

Arista 7170 Multi-function Programmable Networking

Over the last decade, the industry has seen broad adoption of modern data center principles to support increasingly distributed and elastically scalable applications. Highly efficient leaf-spine network designs and improvements in server scale and connectivity have removed many bottlenecks and enabled a substantial increase in the scale of compute resources, the volume of data carried through the network and the number of connected devices.

These optimizations have in turn exposed other weaknesses, particularly in the area of higher layer appliances such as load balancers and firewalls often deployed to perform more advanced traffic handling. Where previously such platforms operated at natural network choke points and throughput was gated by upstream connectivity such as WAN interfaces, today they have been overtaken by surrounding technologies and administrators are forced to build large clusters of expensive platforms to keep up.

In parallel, many of the rich features of traditional appliances are no longer necessary; modern applications are natively designed for secure connectivity and distributed deployment, requiring only basic application services to be delivered at very high speed.

Arista's 7170 Series is a unique family of fully programmable platforms that, for the first time, provides the ability to implement higher layer functionality directly into the network data-plane while providing switch-like levels of performance and economy. The 7170 Series provides new tools for addressing traditional networking challenges that are perfectly aligned with today's data rates and result in substantial cost, power and space savings.

Multi-function Programmable Platform

The Arista 7170 Series, based on the Intel Tofino family of P4 programmable switch chips, are uniquely multi-functional 100GbE systems built to solve complex networking challenges. With a highly flexible and reconfigurable architecture, the 7170 series with Arista EOS® (Extensible Operating System) enable forwarding profiles to customize the hardware data plane to suit an array of network roles for cloud, enterprise, service provider and content delivery networks. The forwarding profiles, written in the open-source P4 language, can address advanced telemetry, network security, and flexible tunnel termination and address translation as well as many other possible use cases enabling a new class of in-network functionality and the ability to customize packet handling for specific use-cases.

The 7170 series supports a unique completely programmable pipeline with the flexibility to allow the entire packet forwarding pipeline to be changed to add new functionality and vary lookup table size and operations. This provides phenomenal investment protection along with opening the possibility to support highly personalized use cases to solve traditional and new network challenges which are otherwise not feasible with traditional switching platforms.

Arista EOS' comprehensive set of features and programmability for cognitive data and control lets the same platform be used in multiple roles, each with its own profile, ensuring consistent management and provisioning.

The following are examples of advanced features enabled by the 7170's programmable pipeline:

- **Network Services Offload** - large scale HW VTEP designed for a bare metal environment freeing up compute resources and accelerating applications
- **Large Scale NAT** - for address and port translation at scale far beyond typical DC switches
- **Stateless Cloud Load Balancing** - enabling high bandwidth load balancing across large pools of Linux hosts or Kubernetes clusters
- **IPv4-v6 Stateless NAT** - hardware implementation of MAP-T, enabling stateless conversion between IPv4 and IPv6 for v4-free SP backbones
- **Broadcast Media Tools** - enabling features like clean switching directly in the network
- **Customer Designed Profiles** - enables customers with specific use cases to write and deploy their own pipeline programs

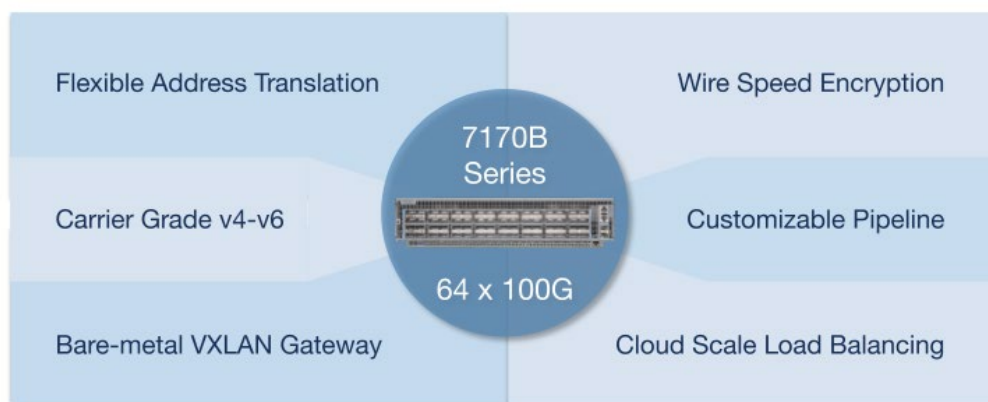


Figure 1: 7170 Series Solutions

Network services such as DDoS filtering, server load-balancing and NAT have traditionally required the use of dedicated appliances, router service modules, and more recently, multiple virtualized software instances.

High-scale network appliances and router service modules tend to be built from large numbers of x86 CPUs or Network Processors (NPU) which offer limited throughput with high latency in large, expensive, power-hungry form factors. In sharp contrast, the Arista 7170 Series replaces expensive, legacy networking appliances and routers by moving critical services directly into the network at wire-speed.

The multi-functional platform may also be deployed to increase server performance by off-loading vSwitch networking functions, such as tunnel termination, security policy enforcement and address translation onto the leaf switch. Embedding network services into the leaf or spine of the network eliminates traditional appliances and enables more efficient network topologies, avoiding the need to hairpin traffic through centralized clusters of traditional appliances. Combining these approaches, the 7170 Series helps improve the efficiency of the compute pool and removes network bottlenecks while lowering equipment, power and cooling costs.

The flexibility of the pipeline architecture further enables advanced users with innovative networking requirements to build their own custom profiles or to collaborate with Arista's engineering team on new use cases. New profiles can be installed and activated on the switch without rebooting or bringing down active links.

The combination of Arista 7170 series and EOS with an open, programmable SDK creates a fungible networking data plane, eliminating costly and disruptive hardware upgrades.

Arista 7170 Series

The 7170 Series is available in various configurations to suit modern datacenter network requirements. The 7170 Series are optimized for deployment, as either the leaf or spine tier, or as centralized service nodes consuming as little as 5W per 100G port.



Figure 2: 7170 Series

Each 7170 model is a fixed configuration system designed around a single high-performance System-on-Chip (SoC) packet processor. The packet processor pipeline comprises a programmable parser, fully programmable match-action engines, a centralized shared buffer and both TCAM and SRAM memory.

The 7170 Series uses the Intel Tofino chip, while the 7170B Series uses the Tofino2 variant. The chips have a common high-level architecture but differ by the number of supported ports and available TCAM, SRAM, buffer and pipeline stages. Figure 3, provides a high-level overview of the two chips:

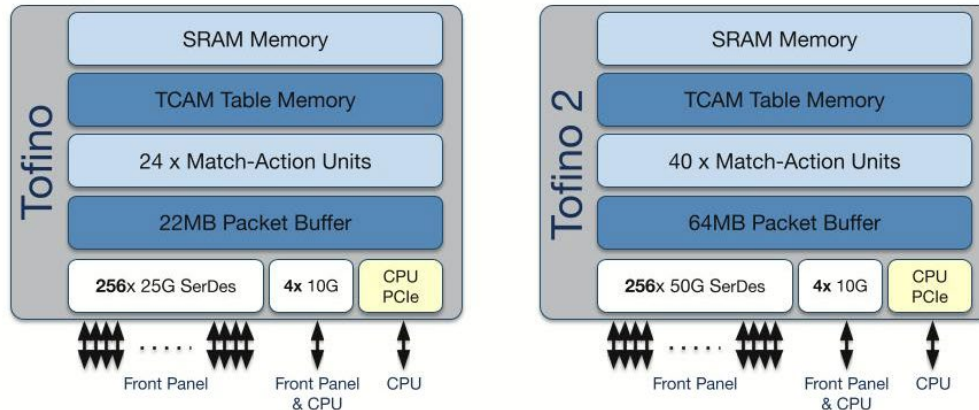


Figure 3: Intel Tofino and Tofino 2 Packet Processors

The Table below provides system details of each model of the 7170 Series:

