Arista 7060X & 7260X Performance

Over the last decade, the industry has seen broad adoption of compute infrastructure based on commodity x86 hardware. In recent years, the power and performance offered by these servers has increased dramatically, to the point where it is no longer unusual to see hosts capable of generating above 10Gb of traffic per second. In parallel the density of servers available within a single rack is increasing, and half RU servers becoming increasingly commonplace. With the overall network traffic growing in line with server performance and density, the demand on the network for additional capacity has never been greater.

In response Arista Networks has developed the 7060X, 7060X2 and 7260X series of high performance data center switches. Designed from the ground up, specifically to deliver increased performance and port density. The 7060X, 7060X2 and 7260X series represent the optimum solution for data center architects either planning a migration to 25/50/100 Gigabit Ethernet or looking to ensure that the infrastructure is 25/50/100 Gigabit Ethernet ready.

Switch Overview

The Arista 7060X, 7060X2 and 7260X Series are purpose built, flexible Data Center switches in compact and energy efficient form factors with wirespeed layer 2 and layer 3 forwarding, combined with advanced features for software defined cloud networking. A family of fixed configuration 1RU and 2RU systems, the 7060X, 7060X2 and 7260X series support a variety of interface form factor and speed options.

Designed precisely to meet the challenges of dense 25, 50 and 100 Gigabit Ethernet switching, the 7060X, 7060X2 and 7260X series feature a flexible combination of 10GbE (SFP+), 40GbE (QSFP+, QSFP100) and 25GbE/50GbE/100GbE (QSFP100) interfaces, supporting up to 64 ports of 100GbE in a single system. The table below highlights the interface and speed combinations.

Table 1: Arista 7060X & 7260X Series overview

Switch Model / Characteristic	7060CX-32S	7260QX-64	7260CX-64	7260CX2-32
Switch Height	1RU	2RU	2RU	1RU
Ports	32 x QSFP100 2 x SFP+	64 x QSFP+ 2 x SFP+	64 x QSFP100 2 x SFP+	32 x QSFP100 2 x SFP+
Max. 10GbE Density	130	2	258	130
Max. 25GbE Density	128	-	256	128
Max. 40GbE Density	32	64	64	32
Max. 50GbE Density	64	-	128	64
Max. 100GbE Density	32	-	64	32
Max. I/O Rate (Tbps)	6.4Tbps	5.12Tbps	12.8Tbps	6.4Tbps
Max. Forwarding Rate	3.3Bpps	3.3Bpps	9.52Bpps	3.3Bpps
Latency	450ns	550ns	500-1500ns	450ns
Packet Buffer Memory	16MB		64MB	22MB
Airflow Direction	Front-to-Back or Back-to-Front			

The Emergence of 25GbE

As CPU, memory and bus technologies continue to evolve, the capacity for both virtualized and physical servers to generate more than 10Gbps of traffic is becoming increasingly commonplace. This level of performance is no longer limited to high-end custom servers, but similarly applies to many of the commodity-based products available on the market today.

Until now, this requirement has presented a significant challenge to Network Architects attempting to not only exceed the current server requirements, but also without multiplying the cost of network infrastructure. This was particularly difficult as the multiplicative cost of the network is a direct function of how higher speed technologies have traditionally been delivered.

10GbE technology is a single lane solution, in which data is run over either a single pair of fiber or a copper cable. Due to technological limitations, clocking that single lane at rates higher than 10Gbps has previously not been possible, therefore existing technologies have delivered higher speeds through parallel cables or optics. Fundamentally a 40GbE interface is comprised of 4x10GbE lanes, whether the transmission medium is a parallel cable, multiple wavelengths or multiplexed 10G signals. Whilst 40GbE technology has seen a cost reduction over recent years, the overall cost continues to be roughly 4X the price of a 10G technology, due to the nature of its parallel design. In order to meet the next generation technical and commercial challenges head on the industry needed a technology that was based on a higher clock rate, facilitating higher data rates over a single lane.

