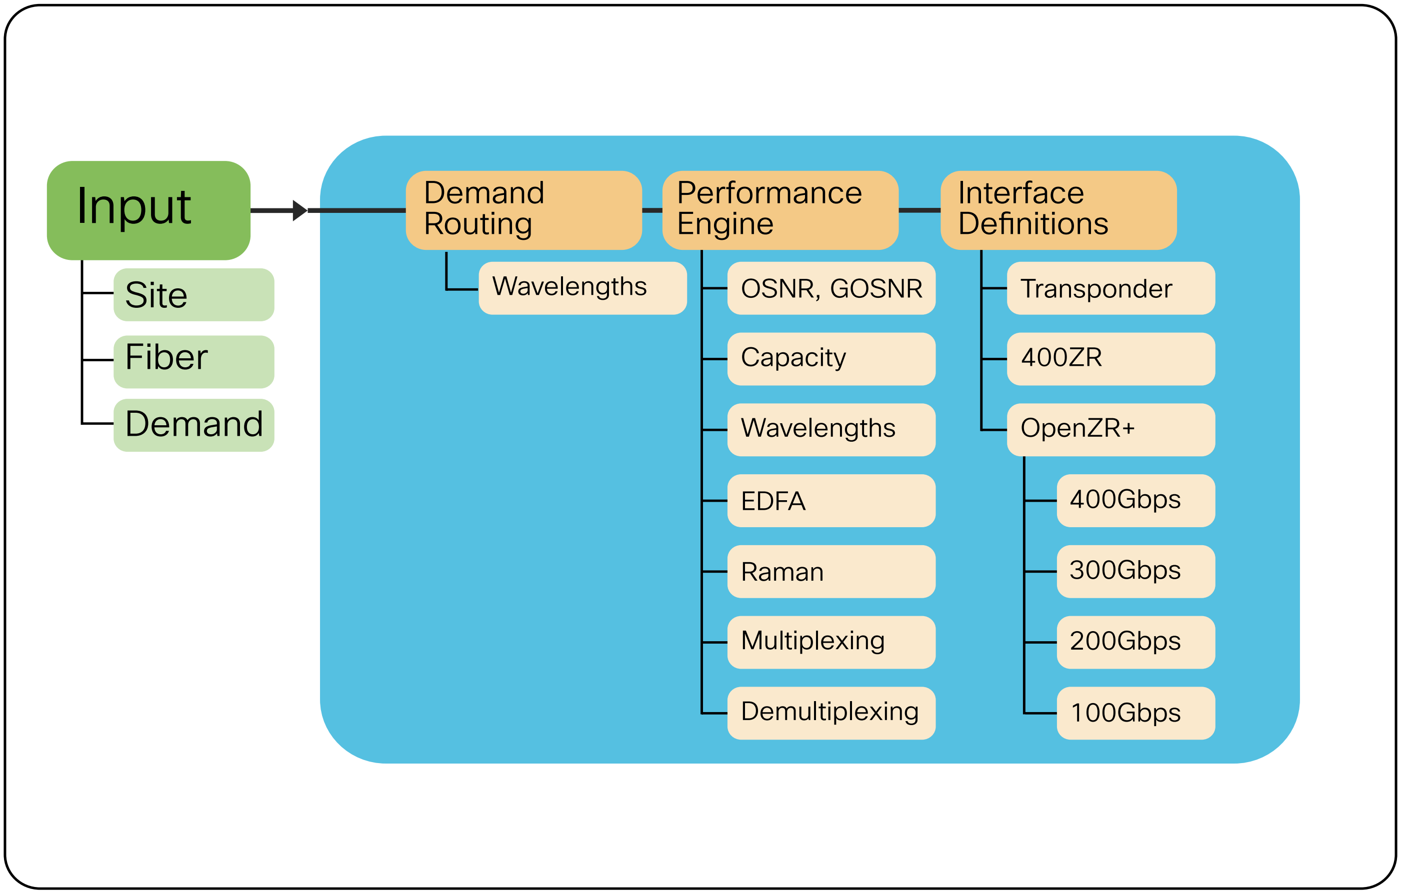

Figure 7. A network modeling approach

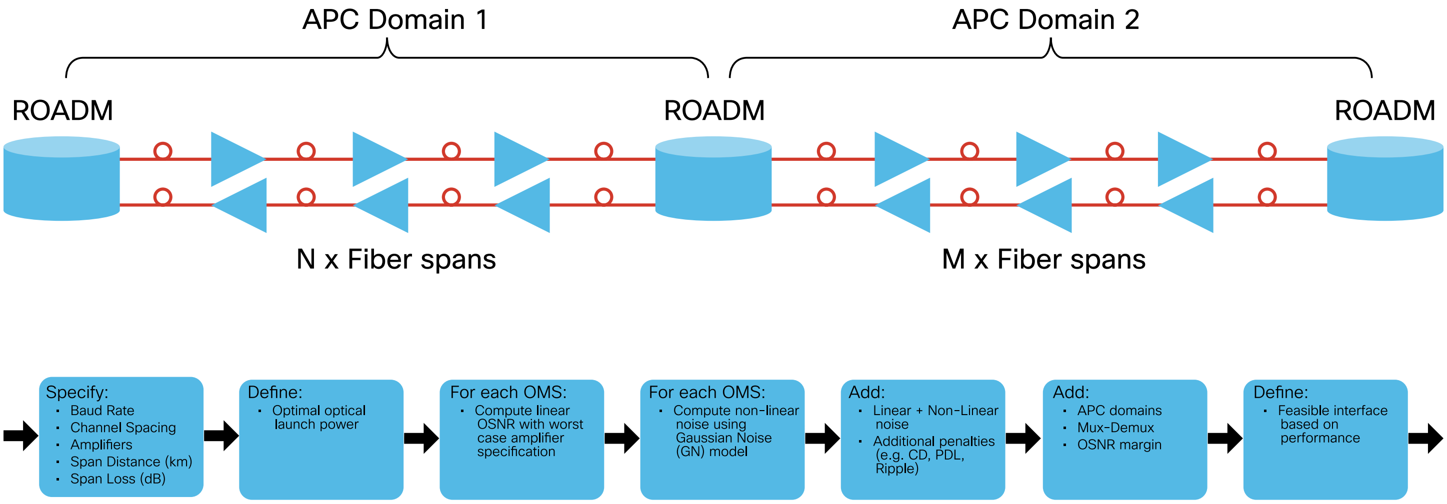

Figure 8. A DWDM network modeling workflow for optical channel feasibility analysis

