

Arista 7160 Series Switch Architecture

Today's highly dynamic cloud datacenter networks continue to evolve with the introduction of new protocols and server technologies such as containers, bringing with them ever increasing bandwidth demands. Server workloads and workflow mobility are characterized by traffic patterns of intensive server-to-server communication from hundreds and even thousands of servers that creates new networking challenges.

Designing and deploying modern datacenter networks creates unique challenges to network architects. They need to select the best in class technology today that delivers on performance, is fully programmable for operational efficiency and offers support for the latest set of modern protocols. Architects also need to ensure that the network includes the flexibility to evolve as new emerging standards become available to increase ROI, extend the useful life cycle and allow deployment of a new service in weeks, not months or years. The selection of leaf and spine network switches that meet the performance requirements of modern applications, while also providing the capability to add support for new protocols, offer easy network migration and that achieve these goals cost effectively is critical.

Many datacenter switches deliver high performance, but they have fixed and rigid pipelines for packet processing, fixed allocations of logical resources and packet buffers, that cannot be adapted for different use cases.

The Arista 7160 offers true innovation for high performance datacenter switches delivering support for a programmable forwarding pipeline, flexible profiles, up to 6.4Tbps of forwarding with 24MB of dynamic packet buffer and a choice of interface types in a set of power efficient compact systems.

- The 7160 Series Flexible Profiles offer a range of custom profiles allocating common resources optimally for the desired network topology with lower complexity for both cloud and enterprise customers who frequently hit limitations with more rigid systems.
- The unique programmable pipeline accommodates support for the addition of new encapsulations and overlay technologies allowing the same system to be used in multiple roles, expanding the life-cycle and allowing rapid deployment.
- Arista AlgoMatch™ enables a more flexible and scalable solution for access control, secure cloud networking and telemetry at 2-6x the scale of traditional methods with up to 50% lower power.
- All ports have access to a fully shared 24MB packet buffer. Dynamic buffer management is designed for speed changes, microbursts or sustained network congestion.
- Native 25GbE that offers the benefits of enhanced compute and storage efficiency at a long-term cost structure similar to 10GbE.

The 7160 Series 25G and 100G ports are fully compliant to both the 25G Consortium specification and the IEEE 802.3by 25GbE standard and provide an easy migration path for 10G networks.

Arista 7160 Series Model Choice

The 7160 Series is available in a choice of three models to allow flexibility of interface type and density.



Figure 1: Arista 7160 Series

The table below provides details on the 7160 Series models.

Table 1: 7160 Series - System Specifications			
7160 Models	7160-32CQ	7160-48YC6	7160-48TC6
Switch Height (RU)	1	1	1
10GBASE-T	--	--	48
10G SFP+	4	--	--
25G SFP	--	48	--
40G QSFP+	--	--	--
100G QSFP	32	6	6
Maximum Density 10GbE ports	128	72	72
Maximum Density 25GbE ports	128	72	24
Maximum Density 40GbE ports	32	6	6
Maximum Density 100GbE ports	32	6	6
Maximum HW System Throughput (Tbps)	6.4	3.6	2.16
Maximum Forwarding Rate (Bpps)	2.4	2.4	2.4
Latency	From 2 usec	From 2 usec	From 3 usec
Packet Buffer Memory	24MB	24MB	24MB

Arista 7160 Series Deployment Scenarios

The Arista 7160 Series offers a wide selection of connectivity options delivering the flexibility to build large-scale leaf and spine datacenter designs in combination with other fixed and modular Arista 7000 Series switches. The operational advantages provided by both flexible profiles and the programmable pipeline ensures suitability for a range of deployment scenarios. The following are examples of just a few of the typical use cases.

- **Dense top of rack** - For server racks with both 10GbE and 25GbE systems the 7160 series offer high performance with a choice of 1/10/25/40/50/100GbE.
- **10GbE to 25GbE Migration** — IEEE 802.3by 25GbE standard and 25G Consortium compliant for seamless transition to the next generation of Ethernet performance as the 7160 provides investment protection using existing cabling and consistent leaf-spine designs.
- **High Performance Computing** — Cost effective and power efficient systems to enable non-blocking or minimal over-subscription for 10G and 25G Servers with the ability to accommodate server and storage capacity expansion with under subscribed uplink designs and efficient flow balancing.
- **Data Center Leaf-Spine** — Open standards based L2 and L3 with rich telemetry and open network monitoring using consistent designs and open APIs.
- **Secure Cloud Environments** — With support for up to 48K ACLs and policy enforcement aligned to business intent that scales for both IPv4 and IPv6.
- **100GbE Scale Out Designs** — Small to medium locations requiring power efficient high density compact systems in a fixed form factor with high performance and a programmable pipeline to accommodate next generation protocols and encapsulation.
- **Software Defined Networking** — With support for CloudVision, VXLAN and eAPI and a flexible and programmable pipeline that allows quicker deployment of new solutions avoiding costly upgrades.

- **ECMP designs up to 128-way** — Cost-effective multi-pathing using open protocols in combination with the Arista 100GbE modular spine switches and routers.
- **Large scale L2 environments** — Flexible resource allocation for maximum L2 scale without inefficiency associated with traditional systems, using all available resources.

