# Exploring the 7150S Family

Demanding networking environments like financial trading systems, high performance computing (HPC) clusters, virtualized large-scale data centers and cloud hosting and application platforms have forced IT managers to reconsider the functionality required from their network platforms. In these deployments, traditional general-purpose network devices have fallen short in meeting the performance, topological, and analytical requirements.

For early adopters it was often necessary to balance the limitations of platforms that focused solely on latency or performance with more traditional enterprise hardware and expensive overlay infrastructure to support granular monitoring, analytics, and time synchronization. The Arista 7150S provides a solution to this dilemma, providing both deterministic ultra low latency, line rate performance and a full suite of tools for high performance environments, including network address translation, precision timing, application performance monitoring and rich management capabilities. This paper describes the Arista 7150S and its unique features.

The Arista 7150S is family of low-latency 1/10/40Gb Ethernet switches packaged in a compact 1-RU form factor. With three available models – 24-, 52- and 64-port – the 7150S sets the benchmark for the new class of data center switches, providing the lowest latency performance while adding additional functions typically seen in larger modular network switches. It provides an extensive toolkit for instrumenting and managing the highest performance environments.

The 7150S is a new breed of mature, full-featured data center devices that can be deployed into many areas in the network – at the edge of the data center, the spine of a mid-sized data center, in ultra low latency (ULL) trading and High Performance Compute (HPC) environments, or in a virtualized cloud facility.

*Leveraging Arista's next generation cut-through architecture, the 7150S provides low and consistent latency and jitter between all ports in the system and regardless of packet size and traffic type.*

## Latency and Cut-Through Technology

Routing and switching performance forms the foundation for the 7150S family. Leveraging Arista's next generation cut-through architecture, the 7150S provides low and consistent latency and jitter between all ports in the system and regardless of packet size and traffic type.

Minimizing jitter, the variance in latency between multiple packets in a stream of data, goes hand in hand with true deterministic performance – a crucial requirement for demanding modern applications where applications are increasingly sensitive to variability and packet loss.

In applications where multicast data sets make up a substantial proportion of traffic loads, there is a further dynamic. It is crucial that packets are not just forwarded quickly but also that they are replicated in a fair and consistent fashion to all subscribers. Traditional iterative replication technologies do not address this need - instead introducing a skew in the generation of each packet copy, which provides a performance advantage to certain ports over others. The innovative zero-copy architecture eliminates this unfairness and ensures that all recipients receive their data simultaneously.
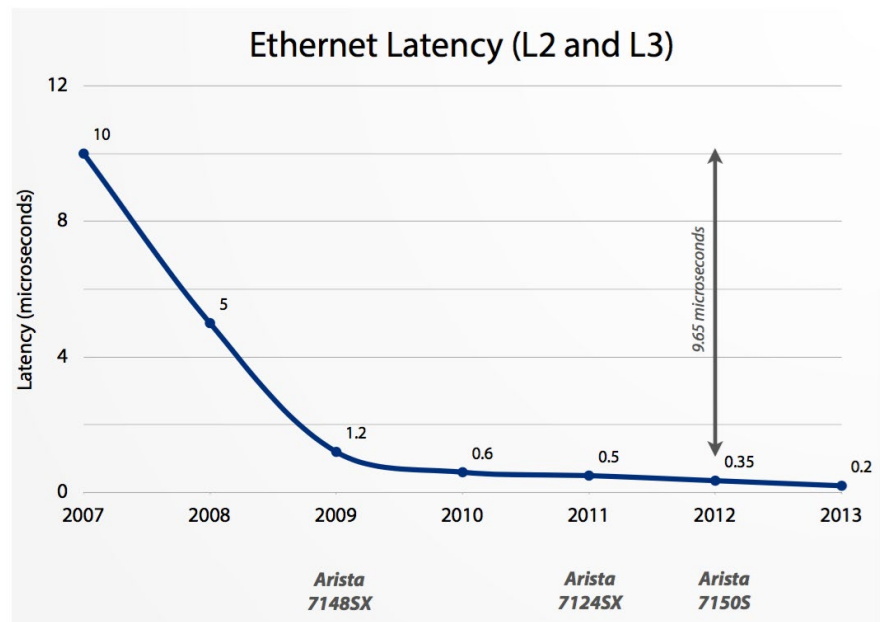


*Figure 1: Lowering latency with cut-through technology*

Figure 1 shows how cut-through forwarding has decreased the latency bar for high-density 10G switches. Starting in 2009, where a regular switch might introduce 10us per hop, Arista Networks 7148SX reduced this to 1.2us – an 88% reduction for 48 x 10GbE solutions.

The 7150S builds on this by increasing density to 64 x 10GbE while lowering latency a further 70% to 350 nanoseconds.

## Latency and Cut-Through Technology

The 7150S is the first data center platform to integrate hardware-enabled precision timing capabilities directly into the forwarding plane, enabling a host of advanced

functionality to support high performance applications and the integration of monitoring features that have traditionally required additional tiers of expensive, proprietary devices.

Central to precision timing is the ability of the ASIC to apply a timestamp to every packet entering the device with nanosecond granularity, at wire speed with no latency or performance penalty.

To ensure the most accurate timing, a high- precision oscillator typically found in dedicated time distribution platforms provides an accurate and stable time signal to the 7150's subsystems. This high precision oscillator minimizes exposure to variation through environmental influences such as voltage and temperature fluctuations and ensures excellent holdover characteristics in the event of a loss of upstream synchronization.