Table 1: 7170 Series – System Specifications				
7170 Models	7170B-64C	7170-64C	7170-32C	7170-32CD
Switch Height (RU)	2	2	1	1
10G SFP+	2	2	2	2
25G SFP	--	--	--	--
40G QSFP+	64	64	32	32
100G QSFP	64	64	32	32
Maximum Density 10GbE ports	130	258	130	130
Maximum Density 25GbE ports	128	256	128	128
Maximum Density 40GbE ports	64	64	32	32
Maximum Density 50GbE ports	128 (50G-1 or -2)	128	64	64
Maximum Density 100GbE ports	64	64	32	32
Throughput (FDX)	6.4 (12.8) Tbps	6.4 (12.8) Tbps	3.2 (6.4) Tbps	3.2 (6.4) Tbps
Maximum Forwarding Rate (Bpps)	3.2	4.8	2.4	2.4
Latency	From 1 usec	From 800 nsec	From 800 nsec	From 800 nsec
P4 Programmable Stages	80	24	48	24
Packet Buffer Memory	64 MB	22 MB	22 MB	20 MB

Arista 7170 System Architecture

Each member of the Arista 7170 Series is designed with a high-performance multi-core x86 CPU connected to internal flash, boot-flash, power supplies, fans, management I/O and peripherals as presented below.

The x86 CPU is further connected over both PCIe and 10GbE to the Tofino or Tofino2 System-on-Chip that performs all the data plane forwarding and connects directly to all the front panel ports.

As the 7170 family is designed for innovative use cases, each system features a robust multi-core x86 processor along with up to 64GB of memory and up to 120GB of SSD storage.

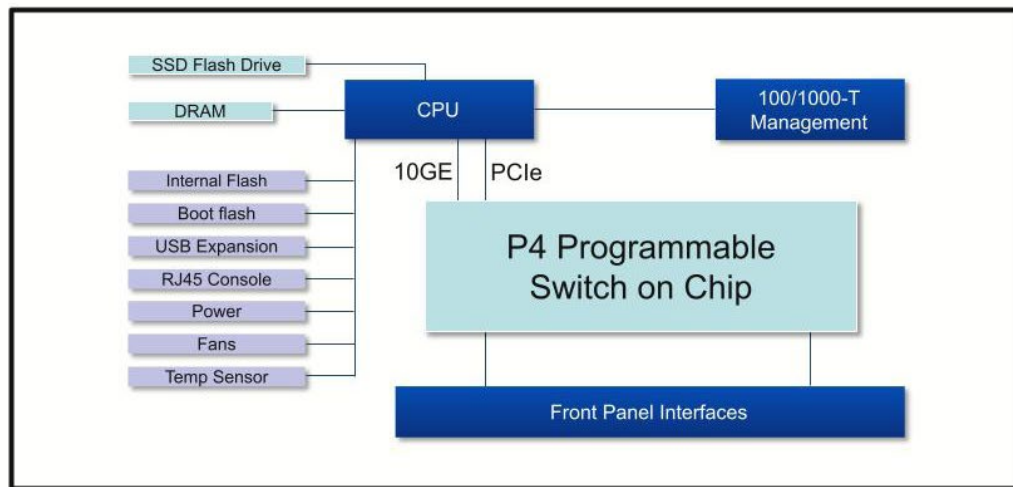


Figure 4: 7170 Series Platform Architecture

Arista 7170 Model Overview

The 7170 Series consists of 4 different models based on a common high performance fully programmable packet processor, combining broad forwarding resources with a flexible packet pipeline and rich EOS features. For each port density option, there are both high and regular scale variants. High scale variants have more pipeline stages available for more complex features and larger table scaling.

Arista 7170B-64C

The 7170B-64C is a 2U 64-port QSFP100 switch with 80 programmable stages and up to 3.2 billion packets per second with 6.4 Tbps of throughput.

All ports are 100G and 40G capable, accept a choice of optics or copper cables. Half of the QSFP ports can be configured for breakout to 4 x 10/25G. An additional 2 SFP+ ports allow for individual 10G dedicated connections. As shown below in Figure 5, additional front panel ports include a 100/1000 Ethernet interface for out-of-band management, a serial RJ45 connection and a USB port.

- 2RU with 64 ports of 100G QSFP and 2 ports of SFP+
- Intel Tofino2 12.8T chip used in 6.4T configuration
- Offers a choice of port combinations with QSFP+ and 100G QSFP optics and cables
- Flexible interface combinations - 64x 100G/40G, 128x 25/10G, 128x 50G
- 100G QSFP ports for 5 speeds for flexible 10GbE, 25GbE, 40GbE, 50GbE and 100GbE
- MACsec support on all QSFP ports
- 6.4Tbps of wire speed performance with 64MB of buffer (12.8Tbps full duplex)
- 80 stage programmable pipeline
- Front to rear cooling, AC and DC power supply options

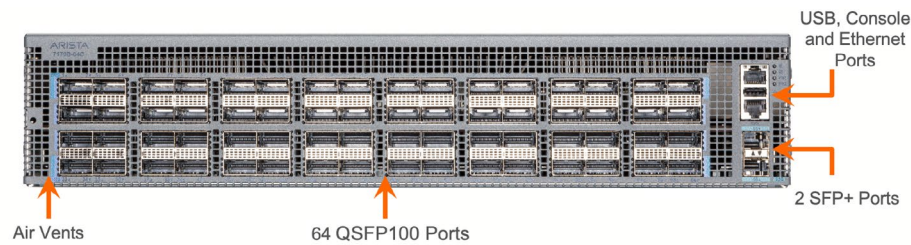


Figure 5: 7170B-64C Switch Front and Rear

The 7170B-64C is available with front-to-rear airflow and has 4 hot-swap fans and dual AC or DC power supplies. Figure 6 below illustrates the 7170B-64C system architecture.