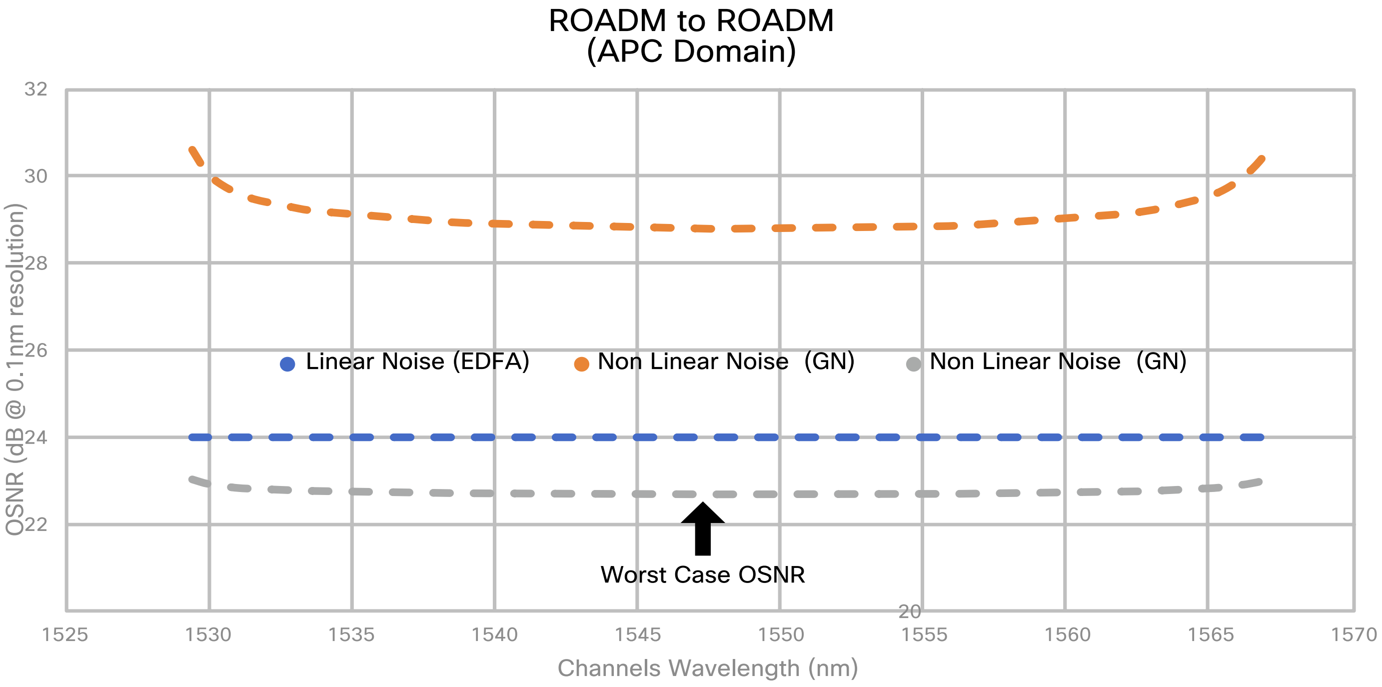

Figure 9. Modeled data flow for optical compute/feasibility and sample output across the C-Band

Supported DWDM line side data rates vary depending on the 400Gbps DCO specification with OIF's 400ZR operating at a fixed 400Gbps and OpenZR+ operating from 100 through 400Gbps in 100G increments at ~60GBaud.

The network modeling should always use the best available data from the real network to avoid over-engineering the solution or overshooting the analysis. Remember that in many cases the goal of network modeling is to get a theoretical validation before more effort is spent in the next step that can be running a PoC in a lab or a field trial in a live network.

Management of DCO pluggables in brownfield deployments

Management of heterogeneous networks that combine multiple technology domains and vendors has been a challenge for decades with no easy solution, especially for legacy networks. One of the reasons for this is that transport networks have historically been designed with tighter integration between hardware and management software, often leveraging proprietary features and interfaces.

Until recently, integration of multiple technology domains or vendor domains has been considered a problem to be solved at upper software layers using northbound interfaces to connect Element Manager Systems (EMS) and Network Management Systems (NMS) to centralized software functions, typically for collecting events like alarms and syslog messages.

Fast forward to today's networks and with the increasing adoption of Software Defined Networks (SDNs) the industry has pushed for open systems combining hardware and software components of multiple vendors. To bring open systems to life and enable multi-vendor management and automation architectures, the collective industry has gathered to define open APIs, data models, and more efficient management protocols, supported by major standards bodies and open industry initiatives including IETF, OpenConfig, OPENROADM, and others.

DCO pluggable transceivers are generally supported by modern routing and switching platforms that are fully aligned with those industry innovations that are taking place in the management software stack, supporting these open interfaces and data models to simplify the integration in modern software architectures.

Even though the current industry scenario is very promising for achieving multi-vendor management using modern open software stacks, the challenge again is how do we address the legacy and brownfield systems that only support proprietary interfaces. There are at least three approaches we can take to solve this engineering problem:

a) Upgrading existing management software infrastructure to support the new DCO transceivers natively

b) Using a hierarchical architecture where the integration happens at an upper network management layer

c) Adding a new, open model-driven telemetry solution to manage the end-to-end network

Let's go through each of these potential solutions and discuss them in more detail.

