QSFP-40G Universal Transceiver

The Arista Universal transceiver is the first of its kind 40G transceiver that aims at addressing several challenges faced by today's datacenters.

Increased adoption of 10 Gigabit Ethernet servers coupled with applications requiring higher bandwidth is accelerating the need for dense 40 Gigabit Ethernet switching. High-density 40G switches, such as the Arista 7500E, 7300X and 7050X Series, are available at an attractive price/port to enable the 10G to 40G migration. However, the widely available optical transceivers, that are required to support the migration to 40G, require re-designing the datacenter layout due to reach limitations and/or require major upgrades to the fiber infrastructure.

The Arista QSFP-40G Universal transceiver is a pluggable optical transceiver in an industry standard QSFP+ form factor that can operate with both duplex multi-mode and single-mode fiber. It has four channels of 10G multiplexed inside the module to transmit and receive an aggregate 40G signal over a single pair (2 strands) of fiber. The Arista Universal transceiver offers a very cost effective connectivity solution and unique value proposition for datacenters to migrate from 10 Gigabits/sec to 40Gigabits/sec with minimal disruption with existing multi-mode or single-mode infrastructure.

Existing 40G transceiver solutions like the industry standard 40GBASE-SR4 transceiver (100m over OM3 parallel fiber) and Arista QSFP-40G-XSR4 (300m over OM3 parallel fiber) for short reach applications require a total of 8 fibers per link; four each for transmit and receive. This is four times more fiber per link than is needed for a 10G short reach link. The Universal transceiver avoids this by requiring just one pair and is compatible with existing fiber patch panels and trunk systems.

Arista Universal transceiver advantages:

- Uses existing duplex fiber infrastructure for 40G
- Extend 40G reach to 150m over OM3 and OM4 fiber (150m reach covers 95%+ of Datacenter links)
- Identical transceiver for both multi-mode and single-mode fiber for simplified operations and investment protection
- Support for Digital Optical Monitoring (DOM) and passive Network Taps for link quality monitoring and passive data analysis
- 100% compatible to IEEE standards without any proprietary lock-in
- Optically interoperable with IEEE 40GBASE-LR4 and 40G-LRL4 for easy connection to third party routers and switches in existing networks
- Supported on all QSFP+ ports on all Arista switches without restrictions.

Further details of these challenges and how the Arista Universal transceiver addresses them are explained in rest of this document.

2-Fiber vs. 8-Fiber

Existing 40G transceivers for short reach, QSFP+ SR4 and the extended reach QSFP+ XSR4, utilize four independent 10G transmitters and receivers for an aggregate 40G link. These QSFP transceivers use an MPO-12 connector and require a parallel multi-mode fiber (OM3 or OM4). This is four times more fiber than is required for 10G short reach links (Figure 1). The Arista Universal transceiver also uses four transmitters and four receivers but has built in optical multiplexing and de-multiplexing, which results in a duplex connector and hence operates over the same duplex fiber infrastructure as 10GBASE-SR.

