

Cloud Interconnect: DWDM Integrated Solution For Secure Long Haul Transmission

The phenomenal growth in mobile, video streaming and Cloud services is driving the need for higher bandwidth within datacenters. Global IP traffic is estimated to grow at a compound annual growth rate (CAGR) of 23 percent from 2014 to 2019. These same applications are also driving the need to interconnect geographically dispersed datacenters to maintain seamless content delivery and to provide application agility without compromising on customer data security from external threats. The implications of poor data protection can be detrimental for an organization, ranging from government fines to costly lawsuits. To keep up with the global demand for higher bandwidth, datacenter operators need an interconnect solution that is dense, economically efficient and secure.

This white paper will provide an overview of the underlying technology of Coherent DWDM and the advantages of the Arista 7500E Series DWDM solution.

Arista Coherent DWDM Solution

Arista's 100G DWDM solution is a 6 x 100G Coherent DWDM line card for the 7500E series with integrated wire-speed encryption and analog coherent CFP2 optical interfaces (Figure 1). The line card delivers error-free performance up to 3000kms of fiber and consumes less than 140W per 100Gbps, which includes Layer 2/Layer 3 switching, wire-speed encryption and coherent DWDM. The standards-compliant coherent CFP2 optics are software tunable and support up to 96 channels in the C-band. Similar to any other Arista platform, the DWDM line card utilizes the single binary image of Arista's extensible operating system (EOS).



Figure 1. DWDM Line card and Analog Coherent CFP2 optic

Cloud Datacenter Interconnect

With most of the world's data residing at mega datacenters around the globe, the operators of these facilities require a tremendous amount of bandwidth and many fiber links to connect these geographically distributed datacenters. Independent market reports forecast the datacenter interconnect market to grow from \$300 million in 2014 to \$4 billion by 2019.

This transition to cloud-based datacenters is driving a fundamental shift in how data is transported at scale. Legacy low-density DWDM transport systems no longer meet the requirements of Cloud Datacenter Interconnect (DCI). While increasing the available capacity and providing seamless services to the users, cloud datacenter operators are looking to reduce capital and operational expenses to keep on track with increasing demand for content at constantly decreasing prices.

Data Security Concerns

Data security issues remain a top priority for cloud datacenter operators, and protecting data from passive wire-tapping after it leaves the physical boundaries of the datacenter is critical to protect against unknown tapping and data snooping. Strong and efficient data encryption is required to tackle a range of ever-changing vulnerabilities and threats. Due to the clear need to encrypt all data, combined with the explosive growth in the volume of data, a highly scalable and cost-efficient data encryption solution is required for DCI networks.

IEEE 802.1AE Media Access Control Security (MACsec) is an industry standard security technology that provides secure communication over Ethernet links that ensures data confidentiality. The Arista 7500E Series 100G DWDM line card integrates support for MACsec encryption at wire-speed on every port. It provides line-rate frame encryption and authentication for traffic from the 7500E that is passed over a DWDM network. This secure encryption capability is optional but removes the need for additional intermediate encryption devices to ensure confidentiality and anti-replay protection.

The MACsec specification uses long-term keys to derive session keys that are used for encryption utilizing the MACsec Key Agreement Protocol per IEEE 802.1X-2010. Long-term keys can either be statically defined or derived via remote authentication dial-in user service (RADIUS) servers. Data is then encrypted using the 128-bit or 256-bit GCM-AES-XPB block cipher suite. MACsec encryption can also be used in combination with other security protocols like IPSec and SSL.

Advanced Modulation and Coherent Detection

As datacenter interconnect bandwidth demands increase, datacenter operators need solutions that can increase the capacity on existing network infrastructure. Adding or leasing new fiber for long distance transmission can be cost prohibitive. Traditional optic solutions that operate with “on/off” modulation techniques do not have the required spectral efficiency at higher speeds to fit into the defined 50GHz DWDM grid that is deployed widely. As a solution to this, the adoption of advanced modulation techniques, in combination with coherent detection and digital signal processing, has proven to be the most viable solution for long distance transmission.