Upgrading existing management software infrastructure to support the new DCO transceivers natively

Whenever possible, upgrading an existing management software to natively support the new DCO transceivers provides a solution with the least impact to network operations. Considering that the DCO transceiver will likely be sitting in a router or switch, this solution makes more sense for integrated management systems (e.g., those that manage both the optical and the packet switching layer). All alarm events for fault management and the various interface counters for performance management for the DCO transceiver will be handled by the integrated management, which can also be extended to provide an end-to-end view of the optical channel trails to the network operator in a similar way to how it's done for DWDM transponders.

An alternative to the integrated management solution upgrade is to leverage software upgrades in separate management systems. A management platform that provides EMS/NMS capabilities only for the packet layer (e.g., IP/MPLS network) can be upgraded to support DCO pluggable transceivers while the optical network will manage the DWDM wavelengths as optical channel trail connections for the "alien wavelength" or "black link", such as at the photonic level. This solution would keep operational separation between the packet and transport domains, while the integration if required is done at another layer of the software architecture as will see in the next scenario.

Using a hierarchical architecture where the integration happens at an upper network management layer

Using hierarchical software architecture is a well-known, industry-proven approach to integrate disparate management systems and more recently SDN controllers and orchestrations platforms. It has been widely used for collecting and managing network events like alarms, creating network wide dashboards for KPIs and network performance, and for creating centralized network inventories and automation frameworks.

In the context of DCO pluggable transceiver management, two approaches are possible.

The first approach is to get data directly from the DCO transceiver host feeding the higher-level management system using modern APIs and device instrumentation to collect operational data. For instance, the network operator could use Model Driven Telemetry (MDT) frameworks to collect operational data modeled using YANG language via the Google Network Management Interface (gNMI). The operational data is then published for higher level management software access. This approach is gaining significant traction in the industry given its simplicity, scale, and performance properties as it allows to receive data based on event triggers or at pre-defined intervals up to the seconds range, compared to typical SNMP-based implementations that work typically with 15 minutes or higher intervals.

Talking about SNMP, that's also an option for integration with upper layer management systems. There's no caveat per se in using SNMP – it's the predominant management interface in many networks. However, model driven telemetry is a much better approach, it's becoming widely available and if it's an option, it's highly recommended.

To simplify the implementation and improve the scale properties of the model driven telemetry approach, the solution could implement a data collector as an intermediate step. The collector receives all the data published by the different network devices and if needed or programmed to do so it can filter it to control the volume of data exported. The data is then published to the northbound components that can subscribe to it in a publisher/subscriber (pub-sub) architecture. In other words, the data collector provides a single interface to the upper management components, avoiding multiple systems talking to the network devices directly.

