

Arista DWDM Solution and Use Cases

The introduction of the Arista 7500R Series DWDM solution expands the capabilities of the Arista 7000 Series with a new, high-density, high-performance, DWDM line card.

Until now, providing secure high-speed Metro and Long haul, high-bandwidth Ethernet connectivity between customer locations required separate layer 2 and 3 IP networking devices, passing each connection through dedicated encryption devices and provisioning optical channels over active DWDM solutions.

With its high port density, scalable performance and deep buffer capabilities, the 7500R makes the ideal choice for a network spine for next generation datacenters, meeting and exceeding the needs of the bandwidth-hungry applications.

With a history of innovations, the Arista 7500R Series 200G DWDM line card extends the capabilities of the Arista 7500R chassis beyond a highly dense 10/25/40/50/100G Ethernet switching platform with full support for standard Layer 2 and Layer 3 protocols.

The introduction of a next generation Metro and Long Haul Ethernet solution combines full featured high-performance layer 2 and 3 switching/routing, inline 802.1AE MACsec encryption and fully tunable coherent DWDM optics in a single line card for the Arista 7500R Modular Chassis for both Metro and Long Haul environments.

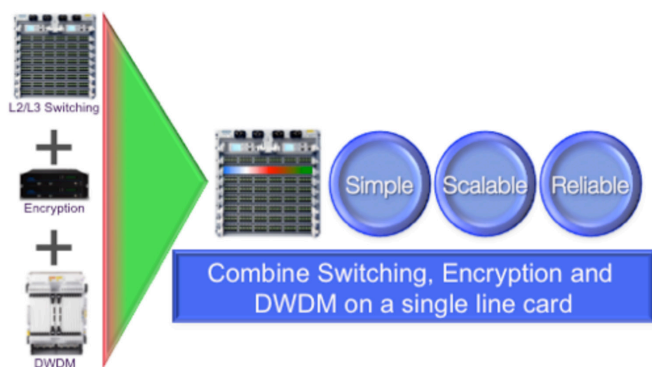


Figure 1: Arista L2/3, Encryption & DWDM Solution

Introduction

Due to cost and bandwidth limitations many enterprises traditionally located their datacenters within the same location as their users. This in turn presented a potential issue in that if the building became unusable, how would they continue to operate.

This gave rise to having separate dedicated datacenters, whereby companies would have a smaller Disaster Recovery (DR) datacenter at a remote location where they could continue to operate should the need arise. Other than for periodic DR testing these locations were rarely used and considered a huge cost with little advantage (until a disaster happened). DR sites also presented a new problem. How do companies get all the data from their primary to DR datacenters? Nightly and weekly backups were created and shipped off-site so that they could be used in the event of a disaster and for DR testing.

When performing DR testing it became clear that the amount of time it took to restore systems from tape backup was slow and would leave the business unable to operate until the systems were back in production, which depending on the amount of data required to restore, could take what seemed like an eternity. In addition the data restored was always stale and did not represent the latest state just prior to the disaster, which could also lead to business issues due to the missing data if DR was ever evoked.

This challenge drove the need for solutions to synchronize data between locations, requiring dedicated network links. Companies also wanted to make better use of their DR facilities, which led to them becoming secondary DC rather than purely DR, allowing some applications to run in a production mode, which in turn led to the need for increased bandwidth between the primary and secondary DC's. Further a need to separate production traffic from back-up traffic drove the need for even more bandwidth and additional dedicated links for each type of traffic.

For many enterprises a DWDM network solution seemed the best course of action. Although expensive, it was more cost effective than provisioning multiple separate links, providing a scalable and elastic solution for the enterprise that offered high bandwidth and data separation, for types of data as well as different business units. In time, these DWDM networks have scaled from low bandwidth rates of 1G to higher bandwidth 10G and 40G.

The Anatomy of DWDM Connections

Let's take a look at a typical 1G and 10G Ethernet over DWDM solution. For intra-datacenter solutions, an endpoint connection uses one of two types of optics, multimode (850nm) for short ranges and single mode (1310nm) for longer ranges. The DWDM node converts this local connection to a channelized frequency or wavelength (λ), which is then multiplexed with other λ s and transmitted over a single fiber connection.