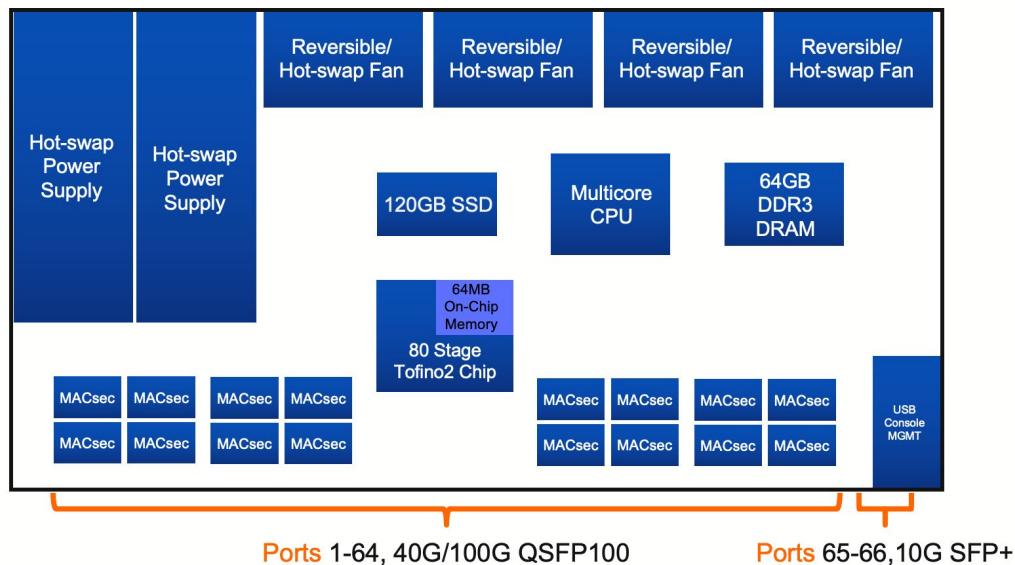


Figure 6: 7170B-64C System Architecture

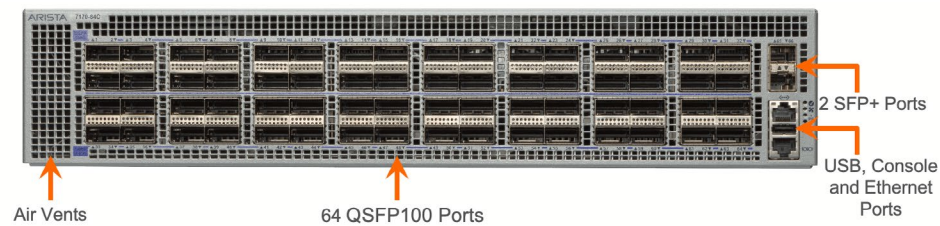
Arista 7170-64C

The 7170-64C is a 2U 64-port QSFP100 switch with 24 programmable stages and up to 4.8 billion packets per second with 6.4 Tbps of throughput.

All ports are 100G and 40G capable, accept a choice of optics or copper cables, and can be configured for 10/25/40/50/100G. An additional 2 SFP+ ports allow for individual 10G dedicated connections. As shown below in Figure 7, additional front panel ports include a 100/1000 Ethernet interface for out-of-band management, a serial RJ45 connection and a USB port.

- 2RU with 64 ports of 100G QSFP and 2 ports of SFP+
- Intel Tofino 6.4T chip used in 6.4T configuration
- Offers a choice of port combinations with QSFP+ and 100G QSFP optics and cables

- Flexible interface combinations - 64x 100G/40G, 256x 25/10G, 128x 50G
- 100G QSFP ports for 5 speeds for flexible 10GbE, 25GbE, 40GbE, 50GbE and 100GbE
- 6.4Tbps of wire speed performance with 22MB of buffer (12.8Tbps full duplex)
- 24 stage programmable pipeline
- Front to rear or rear to front cooling, AC and DC power supply options



64 x QSFP100 / 2 x SFP+: 64x 40/100G to 256x 25G

Figure 7: 7170-64C Switch Front and Rear

The 7170-64C is available as either front-to-rear or rear-to-front cooled and has 4 hot-swap fans and dual AC or DC power supplies. Figure 8 below illustrates the 7170-64C system architecture.

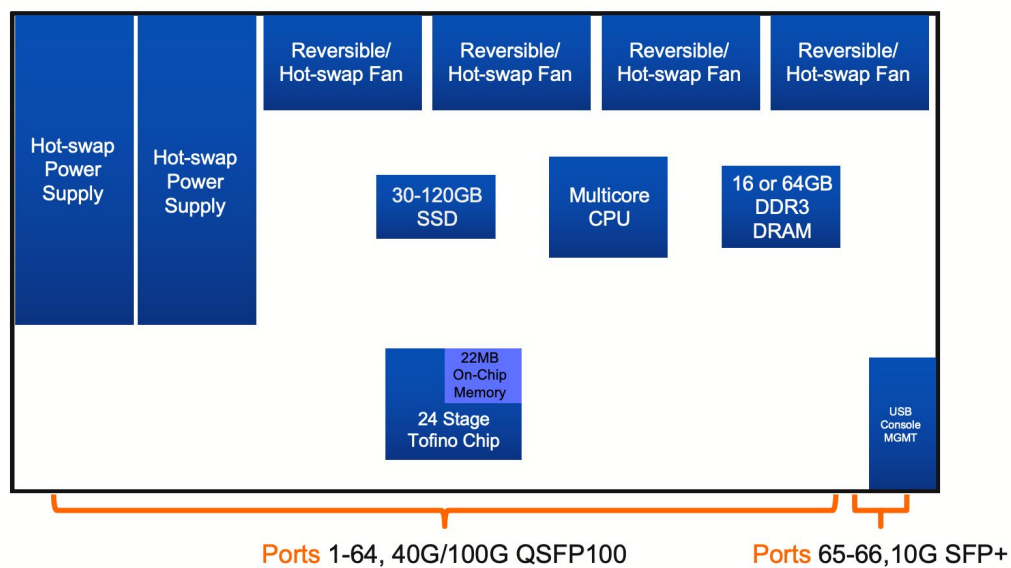


Figure 8: 7170-64C System Architecture

7170-32C

The Arista 7170-32C is a 1RU 32-port QSFP100 switch with 48 programmable stages and up to 2.4 billion packets per second and 3.2 Tbps of throughput. The 7170-32C offers a higher number of programmable pipeline stages for more complex or higher scale profiles.

As with the 7170-64C, all ports support a combination of optics and cables for interface speeds from 10G to 100G. The 7170-32C have two SFP+ ports for 1/10G dedicated connections on the front. At the rear is a 100/1000 Ethernet interface for out-of-band management, a serial RJ45 connection and a USB port.

- 1RU with 32 ports of 100G QSFP and 2 ports of SFP+
- Intel Tofino 6.4T chip used in 3.2T configuration
- Offers a choice of port combinations with QSFP+ and 100G QSFP optics and cables
- Flexible interface combinations - 32x 100G/40G, 128x 25/10G, 64x 50G
- 100G QSFP ports for 5 speeds for flexible 10GbE, 25GbE, 40GbE, 50GbE and 100GbE
- 3.2Tbps of wire speed performance with 22MB of buffer (6.4Tbps full duplex)
- 48 stage programmable pipeline
- Front to rear or rear to front cooling, AC and DC power supply options

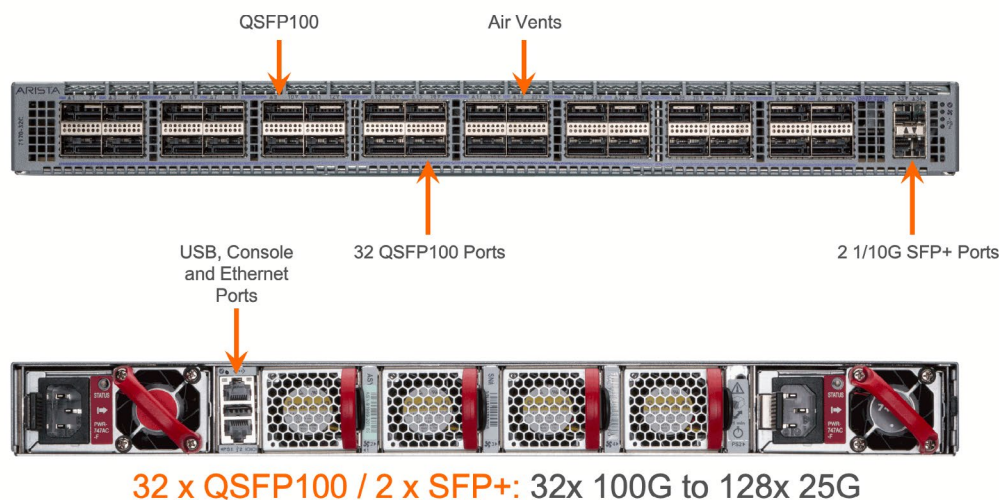


Figure 9: 7170-32C Switch Front and Rear

The 7170-32C models, seen above in Figure 9, are available as either front-to-rear or rear-to-front cooled, and have 4 hot-swap fans and dual AC or DC power supplies. Figure 10, below, presents the 7170-32C systems architecture.