The precision timing capabilities of the 7150S provide the foundations for numerous features, including Precision Time Protocol (IEEE 1588- 2008), LANZ and the hardware time-stamping capabilities integrated with EOS Advanced Mirroring and EOS TAP Aggregation.

### Enhanced Functionality

Many latency-focused devices make significant compromises in functionality, flexibility and feature parity to achieve low latency, with some specialized solutions providing only a fraction of the feature-set or scalability of a traditional enterprise device. For many applications, simply removing functionality to get lower latency is a false economy because those functions must still exist somewhere in the environment, typically at a point that undermines the latency savings made at the network edge.

The foremost goal is to remove the need to choose between performance, determinism and rich functionality – to unify all of these requirements in one device.

To fully address the needs of enterprises seeking platforms to address emerging applications, the device must deliver on five fronts:

- Low consistent latency
- High port density
- Enterprise scalability
- Open, software defined programmability
- Precision instrumentation

The 7150S addresses all of these requirements, providing support for many different functions that would otherwise require discrete devices. In one rack unit, the 7150S provides 64 10GbE ports with ultra low latency, a rich Layer 2-4 feature set and advanced monitoring functionality that enables enterprises to standardize on a common product family, reducing total device count and the operational costs associated with traditional enterprise platforms.

*With the 7150S Series, users don't have to make this compromise – providing much better latency and a selection of interface configurations up to 64 x 10GbE (or 16x40GbE).*
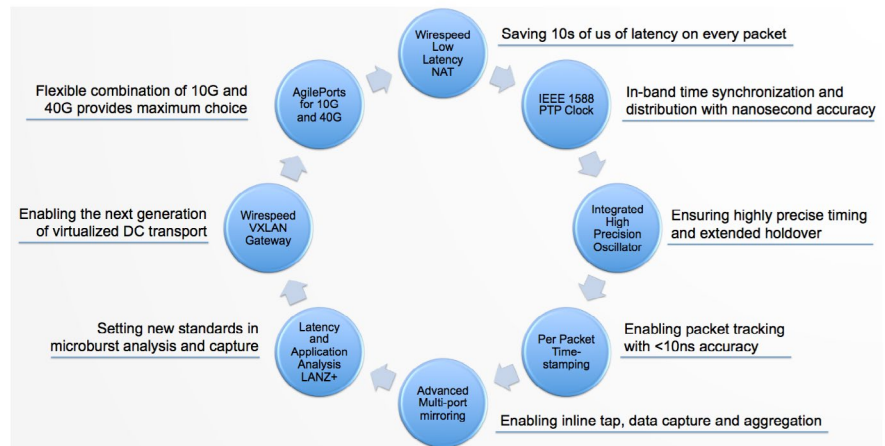


*Figure 2: 7150S Advanced Functionality*

## Lower Latency

The 7150S delivers 350 to 380 nanoseconds latency (350 for the 24-port variant and 380 for 52- and 64-port variants). 350 nanoseconds represents a 25% reduction in latency from the Arista 7124SX with an increase in functionality. The 380 nanoseconds in the 52- and 64-port variants is a 70% saving over the outgoing generation of 48-port devices.

Figure 3 shows the very consistent latency attributes that the 7150S family provides for all packet sizes. By contrast, the red line at the top shows the typical latency profile for a general- purpose data center device – as the size of the packets grows, the latency increases.

This is typical for many comparable products – users must make a trade-off between very good performance but typically lower port density and much less deterministic, worse latency. The lower three lines show that with the 7150S Series, users don't have to make this compromise – providing much better latency and a selection of interface configurations up to 64 x 10GbE (or 16x40GbE).
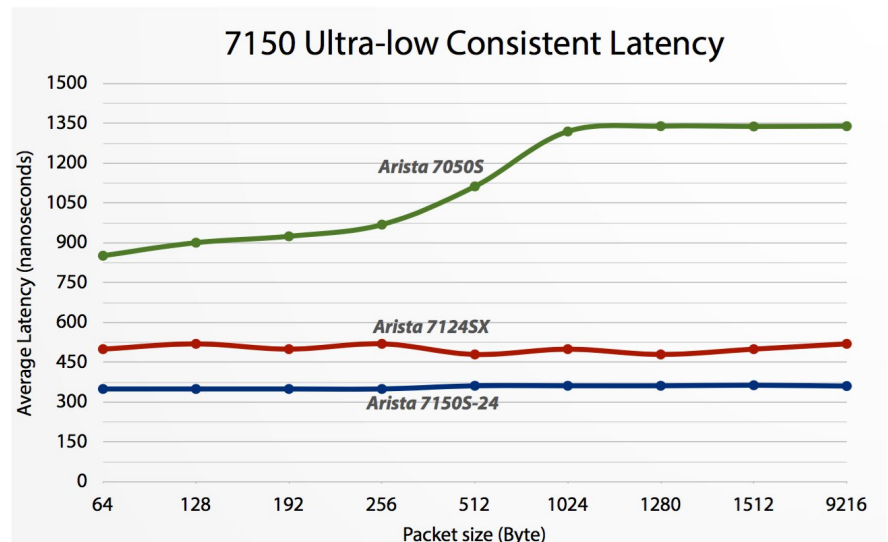


*Figure 3: Comparative latency profile for Arista 7000 Series*

*With the 7150S taking over more complex functions it is possible to drastically reduce the network footprint, removing larger modular devices and recovering the physical space and high power consumption that they would normally require.*

## Scability and Density

Traditional chassis-based switches are expensive to purchase, operate and upgrade but have typically been necessary to support the protocol scaling requirements of larger networks with larger tables and unique L3/4 and encapsulation capabilities.