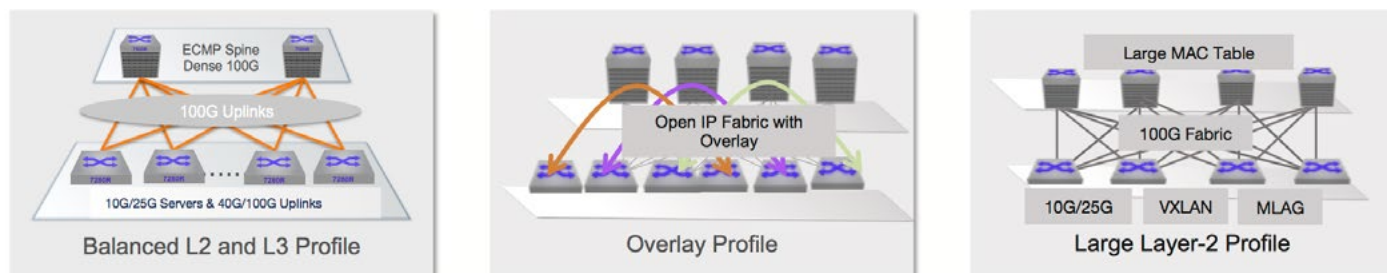


Figure 2: 7160 Deployment Scenarios

The 7160 Series has a number of differentiators compared to other fixed configuration systems that make it ideal for next generation datacenter networking.

- Standards-based high density 100G switch ideal for future proof 10/25/40/50/100G designs and high bandwidth needs such as in High Performance Storage or Content Delivery Networks (CDNs)
- Dynamic and fully shared packet buffers makes for an ideal top of rack switch for lossless performance with in data analytics and IP storage environments
- Directly connected 25GbE, 40GbE and 50GbE attached storage systems, requiring high performance and predictable latency
- Streaming network state for advanced analytics with Arista CloudVision
- Network-wide virtualization platform for next generation cloud bursting with wire-speed VXLAN
- Comprehensive L2 and L3 feature set and open APIs for multi-vendor networks with no proprietary lock-in
- Scalable L2 and L3 table resources allow deployment flexibility in both large L2 and L3 environments without compromising system efficiency and scale
- Latency Analyzer (LANZ) - Microsecond granularity on port utilization using buffering watermarks to provide immediate feedback and precision monitoring
- Network-wide approach for workload orchestration and workflow automation as a turnkey solution for cloud networking with CloudVision

Arista 7160 Switch Architecture

All of the 7160 Series share a common system design built around a high performance x86 CPU and 8GB of system memory for the control plane. The CPU is connected to internal flash, bootflash, power supplies, fans, management I/O and peripherals.

The x86 CPU is also connected over PCIe to the 3.2Tbps Switch on Chip that runs all the data plane forwarding and has all the directly connected front panel ports.

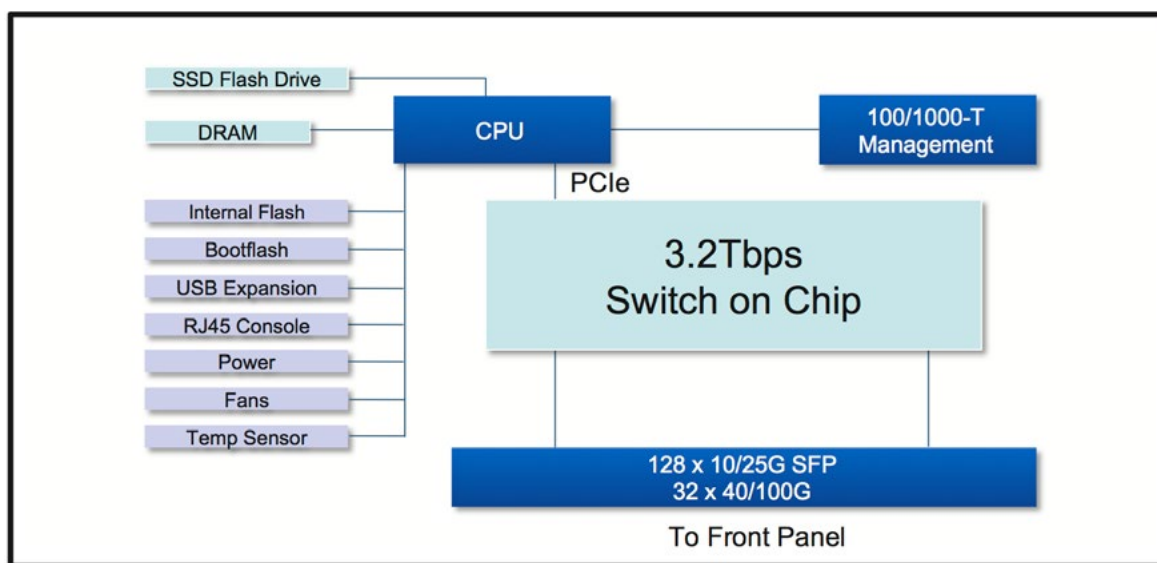


Figure 3: 7160 Architecture

Arista 7160-32CQ

The 7160-32CQ is a 32 QSFP port based switch. All ports are 100G and 40G capable, accept either 100G or 40G QSFP optics and cables and can be configured for a choice of 10/25/40/50/100G depending on the transceiver. An additional 4 SFP+ ports allow for individual 1G and 10G connections and are shared with the first QSFP port. Other ports on the front panel include a 100/1000 Ethernet interface for out of band management, a serial RJ45 connection and a USB port.




