

Arista 7150S Series: Q&A

Contents

Product Overview	1
Advanced Features	3
Precision Time Protocol	3
Port Mirroring	4
Latency Analyzer (LANZ)	5
Time-Stamping	6
Network Address Translation (NAT)	6
Wire-speed VXLAN Gateway	7
AgilePorts	8
General Information	9

Product Overview

Q. What are the unique advantages of the Arista 7150S Series?

The Arista 7150S Advantages:

- First SDN Switch architected for leading-edge applications including Big Data, Cloud Networks, Financial Trading, HPC and Web 2.0 environments
- Industry leading accurate and predictable performance in a range of densities
- Unprecedented balanced resources and deployment flexibility for "Any Application" suitability
- First network-wide virtualization platform for next generation cloud bursting support with wire-speed VXLAN hardware-based Tunnel Endpoint termination
- First platform to support sub-microsecond Network Address Translation (NAT)
- First deterministic latency and high-performance 10GbE/ 40GbE switch and IEEE 1588 Platform
- Advanced Forensics and Monitoring capabilities – redefining instrumentation, automation and analysis of high-end infrastructure

Unique Features:

Feature	Benefits
Wire-speed Low Latency Network Address Translation	Reduce NAT latency by 10s of microseconds vs. traditional high latency solutions
IEEE 1588 Precision Time Protocol (Boundary and Transparent modes)	Provides hardware-based timing for accurate in band time distribution with nanosecond accuracy
Integrated High Precision Oscillator	Ensures highly accurate timing with extended holdover

Latency and Application Analysis LANZ+	Detect, capture, stream microbursts and transient congestion at microsecond rates
Advanced Multi-port Mirroring Suite	Avoid costly SPAN/TAP aggregators with in-switch capturing, filtering and time-stamping
Wire-speed VXLAN Gateway	Enabling next generation Data Center virtualization
AgilePorts	Adapt from 10G to 40G without costly upgrades

Q. What are the different models in the Arista 7150S Series?

The 7150S Series provides three models with 24, 52 and 64 10Gb ports. Each model supports wire-speed performance at both layer 2 and layer 3. Both the 7150S-24 and 7150S-52 support SFP+ on all ports. The 7150S-64 has 48 SFP+ and 4 QSFP+ ports. The SFP+ ports support 1G and 10G operation and also 100Mb support with 100/1000-TX transceivers. The QSFP+ ports support 40GbE and 4x10GbE modes.

7150S Switches			
	7150S-24	7150S-52	7150S-64
Description	7150 Switch 24 Port SFP+	7150 Switch 52 Port SFP+	7150 Switch 48 Port SFP+ 4 Port QSFP+
Maximum 10GbE	24	52	64
SFP+ Ports	24	52	48
40GbE Agile Ports	4	13	16
L2/L3 Throughput	480Gbps	1.04Tbps	1.28Tbps
L2/3 PPS	360Mpps	780Mpps	960Mpps

Q. What are the 7150S improvements over existing 7100 Series devices?

The 7150S Series maintains feature consistency with the 7100S and 7100SX Series and provides a number of enhancements:

- Reduced latency by as much as 66% compared to the 7148SX
- Increased port density up to 64 ports of 10G and support for 40G interfaces
- AgilePorts to enable 4 x SFP+ to be grouped into a single 40GbE interface to support mixed 10G and 40G
- Single ASIC architecture for consistent latency between any ports in the system
- Larger table sizes for greater scalability and wider applicability in large data center networks
- Larger dynamic buffers available to congested interfaces
- Hardware support for many new features (NAT, PTP, VXLAN, Advanced Mirroring, Time-stamping)
- Support for a flexible programmable pipeline that enables new features like VXLAN to be rapidly released

Q. What are the focus markets for 7150S Series?

The 7150S is a robust any-purpose data center switch. Key markets that will realize value from its enhanced feature set are:

- Financial Services: Lowest latency, high density, large tables, NAT, precision timing and application monitoring
- Big Data: High performance, high density, flexible 1G/10G/40G, advanced monitoring and time-stamping
- HPC (High Performance Compute) and Research: Lowest latency, high density, precision timing, precision monitoring, and support for flexible 10G and 40G.
- Virtualized and Large Scale Data Centers: High density, large tables, VXLAN, and line rate performance with L2 and L3 features enabled, multi-port mirroring for increased visibility and monitoring of applications.