The 7150S achieves the combination of low latency with the scalability and density expected, in a significantly more cost effective and compact format. Its well-balanced resource allocation, which represents up to a 12X table scale improvement for Layer 2 and Layer 3 unicast and multicast features, provides users significantly more headroom and the ability to deploy these more compact devices in the network where traditional modular systems may have previously been required. (Figure 4)

With the 7150S taking over more complex functions it is possible to drastically reduce the network footprint, removing larger modular devices and recovering the physical space and high power consumption that they would normally require.

## Agile Ports for Investment Protection

While deploying high performance infrastructures, one must be mindful of the increasing requirement for 40GbE connectivity. In addition to the network hardware itself, 40GbE brings a number of new challenges to the data center as its deployment typically requires considerable investment in new cable plant design.

The 7150S with Arista's AgilePorts capability is designed to address device and cabling life cycles, providing a transparent method to upgrade from 10GbE to 40GbE without replacing cabling or transceivers (Figure 4).

AgilePorts provides a means for customers to protect their investment and use flexible combinations of interfaces and interface speeds. AgilePorts removes the need to invest in multiple devices with a fixed number of 1, 10, or 40GbE ports; instead offering a seamless path to migrate from 1GbE to 40GbE in a single platform.

AgilePorts combines four existing SFP+ ports into one native 40GbE interface, instead of requiring a dedicated QSFP+. Simple EOS CLI commands enable users to quickly provision native 40GbE interfaces on existing hardware and parallel fiber-optic or direct-attach cables.



*Figure 4: Agile Ports*

AgilePorts interoperates with another device with similar functionality or, with the right cabling a native 40GbE QSFP+ device.

Through the AgilePorts feature the 64-port 7150S can support up to 16 40-Gigabit ports. With one device the customer maintains the investment in cabling, the switch, and the transceivers - avoiding the need to replace the cabling infrastructure, which in many cases is more expensive than the switch.

| Table 1: 7150S Scaling Data | |
| --- | --- |
| **Feature** | **Scale** |
| 40 Gb Interfaces | 4, 13, 16 |
| Switching Capacity | 1.28 Tbps |
| Forwarding Rate | 960 Mpps |
| L2 Table Size (unicast / multicast) | 64K / 36K |
| L3 Table Size (unicast / multicast) | 84K / 23K |
| ACL Scale | 20K |

## Optional Integrated SSD

The 7150S supports the use of an optional integrated 50GB SSD drive that facilitates a number of unique functions including:

- Storing DANZ data, packet captures and system logs on the switch

- Installing custom packages and applications

- Storing extended system logs

- Storing performance data

- Provisioning networks services, tools, and deployment applications like DHCP , Chef/Puppet, Wireshark, and NTP at the top of rack

## Software Defined Networking and Advanced Data Center Technologies

### Open programmability

The open programmability of the 7150S is founded in EOS, a Linux-based operating system that brings the power, stability, and modularity of Linux to the network device. EOS is inherently extensible with is modular agent architecture and separation of process and state. This architecture provides a platform for open integration with third party agents, rapid inclusion of new technologies or features, and operational flexibility by applying the same modular software image to form factors from a VM to a large modular switch.

Many organizations want to manage their data center environments without having to invest in proprietary and costly vendor specific tools, and integrate their network into the open application platforms that they already use for orchestration. The expectation is to apply the same open implementation and operational process and policies to the network as they do servers, virtual machines, and applications, avoiding getting "locked-in" to specific architectures. The 7150S provides the open programmability required to enable the use of a broad range of open source and commercially available tools that interact directly with the switch, or are embedded directly into EOS.

### VXLAN

It is increasingly important that networks not be vertically integrated, as current trends have moved away from single-vendor stacks and one application per server, and are moving toward an age of applications that are experiencing massive consolidation. In the virtualization-heavy environments, traffic behavior in the

*The 7150S is the first network device to provide support for VXLAN functionality in hardware and provides EOS-enabled extensions that directly link to VMware for insight into the VXLAN environment.*

network is increasingly "east-west" between application servers, databases and network storage rather than "north-south" between clients and servers. Network architectures are forced to evolve to effectively address the east-west performance and scalability.

VXLAN is a forthcoming standard for improving scaling and virtual machine mobility. It is co- authored by industry leaders including Arista and VMware. The 7150S is the first network device to provide support for VXLAN functionality in hardware and provides EOS-enabled extensions that directly link to VMware for insight into the VXLAN environment.

Modern highly virtualized environments face a challenge – they are highly reliant on the underlying network infrastructure's topology and addressing schema, while generally decoupled from the actual provisioning and reporting of the network.

The ability to create virtual machines and move them around the infrastructure is a sophisticated process from a software perspective. Many of today's orchestration tools rely heavily on the network infrastructure to have large numbers of VLANs deployed spanning the network at all points – a design that itself is rooted in legacy network thinking and presents many serious challenges for both scalability and management.