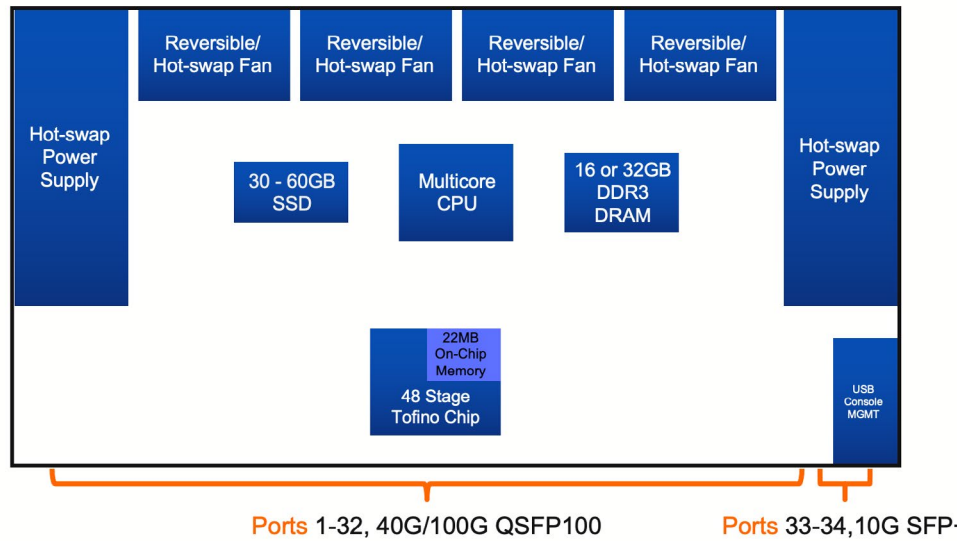


Figure 10: 7170-32C System Architecture

7170-32CD

The Arista 7170-32CD is a regular scale variant of the 7170-32C with 24 pipeline stages. It offers the same port count and interface flexibility as is detailed below.

- 1RU with 32 ports of 100G QSFP and 2 ports of SFP+
- Intel Tofino 3.2T chip used in 3.2T configuration
- Offers a choice of port combinations with QSFP+ and 100G QSFP optics and cables
- Flexible interface combinations - 32x 100G/40G, 128x 25/10G, 64x 50G
- 100G QSFP ports for 5 speeds for flexible 10GbE, 25GbE, 40GbE, 50GbE and 100GbE
- 3.2Tbps of wire speed performance with 20MB of buffer (6.4Tbps full duplex)
- 24 stage programmable pipeline
- Front to rear or rear to front cooling, AC and DC power supply options

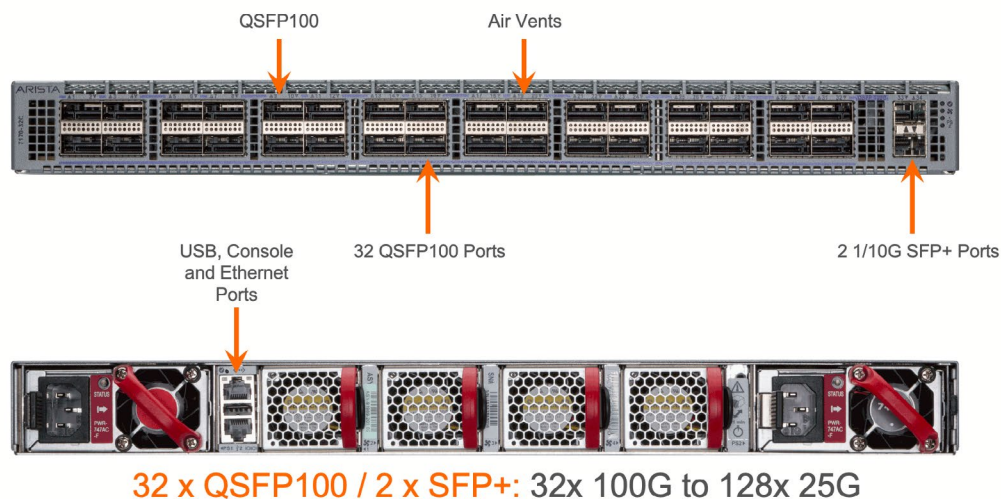


Figure 11: 7170-32CD Switch Front and Rear

The 7170-32CD models, seen above in Figure 11, are available as either front-to-rear or rear-to-front cooled, and have 4 hot-swap fans and dual AC or DC power supplies. Figure 12, below, presents the 7170-32C systems architecture.

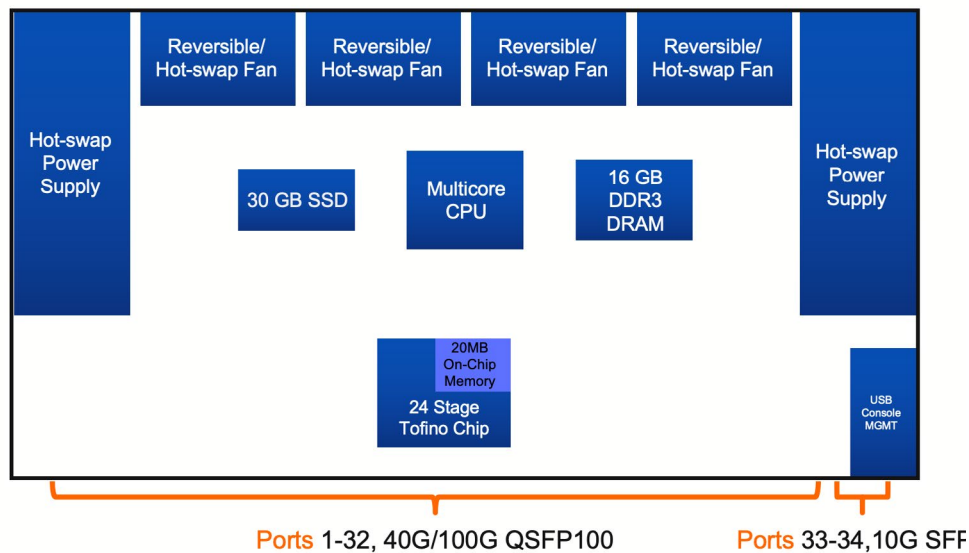


Figure 12: 7170-32CD System Architecture

Data Center Grade Availability and Redundancy

Arista 7170 Series switches are designed for continuous operations with system-wide monitoring of hardware and software components, simple serviceability, and provisioning that prevent single points of failure. Key high availability features include:

- 1+1 hot-swappable power supplies along with four hot-swap fans provide dynamic temperature control combined with N+1 redundancy as illustrated in Table 2,
- Color coded PSU's and fans that deliver platinum-level power efficiency,
- Both rear-to-front and front-to-rear airflow options,
- AC and DC power options,
- Live software patching,
- Self-healing software with Stateful Fault Repair (SFR), and
- Flexible pipeline profiles that do not require a system reload.

Table 2: 7170 Series - Power and Fan Redundancy

7170 Models	7170B-64C	7170-64C	7170-32C	7170-32CD
Power Supplies (Redundancy)	2 hot swappable (1+1)		2 hot swappable (1+1)	
AC / DC Supplies	AC and DC options		AC and DC options	
Fans (Redundant)	4 hot swappable (N+1)		4 hot swappable (N+1)	
Airflow	Front to rear	Front to rear & rear to front	Front to rear & rear to front	

Scaling Control Plane Services

The central CPU complex on the 7170 Series switches, as presented below in Table 3, is used for all control-plane and management functions while data-plane forwarding logic occurs in the programmable packet processor. Arista EOS®, the control-plane software for all Arista switches, executes on multi-core x86 CPUs with 16-64GB of DRAM. As EOS is multi-threaded, runs on a Linux kernel and is extensible, the large system memory and fast multi-core CPUs support an efficient control plane with headroom for running third-party software, either within the same Linux instance as EOS or within a guest virtual machine or container.