Q. Which software licenses are available and which features require licenses?

There are three licenses available for the 7150S series, with options for the smaller density 24 port and the larger density 52 and 64 port systems. The 7150S Series use the same license structure as the 7100 Series and the 7050 Series.

- The Enhanced Routing License enables dynamic unicast and multicast routing protocols - OSPF, BGP (v4 and v6) and PIM as well as Network Address Translation (NAT). RIPv2 is supported without the Enhanced License.
- The VMTracer License is required for both the VMTracer feature and future VXLAN features.
- The Network Monitoring and Provisioning (Z) License is required for ZTP, LANZ and LANZ+ and Advanced Monitoring (future) features.

Advanced Features:

Precision Timing Protocol - PTP (IEEE 1588)

Q. Why is supporting PTP (IEEE1588) in the switch important?

Precision Time Protocol (IEEE1588) is an alternative to the popular Network Time Protocol (NTP), a means of using the existing IP network infrastructure to distribute highly accurate time-of-day enabling multiple devices to synchronize their clocks to a central source (the Grand Master (GM)).

Traditionally, high precision environments (e.g. HPC, Finance) have deployed dedicated time distribution networks, which consist of an overlay network of co-axial cabling and dedicated hardware, required in each client machine. These networks are expensive to scale and require significant additional cable infrastructure.

In financial exchange co-location facilities, the limited availability of accurate timing leads to each firm building their own time distribution infrastructure, usually requiring a roof mounted GPS antenna (per client) which is also expensive to maintain and scale.

The ability of PTP to offer scalable, hierarchical in-band time distribution is very attractive to simplify deployments, lower costs and limit GPS antenna sprawl.

Q. What is different about Arista's Hardware PTP implementation?

While software implementations of PTP can provide improvements over NTP, it is with hardware support that substantial accuracy gains are made. Deployment of PTP in hardware may improve upon NTP's accuracy by several orders of magnitude - from 10-100 μ s down to 100s of nanoseconds.

The 7150S Series uses ASIC based timing capabilities to support both Transparent Clock and Boundary Clock modes. The unique onboard high precision oscillator further improves accuracy in several ways:

- 1000x Better Frequency Stability - a better quality, more consistent tick ensures better accuracy, less variation and allows downstream devices to minimize their offset from the master clock
- 3000x Better Resistance to Environmental Fluctuations which lowers the sensitivity to voltage and temperature variations
- 1000x Better Hold-over to ensure that when no upstream master clock is available the onboard clock keeps "ticking" consistently accurate to maintain an accurate time.

Q. Why is this different from Switches with PTP offerings?

Outside of special purpose Industrial Ethernet solutions, there are no high-performance 10G/40G devices supporting hardware based Transparent Clock and Boundary Clock operation. Current solutions use limited software solutions that are unable to achieve the very high levels of accuracy demanded by both Financial, Scientific and Federal customers. The Arista 7150S is the first to offer both native hardware PTP and a high precision oscillator in a 10G/40G device.

Q. Do the PTP and time-stamping features require the high precision clock module?

Both PTP and the time-stamping features will work without the use of the clock module, which is an optional item on the 7150S-24. The high precision clock improves the accuracy of the timing features over time, and reduces sensitivity to temperature variations.

Port Mirroring

Q. How many SPAN sessions are supported on the 7150S Series?

The 7150S Series supports 4 discrete mirroring sessions, each able to support TX, RX or both with multiple source interfaces.