Recognizing these clear technology drivers Arista Networks was a founding member of the 25 Gigabit Consortium. The 25G Consortium is an open organization for members to enable the transmission of Ethernet frames at 25 or 50 Gigabit per second and to promote the standardization and improvement of the interfaces for applicable products.

The 25GbE technology provides an alternative to based parallel options with a more cost effective serial approach, enabling the delivery of significantly more packets per second, over a single lane. This development ensures 25GbE delivers 2.5X the performance at the same cost as 10GbE.

25GbE also provides a convenient, cost effective upgrade path to higher speeds such as 50GbE and 100GbE. Leveraging the same parallel approach that enabled 40GbE technologies, 50GbE and 100GbE are also parallel distributions of two and four 25GbE lanes respectively. By grounding this parallelism upon a higher clock rate of 25Gb per lane, 50 GbE and 100GbE provide architects with higher performance, higher density and reduced cost points. This approach enables a system based on 25Gb technology to deliver 2.5X the I/O capacity or 2X the interface density in the same form factor and at similar price points.

Data center operators looking to maximize server throughput, minimize per node connectivity costs while retaining maximum network scalability are evolving all network requirements to 25GbE based on these fundamental concepts.

In response to this demand, the Arista 7060CX-32S, 7060CX2-32S and 7260CX-64 high performance systems have been designed to facilitate 25GbE, 50GbE or 100GbE connectivity in compact form factors, providing 32 or 64 100GbE interfaces in 1RU and 2RU respectively. Each physical port can be individually configured as a single 100GbE, dual 50GbE or quad 25GbE interfaces. Based around the latest QSFP100 optic they deliver high density without compromises on system performance, scale or functionality.

Investment Protection

Many large organizations perform infrastructure migration as a phased process. There are a variety of motivating factors behind this, including operational restrictions, ongoing CapEx budgets or in many cases the sheer scale of work required. Regardless of justification, the industry has a clear requirement for investment protection and interoperation between generations of architecture.

The QSFP100 interfaces in the 7060CX-32S, 7060CX2-32S and 7260CX-64 platforms are ideally suited to meeting this requirement. While Arista QSFP100 based interfaces naturally support a versatile choice of 25, 50 and 100GbE options, the exact same interfaces also support current 10/40GbE technologies, requiring only a change in the pluggable optic. This empowers architects to design and deploy an Arista 7060X, 7060X2 or 7260X infrastructure based on 10/40GbE technologies today, then evolve over time to 25/50/100GbE connectivity. Migrating entire racks, or individual hosts as the application or business demands it, without requiring a replacement of physical switching infrastructure.

Despite clear benefits, this modular approach to migration has traditionally being limited by the business impact of migrating live services. Legacy architectures typically require a full system reload when changing the speed of an interface. This forces operations teams towards the batch migration of services in order to limit the number of system outages and required number of maintenance windows.

The Arista 7060X, 7060X2 and 7260X entirely negate this limitation by optimizing the programming of the underlying packet processor. This provides a great number of architectural and performance advantages, including the ability to allow the migration of an interface between any combination of speeds without the requirement for either a system or process reload, ensuring that the speed migration of one application has no functional or operational impact on any other services.

The 7060X Series at a Glance

As networks continue to diversify, a single location often comprises a multitude of technologies and interface speeds. The 7060X and 7260X series address that challenge by providing flexibility in both scale and connectivity. With interfaces supporting 10G, 25G, 40G, 50G and 100G Ethernet speeds, the 7060CX and 7260CX series present a wide range of options for interconnecting a variety of current or next generation infrastructures into a common leaf/spine network.

Figure 1: Arista 7060X & 7260X series overview

At their core the 7060X, 7060X2 and 7260X series are specifically developed to address the requirements for increasing server and storage interface speeds. Faster host connectivity gives rise to a network requirement for the forwarding of a significantly greater number of packets per second. In order to meet this requirement the 7060X, 7060X2 and 7260X feature an optimized forwarding pipeline supporting wire-rate forwarding even under extreme loads.