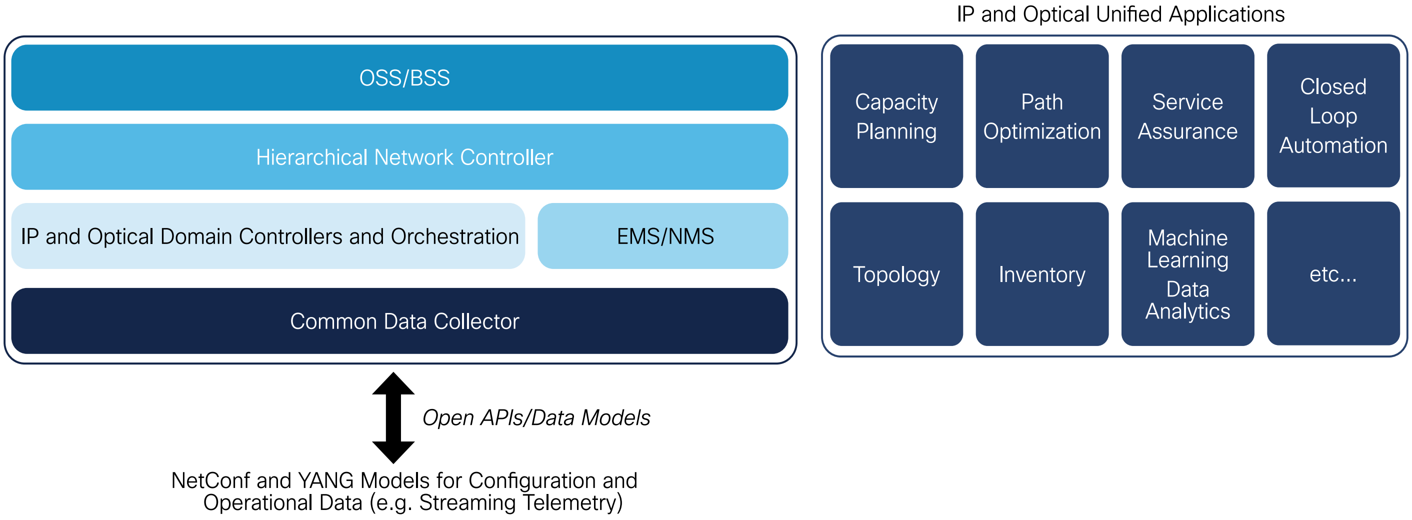

Figure 10. Modern model-driven telemetry-based management architecture for DCO

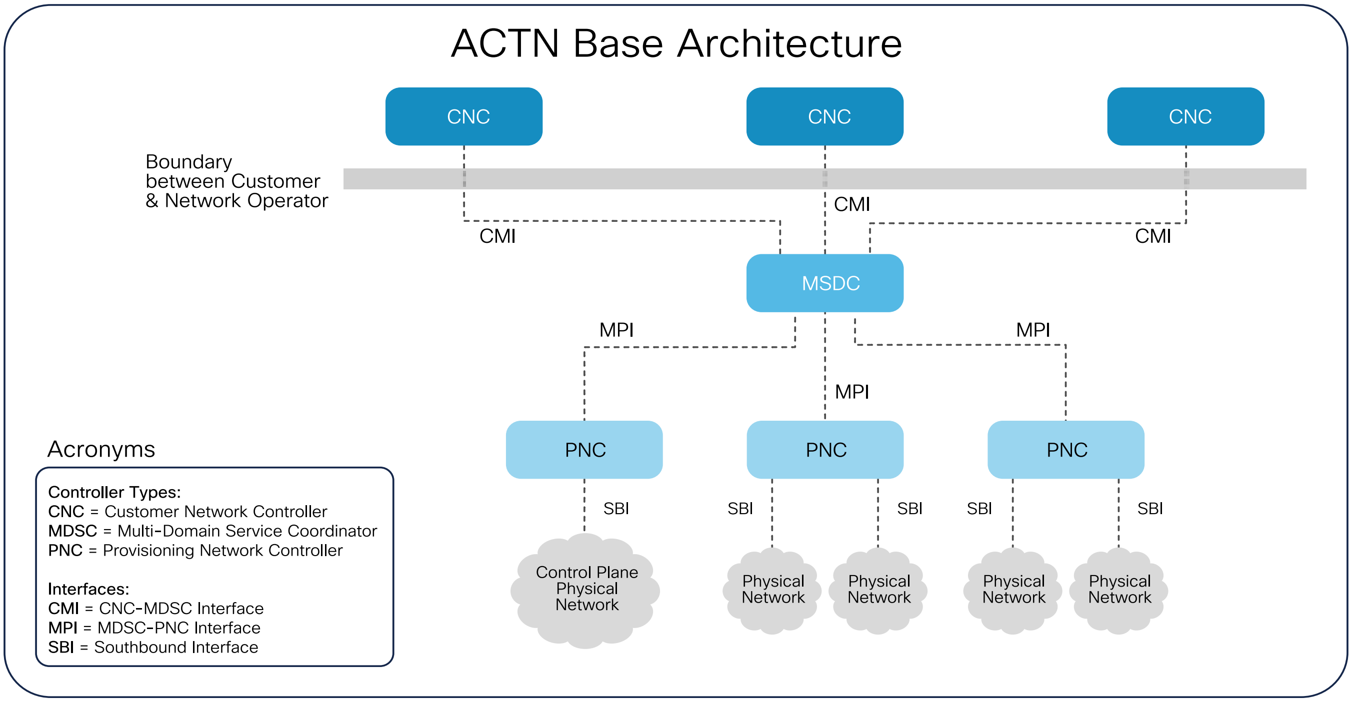

Figure 11. IETF RFC 8453 – Framework for Abstraction and Control of TE Networks (ACTN)