Out-of-band management is available via a serial console port and/or the 10/100/1000 Ethernet management interface. The 7170 Series also offer USB2.0 interfaces that can be used for various functions, including the transfer of images or logs.

Table 3: 7170 Series CPU Complex

7170 Models	7170B-64C	7170-64C	7170-32C	7170-32CD
CPU	Multi-Core x86	Multi-Core x86	Multi-Core x86	Multi-Core x86
System Memory	64GB	64GB or 16GB	32GB or 16GB	16GB
SSD	120GB	120GB or 30GB	60GB or 30GB	30GB

Programmable and Flexible Packet Processor

Each of the 7170 Series systems is built using a single System-on-Chip packet processor that provides both the ingress and egress forwarding for all packets that arrive on or are destined to ports located on the packet processor. All stages of the packet forwarding are performed entirely in the hardware/data plane. The 7170 architecture enables a programmable parser and pipeline along with flexible resource allocation through a series of profiles.

Figure 13 shows key components of the packet processor and the functionality provided by each component. Each switch chip is constructed from multiple parallel identical pipelines, interconnected via the centralized shared buffer. The number of pipelines and pipeline stages varies between each 7170 model, depending on the target use case and model of switching silicon.

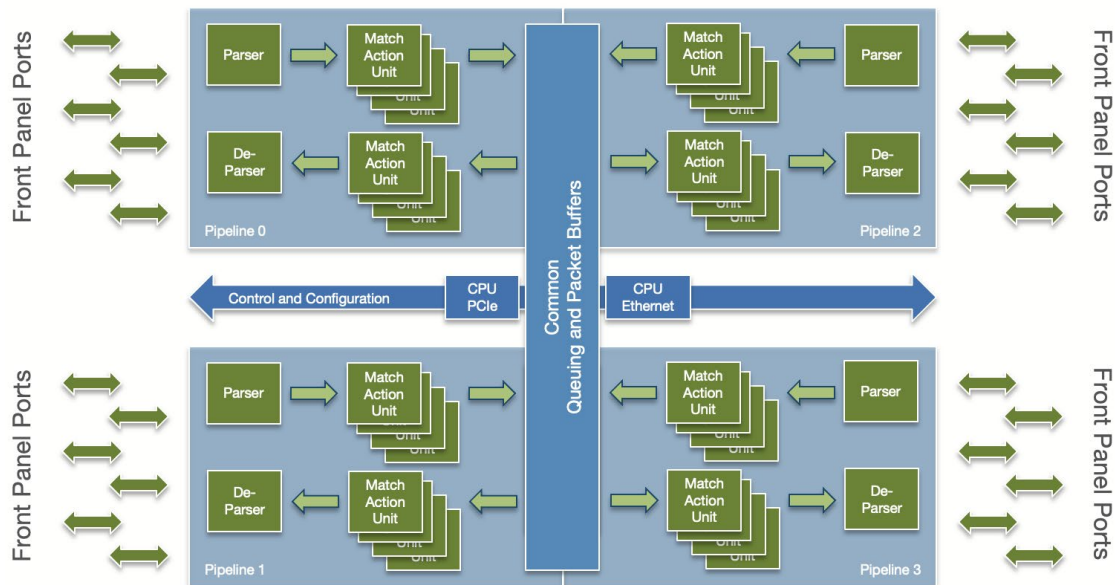


Figure 13: 7170 Series Packet Processor Architecture

Each pipeline supports a number of front panel ports, up to 32 x 100G in the 7170B, while in certain models, additional internal pipelines are present which can be used by pipeline profiles to deploy additional functionality and scaling.

Packets arriving at the ingress MAC are processed by the corresponding ingress match-action pipeline and then enqueued in a common shared memory block which connects the ingress and egress ports. On being scheduled for transmission, packets are passed through the egress match-action pipeline to the transmit MAC.

The following section, as illustrated in Figure 14, provides a detailed look at packet processing through the pipeline.

First, the parser analyzes incoming packets, identifying fields of interest (e.g., MAC/IP addresses) and maps them to locations within per-packet metadata labels called Packet Header Vectors (PHVs). PHVs are then passed through a sequence of match-action units.

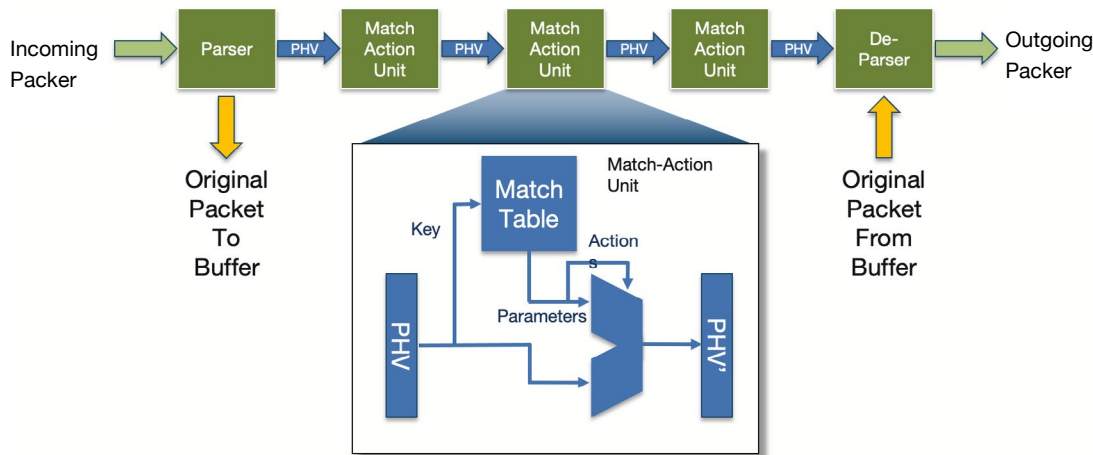


Figure 14: 7170 Series Match Action Unit

The MAUs contain Arista-defined logic and lookup tables that are used to make forwarding and packet rewrite decisions. The series of match-action logic, order, operation, and size and structure of the memory tables are defined in the data plane profile. As PHVs pass through match-action units, keys are extracted from the set of packet fields. Each match-action unit has dedicated SRAM and TCAM banks, which can be assigned to specific functions. Table hits can trigger counters, policers, PHV writes and memory updates. As PHVs are passed through the match-action units, the pipeline state from one table can be used as the key to another. Finally, the de-parser combines the PHV metadata with the original packet header to assemble the modified packets for transmission.

7170B-64C

The 7170B-64C series is based on a 4-pipeline Tofino2 chip, presenting 64 x 100G interfaces to the front panel from two of the pipelines and retaining two further pipelines internally to provide support for more complex or higher scale pipeline profiles.

Each pipeline allows up to 20 processing stages on ingress and egress, enabling the 7170B-64C to offer a total of 80 processing stages. Figure 15 below illustrates 7170B-64C Series packet processor architecture.