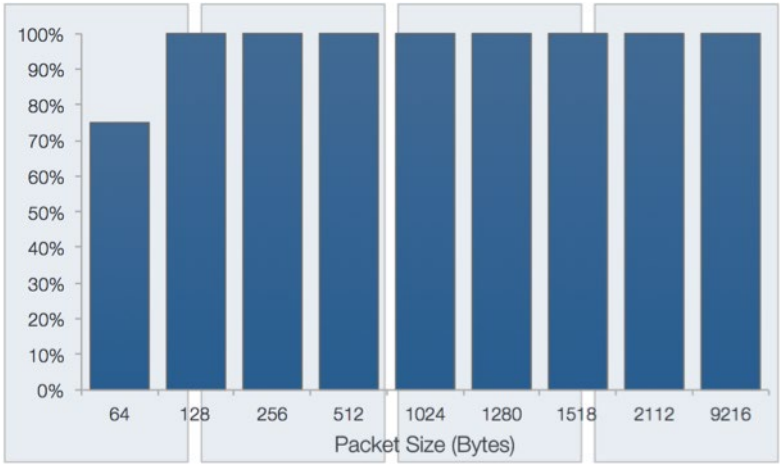

Figure 2: Arista 7060X & 7260X throughput

Numerous additional pipeline enhancements provide a variety of benefits, including low latency from 450ns and consistently low jitter under varying loads.

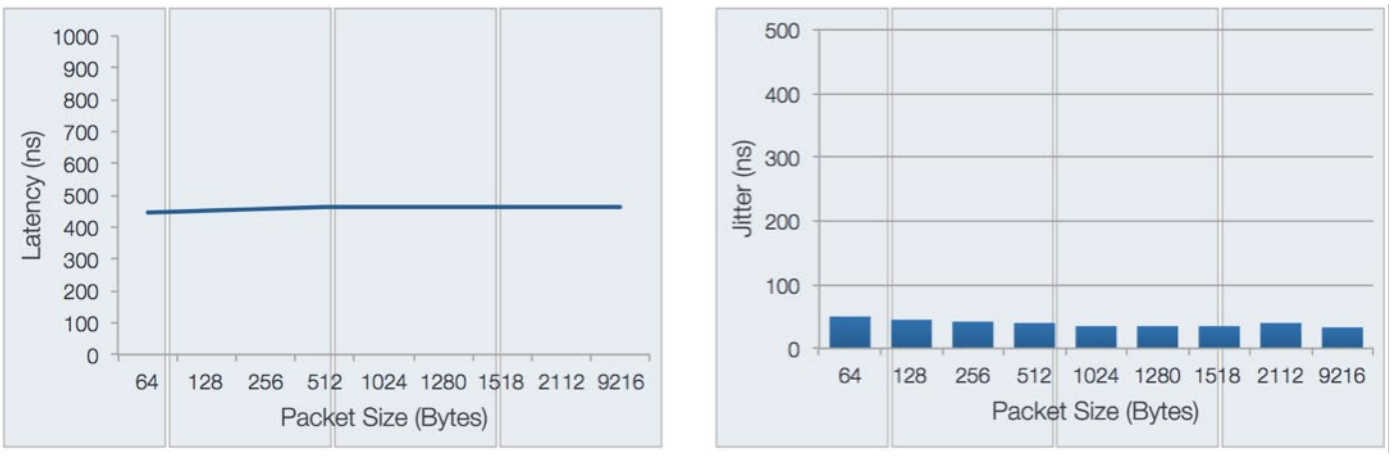

Figure 3: Arista 7060X & 7260X forwarding latency and jitter

The 7060CX-32S and 7260QX-64 offer 16MB of packet buffer, allowing the system tolerance for microbursts, avoiding one of the primary triggers of application performance degradation, dropped packets. While the 7260CX-64 has a total of 64MB of packet buffer, shared across 4 port groups.