First, scaling can be an issue – not all third party switches support full VLAN scalability, and even if they do, large multi-tenant environments will consume large numbers of VLANs with the risk of exceeding the maximum number of 4096 VLAN IDs.

Second, large numbers of virtual machines in Layer 2 networks requires each switch to learn a substantial amount of MAC addresses. This has led to a class of devices that have unbalanced table capacities. They can store a large number of MAC addresses, but are unable to provide L3 routing services to all of the downstream devices because the network resources are biased towards a pure Layer 2 environment.

Finally, larger Layer 2 networks are best avoided due to the risk of fault propagation and inefficient flooding and broadcasting of data. This can be partially mitigated by intelligently pruning idle VLANs with functionality such as Arista's VMTracer but needs to be addressed from the ground up to expand further and enable the next generation of scalability.

VXLAN provides a robust solution to all of the issues above. It abstracts the traffic flowing between virtual machines away from the L2 network topology (Figure 5/6). Instead of requiring Ethernet bridging between virtual machines in the same domain for mobility services, the bridged traffic is encapsulated in IP datagrams and unicast between virtual switches or other tunnel endpoints, which can be software devices such as hypervisor vSwitches or physical endpoints such as the Arista 7150S.

The 7150S provides the capability to provision a VXLAN and bridge it to a well understood 802.1Q VLAN to provide backwards compatibility with devices that don't support VXLAN. This enables seamless interaction at line rate for existing

*The Arista 7150S is the first hardware device that can translate between the VXLAN environment and the VLAN world. [...] This provides investment protection as physical servers or devices that haven't been virtualized but need to communicate with the virtualized environment can be integrated via VXLAN tunnels provisioned directly to the top-of-rack switch.*

firewalls, routers, load balancers, physical servers, and storage that do not support VXLAN natively.

As VXLAN encapsulates and forwards virtual machine traffic over a standard Layer-3 IP network, data centers can scale significantly beyond traditional Layer 2 environments. VXLAN also leverages well-understood routing technologies and Equal Cost Multi Path (ECMP) routing to enable large-scale two tier environments that support tens of thousands of interfaces. Furthermore, VXLAN provides a standards-based mechanism to forward inter-VM traffic between data centers and globally diverse facilities without having to extend the Layer 2 network between those sites.

The Arista 7150S is the first hardware device that can translate between the VXLAN environment and the VLAN world. It can do this at wire speed across 64 ports. This provides investment protection as physical servers or devices that haven't been virtualized but need to communicate with the virtualized environment can be integrated via VXLAN tunnels provisioned directly to the top-of-rack switch. For non-virtualized environments, VXLAN provides a means to transport traffic between sites without having to extend traditional 802.1Q VLANs.
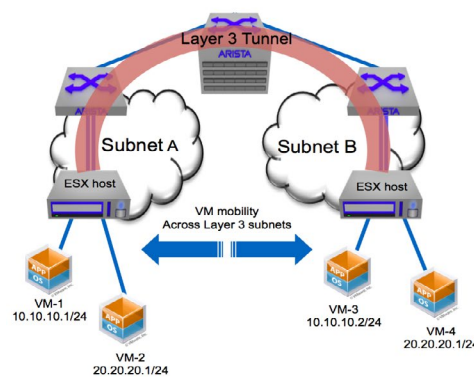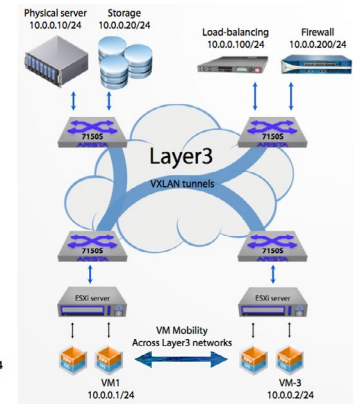


*Figure 5/6: VXLAN provides L2 overlay functionality*

**Network Address Translation (NAT)**

Network address translation (NAT) is typically used to mask an internal network or to overcome a limited IP allocation that may be too small for the number of hosts you wish to use. This is common in financial venues and exchanges where collocation participants are limited to a small range of IP addresses allocated by the exchange or venue. For example, if the venue gives the trading entity 30 addresses and there are 60 servers, a way around this is needed either by obtaining another 30 addresses from the exchange, if available, or using NAT to mask the server addresses. In another example, an exchange may be using a well-known multicast address to publish a feed and need to change the sources or even the group internally without modifying the externally well known S,G information. Typically NAT has been a function of modular chassis devices, firewalls, or software- based routers. All of these devices introduce latency, typically from10μs up to 100s of μs in the worst case.

Arista Networks is the first to offer NAT at wire speed in an ultra low latency platform. The Arista 7150S offers a leap in performance by bringing NAT functionality into a

*PTP enables a precision timing implementation that can be deployed on a standard Ethernet network, either in band or out of band of the standard data plane. This removes the requirement for dedicated timing infrastructure while providing similar or increased precision.*

*Figure 7: Hardware NAT results in substantial latency savings*

high performance, compact data center switch. It implements NAT at the same low latency as standard L2 bridging, L3 routing, or L4 inspection.