Figure 4: 7160-32CQ Switch

7160-32CQ - 1RU with 32 Port 100GbE QSFP

- Offers a choice of port combinations with 40G and 100G QSFP optics and cables
- Flexible interface combinations - 32 x 40G, 128 x 10G, 32 x 100G, 128 x 25G, 64 x 50G
- IEEE 25GbE and 25G Consortium specification support
- 5 speeds for flexible 10GbE, 25GbE, 40GbE, 50GbE and 100GbE with optics or cables
- 6.4Tbps of wire speed performance with 24MB of buffer

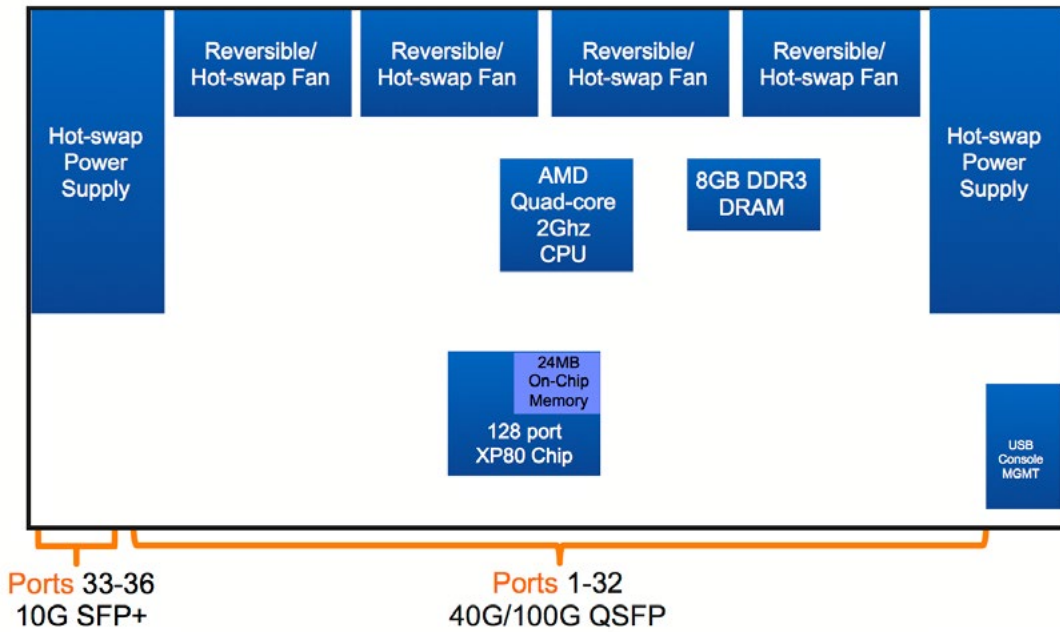


Figure 5: Arista 7160-32CQ Architecture Block Diagram

Arista 7160-48YC6

The 7160-48YC6 provides customers with a 10/25GbE solution with 48 x 10/25GbE SFP ports and 6 x 100G QSFP ports for connection to spine switches. The QSFP ports will accept either 100G or 40G QSFP optics and cables and can be configured for a choice of 10/25/40/50/100G depending on the transceiver.