The second option is to use the Northbound Interfaces (NBIs) from an existing network management software for the host equipment to do the integration to the hierarchical management. As we discussed earlier, an existing EMS/NMS can be upgraded to add capabilities to manage DCO transceivers that are deployed in routers or switches. Many upgrade scenarios will end up with different management systems for the packet and the optical network. In this context, the hierarchical manager will centralize the information from both networks, and potentially other networks too, to provide a global view of the network and KPI dashboards.

Whereas in the first approach the interface to the hierarchical manager is at the device directly or through a data collector, here the integration is done via the NBIs provided by the EMS/NMS systems. Multiple interface options exist for this implementation, for example SNMP, and some management systems are highly flexible to integrate even with proprietary systems through the development of software adapters or connectors.

Currently there are multiple solutions available in the market that allow for a hierarchical software architecture by leveraging northbound and southbound APIs for integration. For instance, IETF RFC 8453 – Framework for Abstraction and Control of TE Networks describes a hierarchical architecture for controlling multi-domain networks, where domains can represent network elements that are grouped based on different technologies, vendor profiles, or geographical proximity.

Leveraging the IETF RFC 8453 framework, network controllers with a variety of standards based and open APIs can be used for integration.

Another example of an industry effort that enables the hierarchical software architecture is Open Networking Foundation (ONF)-defined Transport APIs (T-API), which support a variety of applications for both monitoring and configuring transport network devices. Using T-API interfaces between the domain-specific network controllers and the hierarchical controller results in an open solution for multi-vendor deployments and is a more future-proof architecture.

Adding a new, open model-driven telemetry solution to manage the end-to-end network

This option represents a modern approach for management and control of networks, and it’s based on the adoption of MDT frameworks with standard APIs and open data models for both the packet layer and the DWDM network.

In such deployments, the network operator may already be leveraging the scale and performance attributes of MDT to manage each network using a classic, DWDM transponder-based approach. To manage the DCO transceivers, new data models are required that expose the configuration and the operational data for these devices.

OpenConfig (OC) is an industry consortium composed mainly by network operators, traditional service providers and web-scale service providers that has been writing open data models for different types of networks and network devices using the IETF YANG language. OC has become very popular in the industry, and it's considered the de-facto standard for data models.

With the introduction of DCOs in the market, OC extended the optical-transport data models to cover 400ZR and 400G ZR+ modules [5].

Conclusion

DCO is a breakthrough pluggable transceiver technology that is set to disrupt the DWDM industry and 400Gbps implementations over QSFP-DD represents its state of the art. To collect immediate benefits from this technology, service providers, network operators, and companies with high traffic demands must understand the technology and design engineering solutions that allow 400G DCO pluggable transceivers co-existence and interoperability with brownfield DWDM networks.

Solutions are readily available that cater to different deployment options and operational preferences. In this paper we described the key characteristics of 400Gbps DCO over QSFP-DD, supporting industry standards and interoperability and integration scenarios for DCO over different brownfield DWDM options including various ROADM options and DWDM terminals. We also highlighted some of the technical challenges at hand. We looked at different management options for DCO when it's deployed as a new technology that sits in a different management domain outside the traditional DWDM network.

Finally, the industry is actively working to extend the open specifications and standards to support DCO technology both at hardware and software levels, including OIF 400ZR, OpenZR+ MSA, and OpenConfig data models.

Learn more

Take a closer look at some of the products and solutions covered in this paper such as Cisco Routed Optical Networking, Cisco Optical Networking, and Acacia (now part of Cisco) optical interconnects.

Citations

(1) ACG Research: The Economic Benefits of IP Transport at 400G

(2) OIF 400ZR Technical Specification (March 2020)

(3) OpenZR+ MSA Technical Specification – version 1.0 (September 2020)

(4) Cisco Network Convergence System 1004 C-Band 1.2T Transponder Line Card Data Sheet

(5) OpenConfig GitHub