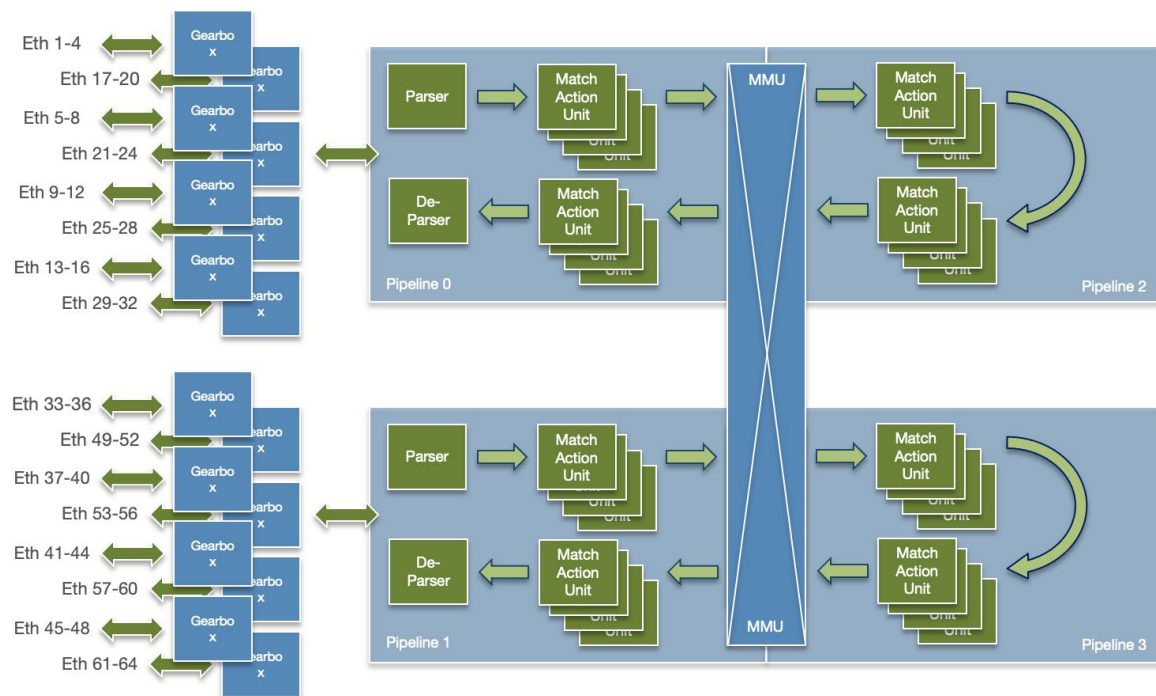


Figure 15: 7170B-64C Series Packet Processor Architecture

7170-64C

The 7170-64C series is based on a 4-pipeline Tofino chip. It presents all four pipelines to the front panel, each with 16 x 100G ports for a total of 64 x 100G ports. Each pipeline allows up to 12 processing stages on ingress and egress for a total of 24 processing stages. Figure 16 below illustrates 7170-64C Series packet processor architecture.

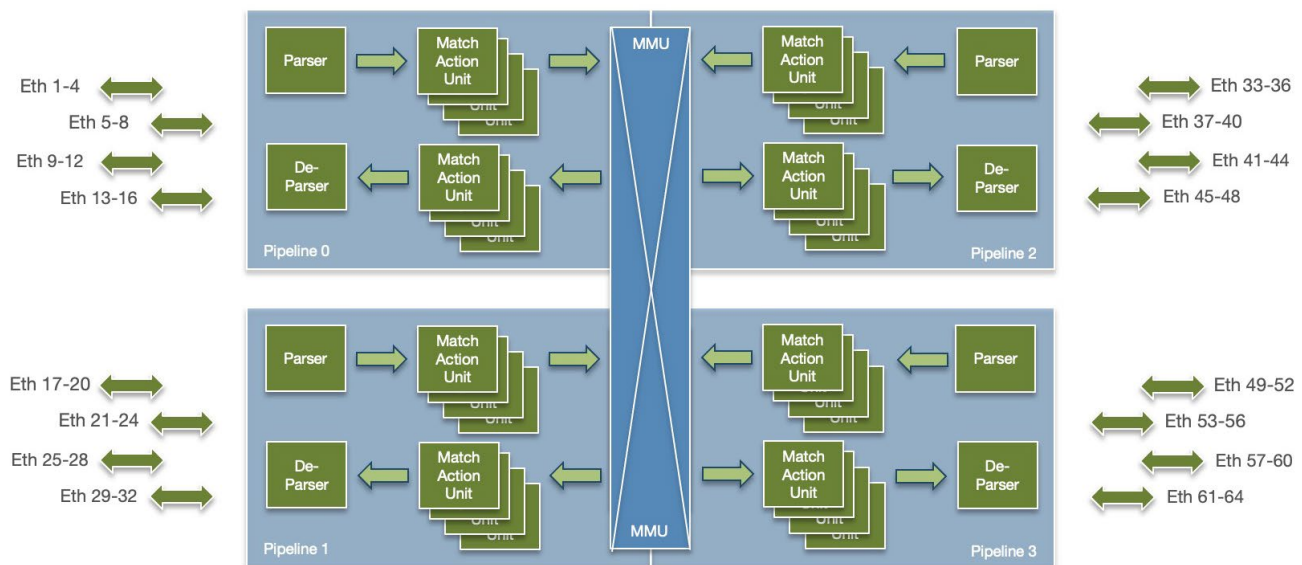


Figure 16: 7170-64C Series Packet Processor Architecture

7170-32C

The 7170-32C also has a four-pipeline architecture, with two pipelines assigned for the 32 x 100G ports. Unlike the 7170-64C, two additional internal pipelines are available for extra processing, doubling the total number of processing stages to 48. Packets are forwarded to the extra stages as needed, before being scheduled to the egress match-action pipeline. Figure 17 below shows the 7170-32C Series packet processor architecture.

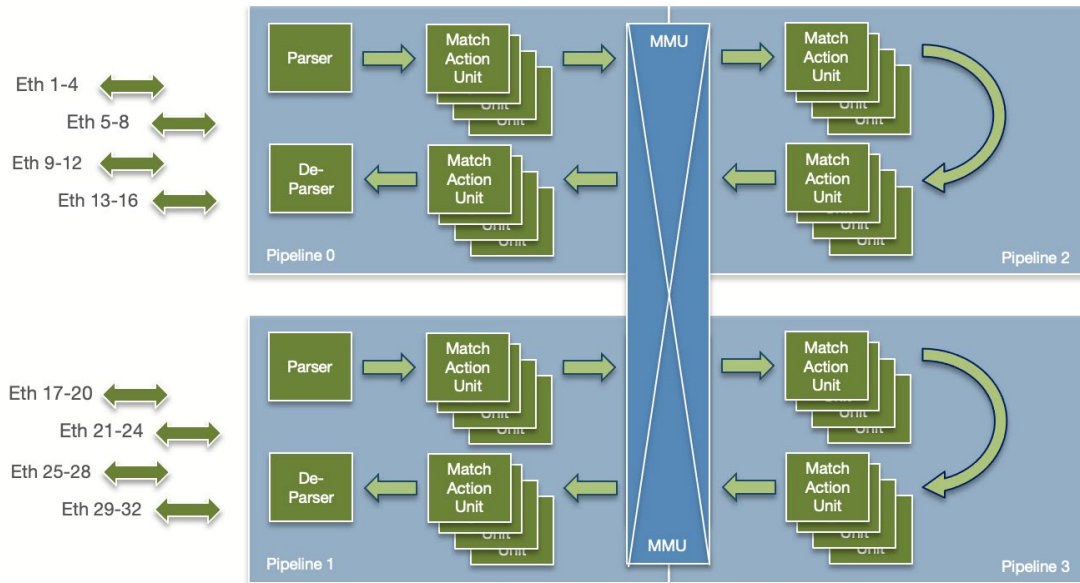


Figure 17: 7170-32C Series Packet Processor Architecture

7170-32CD

The 7170-32CD has a two-pipeline architecture, with two pipelines assigned for the 32 x 100G ports. Unlike the 7170-32C, it does not offer additional internal pipelines available for extra processing and therefore provides the equivalent 24-stage processing capability as the 7170-64C but in a half-sized form factor. Figure 18 below shows the 7170-32CD Series packet processor architecture.