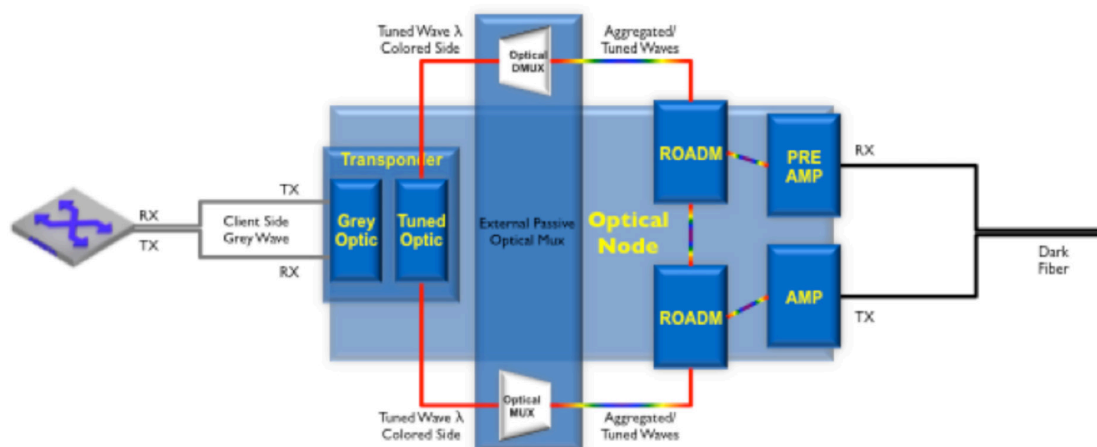


Figure 2: Basic DWDM System

The “Transponder” in a DWDM node performs the conversion from the client side grey wavelength to a channelized lambda or colored wave. The colored wave is then connected to an Optical Mux/DeMux port matching the tuned wavelength. Often external to the DWDM node the Mux/DeMux multiplexes multiple colored wavelengths together onto an aggregated signal and over a single fiber. The aggregated signals are then passed into a ROADM (Reconfigurable Optical Add-Drop Multiplexer), which is configured to add specific lambdas from the node to the dark fiber line. In a multi-node or multi-degree network, the ROADM is responsible for adding, dropping or passing wavelengths on, through an express connection, based on specific configurations for each lambda.

The aggregated signal from the ROADM is then passed on to an amplifier to boost the signal on the outbound connection. The node may also employ a pre-amplifier that boosts an incoming aggregated signal as it comes into the node from another location, prior to passing on to a ROADM to drop or pass signals at that node. The basics of DWDM systems are the same, though some pieces are not present in all networks.

Arista 7500R 200G DWDM Line Card

As a best of breed solution, with a proven history of innovation, investment protection, high port density for 10G, 25G, 40G, 50G and 100G, low latency, scalable performance, deep buffers, fully featured Layer 2 and Layer 3 capabilities, advanced network virtualization, open monitoring and network analysis features, the Arista 7500R is the ideal choice for the network spine for two tier datacenters.

Arista’s new 7500R 8-port 200G Coherent DWDM line card expands on previous innovations by introducing a low-cost, high-density, secure and reliable DWDM connectivity solution with no loss of features or performance.

Integration Simplicity

With the 8-port DWDM line card, the Arista 7500R extends the capabilities of a full service spine switch to provide the perfect resolution for high bandwidth Metro and long haul DCI (Datacenter Interconnect) solutions, for Ethernet over DWDM (EoDWDM).

To migrate from an existing 10G DWDM solution to new 100G-transceiver / transponders and 100G Coherent Line side modules, significant investment in transport hardware would be required, making it cost prohibitive to upgrade an existing DWDM system to provide the levels of bandwidth necessary for next generation DCI’s.

Arista has taken a different approach; rather than requiring customers to upgrade the DWDM equipment to provide 100G DWDM long haul capabilities, we have integrated the DWDM capabilities directly into a 200G DWDM line card, mitigating the need to upgrade any portion of a customers DWDM install base. The native Ethernet speed supported by this line card is 100GbE but the integrated Digital signal processor combines 2x100GbE to transmit 200G over a single wavelength.