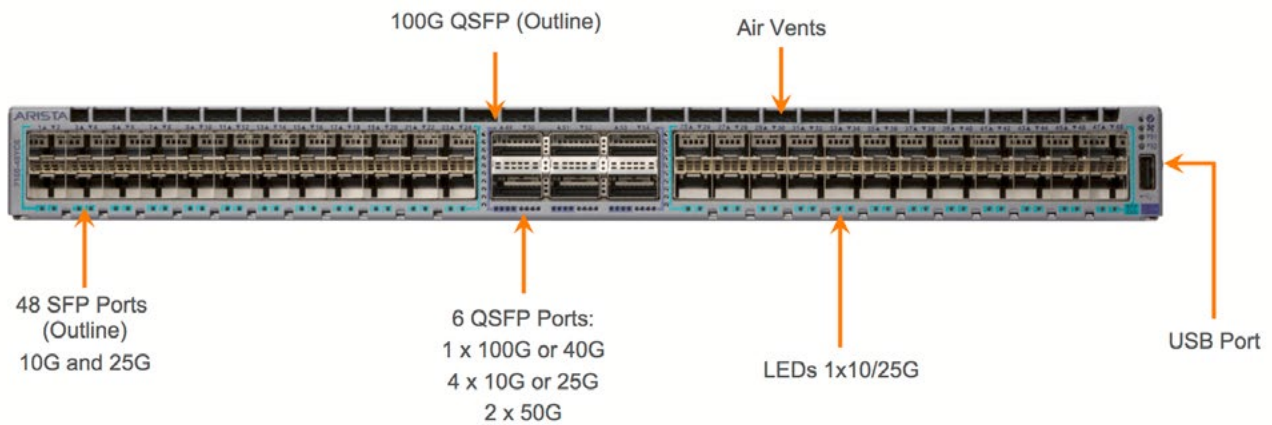


Figure 6: 7160-48YC6 Switch

7160-48YC6 - 1RU with 48 Port 25GbE SFP and 6 Port 100GbE QSFP

- Offers 48 wire speed 25GbE ports with six 40/100G QSFP ports for up to 72 total 25G or 10G ports
- IEEE 25GbE and 25G Consortium specification support
- Easy migration from 1/10G to 1/10/25G using familiar SFP connections.
- Six 40/100G QSFP ports for flexible combinations of 40G and 100G QSFP
- 3.6Tbps of wire speed performance with 24MB of buffer

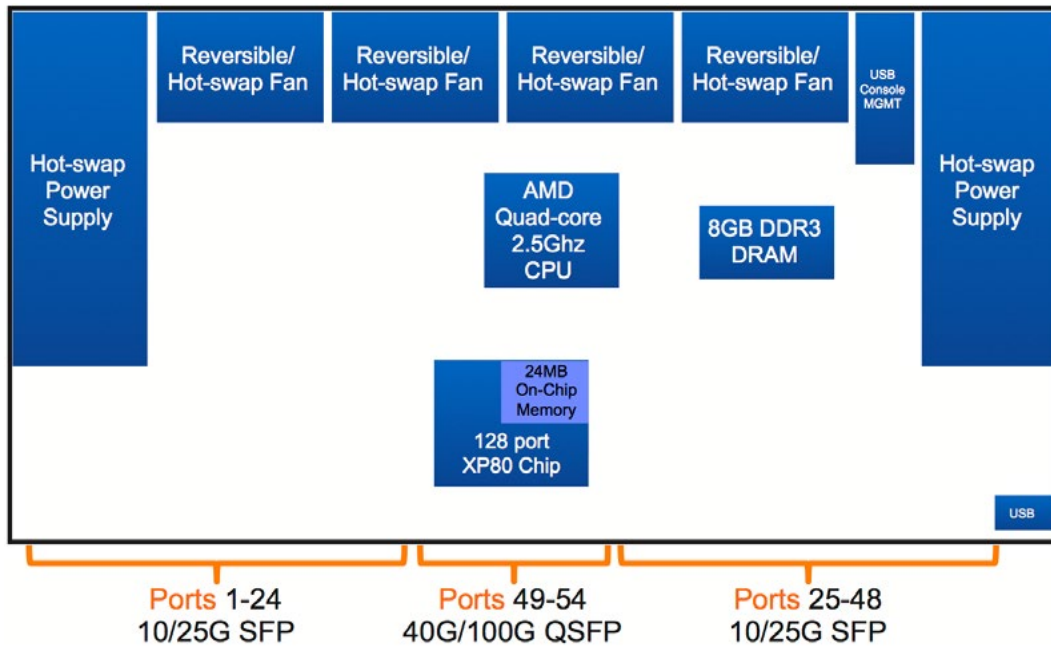


Figure 7: Arista 7160-48YC6 Architecture Block Diagram

Arista 7160-48TC6

The 7160-48TC6 provides customers with a 10GBASE-T solution with 48 x 1GbE/10GbE RJ45 ports and 6 x 100G QSFP ports for connection to spine switches, located in the center of the front panel. The QSFP ports will accept either QSFP100 or QSFP+ optics and cables and can be configured for a choice of 10/25/40/50/100G depending on the transceiver.