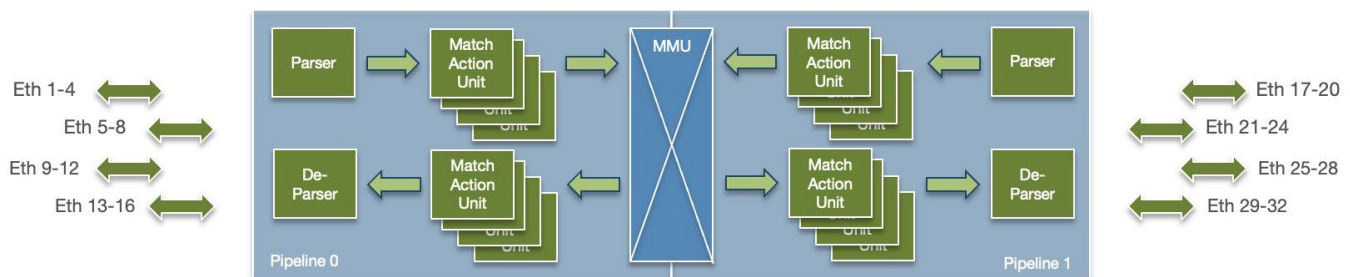


Figure 18: 7170-32CD Series Packet Processor Architecture

Shared Buffer Architecture

The 7170 Series incorporates an advanced traffic manager with up to 64MB of packet buffer that is fully shared across all ports, making it an excellent choice for data center applications. Unlike legacy architectures where the total buffer is statically allocated to a port or group of ports, or the packet buffer memory is segmented into slices, the 7170 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.

Flexible Switch Forwarding Profiles

As with all products in the Arista portfolio, a single common EOS image supports the 7170 series, ensuring a common control plane for feature consistency across multiple system architectures to deliver higher quality and simplified network operations.

To leverage the programmable capability of the 7170 Series pipeline, the 7170 series introduces a series of forwarding profiles that are packaged with EOS. Profiles are used to define the data plane. Each profile is optimized to provide either a unique feature-set or to maximize the set of forwarding services and scale for a specific use case.

New profiles are introduced through EOS upgrades or by loading a new RPM directly to the system. Toggling between profiles is hitless, with just a brief pause in packet processing without rebooting or restarting the system.

While profiles are defined in P4, and used to program the data plane, EOS is still the operating system for control and management planes, with a set of open programmable interfaces and SDK.

Changing Profiles

Changing profiles is a simple operation and requires the use of just a single command as shown below in Figure 19. The “default” profile is replaced by the “baremetal” profile as the active profile although both remain available.

```
al108(config)#show platform barefoot profile
  Profile    Chips
  -----
  baremetal
  default    Bfn0
  -----

al108(config)# platform barefoot profile baremetal Bfn0

al108(config)#show platform barefoot profile
  Profile    Chips
  -----
  baremetal   Bfn0
  default
  -----
```

Figure 19: 7170 Series Profile Change

Multi-Tenant Bare Metal Hosting

Bare-metal hosting is a public cloud service that lets customers rent hardware resources from a service provider. Flexibility is a key benefit that bare-metal cloud provides to customize a bare-metal system to meet a customer’s unique requirements without interference from neighboring virtual machines (VMs). Bare metal environments are ideal for big data applications or high-transaction workloads that do not tolerate latency.

In a bare-metal environment, additional segmentation, security, encapsulation/tunnel and overlay services are enabled in either dedicated hosts or appliances with all traffic directed through these systems, adding to the latency and lowering performance. The Arista 7170 bare metal profile, however, offloads the virtualized services into the network, ensuring consistent performance and comprehensive services.

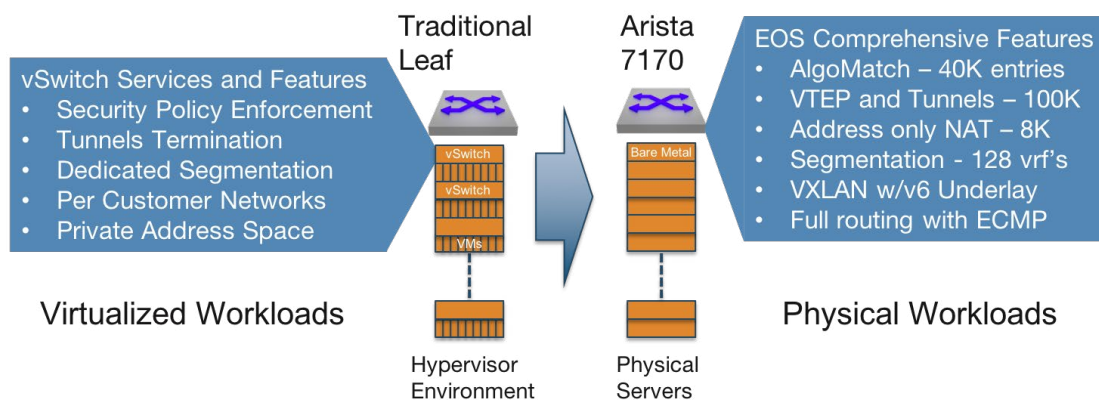


Figure 20: Multi-tenant Bare Metal Profile

Cloud Load Balancing

Traditional Application Delivery Controllers (ADC) offer a rich set of processing logic and offload capability for homogenous data center environments, where applications and services have not been designed with scale-out in mind. The extra functionality comes with the penalty of complex products with limited throughput and higher CAPEX and OPEX costs.

Modern high-bandwidth scale-out applications such as web farms, cloud infrastructure and microservices are already intrinsically designed for elastic growth and end-to-end security. Much of the functionality provided by traditional ADCs is not required, instead cost effective scaling, higher throughput and advanced redundancy models are key.

Operating as a Load Balancer, the 7170 can focus on these scale-out use cases, providing a unique platform for multi-terabit load balancing in a compact form factor and at a fraction of the cost, size and power of traditional ADCs. Combined with Arista EOS's industry leading reliability and flexible programmability to integrate smoothly into large scale highly automated environments.

IPv4 to v6 Transition

As IPv4 address space becomes more scarce and more expensive to procure, optimal deployment and reuse of existing allocations is critical. For Service Providers, options include removing IPv4 entirely from core networks and leveraging large scale Carrier Grade NAT (CGNAT) solutions to provide access to the predominantly IPv4 Internet while conserving IPv4 addresses.

CGNAT solutions that support the required millions of concurrent sessions may help to alleviate IPv4 exhaustion but introduce several challenges for Service Providers. Scaling to the very large data rates required for triple-play services, assuring resilience and accurately recording every user session for audit and lawful intercept purposes requires a large fleet of equipment and an onerous level of data collection.

As an alternative to stateful address translation, the 7170 Series is ideally suited for wire-speed stateless translation between IPv4 and IPv6. Stateless translation removes the need for extensive logging (as all translations are predictable) and enables multi-level fault tolerance without inter-device synchronization.

The 7170 is capable of supporting both of the predominant stateless translation mechanisms (MAP-T and MAP-E), offering up to 6.4Tbps of translation capacity per device in a fraction of the footprint of CGNAT solutions.