Delivering address translation in the same device not only reduces latency, it also enables device consolidation, which brings significant reductions in CapEx and OpEx. (Figure 7).
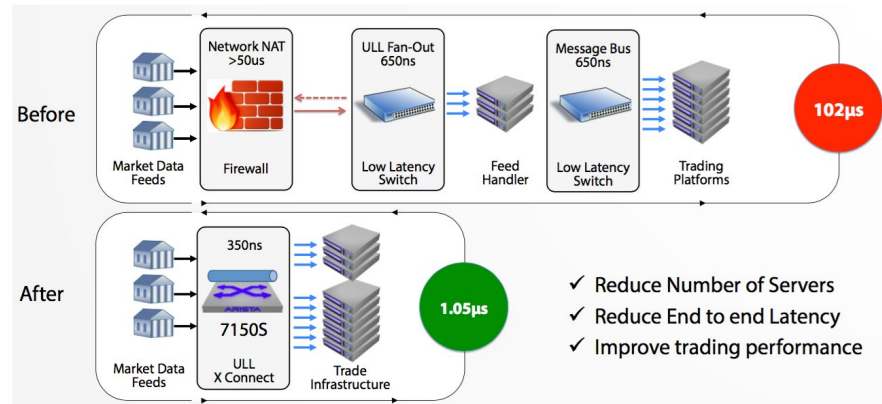


Multicast NAT is the ability to translate the addresses of multicast traffic for trading and video networks, and is also supported by the Arista 7150S. Multicast NAT is a useful capability for those who are subscribing to or publishing data, and can be used to transform traffic so it appears that it came from a single source. It also enables translating multicast groups and destinations to avoid conflicts in the IP infrastructure, or differentiate identical traffic being received at multiple points. Support for multicast NAT provides techniques for simplifying and enhancing PIM configurations and enables providing unique services that were restricted by collisions of address space or "S,G" proliferation. The 7150S provides versatility through the unique delivery of these features with the ultimate in control for multicast data.

**Precision Time Protocol (PTP)**
IEEE 1588 is a standardized precision timing protocol (PTP). PTP is designed to provide precise time distribution over a data network, as opposed to other timing solutions requiring a discrete physical infrastructure within the data center. Traditionally, where synchronized precision timing is critical between individual hosts or networks, a dedicated coaxial network directly connecting to a costly grand master clock with GPS synchronization is implemented. This requires dedicated cabling and dedicated timing cards in every server, along with GPS antennas for GPS synchronization and atomic clock options for holdover stability.

PTP enables a precision timing implementation that can be deployed on a standard Ethernet network, either in band or out of band of the standard data plane. This removes the requirement for dedicated timing infrastructure while providing similar or increased precision.

With the expansion of low latency applications, precision timing has become increasingly important. This is especially critical in HPC environments for measuring events with very high accuracy, or in financial networks where the need exists to

Product Brief

*The Arista 7150S delivers hardware-based functionality that allows the switch to participate in the PTP protocol in two ways: by improving the scaling of timing environments and by providing an accurate calculation of the time taken by PTP packets to traverse the device or network.*

instrument transactions among several devices with nanosecond precision.

The Arista 7150S delivers hardware-based functionality that allows the switch to participate in the PTP protocol in two ways: by improving the scaling of timing environments and by providing an accurate calculation of the time taken by PTP packets to traverse the device or network. The boundary clock function allows end devices (ordinary clocks or "slaves") to be served by the local switch, offloading processing from the grand master, and distributing time over VLAN or routed boundaries. Transparent clocking instruments how long PTP packets spent in the network or device so that slaves can synchronize quickly and accurately with the time being distributed by the master clock.

The PTP implementation on the 7150S is a two- step process that is hardware-assisted for higher accuracy. On ingress and/or egress of prescribed PTP messages, the 7150S ASIC records the arrival or transmission time with better than 10ns resolution. This enables a high precision PTP boundary clock and tightly integrated transparent clocking functionality and provides nanosecond time synchronization performance. The 7150S also provides an onboard, high precision "Oven- Controlled Crystal Oscillator" (OXCO) that delivers an exponentially better cycle-to-cycle jitter and holdover performance than a standard system clock.

*Boundary Clock*

A factor to consider in the design of a PTP system is how many slaves or network segments a grandmaster clock can service while maintaining a defined level of accuracy. The scale is typically limited by the amount of PTP sync packets a master can generate, delay messages it can receive, or interfaces it supports without delaying PTP messages and introducing unintentional clock skew. The PTP boundary clock function addresses this scale limit acting as an intermediate clock that acts as a slave to the master and performs master functions for multiple slaves downstream, while it maintains a high level of accuracy.
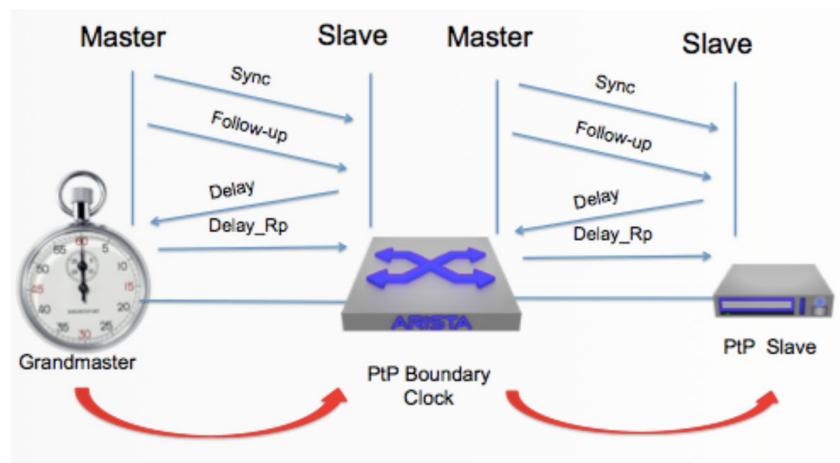