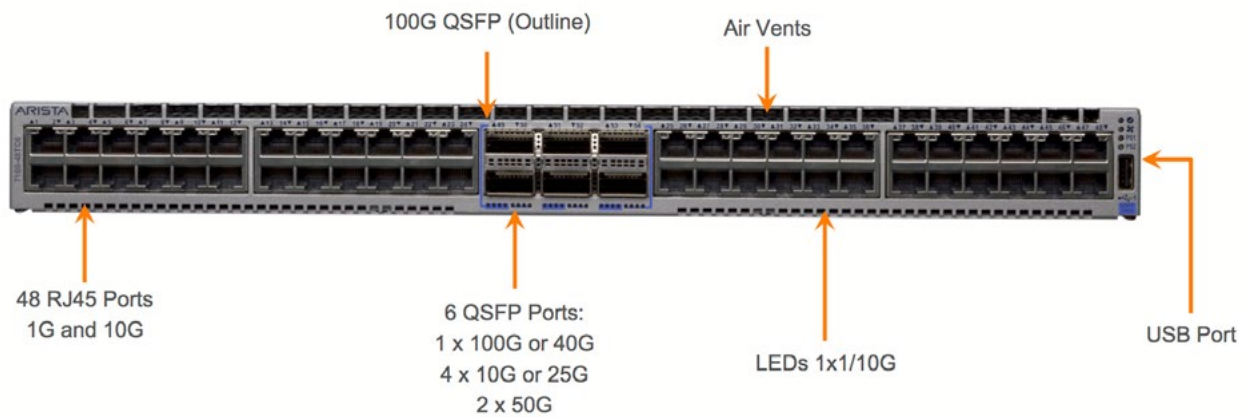


Figure 8: 7160-48TC6 Switch

7160-48TC6 - 1RU with 48 Port 10GBase-T and 6 Port 100GbE QSFP

- Offers 48 wire speed 10G ports with six QSFP100 for up to 72 total 10G ports
- Six QSFP100 ports allow choice of 6 x 100GbE, 24 x 25GbE, 6 x 40GbE or 24 x 10GbE
- IEEE 25GbE and 25G Consortium specification support on 25G breakout ports
- 5 speeds for flexible 10GbE, 25GbE, 40GbE, 50GbE and 100GbE with optics or cables
- 2.16Tbps of wire speed performance with 24MB of buffer

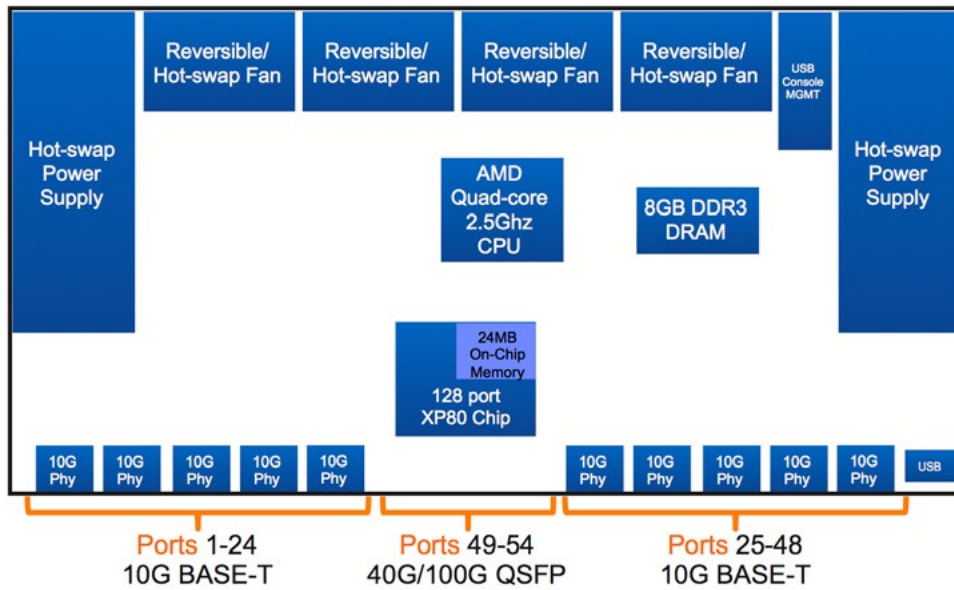


Figure 9: Arista 7160-48TC6 Architecture Block Diagram

Datacenter Grade Availability and Redundancy

The Arista 7160 switches are designed for continuous operations with system wide monitoring of hardware and software components, simple serviceability and provisioning to prevent single points of failure. Key high availability features include:

- 1+1 hot-swappable power supplies and four hot-swap fans provide dynamic temperature control combined with N+1 redundancy
- Color coded PSU's and fans that deliver platinum level power efficiency
- Live software patching
- Self healing software with Stateful Fault Repair (SFR)
- Smart System Upgrade (SSU) and Accelerated Software Update (ASU)

Table 2: 7160 Series - Power and Fan Redundancy

7160 Models	7160-32CQ	7160-48YC6	7160-48TC6
Power Supplies (Redundancy)		2 hot swappable (1+1)	
Fans (Redundant)		4 hot swappable (N+1)	
Airflow		Front to rear and rear to front	



Figure 10: Arista 7160-48YC6 and 7160-48TC6 Switch Rear View

Scaling The Control Plane

The CPU complex on the 7160 Series switches is used exclusively for control-plane and management functions; all data-plane forwarding logic occurs at the packet processor level.

Arista EOS®, the control-plane software for all Arista switches, executes on multi-core x86 CPUs with multiple gigabytes of DRAM. As EOS is multi-threaded, runs on a Linux kernel and is extensible, the large RAM and fast multi-core CPUs provide for operating an efficient control plane with headroom for running 3rd party software, either within the same Linux instance as EOS or within a guest virtual machine.

Out-of-band management is available via a serial console port and/or the 10/100/1000 Ethernet management interface. The 7160 Series also offer USB2.0 interfaces that can be used for a variety of functions, including the transferring of images or logs.

Table 3: 7160 Series CPU Complex

7160 Models	7160-32CQ	7160-48YC6	7160-48TC6
CPU	Quad-Core x86		
System Memory	8GB		
Flash Storage	4GB		

The Next Generation Programmable and Flexible Packet Processor

The 7160 Series is built using a single Switch on Chip (SoC) silicon. All stages of the packet forwarding are performed entirely in the hardware/data plane. The packet processor is capable of providing both the ingress and egress forwarding for all packets that arrive on or are destined to ports located on the packet processor.