The new 7500R Series DWDM line card provides eight 200G ports, with coherent 200G tunable optics allowing customers to connect directly into existing Optical Mux/DeMux without the need for additional external transceiver modules. The Coherent optics use the standard ITU wavelength schema (C-band) and spacing providing a cost efficient solution for up to 96 channels of 200Gb over a single dark fiber pair. The line card provides the flexibility to operate in 100G, 150G and 200G modes to optimize reach and bandwidth based on network requirements. Below table shows the various modulation formats and the bandwidth and reach capabilities of each mode.

Table 1: Programmable modulation formats of Arista DWDM solution for 7500R

Operating Mode	Max Line Card Bandwidth	Modulation	Reach
100G	800GB	QPSK	> 5000km
150G	1.2TB	8QAM	2000km
200G	1.6TB	16QAM	1000km

In recent years there has been a great deal of discussion in regard to passive sniffing of fiber optics cables without the need to splice into them. To address this the 7500R Series DWDM line card integrates 802.1AE 256 bit MACsec encryption capabilities that allows all data to be encrypted before it is transmitted over any optical network, ensuring that data being transmitted cannot be compromised.

This high level of integration of the 7500R Series Coherent DWDM line card dramatically reduces the CapEx costs associated with upgrading an optical network to 100G, simplifying connectivity and management. In addition, the integrated support for wire speed encryption eliminates external dedicated devices and provides an all-in-one solution for L2/3 switching MACsec encryption and DWDM transmission.

Use Cases

Several use cases for the Arista 7500R Series DWDM card in multi-site datacenter networks exist. The following sections identify a number of common ones.

Use Case 1 - Less Than 80KM Dark Fiber Connection

For a simple solution of less than 80km the 7500R Series DWDM line cards can directly terminate a dark fiber connection providing a point-to-point connection between two locations. The 200G mode provides the highest density and lowest cost/bit and is best suited for such applications.

Many enterprise deployments will use far less than 96 channels, but typically require more than one; passive multiplexers allow for simple scaling by aggregating multiple wavelengths without additional equipment complexity.

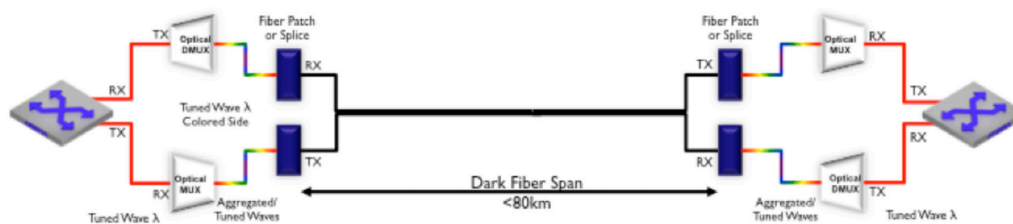


Figure 3: Passive Metro Ethernet Solution

Use Case 2 - Greater Than 80KM But Less Than ~ 150KM

Due to losses incurred in the light signal when passing through fiber patch panels, splices and the fiber itself, there is a need to amplify the signal when extending the distances beyond 80km. To obtain these greater distances, EDFA's or Erbium Doped Fiber Amplifiers are employed to give the aggregated wave a boost.

An EDFA takes the aggregated wave signal and passes it through a multiplexer, where it adds a pump laser. This signal is then passed through an Erbium Doped fiber spool, which excites the photons and boosts the signal. The boosted signal is then passed through a filter to remove the pump laser signal and the amplified aggregate wave is passed on as shown in Figure 4 below.

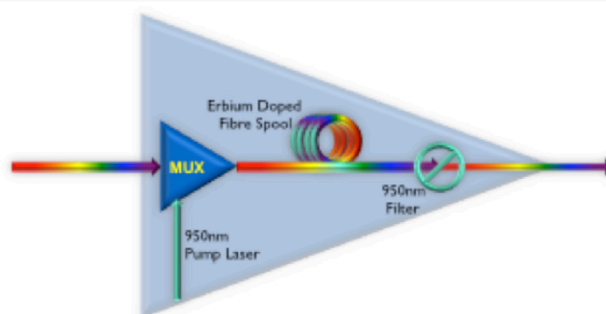


Figure 4: EDFA (Erbium Doped Fiber Amplifier)

By introducing EDFAs to the transmit side of each end of the dark fiber link, the signal can be boosted to achieve distances of up to 150km. Exact distances will be dependent on the number of patches, fiber splices and quality of the fiber.