Coherent optical solutions have the required spectral efficiency, increase noise tolerance and compensate for several optical impairments without requiring any regeneration or dispersion compensation of the signal on the link. As data travels over long distance fiber, Amplifiers in the path (typically Erbium doped Fiber Amplifiers, EDFA) degrade the optical signal-to-noise ratio (OSNR) of the system. Additionally, fiber impairments such as Chromatic Dispersion (CD) and polarization mode dispersion degrade the Bit Error Rate (BER) of the channel. The Soft-Decision Forward Error Correction (SD-FEC) incorporated to digital signal processors compensates for the fiber impairments and recovers error bits to deliver an error-free signal.

The Arista 100G DWDM solution has digital signal processor (DSP) integrated to the line card and utilizes advanced Analog Coherent Optics in a CFP2 form factor to deliver the 100G Coherent solutions. The integrated Digital Signal Processor applies SD-FEC encoding and delivers 4x32GB NRZ signals to the Analog Coherent optics. The Coherent optics takes the 4x32GB electrical signals and generates 100G Polarization-Multiplexed Quad Phase Shift Keying (PM-QPSK) signals. At the receiving end, the Coherent optic receives the incoming optical signal and, using a local Oscillator (LO), sends the signal to the DSP on the line card. The DSP receives the polarization and amplitude information from the Coherent optic and processes the signal mathematically to compensate for fiber impairments and delivers an error-free signal as a result (Figure 2).

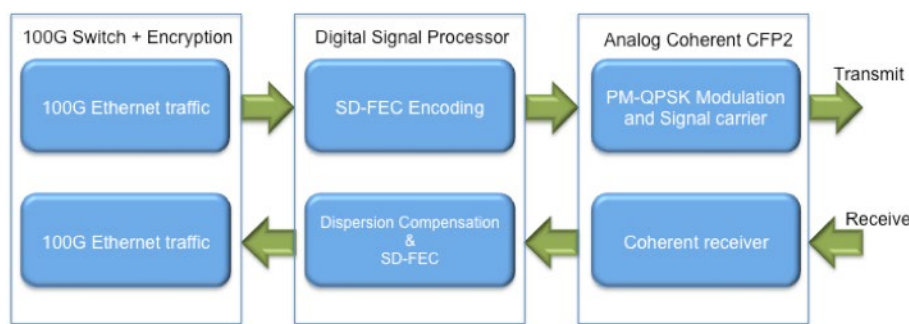


Figure 2. Functional block diagram of the Arista 100G DWDM solution

DWDM Pluggable Optics

The use of pluggable optics has gained tremendous popularity over the years and almost all common client interface types are predominantly addressed by a selection of pluggable optics. The key value offered by pluggable optics is the ability to easily replace or repair a link without having to replace the entire system. Pluggable optics also provide the flexibility to use various physical media devices (PMDs) depending on fiber type and reach requirement. Additionally, the disaggregation of optics from the rest of the system enables a pay-as-you-grow model without having the large upfront CapEx of a fully populated system.

Looking back at Ethernet optics, the “PHY” layer functionalities are incorporated inside the pluggable optics in the early phases of any technology. Eventually the “PHY” is moved out to the host board to reduce both the size and cost of transceivers. The 10G Ethernet evolution of pluggable transceivers started with XENPAK and eventually settled with SFP+, which had moved all PHY functions to the host. For 40G the early CFP evolved to QSFP+, and 100G CFP evolved to CFP4 and QSFP100 followed similar evolutions.

Until recently, long distance capable DWDM optics for 40Gb and 100Gb have been fully integrated to host systems due in part to signal integrity concerns between the optics and digital signal processor. With the standardization of 25G electrical interfaces and advancements in coherent optics, the time is now right for pluggable 100G DWDM applications. The Optical Internetworking Forum (OIF) has taken the initiative and standardized Analog Coherent Optics, as pluggable optics for 100G DWDM. Arista's 100G DWDM separates the DSP out from the optics and provides an analog interface to the pluggable Coherent optics, offering the benefits of pluggable optics without compromising the system performance.

Arista Cloud Interconnect Integrated Solution

Datacenter operators commonly use total cost per gigabit of bandwidth as a measure of the cost-effectiveness of a given solution. A traditional DCI solution with the spine tier connecting to an edge router, which in turn connects to encryption devices and a dedicated DWDM system, does not offer an efficient cost model. This has led to a series of pizza-box type DWDM solutions that have improved density and lowered CapEx compared to legacy systems. However, the pizza-box DWDM model still has many devices in the link, which increases CapEx and OpEx as well as requiring multiple client optics to connect between the devices. Figures 3 and 4 below show these two models.