*Figure 8: PTP Boundary clock implementation*

As an example, if a master clock is capable of supporting 100 devices, and each boundary clock can support an additional 100 devices, the PTP implementation can scale by 100X by implementing a single tier of boundary clocks in the network,

*The 7150S is the first product in the data center switch segment to offer complete hardware support for PTP boundary and transparent clock functions with the highest levels of precision and scalability in the market.*

achieving a much larger network than with the grand master clock alone. This also prevents having to provision numerous grand masters at different points in the network. As the boundary clock becomes the authoritative source of time to slaves downstream, the on-board high precision clock is absolutely critical to maintain precision in the case where synchronization with the master clock is lost.

The boundary function is useful in deployments where a grandmaster is installed at a fixed location with GPS access and the time distribution network encompasses multiple network devices with many slaves, or even multiple sites. Deploying multiple boundary clocks in a remote site is useful for providing redundancy in a PTP deployment. Through the Best Master Clock (BMC) election process, the boundary clock ensures the continued synchronization for downstream slaves if a single master fails.

The 7150S boundary clock implementation also addresses a problem seen in many of the popular financial colocation data centers, which is a limitation on provisioning unique GPS antennas for every customer. As opposed to dedicated roof access and serial cabling infrastructure, data center providers would like to be able to deliver precision timing in-band with lower cost and better scalability. The consumers of the time signal want an accurate, precise sync without having to buy additional infrastructure or manage off-premise wiring. The 7150S boundary clock implementation addresses these desires and provides a means of distributing precise time over the Ethernet cross connect to each participant, while providing network isolation and guaranteed services via routing and data plane filtering.

*Transparent Clock*

PTP relies on instrumenting the delay between the master and slave to correctly instrument offset. The less variability in the delay, the more accurate the synchronization will be maintained. In some PTP deployments, the slaves may be a number of network hops away from the master clock. Each network hop has the potential to introduce a non-deterministic amount of latency based on queuing or congestion of the interface at the time the PTP packet is transmitted. The PTP transparent clock functionality eliminates queuing and processing delay for PTP messages. Each transparent clock in the path to the slave measures the residence time (RT) of a PTP packet in its queue and modifies the RT field in the PTP messaging to indicate if the message was delayed. The slave receives the PTP message and uses the RT data to calculate and remove the jitter the network introduced, thereby maintaining lock with the master based on a consistent delay.

In comparison with other offerings in the data center Ethernet market, there are very few transparent clock-capable devices, and most boundary clocks were traditionally implemented in software-powered metro Ethernet routers. The 7150S is the first product in the data center switch segment to offer complete hardware support for PTP boundary and transparent clock functions with the highest levels of precision and scalability in the market.

*Arista has focused on delivering a new class of high precision instrumentation; bringing management tools up to speed and aligned with the visibility needs of high performance environments for 10GbE speeds and beyond.*

## Precision Data Analysis

### EOS Precision Data Analyzer (DANZ)

Network performance monitoring and capacity planning technologies haven't changed significantly for over a decade. Many customers are still using SNMP for capturing interface utilization statistics – tools that were initially designed to work on 10Mbps and 100Mbps networks and that work on single- or multi- second granularity. As a single 10-GbE interface can pass nearly 30 million packets in one second it is clear that traditional tools are no longer sufficiently reporting resolution and granularity.

Arista has focused on delivering a new class of high precision instrumentation; bringing management tools up to speed and aligned with the visibility needs of high performance environments for 10GbE speeds and beyond.

EOS Data Analyzer (DANZ) tool-set provides a broad range of features for fine and coarse grained troubleshooting in three main categories: Advanced Mirroring, Precision Capture and Event Driven Programmability leveraging core technologies including Hardware Time-stamping, TAP Aggregation, Advanced Event Management (AEM) and sFlow.
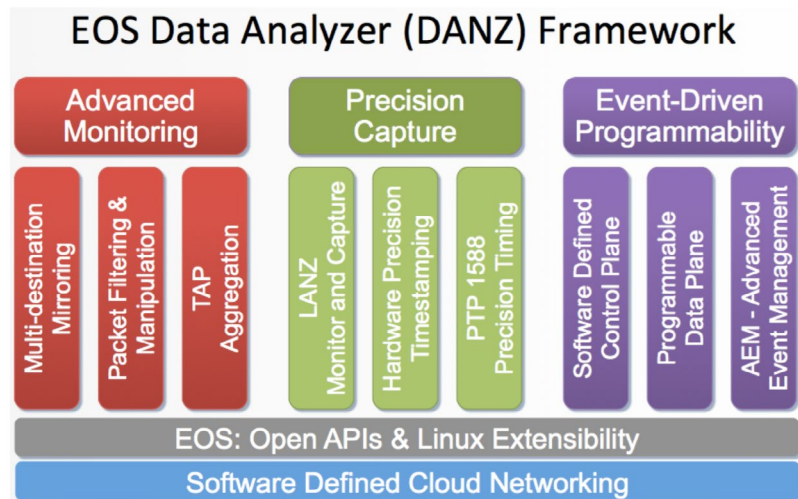


*Figure 9: EOS Data Analyzer Framework*

### Advanced Mirroring and TAP Aggregation

In many high performance environments, the need to accurately measure application performance and collect high precision traffic data has led to a proliferation of specialized filtering and aggregation products like matrix switches and TAP aggregators. These products often add considerable expense to a network deployment and increase the number of different platforms under administration.