The 7160 forwarding architecture enables a programmable pipeline along with flexible resource allocation made available through custom forwarding profiles.

Below is the block diagram showing key components of the packet processor and the functionality provided by each component.

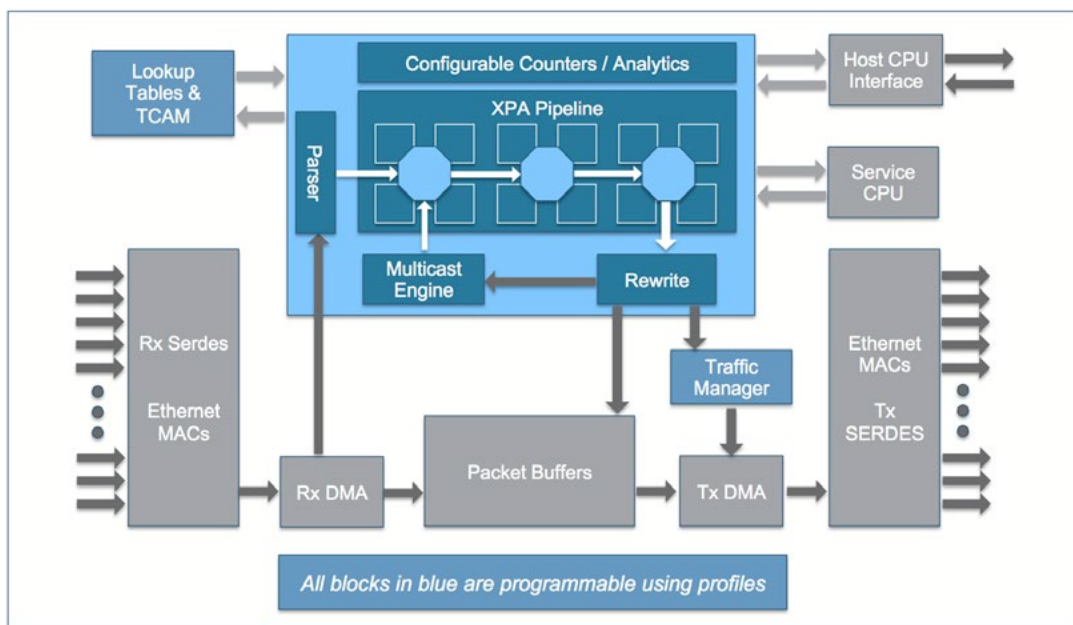


Figure 11: 7160 Series Packet Processor

Programmable Pipeline

The 7160 programmable packet pipeline allows for the addition of new protocols, encapsulation and tunneling features to the packet processor through simple software upgrades without changes to the underlying hardware. This flexibility avoids the needs to wait for new switch cycles, costly replacements or major upgrades. Recent examples of new encapsulations in the datacenter environment include VXLAN and GENEVE, both of which required new silicon technology, delaying deployments and innovation.

Flexible Resource Profiles

The 7160 Series provides unique flexibility on how memory resources are allocated among the different forwarding tables and features. A forwarding profile is a pre-defined allocation of the lookup memory resources on the 7160 to assign specific table sizes for different features. A forwarding profile allows the available resources to be optimally allocated for a specific feature or set of features as needed for a given use case. This allows for deployment of the 7160 Series across multiple use cases each configured with a different forwarding profile to suit the unique requirements.

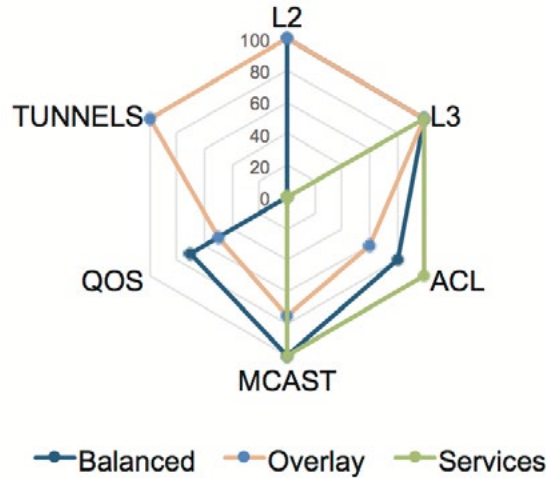


Figure 12: Arista 7160 Flexible Profiles

This flexibility allows customers to standardize their data center switching with 7160 Series across multiple use cases, each most efficiently using all the switch resources available. A few example profiles are described below.

Balanced L2 and L3 Profile

This profile is suitable for typical data center leaf/spine deployments supporting virtualized and non-virtualized workloads with balanced Layer and Layer 3 feature scale.

Table 4: 7160 Series Balanced Profile Example

Interfaces	MAC Address	VLANs	VXLAN VNIs	VXLAN VTEPs	Multicast Groups	VRRP Instances	IPv4 Routes
10G – 100G	128K	4K	4K	4K	128K	255	128K

Large L3 Profile

This profile is applicable for both spine and WAN edge roles supporting large routing tables in combination with ACLs.

