

Arista 7368X4 Series: Q&A

What are the 7368X4 Series?

The expansion of applications for machine learning (ML) and Artificial Intelligence (AI) driven by faster CPUs, flash storage and server-less compute is driving intra-datacenter traffic levels to new highs. The availability of high density 100G and 400G systems based on single chip packet processors addresses the need for increased performance, lower cost and more power efficient while retaining all the advantages of high radix network architectures with reduced leaf-spine tiers.

The Arista 7368X4 Series are purpose built 100G and 400G data center switches in a highly compact and energy efficient form factor with wirespeed layer 2 and layer 3 features, combined with enhanced traffic management and monitoring features optimized for modern cloud environments.

The 7368X4 Series provide a 2X density improvement, a 60% reduction in power per port and collapse multiple network tiers into one tier, leveraging a single 12.8Tbps packet processor, lowering end to end latency and powering the next generation of distributed storage, hyperscale cloud, AI and ML networks based on 100G and 400G Ethernet.

The Arista 7060X Series comprises the 7060X, 7060X2, 7260X, 7260X3, 7060X4 and 7368X4 purpose-built 10G/25G/40G/100G/400G data center switches.

What switch models are available in the 7368X4 Family?

The 7368X4 Series comprises the 7368X4 system and the set of IO modules:

7368X4 - 4RU SYSTEM 8 I/O MODULE SLOTS, REDUNDANT POWER AND COOLING

- Flexible interface combinations - 32x 400G, 128x 100G or 40G, 64x 200G
- A full range of 400G OSFP or QSFP-DD and 100G QSFP optics and cables
- IEEE 400G and 100G specification support
- 12.8Tbps of wire speed performance with 64MB of shared packet buffer
- Front to rear and rear to front airflow
- Low Latency and Power Efficient

7368-16C – QSFP BASED 100G MODULE

- 16 ports supporting 40/100G mode on all ports and up to 8 ports of 200G using alternate ports
- Wide range of standard optics and cables
- Hot-swap

7368-16S – SFP28 BASED 10/25G MODULE

- 16 ports supporting 10/25G mode on all ports
- Wide range of standard optics and cables
- Hot-swap

7368-4D – QSFP-DD BASED 400G MODULE

- 4 ports supporting 400G, 4x100G or 2x200G mode on all ports
- Wide range of standard optics and cables
- Hot-swap

7368-4P – OSFP BASED 400G MODULE

- 4 ports supporting 400G, 4x100G or 2x200G mode on all ports
- Wide range of standard optics and cables
- Hot-swap

The Arista 7368X4 switches were designed for continuous operations with system wide monitoring of hardware and software components, simple serviceability and provisioning. Hot-swappable power supplies and five hot-swap fans provide dynamic temperature control combined with N+1 redundancy.



Figure 1: 7368X4




Figure 2: 7368X4 I/O Modules

What are the key use cases for the 7368X4 Series?

Key features of the 7368X4 Series include support for high density and cost effective 100G and 400G with enhanced ECMP, traffic aware load balancing and scheduling, cut-through forwarding and a shared packet buffer that improves application latency for AI and storage workloads. These characteristics are suitable for a number of solutions:

Cloud Scale Data Centers

With the adoption of 25G and 50G on servers an increase of the spine network capacity is needed to accommodate the increased traffic. As a result, the primary market for the 7368X4 series is for cloud-based spine networks, when lowest cost per gigabit of bandwidth is critical, along with reductions in space, improved power efficiency and reducing the number of devices. Enhanced traffic management and multi-pathing ensure optimum performance of large scale ECMP environments.

25G and 100G HPC Clusters

In addition to 25G the emergence of 50G and 100G for host connectivity in the largest scale high performance compute (HPC) systems is driving the need for higher density 100G systems with scalable throughput, low latency and power efficiency. The 7368X4 and the Arista 7060X Series deliver high density non-blocking performance in both spine and leaf systems, with a wide range of port configurations, that allow scale-out of high performance compute for applications such as machine learning and analytics to achieve maximum performance. In these HPC environments the elephant flow detector automatically identifies larger bandwidth long lived flows and assigns to lower priority queues, allowing shorter duration mice flows to remain uncongested.

GPU clusters for deep learning, automation and machine learning

Compute clusters designed for deep learning and automation equipped with GPUs require high speed networks to shuttle data between nodes and reduce latency to maximize the performance. Leaf and Spine networks based on the 7368X4 reduces the end to end latency with a large integrated buffer to maximize the buffer available for congested ports, reducing the job-critical tail latency. Deploying high density 100G and 400G systems is a cost-effective option for the removal of network tiers, aggregating bandwidth between network tiers and migration to 400G.