Q. Does the 7150S capture real TX frames when TX mirroring is enabled?

Unlike the 7100 Series, the 7150S captures the actual TX packet when using TX mirroring.

Q. Is it possible to mirror a set of sources to multiple monitor destinations (i.e. to multiple ports)?

The 7150S Series supports this capability and the feature is planned for a future release of EOS software.

Q. Is it possible to have a single monitor source mirrored to multiple destination sessions?

The hardware supports this capability and will be offered in a future release.

Q. What is the Advanced Mirroring suite?

Advanced Mirroring is a group of expanded mirroring features that include the following:

- Enhancements to regular SPAN
- Filtered/ACL based mirroring
- Packet truncation on mirror egress
- Hardware packet time-stamping of mirrored traffic
- Capability to act as a SPAN/TAP aggregation device – expanding on the current capability to operate a 7124SX as a SPAN/TAP aggregator: <https://eos.aristanetworks.com/2011/05/153/>
- Arbitrary N:M mirroring

Latency Analyzer (LANZ)

Q. What is Latency Analyzer (LANZ)?

LANZ (Latency Analyzer) is an advanced monitoring feature available on Arista 7100 (7124S(X), 7148S(X), 71xxT) and Arista 7048/7500 platforms. It has two modes of operation based on target use-cases:

- Arista 7100 series devices LANZ provides real-time monitoring of the depth of queues on an egress interface allowing a user to instrument even very small, very brief congestion events (e.g. microbursts) that might not result in packet loss
- Arista 7048 and 7500 series LANZ provides the ability to track utilization of deep VOQ buffers, providing the user understanding of which destination interfaces are most congested and how much data is being en-queued and its sources.

Q. What is Latency and Application Analyzer Plus (LANZ+)?

LANZ+ refers to the enhanced LANZ features available with the Arista 7150S series of devices. LANZ+ has the following advantages over LANZ on Arista 7100 switches:

- **Higher Granularity** - While all variants of LANZ are trigger driven (rather than sampled or polled), LANZ+ can dynamically adjust the trigger rate depending on load allowing very fast repeat triggers, over 10x faster than LANZ
- **Per Queue per Port Reporting** - LANZ reports on a per port basis, LANZ+ allows reporting and

configuration for each traffic class on every port

- **Accurate Time-stamps** - High granularity timestamps for LANZ+ events are available using hardware assisted timing rather than the system clock
- **Congestion Lifecycle Monitoring** - LANZ+ tracks congestion events end-to-end, providing Start-Time, Duration and Maximum Queue Length
- **Drop Monitoring** - LANZ+ is able to report on lost traffic occurring through severe congestion events
- **Global Buffer Monitoring** - LANZ+ reports on shared buffer usage in addition to interface queues
- **Data Capture** - LANZ+ allows the user to capture packets that were involved in congestion events, with accurate timestamps

Time-Stamping

Q. What is Time-Stamping?

Time-stamping is the act of adding a field to a packet to indicate the time at which the packet was received at the device.

Traditionally network analyzers add time-stamps as they receive copies of traffic from SPAN ports or networks taps. This is inherently inaccurate because there may be tiers of additional devices between the analyzer and the traffic source. The analyzer has no method to measure how much additional delay the intermediate devices and infrastructure have introduced and how this varies due to queuing.

Adding a time-stamp at source is the most reliable method to ensure the analyzer has an entirely accurate view of when the traffic was captured.

This type of functionality has only been available in proprietary, low-density platforms. The Arista 7150S Series are the first to bring this to a switching platform for up to 64 x 10G ports at a fraction of the cost.

Several tool and NIC vendors such as TS-Associates provide support for decoding 7150S timestamps to enhance the accuracy of their applications. Arista networks are committed to open interoperability and are working with third parties to help them integrate time-stamping.

Q. How does the precision time option work?

The precision time option provides a high-resolution clock on the EOS control plane. This allows the applications running in EOS to have a precise time value for instrumentation of application events on the switch and FPGA subsystem.