Table 5: 7160 Series Large L3 Profile Example

Interfaces	IPv4 LPM	IPv6 LPM	Multicast Routes	ECMP	ACL
10G – 100G	1M	500K	48K	128-way	30K

Large Hosts - L2 Overlay VXLAN Profile

This profile is targeted for leaf/top of rack switch supporting large L2 host scale in an overlay network.

Table 6: 7160 Series Large Host Profile Example

Interfaces	MAC Address/IP Host	VLANs	VXLAN VTEPs	IPv4 LPM	IPv6 LPM	Multicast Groups
10G – 100G	300K	4K	96K	64K	32K	8K

Each custom forwarding profile leverages the common set of forwarding resources, and additional profiles enable more use cases.

Arista AlgoMatch

AlgoMatch is an Arista innovation combining software and hardware to implement Access Control Lists for policy control and network telemetry that are more scalable and require lower power use. AlgoMatch utilizes an efficient packet matching algorithm that in turn enables flow matching for access control, policy and visibility. It is more flexible than traditional technology, allowing for multiple actions to be performed on a single packet or flow, and user defined filters for packet classification and custom actions.

AlgoMatch is the combination of a set of hardware and software technologies to implement Access Control Lists in a novel way without the usage of hardware TCAMs.

AlgoMatch delivers the efficiency and flexibility required for next generation architectures in the datacenter to address the demands of increased scale, whether it be dual stack implementations or flow analysis at Layer-4 and above.

Shared Buffer Architecture

The 7160 Series incorporates an advanced traffic manager with 24MB of packet buffer that is fully shared across all ports and is an excellent choice for scalable datacenters and modern intensive workloads. Unlike architectures where the total buffer is statically allocated to a port or group of ports or the packet memory or buffers are formed of multiple slices, the 7160 Series packet buffer is dynamically allocated across all ports with the ability to adjust in real time to the demands of bursty applications, mixed interface speeds and congestion.

To optimize the memory performance, the 24MB is divided into 96K pages of 256 byte cells, organized in 16 banks each of 1.5MB. This ensures an even distribution of page usage over the total memory. A cell can contain a single packet, or a linked list of 256 byte cells is used to buffer packets larger than a single cell.

The packet buffer is carved by software into multiple shared pools to handle high priority traffic and Priority Flow Control (PFC) enabled ports separately from the other traffic. This ensures the switch packet buffer is optimally allocated for the specific use case's unique requirements.

Extensive support for Active Queue Management mechanisms such as WRED, DCTCP and ECN ensure that both high priority flows and lossless storage traffic are handled equally well with the ability to absorb large bursts with extensive counters for visibility and accounting.

QOS Architecture

The 7160 Series supports a sophisticated and advanced traffic manager with 4 levels of scheduling hierarchy that supports traffic shaping along with WRR and SP scheduling. The architecture supports up to 16 transmit queues per port consisting of 8 unicast queues and 8 multicast queues. The Traffic Class (TC), traffic type and egress port select a transmit queue. Each transmit queue can be independently configured for various scheduling parameters including strict priority, WRR, shaping rate etc. In addition, each queue can be configured to apply WRED. The 7160 Series supports 3 trust modes including CoS trusted, DSCP trusted and untrusted, along with CoS and DSCP rewrite capability.

802.3BY IEEE 25GbE and 25G/50G Consortium Specification

The 7160 Series offers full support for the IEEE 802.3by 25Gigabit Ethernet standard, ensuring long term investment protection, and support for the 25G and 50G Consortium specification for backward compatibility to existing 25G devices.

Unrelenting traffic growth and storage capacity expansion is driving demands for higher performance networks. The introduction of 25GbE provides a 2.5X performance improvement over 10GbE while using the same familiar cabling and designs. Support for 10G/25GbE modes allows for future investment protection with the ability to migrate as needed without expensive network upgrades.

Some of the advantages for migrating to 25GbE include:

- Maximize the switch and server throughput and efficiency, by using all available switch bandwidth on high performance systems
- Reduce the capital expense by reducing the number of cables and switch ports compared to multiple 10G ports used to increase bandwidth
- Lower operational expenses by reduced power and cooling compared to a 40G alternative
- Lowest cost per bit of performance since 25GbE provides 2.5X higher throughput using the same switch technology

The 25GbE interfaces are backward compatible with a wide range of 10G SFP+ optics and cables allowing for dual speed support with each port independent of the others. In addition, all 7160 Series 100GbE ports allow for 5 interface speeds, including 10G, 25G and 50G with parallel optics and breakout cables, as well as both 40G and 100G.

Arista EOS: A Platform for Scale, Stability and Extensibility

At the core of the Arista 7160 Series is Arista EOS® (Extensible Operating System). Built from the ground up using innovations in core technologies since our founding in 2004, EOS contains more than 8 million lines of code and over 1000 man-years of advanced distributed systems software engineering. EOS is built to be open and standards-based, and its modern architecture delivers better reliability and is uniquely programmable at all system levels.

EOS has been built to address two fundamental issues that exist in cloud networks: the need for non-stop availability and the need for high feature velocity coupled to high quality software. Drawing on our engineers' experience in building networking products for more than 30 years and on state-of-the-art open systems technology and distributed systems, Arista started from a clean sheet of paper to build an operating system suitable for the cloud era.