The advanced mirroring functionality of the Arista 7150S platform is a suite of enhancements to common port mirroring functionality that enables users to leverage capabilities normally only found in dedicated aggregation devices directly in the data center switch – significantly reducing the complexity, cost and footprint of precision monitoring.

For larger deployments, the 7150S can also be configured as a matrix switch or TAP

*The 7150S is the first device to add dedicated functionality that allows users to standardize on a single product for both standard switching and performance monitoring duties, dramatically lowering both capital and operational costs.*

aggregation device, allowing users to tag, filter and distribute traffic sourced from monitoring devices such as optical taps or mirror ports towards one or more capture and analysis tools.

Performing aggregation, filtering and multiplexing at ultra low latencies insures that a high value analytics tool receives only relevant data with minimal delay for optimal performance.

The 7150S is the first device to add dedicated functionality that allows users to standardize on a single product for both standard switching and performance monitoring duties, dramatically lowering both capital and operational costs. Deploying a consistent platform for both functions also provides additional efficiencies for sparing and software integration – a consistent and mature EOS software platform delivers comprehensive management tooling, open programmability for automation, and a familiar configuration paradigm.

**Latency Analyzer (LANZ)**

LANZ is a unique tool for microburst detection and monitoring interface contention. It allows a network operator to move away from traditional sampled data and interface counters, which typically rely on single- or multi-second increments, as "interesting" events typically occur at nanosecond and microsecond rates in today's applications.

Utilization itself is a very coarse metric. Tools like NetFlow may give you better visibility but they do not scale well to high density, high speed environments. These tools are also incapable of providing early warnings of transient network hotspots before packet loss occurs, making most SLA management reactive rather than proactive.

IT managers need to quickly know how the network is trending and whether hotspots occur at specific times of day or due to specific application behavior. The traditional approach of sampling and software correlation is unable to keep up with faster interface rates and higher densities, requiring substantial CPU resources to manage a single 10G interface alone.
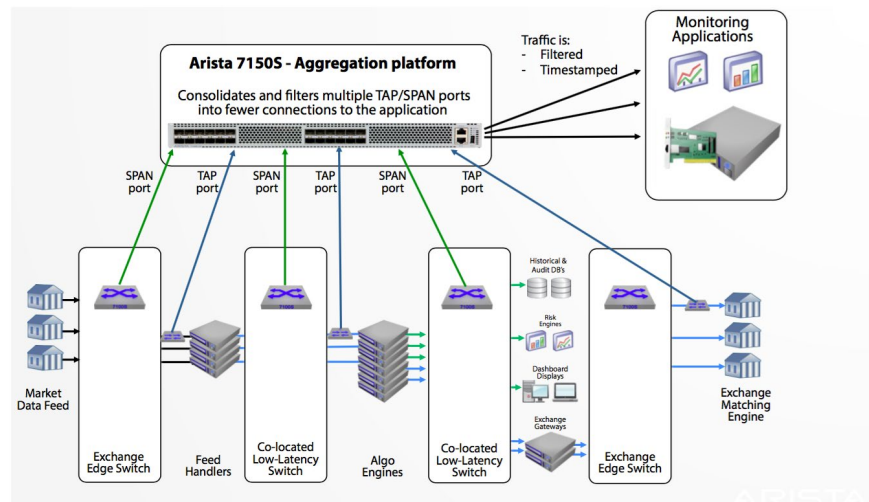


*Figure 10: Common platform for Advanced Mirroring and TAP Aggregation*

LANZ works from the opposite angle, allowing the user to set thresholds for triggered alerts generated by hardware. By removing the inefficiencies of sampling, LANZ works at the byte level, not on volumes of data. It can be configured to alert on the smallest levels of interface contention, being able to instrument to the granularity of just one packet queued behind another.

Such capability is especially important in financial environments where increasing latency can have a profound effect on the success of a trading strategy, reduce fairness, or provide early warning of degradation in connectivity to a venue due to contention for cross-connect bandwidth.

The need to follow burst and congestion trends in the network is not specific to financial applications however; it's also useful in regular data center, service provider, content delivery, and big data networks. LANZ is the only early warning system to instrument and alert on transient congestion events before packets are dropped.