The 7060CX-32S is a 1 RU system with 32 QSFP100 ports and 2 SFP+ ports. Each QSFP100 port can be individually broken out into 4 ports of 10GbE or 25GbE; 2 ports of 50GbE or used as a single port of 40GbE or 100GbE.

Figure 4: Arista 7060CX-32S series

Figure 5: Arista 7260QX-64 series

The 7260QX-64 is a 2 RU system that offers 64 QSFP+ ports and 2 SFP+ ports. The QSFP+ interfaces each support wirespeed 40GbE, enabling high density with a single chip design that ensures latency remains below 550ns regardless of ingress and egress port mappings, while maintaining consistently low jitter.

The 7260CX-64 is a 2 RU system that provides 64 QSFP100 ports and 2 SFP+ ports. Each QSFP100 port can be individually broken out into 4 ports of 10GbE or 25GbE; 2 ports of 50GbE or used as a single port of 40GbE or 100GbE.

Figure 6: Arista 7260CX-64 series

Figure 7: Arista 7060CX2-32S series

The 7060CX2-32S is a 1RU system that re-implements the 7060CX-32 but provides packet memory buffer that is increased by 40% to 22MB to address lossless network environments. Featuring 32 QSFP100 and two SFP+ ports, all of which can be enabled concurrently. Each QSFP100 interface can be individually broken out into 4 ports of 10GbE or 25GbE; 2 ports of 50GbE or used as a single port of 40GbE or 100GbE. The 7060CX2-32S also features full IEEE support for 802.3by (25GbE) including RS-FEC on all ports. This allows interfaces run at 25GbE to support maximum reach over multimode fiber, ensuring that 25GbE is no longer limited to an intra rack technology.

All four products share common system architecture, built on the same underlying system on chip silicon, to ensure consistent performance and feature parity. With typical power consumption of under 7W per 100GbE port on the 7060CX-32S, the 7060X, 7060X2 and 7260X series provide industry leading power efficiency coupled with power supplies that are rated at greater than 94% efficiency at typical loads, which is equivalent to platinum level. All 7060X, 7060X2 and 7260X models feature reversible airflow options providing the flexibility to ensure suitability at the top of rack, middle of row, end of row, or at the network spine layer.

The 7060X, 7060X2 and 7260X series run a single EOS binary image which is common across the entire Arista product portfolio, simplifying certification and compliance, while removing legacy challenges such as release train divergence, or platform specific inconsistency and recurring software defects. Building upon the solid foundations of EOS enables the 7060X, 7060X2 and 7260X Series to deliver advanced features for big data, cloud, SDN, virtualized and traditional enterprise network designs.

Efficient Non-Stop Operations

The latest developments of cloud computing have given rise to a new whole new generation of challenges, many of which revolve around the unparalleled demand for storage and compute capacity. While such requirements are often at the forefront of any design conversation, many more traditional concerns remain equally relevant for Network Architects.

