

Investment Protection with the Arista 7500 Series

Arista Networks award-winning 7500 Series was introduced in April 2010 as a revolutionary switching platform, which maximized datacenter performance, efficiency and overall network reliability. It raised the bar for switching performance being five times faster, one-tenth the power draw and one-half the footprint compared to other modular datacenter switches.

Just three years later, the introduction of the Arista 7500E Series linecard and fabric modules delivers a three-fold increase in density and performance with no sacrifices on functionality, table sizes or buffering with an industry-leading 1,152 x 10GbE, 288 x 40GbE or 96 x 100GbE in the same quarter-rack 11RU chassis.

This whitepaper details the investment protection capabilities of the Arista 7500 E-Series linecards and fabrics and how they utilize the Arista 7500 chassis and interoperate with existing Arista 7500 linecards and fabric modules.

Investment Protection Summary

Where many vendors promise investment protection but fail to deliver (or only provide ‘ships in the night’ which actually means that they “can coexist in the same chassis but not forward between one another”), the Arista 7500 series delivers complete interoperability and investment protection across all components that make up an Arista 7500 series switch:

- There is no change to the chassis or power-supplies
- Upgrading a first-generation Supervisor to a second-generation Supervisor-E is optional (note: Supervisor-E offers a more than three-fold increase in switch control-plane performance, so it may be desirable to upgrade)
- Mixing/matching first and second generation linecard modules with distributed data-plane forwarding is supported
- First-generation linecard modules are forward compatible with second-generation fabric modules. Likewise, second-generation linecard modules are backward compatible with first-generation fabric modules.
- In addition to complete interoperability it’s also supported to perform an in-place upgrade from first-generation to second-generation fabric modules without disruption to data-plane forwarding.



Figure 1: 7504E and 7508E with up to 1,152 10G ports, 288 40G ports or 96 100G ports

The investment protection provides both a flexible migration and upgrade path for Arista 7500 switches and fulfils Arista’s promise of future investment protection without expensive forklift upgrades.

Upgrade Considerations

The following section contains details on interoperability and any steps that are required.

Minimum EOS Software Release

A minimum software release of EOS 4.12.0 must be running on a switch in order to support a second-generation linecard or fabric module.

The minimum supported software release required for a Supervisor-E is EOS 4.12.0.

Supervisor-E ships with EOS 4.12.0 or later installed.

A minimum software release of EOS must be running on the switch in order to support interoperability of first/second generation linecard modules/fabrics. Refer to Arista EOS release notes for details.

Arista 7500 Series Chassis

A minimum software release of EOS 4.12.0 must be running on a switch in order to support a second-generation linecard or fabric module.




Figure 2: Arista 7508 and Arista 7504 Chassis

There is no change to the 7500 Series chassis. Designed with future proofing in mind, the existing Arista 7508 11RU and 7504 7RU chassis are capable of supporting the three-fold increase in performance and density enabled by second-generation 7500E linecard and fabric modules.

Arista 7500 Power Supplies



Figure 3: Arista 7500 Series Power Inputs and Power Supply

There is no change to the existing 7500 Series 2900W AC Power Supplies.

Arista 7500 second-generation linecard and fabric modules provide a three-fold increase in the performance and density of ports within a single system while remaining within the same system-level power envelope. The end result is a two-thirds reduction in the power/port on what was already an industry-leading power consumption/port metric.

The existing hot-swappable 2+2 redundant 2900W Gold Climate-rated Power Supplies (92% efficient) remain the same.

Arista 7500 Supervisor-E

The 7500 Series Supervisor-E provides a more than three-fold increase in control-plane performance compared to the original Arista 7500 Supervisor (increased from a dual core to a hyper-threaded quad core), and more memory (from 4GB to 16GB). Supervisor-E also adds dual out-of-band management interfaces, a PPS clock input and an optional enterprise grade SSD for extensibility in addition to the capabilities of the 7500-Supervisor (or Supervisor-1).



Figure 4: Arista Supervisor-E

While upgrading a Supervisor-1 to Supervisor-E is optional, Arista recommends that Supervisor-E be used if expanding a chassis beyond the original 384 ports of 10GbE supported on first-generation systems.

[Note: It's not possible to run a system with Sup-1 and Sup-E in the same chassis and Supervisor upgrades will be disruptive.]

Mixing First and Second-Generation Linecard Modules

Mixing both first and second-generation linecard modules in the same system and forwarding between the different generation modules is supported. As the forwarding engine on second-generation linecard modules has larger tables and more features compared to the forwarding engine on first-generation linecard modules, Arista EOS needs to know whether to limit the system capabilities to those of first-generation or not. This is controlled using the following commands:

```
switch(config)# [no] platform sand compatibility forwarding [PetraA | Arad]
switch(config)# write memory
switch(config)# reload
```

A reboot is required after making changes to this setting for it to take effect.