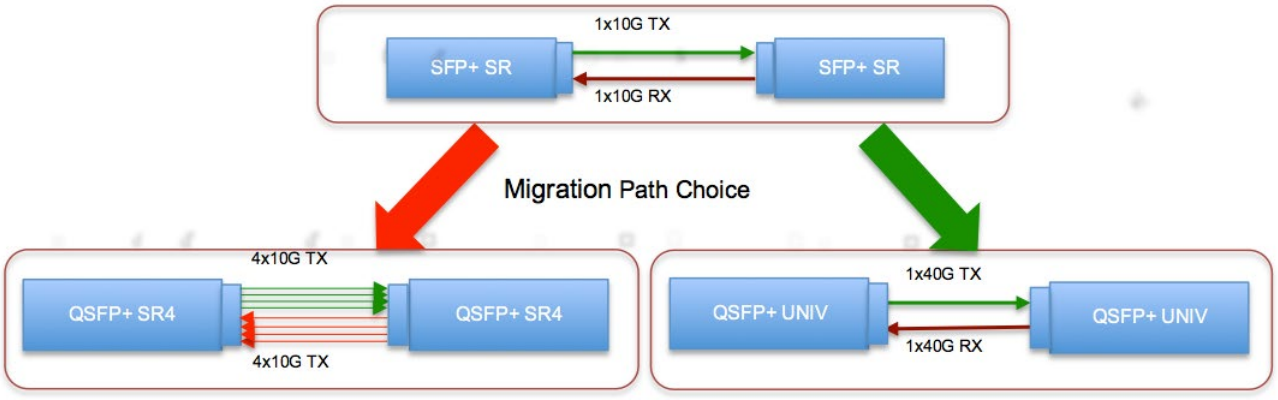

Figure 1: Fiber requirements for a single link using SR4 or UNIV

In addition to the reduced number of fibers per 40G link, the Arista Universal transceiver offers significant cost savings in the overall fiber cable infrastructure. Customers retain the existing structured cabling system as is for 10G to 40G migration, but they have to change the patch cables and patch panel infrastructure to use a QSFP+ 40GBASE-SR4. Figures 2, 3 and Table 1, 2 summarize the cost savings that can be achieved by using the universal transceiver in place of QSFP+ 40GBASE-SR4.

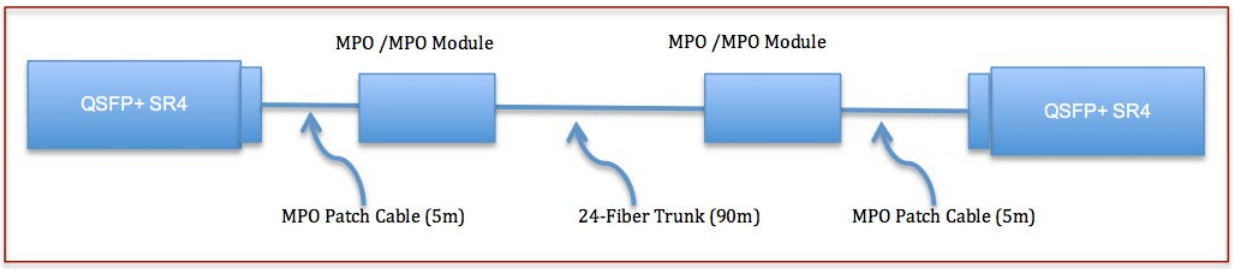

Figure 2: Typical 1-trunk cable infrastructure for QSFP+ SR4

Table 1: Cabling cost associated with 40G link as in Figure 2			
Item	Unit Price"	Extended Price Per 40G Link"'	Notes
MPO patch cable (5m, OM3)	$448	$896	2 patch cables required per link
MPO/MPO Module (OM3)	$1185	$790	2 modules are required which can be used for 3 links (per Link = $1185*2/3)
24-Fiber Trunk (90m, OM3)	$3198	$1066	QSFP SR4 requires 8 fibers in the trunk per link
Total		$2752	

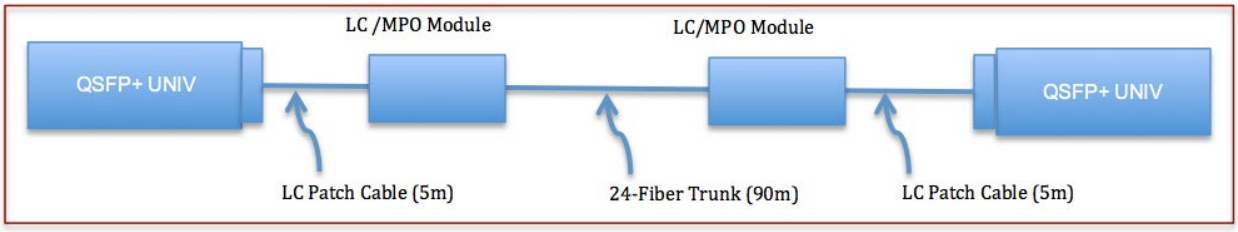

Figure 3: Typical 1-trunk cable infrastructure for QSFP+ UNIV

Table 2: Cabling cost associated with 40G link as in Figure 3			
Item	Unit Price[2]	Extended Price Per 40G Link[3]	Notes
Duplex LC patch cable (5m, OM3)	$64	$128	2 patch cables required per link
LC/MPO Module (OM3)	$482	$161	2 modules are required which can be used for 6 links (per Link = $482*2/6)
24-Fiber Trunk (90m, OM3)	$3198	$267	QSFP Univ requires 2 fibers in the trunk per link ($3198*2/24)
Total		$556	