At its foundation, EOS uses a unique multi-process state-sharing architecture where there is separation of state information from packet forwarding and from protocol processing and application logic. In EOS, system state and data is stored and maintained in a highly efficient, centralized System Database (SysDB). The data stored in SysDB is accessed using an automated publish/subscribe/notify model. This architecturally distinct design principle supports self-healing software resiliency, easier software maintenance and module independence, higher software quality overall, and faster time-to-market for new features that customers require.

Arista EOS contrasts with the legacy approach to building network operating systems developed in the 1980's that relied upon embedding system state held within each independent process, extensive use of inter-process communications (IPC) mechanisms to maintain state across the system, and manual integration of subsystems without an automated structured core like SysDB. In legacy network operating systems, as dynamic events occur in large networks or in the face of a system process failure and restart, recovery can be difficult, if not impossible.

Arista took to heart the lessons of the open source world and built EOS on top of an unmodified Linux kernel. We have also maintained full, secured access to the Linux shell and utilities. This allows EOS to preserve the security, feature development and tools of the Linux community on an on-going basis, unlike legacy approaches where the original OS kernel is modified or based on older and less well-maintained versions of Unix. This has made it possible for EOS to natively support things like Docker Containers to simplify the development and deployment of applications on Arista switches. Arista EOS represents a simple but powerful architectural approach that results in a higher quality platform on which Arista is faster to deliver significant new features to customers.

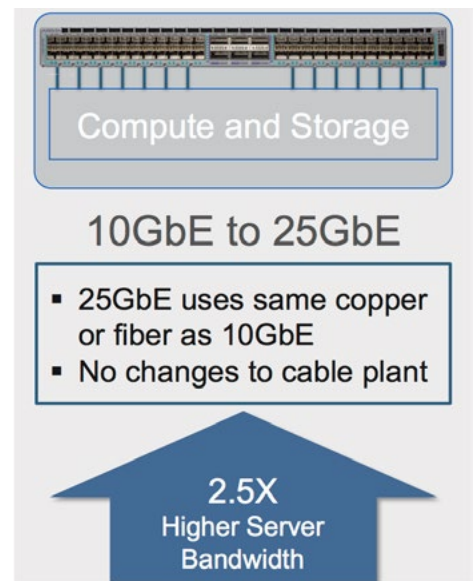


Figure 13: 10GbE to 25GbE Migration Advantages

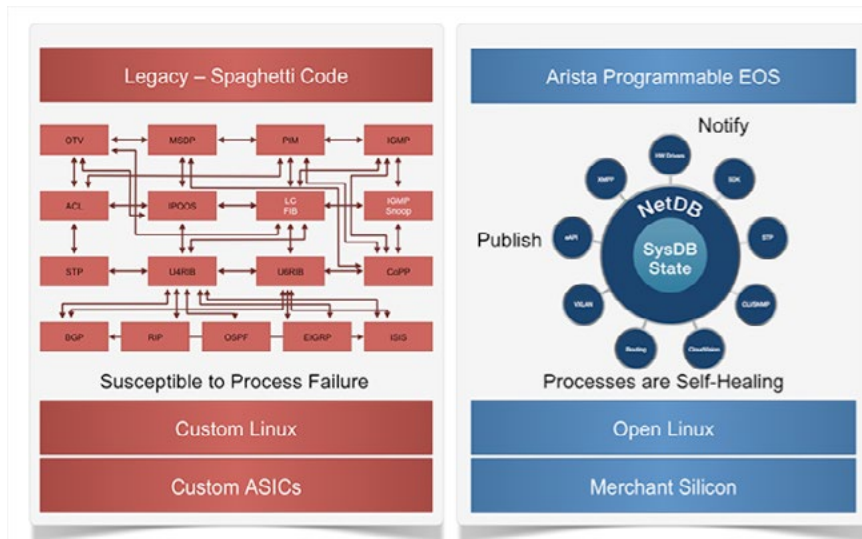


Figure 14: Legacy approaches to network operating systems (left), Arista EOS (right)

EOS is extensible at every level, with open APIs at every level: management plane, control-plane, data-plane, services-level extensibility, and application-level extensibility. With access to all Linux operating system facilities including shell-level access, Arista EOS can be extended with unmodified Linux applications and a growing number of open source management tools to meet the needs of network engineering and operations.

Open APIs such as EOS API (eAPI) and OpenConfig and EOS SDK provide well-documented and widely used programmatic access to configuration, management and monitoring that can stream real-time network telemetry, providing a superior alternative to traditional polling mechanisms.

Conclusion

The Arista 7160 series offers a selection of leaf and spine network switches that meet the performance requirements of modern applications, while also providing the capability to add support for new protocols, offer easy network migration and achieve these goals cost effectively. Combined with Arista’s EOS and the flexibility and programmability offered by the architecture, the Arista 7160 series is a great addition to the Arista product portfolio, providing comprehensive datacenter networking solutions.

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



Copyright © 2016 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document. Dec 5, 2016 02-0067-01