The 200G mode of operation is recommended for this use case as it provides the highest density and lowest cost/bit.

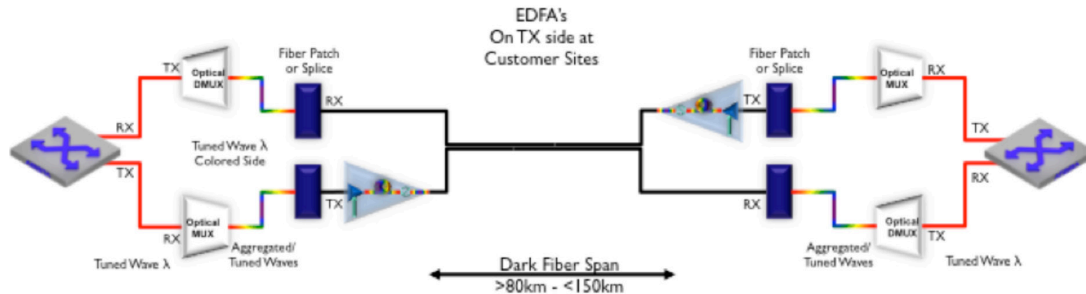


Figure 5: Passive Metro Ethernet with Amplification Solution

Use Case 3 - Greater Than 150KM But Less Than 3,000KM

As with the last example, employing additional EDFAs can further extend the distance of the link well beyond 150km. Installing EDFAs at a spacing of approximately 80km along the fiber route allows the length of a connection to be extended to over 3,000km. As shown below, EDFAs are used on both paths to boost the signal.

The use of 200G mode on the Arista DWDM line card provides a reach capability of up to 1000km. For distances between 1000km and 2000km, the 150G mode is recommended for an aggregate bandwidth of 1.2Tbps per line card and a maximum reach of 2000km. For any distance beyond 2000km, the 100G mode of operation is recommended which provides a total bandwidth of 800Gbps per line card.

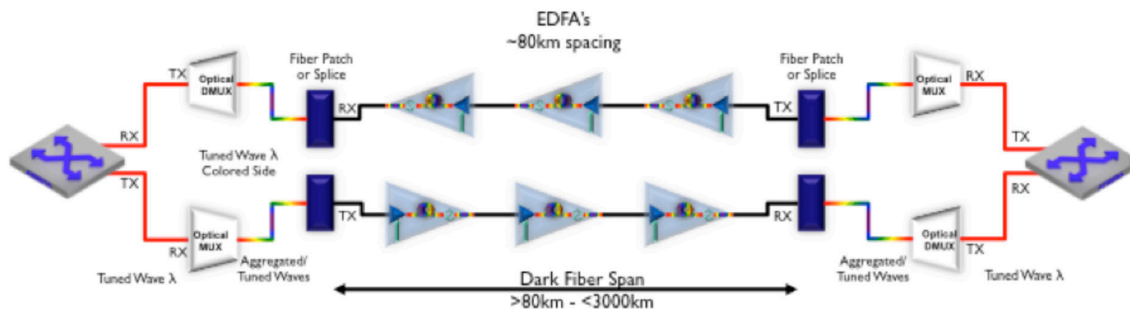


Figure 6: Passive Long Haul Solution

Use Case 4 - Multi-Node Passive Networks with Redundant Connections

Most customers require redundancy within their optical networks, to provide diverse paths between datacenters. Using the Arista 7500R Series DWDM card, a second independent link can be established to run over the redundant fiber connection. These connections can be sourced from different ports on the same card, a separate card in the same system, or a completely separate system depending on the level of isolation and redundancy required. Robust Layer 2 and Layer 3 IP protocol resilience will ensure convergence in the event of link or path failure.

In larger networks, redundant connections may go via a third location or node. To achieve the redundancy required and take advantage of the second path via the third location, some simple patching can be used to pass the signal through that site.

The primary connection between the two sites would be routed through the optical network in the most direct way, and the second link would be routed via the third location. At this location, a simple patch would connect the east and westbound Multiplexers together on the appropriate colored wave, and provide a pass through connection for that site.