In addition, Big Data analysis, Content Delivery, High Tech Enterprise and Manufacturing and SW Development environments all benefit from higher density 100G-based leaf and spine solutions and investment protection when migrating to 400G connectivity.

What are the key enhancements with the 7368X4 Series?

The Arista 7368X4 enhance both the modular 7320X and fixed 7060X series portfolio with the addition of key new technologies and features, with a consistent architecture and proven EOS features for resilient and scalable L2 and L3 networks, open programmability, automation and traffic monitoring.

Dynamic Load Balancing – uses a large number of hash results and allocates (and dynamically re-allocates) flows to new links based on current utilization.

Elephant Flow Detector – allows the tracking of large flows and re-assigns to a lower priority to avoid congestion for mice flows.

Enhanced ECMP – allows for larger scale ECMP networks with better control of traffic over all available paths, improved convergence and immediate recovery from failed links.

Shared Buffer and Traffic Prioritization – provides improved burst absorption with support for multiple lossless classes and advanced congestion controls for end to end latency improvements

The 7368X4 offer support for a consistent set of EOS features that are already supported on other Arista 7060X Series systems including Hitless Speed Changes, Smart System Upgrade, LANZ and Network Telemetry. Maintaining operational and feature consistency lowers the qualification time typically associated with introducing new switches and the 7368X4 systems seamlessly insert into existing networks. *

Is IEEE 400G standard support available on the 7368X4 Series?

The 7368X4 Series offers full support for the IEEE 802.3bs 400Gigabit Ethernet standard ensuring long term investment protection, along with support for the 100G and 200G standards for backward compatibility to existing networks. The 7368X4 supports both OSFP and QSFP-DD interfaces for 400G Ethernet.

The introduction of 400G systems provides a 4X performance improvement over 100G while using the same familiar fiber infrastructure and designs. Support for a mixture of 400G and 100G modules with breakout options allows for the ability to migrate as needed without expensive network upgrades.

What is the difference between the OSFP and QSFP-DD 400G 7368X4 I/O Modules?

The 7368X4 Series offers a choice of OSFP and QSFP-DD 400G modules as there are customers who prefer the QSFP-DD form factor, and some who prefer the OSFP. Both module types leverage the same system architecture, high performance and rich features. The two module types differ only in the type of interface, the OSFP or the QSFP-DD, and both allow for 4x100G and 2x200G mode when used with optical fiber and breakout cables and the use of existing 100G QSFP optics.



Figure 3: QSFP-DD Optics



Figure 4: OSFP Optics

Both OSFP and QSFP-DD are industry standard form factors for 400G pluggable optics and cables, with a similar set of available types. The OSFP is designed for a future higher speed connector, enabling 800G and uses an adaptor to support existing QSFP28 optics. The QSFP-DD leverages the existing QSFP28 form factor to allow the use of existing 100G optics, with no adaptor needed, but cannot easily support the future higher speeds required for 800G. The maximum thermal capacity of the QSFP-DD is approximately 5W less than the OSFP which also limits the ability to support long-reach high power optics. Both allow for shorter reaches with optics and cables, and the same number of interfaces in a system.

For more information on both the 400G OSFP and QSFP-DD optics and cables refer to Arista.com

How many ports do each of the 7368X4 series I/O modules have?

Within the 7368X4 series the I/O modules provide a choice of QSFP, OSFP and QSFP-DD 400G interfaces. The table below summarizes the interface options.

Platform	QSFP	QSFP-DD	OSFP	SFP28
7368-16C	16			
7368-16S				16
7368-4D		4		
7368-4P			4	

What speeds do the 7368X4 series I/O Modules support?

The table below shows the combinations of speeds supported on each switch.

I/O Module	100G	4x 100G Mode	200G	2x 200G Mode	400G	10/25G
7368-16C	All Ports (1-16)	--	Alternate ports	--	--	--
7368-16S	--	--	--	--	--	All Ports (1-16)
7368-4D	All Ports (1-4)	Ports 1-4	Ports 1-4	Ports 1-4	Ports 1-4	--
7368-4P	All Ports (1-4) With Adapter	Ports 1-4	Ports 1-4	Ports 1-4	Ports 1-4	--

Which cables and optics can be used in the QSFP, OSFP, QSFP-DD ports?