Network Address Translation (NAT)

Q. Why do we need NAT?

NAT is a feature commonly found in modular routers and firewalls, so what is the benefit of having NAT capability in a low-latency fixed configuration data center switch?

There are two major use cases for NAT:

Financial Services - Exchange Connectivity

When cross connecting with a financial exchange, it is typical for the exchange to provide a dedicated (private) subnet from which they will accept incoming connections for order entries and repeat requests. The size of the provided subnet varies considerably by venue but may be restricted to /27 or perhaps less.

With only a small number of source addresses available, customers are restricted to the number of hosts that can actively participate in a venue. If the firm is a broker with large numbers of sub-customers, they must use some form of NAT to consolidate and translate their customer subnets to venue friendly addresses.

A further use case exists where a firm wishes to trade on two (geographically distributed) venues simultaneously with the same servers. If the server must have an interface in both venues' private subnets, this typically requires either a stretched VLAN connection or some form of NAT.

The typical solution is to deploy a firewall or high-end router with onboard NAT functionality, however these types of platforms may add up to 100µs of extra latency. The ability of the 7150S to perform NAT represents an enormously compelling latency saving.

Cloud-Bursting

One way of maintaining subnet coherence when turning up additional contingent DC capacity is to leverage NAT to mask the remote DC behind local IP addresses. This allows all local services (SLB, FW, etc.) to treat remote servers in the same way as their existing local pool of devices.

In this case NAT compliments VXLAN for scenarios where the remote capacity may be physical only or may not support VTEP (VXLAN Tunnel End Point).

Wire-speed VXLAN Gateway

Q. What is VXLAN?

VXLAN is a tunneling mechanism that runs between virtual or physical switches and enables applications to be deployed and moved between any server within the Data Center regardless of IP subnet or physical host location.

This enables IT departments to dynamically scale network architectures to support capacity on demand and workload mobility regardless of geography and existing IP addressing. VXLAN also overcomes scalability and segmentation limitations that exist within today's data center protocols such as TRILL, Spanning Tree, and others. Unlike other network virtualization overlay models, VXLAN uses proven IP protocols and requires no change to the underlying IP architecture or existing data center infrastructure.

Q. How does VXLAN work?

VXLAN is a network encapsulation and segmentation protocol enabling applications to be deployed on any server, on any network, at any time. It accomplishes this by encapsulating the application's MAC and IP packets within a UDP header and utilizing IP multicast groups to emulate broadcast domains.

VXLAN encapsulation and de-capsulation is done at the Virtual Tunnel End Point (VTEP) located at the virtual or physical edge of the network. Because VXLAN networks are not bound by IP subnets or L2 boundaries, VXLAN can utilize the existing L3 network topology.

Q. What advantages does VXLAN provide over building large L2 networks?

Large Layer 2 broadcast and failure domains can be eliminated and traded for more stable L3 networks supporting greater scale, better multi-pathing and millisecond convergence. In addition, previous scalability limitations due to MAC address table exhaustion and limited VLAN tags (4K VLANs) are replaced with the VXLAN header allowing for up to 16 million customer segments.

Q. How can I deploy VXLAN?

There are two fundamental ways that VXLAN can be deployed:

- Virtual switch (in software, distributed across physical hosts)
The virtual switch approach is supported by some virtual switches (VMware Distributed vSwitch / vSphere 5.1, Open vSwitch, others..) with between a 10-30% performance impact on the server (ref: <http://www.vmware.com/files/pdf/techpaper/VMware-vSphere-VXLAN-Perf.pdf>)
- Physical switches (hardware VXLAN gateways)

The second example is provided on Arista 7150S series for hardware VXLAN gateway functionality.