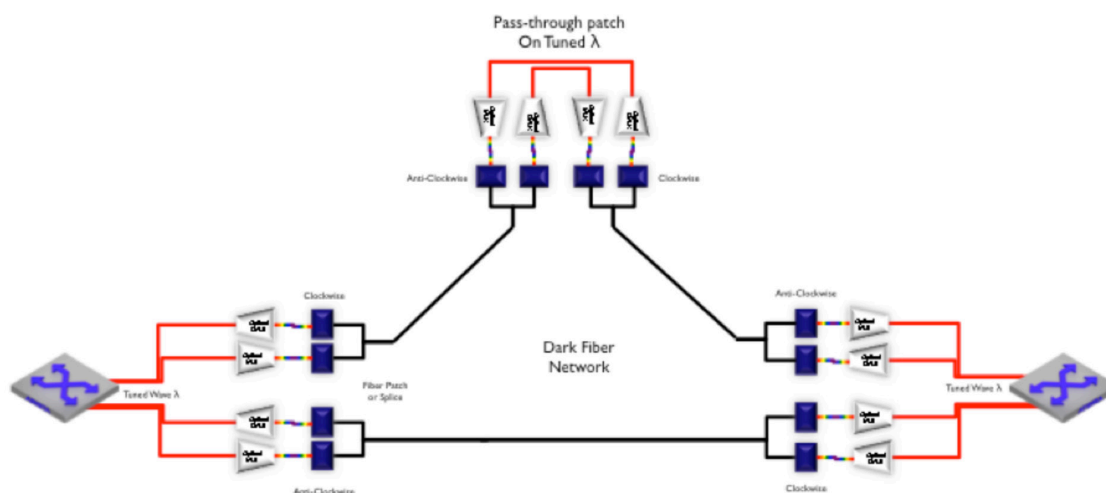


Figure 7: Passive DWDM with L3/L2 and fiber path redundancy

From the switch perspective, these are two regular connections that provide ECMP (Equal Cost Multi-Path) load balanced solution between the two sites. Since traffic would be hashed on a flow basis, the links could be members of a port channel or just simple point-to-point links.

Use Case 5 - Integration with Existing Active Optical Networks

The Coherent optics installed in the Arista 7500R-8CFPX are customer tunable across a full range of 96 channels and adhere to the ITU 50GHz and 100GHz grid for colored waves, frequencies and inter-lambda spacing. They can easily integrate with customers' existing optical systems that also comply with the ITU standards. The ROADMs on the active optical network can be configured to provide a path from A to B between locations as desired on any specific lambda. The tuned wave will then be connected directly to the lambda port on the Optical Mux/DeMux panel. This signal is then fed back into the ROADM, for switching through the active optical network, and is amplified in the same way as a signal generated from a transceiver module.

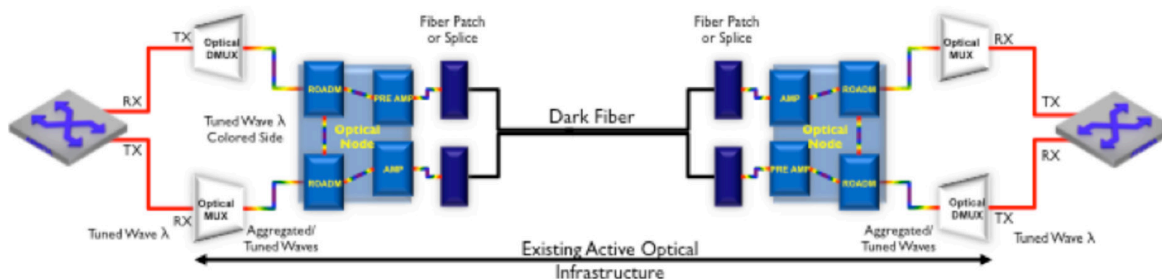


Figure 8: Integration with an Active DWDM system

Use Case 6 - Integration with Existing Optical Network with Redundancy

Similar to the use case above, two circuits can be built from A to B, one via the eastbound route, the other via the westbound route. If a path passes through a transit node, that node will be configured to use the express link between ROADMs and pass the signal on without dropping it at the node.