A wide and comprehensive range of 100G QSFP, OSFP and QSFP-DD 400G transceivers and cables are supported on the Arista 7368X4 series including direct attach copper (DAC) cables, active optical cables (AOC), multimode and single mode fiber for 100G and 400G connections.

There are a wide range of options for using both OSFP and QSFP-DD ports with breakout cables, and optics for single and multi-mode fiber to convert the 400G ports when operating in 4x100G or 2x200G modes.

The 100G QSFP ports accommodate a wide range of 100G and 40G QSFP transceivers and cables to provide support for a wide range of connectivity options from short reach copper and multi-mode fiber, to longer reaches over single mode up to 80km. For more information refer to the Arista Optic Modules and Cables Data Sheet and the Arista 400G Optics Q&A document.

How are the multi-purpose OSFP and QSFP-DD ports on the 7368X4 Series moved between 100G, 200G and 400G modes and what is the default?

The default interface speed is 400G for both OSFP and QSFP-DD ports. All 400G ports allow configuration as 4x100G or 2x200G mode. To migrate the links to a different speed, use the 'speed forced' command on the master interface.

```
7368X4(config)#interface ethernet 7/1
7368X4(config-if-Et7/1)#speed ?
```

What latency can be expected on the 7368X4 series?

The 7368X4 series support both cut-through and store-and-forward capability. The 7368X4 Series are built upon a single system-on-chip, with latency from 700ns between all ports. The port to port latency depends on the forwarding mode of the switch – cut-through or store-and-forward, and the speeds of the interfaces.

What are the maximums for forwarding tables on the 7368X4 series?

The 7368X4 series support comprehensive L2 and L3 resources optimized for hyper-scale cloud, multi-tier compute clusters and HPC deployments:

Resources	7368X4 Series
MAC Addresses	72K
IPv4 Hosts	80K
IPv4 Routes - Unicast	480K
IPv6 Routes – Unicast	300K

* Maximum values are dependent on shared resources in some cases

What EOS licenses are available and what features require them?

The 7368X4 series use the same license structure as the existing 7000 series fixed platforms.

Description	License
Virtualization feature license for Arista Fixed switches (VM Tracer and VXLAN)	LIC-FIX-3-V
EOS Extensions, Security and Partner integration license for Arista Fixed switches	LIC-FIX-3-V2
Network monitoring and provisioning feature license for Arista Fixed switches (ZTP, LANZ, API, Time-stamping)	LIC-FIX-3-Z
Enhanced L3 License for Arista Fixed switches (BGP, OSPF, ISIS, PIM, NAT)	LIC-FIX-3-E
FLX-Lite License for Arista Fixed switches, OSPF, ISIS, BGP, PIM, Up to 256K Routes, EVPN, VXLAN	LIC-FIX-3-FLX-L

NOTE: RIPv2 is supported without the Enhanced License.

For more information on Arista licensing please refer to the official [licensing page](#).

What are the key high availability options?

The Arista 7368X4 Series were designed for high availability from both a software and hardware perspective. Key high availability features include:

- Hotswap of all interface modules
- Removable switch card with no disruption to I/O modules and cabling
- Up to four N+N hot-swappable power supplies
- Five high performance N+1 hot-swap fans
- Live software patching
- Color-coded PSUs and fans
- Self-healing software with Stateful Fault Repair (SFR)
- Smart System Upgrade (SSU) Leaf and Spine *
- Multi-chassis LAG for active/active L2 multi-pathing
- 128-way ECMP routing for load-balancing and redundancy

What is the power draw on the 7368X4 series?

The 7368X4 series feature low power draw, with typical per port power lower than 10W per 100G port. The actual power draw will be dependent on the type and quantity of 100G and 400G optics and cables.

What efficiency rating do the power supplies have?

The AC power supplies are rated at over 93% efficient for typical use, or Platinum rated with a maximum output of 2kW

Do the 7368X4 series support both AC and DC PSUs?

Yes, both AC and DC power options are supported. AC power options are available immediately, and a DC power option will be added.

What are the options for support?

Arista A-Care Service Options are designed to provide you with world-class support. A-Care service offerings are available 24x7x365 with advance replacement options to minimize any network downtime. All A-Care Service options include full access to bug fixes and software downloads. For more information about A-Care Service options go to <http://www.arista.com/en/service>

Where do I get more information on the Arista 7368X4 series?

For more information please go to www.arista.com or contact us at sales@arista.com