- Power budgets continue to restrict the deployment of compute resources far more than physical footprint. The 7060X, 7060X2 and 7260X continue Arista's commitment to produce lower power draw, environmentally efficient products, which are power efficient and generate less heat avoiding costly and complex upgrades to facilities, and allowing for higher server density in each rack.
- Parity errors dominate the list of common difficulties by Operations teams globally. This is due to the entirely random nature of these errors, combined with the significant impact of a device reload (best case at a planned time, worst case in the middle of a business day). The 7060X, 7060X2 and 7260X series, as with all Arista products, includes significant custom health monitoring mechanisms for parity error detection and handling. These ensure that errors can be detected and corrected on the fly, with no discernable impact to the device avoiding system reloads.
- Yellow Zone Upgrades add increasing complexity for Network Operations teams executing maintenance tasks. Increased uptime expectations make it difficult for Network Operations teams to secure maintenance windows, with many organisations offering production time 'yellow zone' upgrades. As a result of this, the demand has transitioned away from functionally complex, device centric, 'zero loss' mechanisms, towards 'low loss' network centric approaches that are operationally and functionally simpler and give predictable and consistent results. Arista has responded to this need with the Smart System Upgrade (SSU) suite. SSU includes discrete mechanisms designed specifically for both highly redundant leaf/spine networks and environments where no redundancy is present, such as single homed servers, typically seen in cloud environments.
  - SSU Maintenance Mode - Provides standards based mechanism to isolate individual, redundant devices from the production data path. Once removed from the data path of operational tasks can be performed on the device, without any risk to the production infrastructure. When used in conjunction with Arista's Cloud Vision Portal (CVP) the operational complexity of such tasks is significantly reduced, minimizing the risk at every step of the change.
  - SSU Leaf - Hitless Upgrade - Provides a platform that allows single homed servers to maintain network connectivity while the directly connected Top-of-rack switch is upgraded. Similarly to other upgrade technologies, SSU Hitless upgrade allows the data plane to continue forwarding based on the last known good state, while the control plane is re-instantiated. However, uniquely SSU Hitless upgrade leverages standards based protocols to inform peer devices of the impending maintenance action. In addition, advanced functionality within EOS enables the data plane to act as a pseudo control-plane during the reload process. This ensures network level intelligence and stability is maintained even while the control-plane is offline. This results in predictable, repeatable upgrades with less than 200ms of application loss.

Scaling the Data Plane

Any system with the level of I/O capacity available on the 7060X, 7060X2 and 7260X series requires appropriately large hardware forwarding tables. Table scale however, is not a one-dimensional value, with individual network design use cases requiring different distributions of hardware resources such as a stronger L2 bias in an HPC or an L3 focus in large-scale cloud environments.

While many industry standard fixed configuration switches statically allocate discrete resources to specific functions with relatively small tables in some areas, such as MAC Address or IPv4 Host route tables, the 7060X, 7060X2 and 7260X series leverage the flexible forwarding table approach introduced with the Arista 7050X series.

Table 2: Arista 7060X & 7260X series Forwarding Table Scale

Table Structure	Maximum Scale
MAC Address Table	136K
IPv4 Host Routes	104K
IPv4 LPM Routes	128K
IPv4 Multicast Routes	52K
IPv6 Host Routes	52K
IPv6 LPM Routes	85K
Packet Buffers	16MB
ACLs	6000 Ingress 4000 Egress

*Roadmap Enhancement (1500/1000) supported at FCS

The use of the UFT empowers architects to define their own data plane scale, by being able to dedicate the entire UFT to expand the MAC address tables in dense L2 environments, or a balanced approach achieved by dividing the UFT evenly between MAC Address and Host route scale. The UFT can also be leveraged to support the expansion of the longest prefix match (LPM tables).

Summary

Simultaneous increases in demand in data center virtualization, social networking activity, streaming video/audio, cloud computing and online gaming are generating massive volumes of network traffic and are driving the need for rapid expansion of the server leaf to 25GbE and the spine to 100GbE. In parallel increasing demands for application availability is generating the need for non-stop operations, with 24 hour uptime, 365 days a year. These two competing challenges exert very different requirements upon the underlying infrastructure. Network Architects and Administrators are in the difficult position of identifying solutions that meet both sets of criteria.

The 7060X, 7060X2 and 7260X series have been purposely designed to solve both challenges. They deliver the raw I/O capacity and system performance required to meet the demands placed upon it by modern cloud applications, while maintaining the hardware and software reliability Arista has become synonymous for.

Through the use of EOS, an Arista solution provides accurate network telemetry data, enabling proactive monitoring of the overall health and operational state of the infrastructure at any point in time, with absolute certainty. Should the network performance or stability be compromised, Arista EOS can proactively notify the Operations staff and take corrective actions to resolve the situation before the end users even become aware an issue existed.

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