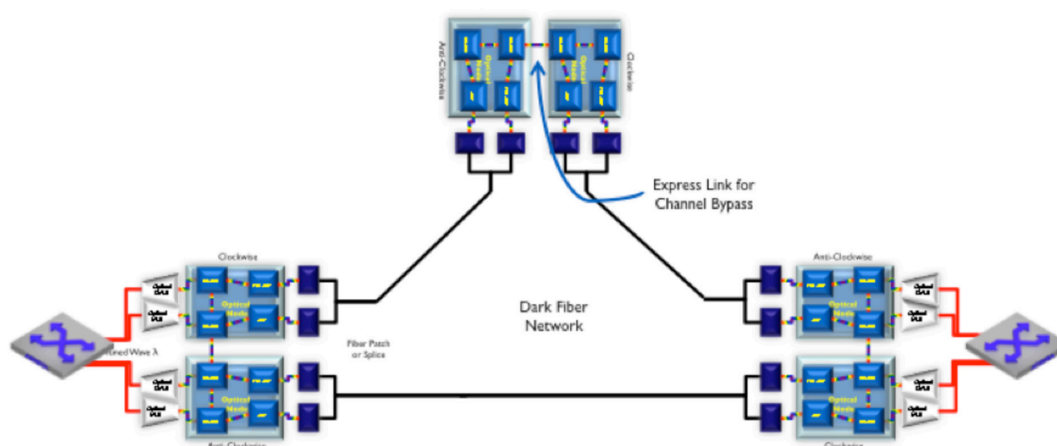


Figure 9: L2/L3 Redundancy integration with an Active DWDM Solution

In this case, the Arista 7500R Series DWDM card will have two ports tuned to the specific lambda's configured and connected to the Optical Mux/DeMux. One port for the east-bound and one for west-bound connections. While it is common to match the east and west Lambda's (for simpler management) this is not necessary providing the signal paths do not overlap each other. The advantage that two connections have over an Optical redundant link is that both paths are active at the same time, providing double the bandwidth over the optical network between two locations. The Arista 7500R Series switches use ECMP load balancing to hash traffic over each link, and any link or path failure would result in a L2/L3 convergence by re-hashing the traffic to the available path with minimal traffic loss.

Use Case 7 - Active Optical Failover Solution

Customers that utilize multiple fiber paths for redundancy, but wish to use an optical solution for rapid failover between paths, whether that is using an active multi-degree ROADM or active failover multiplexer, can still take advantage of the 7500R DWDM system. The optical failover solution is external to the 7500R switch and would be managed by the downstream active devices; the device can detect signal loss on the primary line and switch over to an alternate path automatically as shown below in figure 10.

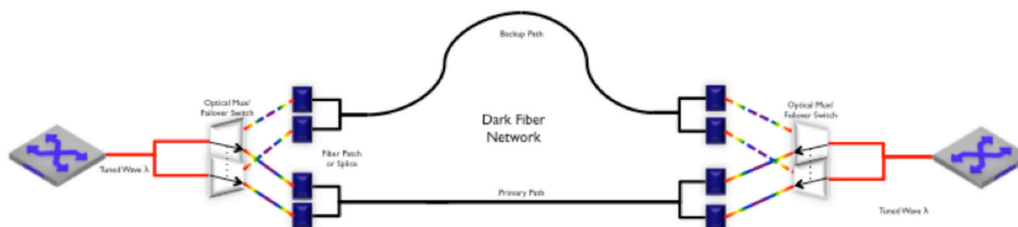


Figure 10: Fiber path redundancy with Active Optical Switching

Conclusion

Extending virtualized services between sites, increasing total capacity both inside and between locations, enhancing network resilience and adopting cloud scale economic solutions are the common drivers for many enterprises. The Arista 7500R has been the foundation of two-tier open networking solutions for cloud datacenters, based on open standards, delivering a feature-rich, highly scalable platform. The Coherent DWDM solution enhances the 7500R system by providing cost-effective, high performance DWDM port density with the same rich features and dedicated secure encryption in compact and power-efficient systems. Enterprises can easily extend and migrate existing metro and long-haul DWDM networks to add new 100G capacities, extend inter-DC L2 and L3 services.

The 7500R Series DWDM solution offers a cost-effective solution for transporting massive volumes of traffic with wire-speed encryption, unmatched density in a proven architecture that is scalable, and resulting investment protection.

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 © 2017 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. July 17, 2017 02-0021-03