Road To Single-Mode Fiber Transition

The number of connectors in a link significantly impacts the supported distance of a transceiver. In the case of double link channels (2 trunks), the insertion loss typically exceeds the allowed budget for connector loss forcing network architects to design to shorter link distances. The Arista Universal transceiver has a loss budget of 2.0 dB. It has been tested to support 1-trunk and 2-trunk structured cabling with up to 4 LC patch cables and 4 MPO connectors. The supported link distance of 150m over multi-mode fiber covers in excess of 95% of the existing fiber optic installed base with 1-trunk and over 80% with 2-trunks (Figure 4). The significant

cost savings realized by re-using 10G fiber infrastructure combined with the extended reach capabilities make the Arista Universal transceiver an ideal choice for 40G leaf-spine connections with existing multi-mode fiber.

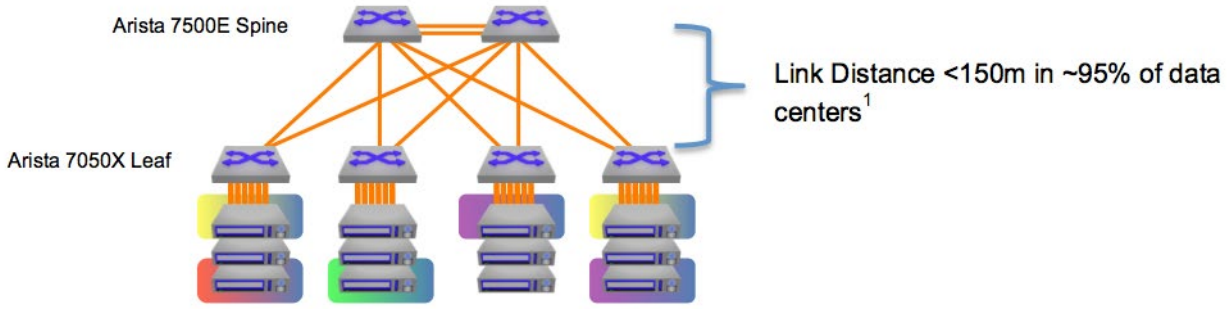

Figure 4: QSFP-UNIV in a leaf-spine architecture with multi-mode fiber

As data rates increase from 40G to 100G and beyond to 400G, there is a strong desire for datacenters to move to Single Mode fiber for cost effectiveness (Figure 5 and Table 3) and to future proof the fiber infrastructure. Due to the limitations of multi-mode transceivers to support existing distances with ever increasing data rates, migrating to 100G and 400G in the future will be simpler with single-mode fiber. However, the major pain point in this transition has traditionally been the optics cost. Single mode transceivers typically cost up to 4X more compared to multi-mode transceivers. Arista has broken this barrier with the universal transceiver:

- Cost of the optical transceiver is less than 2X multi-mode transceiver
- Same optics for multi-mode and single-mode fiber
- Investment protection as the same optics can be used when migrating from MMF to SMF
- Compatible with industry standard 40GBASE-LR4 and the LR4-Lite (1km version)

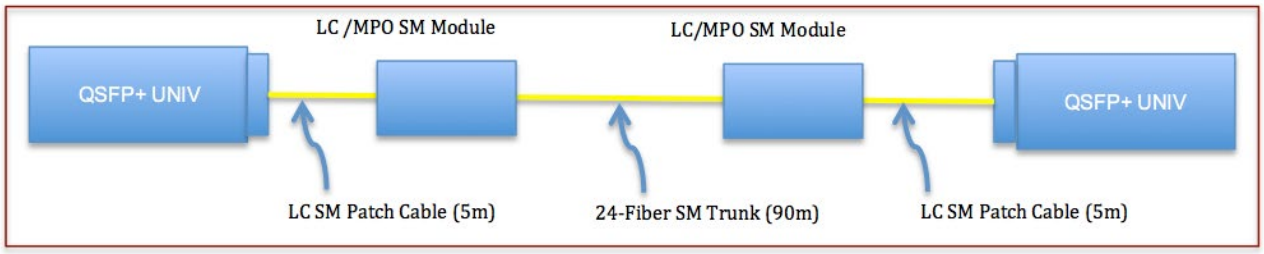

Figure 5: Typical 1-trunk Single-mode Fiber cable infrastructure for QSFP+ UNIV

Table 3: Cabling cost associated with a 40G link using all single-mode fiber infrastructure as in Figure 5