Q. What VXLAN functionality will be available on the Arista 7150S?

Arista 7150S switches will provide Hardware VXLAN Gateway functionality. Specifically this is the ability to gateway a VLAN to/from a VXLAN Tunnel Interface (VTI) at line rate, acting as a Virtual Tunnel Endpoint (VTEP) and participating as a VXLAN Segment with both hardware and software based VTEPs.

AgilePorts

Q. What are "AgilePorts"?

AgilePorts are a unique feature of the 7150S Series that allows the user to configure adjacent blocks of 4x SFP+ interfaces as a single 40GbE link. This allows the 7150S-24 to provide up to 4x40G and 8x10G (using 16 SFP+ for 40G with the remaining SFP+ ports offering 8 x 10G ports). On the 7150S-52 this feature allows up to 13x 40G (all 52 SFP+ ports grouped into 40G ports) and on the 7150S-64 AgilePorts allows the switch

to be deployed with up to 16 native 40GbE interfaces where 4 40G ports are the QSFP+ and the additional 12 40G ports use the 48 SFP+ in 4 port groups.

Q. How do you use AgilePorts?

From the CLI first configure AgilePort mode on a group of interfaces. Then install the same flavor of SFP+ in each required cage and wire all four links in order to another AgilePorts capable device or in the case of 10GBASE-SR(L) via LC-MTP cables to a 40G-SR4 device. 10GBASE-CR can also be utilized in the same way (10G CR back to back or 10G CR to CR4).

Q. Are AgilePorts compatible with normal 40GbE links?

As 40GBASE-SR4 and 40GBASE-CR4 are based on 4 x 10G compatible sub-links, both are compatible with AgilePorts when deployed with 10GBASE-SR or SRL transceivers, and 10GBASE-CR copper cables.

General Information

Q. Is this a store and forward or cut-through switch?

The Arista 7150S switch utilizes an ultra-low latency cut-through switch architecture, which provides deterministic ultra low latency for all packet sizes. The latency does not change even when additional features such as L3 forwarding, L4 inspection, ACL, QoS, Multicast or Port Mirroring functionality are enabled. The 7150S also forwards packets in cut-through mode at 1GbE speeds at low latency for legacy connections.

Q. What is the power draw for 7150S series devices?

The power draw both typical and maximum for the three 7150S Series models is shown below. Typical power is measured at 25°C ambient with 50% load on all ports. Maximum power is specified for all ports running at line rate and the maximum specified operating temperature.

7150S Model	Typical Power	Maximum Power
7150S-24	191W	334W
7150S-52	191W	450W
7150S-64	224W	455W

Q. Does the 7150S have AC and DC power Options?

The Arista 7150S switch supports both AC and DC power options that maximize the deployment options and provides high availability with hot-swappable capability.

Q. What are the high availability options?

The Arista 7150S switch was designed for high availability from both a software and hardware perspective.

Key high availability features include:

- 1+1 hot-swappable power supplies and four N+1 hot-swappable fans
- Color-coded PSUs and fans common to Arista 1RU devices
- EOS Zero Touch Provisioning (ZTP)
- Self-healing software with Stateful Fault Repair (SFR)
- Multi-chassis LAG for active/active L2 multi-pathing
- 32-way MLAG and ECMP routing for all-active L2 and L3

Q. Which cables and optics can be used in the SFP+ ports?

All currently supported transceivers, with the exception of LRM, are supported on the Arista 7150S SFP+ ports. The SFP+ ports accommodate a full range of 10GbE SFP+ and 1GbE SFP 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 40km and DWDM solutions up to 80km. The SFP options include multi-mode and single-mode fiber transceivers, and both 100Mb and 1Gb over copper cabling.

Q. What is the minimum EOS software version for the Arista 7150S-24?

The minimum version of EOS that supports the Arista 7150S is 4.10.0-7150. This image adds support for the 7150S Series. Future EOS releases will also support the 7150S with a single image for all Arista platforms.

Q. 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.aristanetworks.com/en/service>.

Q. Where do I get more information on the Arista 7150S?

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