Figure 3. Legacy DWDM transport solution




Figure 4. Small form factor DWDM transport solution

By taking the long haul DWDM interface and integrating directly into a fully featured Layer 2/Layer 3 spine switch, the Arista solution provides an opportunity to further lower CapEx by both minimizing the number of devices and client optics in the link and also lowering OpEx. With far fewer systems to operate and manage, this also provides space and power savings over discrete systems. With the integration of 100G DWDM to the Arista 7500E Modular Chassis Datacenter, operators have the flexibility to mix and match the client interfaces by using the wide range of complementary line cards available with 10G, 40G and 100G speeds, as shown in Figure 5 below.



Figure 5. Arista Integrated DWDM solution with encryption

Multiple Applications for Metro and Long Haul

The high spectral efficiency of PM-QPSK modulation allows the signal to propagate 100Gb Ethernet traffic through many cascaded reconfigurable optical add/drop multiplexers (ROADMs). With a robust SD-FEC residing on the DSP, a well-designed Coherent system can travel several thousands of kilometers and still deliver an error-free link. For a typical metro link that is less than 80kms, the Arista DWDM solution delivers error-free links without the need for any additional amplification (Figure 6).




Figure 6. A Metro link (~80kms) without any in-line amplification

Both regional and long haul point-to-point links can be achieved with in-line amplification, either EDFAs or Hybrid Raman EDFAs depending on the span length between amplification, as shown in Figure 7. The robust OSNR performance of the Arista DWDM solution allows point-to-point links to travel several thousand kilometers enabling seamless connectivity between datacenters that are separated by multiple time zones. The Arista 100G DWDM solution can also be integrated with a conventional ROADM system for Metro and long distance, as showing in Figure 8.




Figure 7. Regional or Long Haul Point to point link




Figure 8. Long Haul DWDM links with ROADMs

Flew-Grid Tuning Enables Additional Bandwidth

The telecommunication standardization sector of International Telecommunications sector (ITU-T) has defined DWDM wavelengths in the 1550 nm region with 100GHz and 50GHz spacing between channels. The most popular band used in telecommunication is the C-Band (1530-1565 nm), which falls in the wide-band amplification range of Erbium-doped Fiber Amplifiers (EDFA). The Arista 100G DWDM solution supports 96 channels in the C-Band, and the pluggable optical transceivers are software-tunable to any supported wavelength in the region. This gives an effective bandwidth of 9.6Tbps in each direction of transmission.

An approach to increase the total bandwidth without adding or leasing more dark fiber is to squeeze more channels into the existing fiber. Flex Grid Tuning is a viable solution, which increases spectral efficiency and allows for narrower channel spacing within the existing C-Band. By pulse shaping the signal using Nyquist filter, 37.5GHz spacing between channels can be achieved with the Arista 100G DWDM interfaces. This delivers 128 channels and up to 12.8Tbps bandwidth per fiber pair for a 33% increase in bandwidth compared to a regular 50GHz spacing.

Summary

Evolution of Coherent DWDM technology, along with advancements in digital signal processors with soft-decision FEC enables 100 Gbps DWDM to be integrated into a spine switch for long distance transmissions. The Arista Coherent 100G DWDM solution integrates Layer 2/Layer 3 switching with wire-speed MACsec encryption and coherent DWDM interfaces. Industry standard encryption with MACsec protects data integrity and customer confidentiality, which is critical to any organization. The use of standard pluggable coherent optics improves serviceability, and the un-bundling of optics allows for a flexible pay-as-you-grow model.

The Arista solution dramatically lowers both CapEx and OpEx for cloud datacenter operators compared to alternative non-integrated DWDM solutions and provides a simple, reliable and secure way to enable high bandwidth between datacenters for cloud interconnect.

Santa Clara—Corporate Headquarters

5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500

Fax: +1-408-538-8920

Email: info@arista.com

Ireland—International Headquarters

3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office
9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office
1390 Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office

Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office

9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office

10 Tara Boulevard
Nashua, NH 03062



Copyright © 2016 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document. Sep 2016 02-0055-01