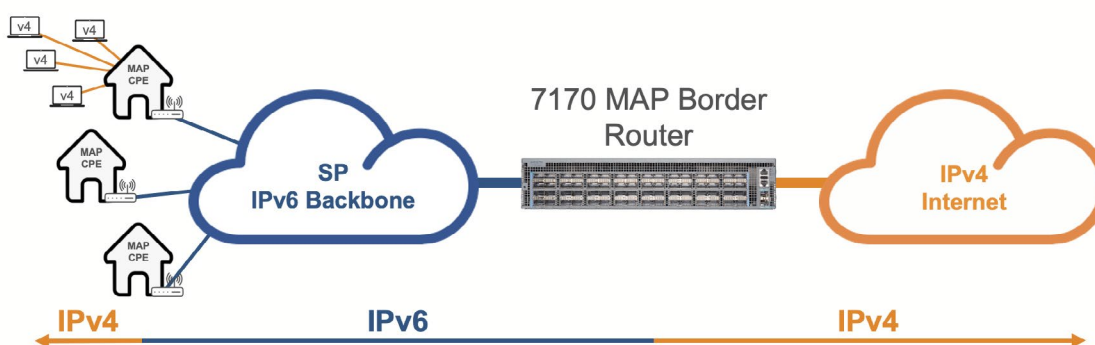


Figure 21: 7170 MAP Border Router for IPv4 over IPv6 only core networks

Example of Cost Advantages

SmartNIC Offload

A typical rack design consisting of servers with 100G capable SmartNICs connecting to a traditional, fixed pipeline top-of-rack switch with 32 ports of 100G is approximately 30% costlier than one with regular 100G NICs connecting to an Arista 7170 series top-of-rack switch. If the comparison is made with 25G capable SmartNICs, the cost difference increases to 45%. The cost difference varies between 40%-50% when using a 64-port 100G fixed pipeline top-of-rack switch with SmartNICs compared to a 64-port Arista 7170 series switch. Cost savings can also be attributed to centralizing the cost of memory and equipment needed to support a programmable networking dataplane per rack using Arista 7170 series, rather than on every server using SmartNICs.

Routers and Services Modules

Many traditional router platforms can deliver additional packet processing through re-circulating a packet through the processor until the set of operations are complete. This is termed a run-to-completion architecture as compared to the pipeline architecture of the 7170 Series. The router throughput linearly reduces with the additional processing steps, and the ability to add new services is limited by the processor specifications and the system vendor, affecting the flexibility.

An alternate approach is to deliver advanced services using “Services Modules” that contain one or more dedicated CPUs, NPUs or FPGAs that are designed to perform intensive packet processing services that cannot be supported by run-to-completion packet processors. However, complex chipsets combined with high power consumption and an inability to provide throughput aligned to data center scale make service modules an expensive option.

A typical service module has less than 100Gbps of total aggregate throughput, despite being a dedicated line card with no physical interfaces. In contrast, the 7170 Series can provide over 64x greater throughput while supporting and accelerating rich services, such as application load balancing, large scale network address translation, tunnel encapsulation, and secure traffic-filtering. Additionally, the 7170 is optimized for the performance requirements of cloud scale data centers with a wide range of physical interfaces in dense form factors and power consumption from less than 5W per 100G port. When comparing the overall cost of a services module, the additional feature licenses and higher operational costs also need to be considered. At best, the service modules offer one tenth the performance for two to three times the cost, resulting in a large cost advantage of the 7170 Series.

Dedicated Appliances

Traditional Application Delivery Controllers (ADC) are complex devices, designed to offer much more than simple load balancing by offering a wide array of services to enhance traditional applications (such as SSL offload) in a transparent fashion. They solve many problems for legacy applications and heterogenous data centers by acting as a centralized resource for enforcing policy and implementing features not supported by the original application or server.

As this extra functionality requires additional processing and the tracking of a significant amount of state information, the typical solution to this problem is to use x86, NPU or FPGA based hardware with large amounts of external memory to track and manipulate all traffic flows. The downside is that each CPU, NPU or FPGA provides limited packet processing throughput, so high bandwidth load balancers tend to consist of large numbers of processors and memory resulting in a large physical footprint, complex software and making them costly to purchase and operate.

For modern, large-scale applications such as web and container farms, the extra functionality is redundant since applications are typically already designed for end-to-end security and are inherently distributed and fault tolerant. Using traditional ADCs for such deployments creates a cost and scaling bottleneck relative to the overall infrastructure.

Switching silicon performance is several orders of magnitude higher than a typical CPU and the flexibility of the 7170 Series Load Balancer solution introduces an alternative stateless load balancing architecture with extremely high throughput that supports up to 6x the capacity of a high-end ADC platform while consuming a fraction of the power and space.

Arista EOS: A platform for scale, stability and extensibility

At the core of the Arista 7170 Series is Arista EOS®. Built from the ground-up using innovations in core technologies since our founding in 2004, EOS contains more than eight 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: non-stop availability, and high feature velocity coupled with high quality software.

Drawing on our engineers' extensive over 30-year experience in building networking products and state-of-the-art open systems technology and distributed systems, Arista started with 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 that separates state information from packet forwarding, protocol processing and application logic. In EOS, system state and data is stored and maintained in a highly efficient, centralized System Database (NetDB). The data stored in NetDB is accessed using an automated publish/subscribe/notify model. This architecturally distinct design principle provides self-healing resiliency in Arista's software, easier software maintenance and module independence, higher software quality overall, and faster time-to-market for new features.

Arista EOS, as seen in Figure 17, contrasts with the legacy approach to building network operating systems developed in the 1980s 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 NetDB. 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.

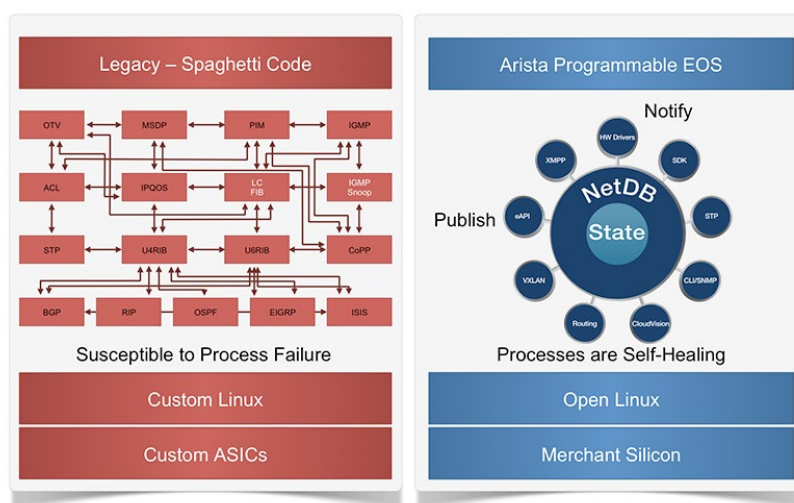


Figure 22: Legacy approaches to network operating systems (left), Arista EOS (right)

Arista took to heart the lessons of the open-source world and built EOS on top of an unmodified Linux kernel maintaining 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 architectural approach lets EOS natively support, for example, Docker Containers that simplify the development and deployment of applications on Arista switches. Arista EOS represents a simple but powerful architectural approach that provides a robust platform to swiftly deliver significant new features.

EOS is extensible at every level, with open APIs at the management plane, control-plane, and data-plane, coupled with services-level extensibility, application-level extensibility, and 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) along with OpenConfig and EOS SDK provide well-documented and widely used programmatic access to configuration, management and monitoring that can stream real-time network telemetry for a superior alternative to traditional polling mechanisms.

Conclusion

The 7170 Series represents a unique solution to many of the challenges of operating large scale applications and networks. Arista's open and extensible EOS combined with the flexibility and multi-function programmability offered by Arista 7170 series, is engineered to address real-world use cases for cloud providers, enterprises, service providers and content delivery networks, eliminating the need for costly, application-specific hardware. Deploying 7170 can help realize substantial CapEx and OpEx savings, while removing bottlenecks and improving the throughput of production traffic by orders of magnitude.

To learn more about the 7170 Series, please contact your Arista representative or visit www.arista.com.

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 © 2022 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. January 20, 2020 02-0079-05