Item	Unit Price[2]	Extended Price Per 40G Link[3]	Notes
Duplex LC patch cable (5m, SMF)	$68	$136	2 patch cables required per link
LC/MPO Module (SMF)	$527	$176	2 modules are required which can be used for 6 links (per Link = $527*2/6)
24-Fiber Trunk (90m, SMF)	$2042	$170	QSFP Univ requires 2 fibers in the trunk per link ($2042*2/24)
Total		$482	

Network Taps And Digital Optical Monitoring

Passive Networks TAPs are the most commonly used devices in a fiber cable infrastructure to analyze the traffic without affecting the live application. TAP modules typically contain an optical splitter that divides the optical signal into two outputs: one for live link traffic and one for monitoring. TAPs are used in Datacenter networks for real time monitoring and analysis of traffic for various reasons listed below. The benefits of network tapping are becoming more profound in other fields such as health care, social networking and cloud computing etc.

- Regulatory Compliance requirements (Banking, High Frequency trading, etc.)
- Security Threat Monitoring
- Performance optimization and availability of the Network

Most of the time the TAP modules are fully integrated into structured cabling and it is expensive to replace or to add more TAPs into an existing purpose built network. The most commonly used TAPs in 10G networks are the 1x2 splitters with a split ratio of 70/30 or 80/20. The Arista 40G Universal transceiver can fully utilize existing 1x2 TAP modules as is, without modification. It offers true "Plug and Play" 10G to 40G migration path with the existing monitoring infrastructure.

One of the challenges of adding a passive TAP to a link is the insertion loss of the TAP module. Table 4 provides a rough estimate of the insertion loss to the live link and monitor link with various split ratios. The Arista Universal transceiver has power margins designed to support existing 10G TAPs. Leveraging the Arista Universal transceiver, customers have the benefit of using the industry standard 40GBASE LR4 on the monitor link if additional receiver sensitivity is needed due to the TAP insertion loss.

Table 4: Insertion loss with various TAP split ratios

TAP Split Ratio	Insertion Loss On The Live Link	Insertion Loss On The Monitor Link
50/50	3.8dB	3.8dB
70/30	2.4dB	5.8dB
80/20	1.8dB	6.6dB

Digital Optical Monitoring (DOM) provides access to real-time operating parameters of the transceiver through a digital interface. It enables pro-active monitoring of the key parameters of the transceiver and helps when performing link troubleshooting. DOM can help network operators to become aware of degrading fiber paths, detect transceiver problems and test splicing/patching work remotely and in a non-intrusive way. The Arista Universal transceiver has full DOM support with access to Transmit, Receive power levels and Temperature, voltage and bias current monitoring.

Standards Based And Interoperable Transceiver

The Arista Universal transceiver specification is based on the industry standard IEEE 40GBASE-LR4 and interoperates with any 40GBASE-LR4 transceivers regardless of form factor (CFP, QSFP, CXP, etc.). The Universal transceiver will also interoperate with 40G LR4-Lite, which is being widely used for 1km datacenter applications. Interoperability with 40GBASE-LR4 and LR4-lite is up to a full 500m over single-mode fiber. The multi-vendor interoperability ensures customers can avoid single supplier and proprietary lock-ins.

Conclusion

The Arista Universal transceiver is yet another example of Arista's innovation and thought leadership in addressing real customer problems and providing seamless migration paths to higher speed networks while maintaining open standards. This transceiver enables datacenters running at 10G today to seamlessly upgrade to 40G without having to re-design or modify the cable infrastructure. It also offers a transition path for customer planning migrations to single-mode fiber in datacenters with a single transceiver that bridges the gap between multi-mode and single-mode optics. The Arista Universal transceiver combined with the high-density Arista 40G switches offers the best in class performance for datacenters of any size.

[1] IEEE802.3ba, Jan. 2008, Flatman_01_0108
[2] Prices Based on IEEE link cost modeling without any discount factor
[3] Transceiver cost not included
[4] IEEE 802.3ba, 40GBASE-LR4 specification

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