With LANZ the reporting resolution is substantially enhanced through the addition of high-precision timing and detailed metrics by interface and traffic class. LANZ reports when congestion occurred, its duration, peak, and any packet losses within the reported event. The LANZ records are instrumented to nanosecond resolution, allowing strict correlation with application logs. The LANZ records can also be streamed to external applications for integration with application monitoring and historical analysis.
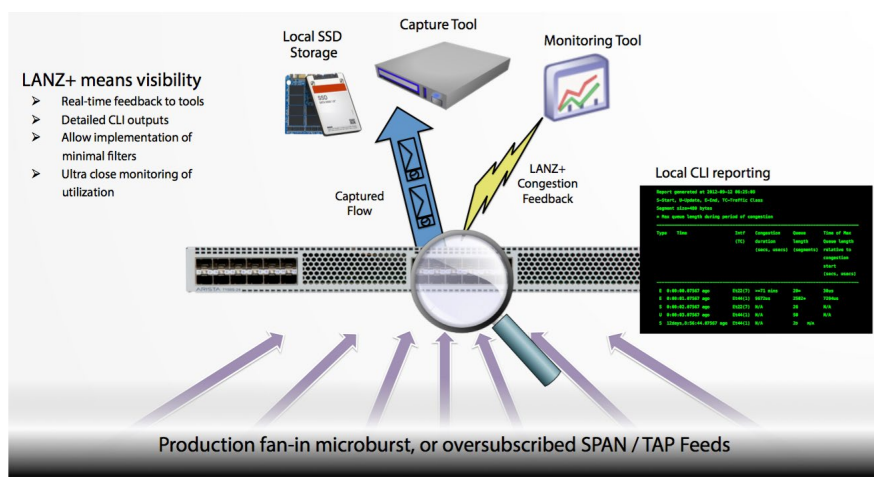
*LANZ is the only early warning system to instrument and alert on transient congestion events before packets are dropped. With LANZ the reporting resolution is substantially enhanced through the addition of high-precision timing and detailed metrics by interface and traffic class.*



*Figure 11: Leveraging LANZ for congestion monitoring*

For forensic purposes, LANZ+ is able to capture a sample of the packets that were queued in a congestion event. This provides application level latency performance monitoring from all angles – always-on proactive, triggered reactive, flow- based, and real time packet sampling.

**Hardware timing**

As interface speeds continue to increase and analytics tools become more precise, it is critical for devices offering mirroring or aggregation services to provide accurate timing information for every packet on the wire.

The 7150S leverages the on board high-precision oscillator coupled with integrated hardware time stamping to apply accurate timestamps to every packet at wire

speed without incurring additional latency.

Previously only available on specialized low- density devices, the integration of hardware time stamping provides new possibilities for accurate measurement of both one-way and round trip traffic latency with basic tools or software (such as a network sniffer).

This time stamping can be leveraged to improve both standard mirror and in-aggregation output. As the device can now instrument with a high degree of accuracy when a packet arrived on ingress, it is possible to remove any additional queuing resulting from the act of aggregating multiple sources to fewer tools. This substantially improves the accuracy and determinism of the captured data.

Arista's open timestamp format is supported by an ecosystem of leading tool vendors and is simple to integrate with in-house developed applications, making highly accurate timing cost effective to deploy network-wide, both in and out of band of the data plane.

**Advanced Event Management**
AEM is a suite of tools that allows flexible localized intelligence to be deployed on EOS platforms and offers the ability to easily schedule events and react to network changes. AEM can execute bespoke scripts or packages to control or reconfigure the device, extract data, automate operational processes, and even customize the behavior of the switch on the fly. It enables the network operator to build a rich set of action/reaction events to proactively manipulate what happens on the switch without human interaction, lowering operational overhead and improving reaction time to issues.

AEM further enhances the traditional approach to troubleshooting. Via AEM, any EOS platform can actively capture a forensic log of key state indicators in a SQL database. This is done on a continuous basis, providing operators a comprehensive view of what's happening in the network both currently and historically. Instead of relying on point-in-time data or external polling, administrators are now able to accurately monitor past events from a topological perspective.

**sFlow**
sFlow is an industry standard analysis tool (RFC3176) that was designed specifically to support very large, high-speed environments. Instead of localized data collection and correlation, sFlow provides a lightweight mechanism to perform centralized, network-wide flow analysis.

sFlow captures samples of traffic and statistics from interfaces in hardware, exporting to one or more collector applications for post processing. sFlow's lightweight implementation means it does not impact forwarding performance or latency, yet when enabled across all infrastructure devices it makes complete network-wide monitoring possible without specialized capture tools.

By collecting statistics and flow data from a large number of devices, sFlow is able to reveal horizontal network activity typically missed by "choke point" monitoring. Instead of trying to correlate data on board each device, exporting direct to a

*The Arista 7150S Series enables construction of scalable topologies that embed tools to instrument and understand your network with the granularity required in high performance, low latency, or virtualized environments.*

collector application allows sFlow to benefit from the broad CPU, memory and storage available to a server to ensure fine-grained data is collected for analysis and can be kept for extended periods for forensic troubleshooting. As sFlow is supported on all Arista EOS powered platforms, complete visibility can be provided across the data center.

## Conclusion

High performance networks need a new breed of switch that combines the highest levels of intelligent forwarding performance with advanced networking and monitoring features. The Arista 7150S Series enables construction of scalable topologies that embed tools to instrument and understand your network with the granularity required in high performance, low latency, or virtualized environments.

The Arista 7150S provides ultra low latency with deterministic performance, open extensibility, and a full suite of management, address translation, and forensic traffic and performance monitoring features. These capabilities reduce device counts, lower capital and operational expenses, and streamline data center operations enabling a new paradigm of proactive, automated manageability at cloud scale and cloud economics.



*Figure 12: 7150S Family*

**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