If forwarding compatibility is set to 'Arad' mode than any first-generation linecard modules will remain powered-down if present in the chassis and only second-generation linecard modules will operate.

If set to 'PetraA' mode then both first and second-generation linecard modules will function together however second-generation linecard modules will be capped to the capabilities of the first-generation forwarding engine as shown in table 1:

Table 1: Comparison of first and second-generation Forwarding Engine characteristics

Characteristic	7500 (PetraA)	7500E (Arad)
MAC Table Size	16K	256K
Max Host Route Prefixes	16K	128K
Max IPv4 Route Prefixes	24K	64K
Max IPv6 Route Prefixes	1K	12K+
Max Multicast Groups	16K	64K
Max Egress Forwarding Entries / Forwarding Engine	7.5K	30K
Max LAG Groups	256	up to 1K
Max LAG Members	up to 16 member ports	up to 32 member ports (initially)
Max ECMP Fan-out	16-way	32-way (64-way)
MLAG Peer-Link physical interfaces	Must be on 7500 linecards	Must be on 7500E linecards
PIM+MLAG	Not supported	Will be supported
MAC Table Size	16K	256K

Upgrading Fabric Modules from First Generation to Second Generation

Fabric modules located at the rear of the systems may be upgraded in-place from first to second-generation without disruption to data-plane forwarding. See figure 5.




Figure 5: Rear of 7508 and 7504 chassis with Fabric/Fan modules

The upgrade procedure is as follows:

1. Remove a single first-generation fabric module by pulling the fan tray/fabric on the module. [Note that there is no need to manually shut down the fabric module, the system recognizes the module is being removed and will gracefully move any traffic from the fabric before it has lost connectivity.]
2. Insert/install the new second-generation fabric/fan module in the open slot where the first-generation fabric module used to be.
3. In order to ensure N+1 redundancy and full fabric bandwidth, wait until the fabric and fan status LED is illuminated green on the newly inserted module before moving onto the next module. Observe the two LEDs on the module, one for the fans, another for the fabric; both should be illuminated green.




Figure 6: Fabric and Fan Status LEDs on Fan/Fabric modules

4. Repeat steps 1-3 across the remaining five fabric/fan modules.

For an in-place upgrade to be non-disruptive, ensure:

- There is only one fan/fabric module removed at a time
- A fabric/fan module slot is not left empty for longer than 60 seconds

Failure to follow these steps will result in a switch shutdown and service disruption.

While its possible for the system to operate in a mixed-fabric mode indefinitely (first and second-generation) Arista recommends that a system should consist of entirely first-generation or second-generation fabric modules as the system performance is based on the lowest-common-denominator of a group of fabric links.

An in-place upgrade from first-generation to second-generation fabric does not enable the full capabilities of the second-generation fabric as the encoding method and multicast expansion groups won't be at the full capabilities of second-generation as highlighted in table 2 below:

Table 2: Comparison of first and second-generation Fabric characteristics			
Fabric parameters for second generation linecard modules	FE600 (Fabric)	In-place upgrade Fabric to Fabric-E	FE1600 (Fabric-E)
In-place upgrade Fabric to Fabric-E	6.25 GHz	11.5 GHz	11.5 GHz
Fabric Links Active Per Linecard	144	192	192
Fabric Encoding Format	8b/9	64b/66	64b/66
Fabric Cell Format	VSC128	VSC128	VSC256
Max Multicast Expansion Groups in fabric	16K	16K	64K

In the majority of cases, the settings highlighted in red don't make any noticeable difference, as the major benefits of second-generation fabrics (more fabric links running at a faster speed with more efficient encoding) will be utilized even with an in-place upgrade. However, if you operate a network with significant amounts of line-rate traffic (>2.2Tbps/slot) and large numbers of multicast groups, then a reboot may be necessary to achieve optimal performance and scale after an in-place upgrade.

Once a system has been upgraded from first-generation to second-generation fabric modules, EOS assumes that you will continue to want to use second-generation fabric modules in future. If you intend to put first-generation fabric modules into the chassis in future then it will be necessary to configure EOS support and 'fe600 fabric compatibility' should be enabled with:

```
switch(config)# [no] platform sand compatibility fabric [fe600 | fe1600]
switch(config)# write memory
```

First Generation Linecard Modules with Second Generation Fabric Modules

First-generation linecard modules operate seamlessly with second-generation linecard and fabric modules by default. No explicit configuration is necessary.

Second Generation Linecard Modules with First Generation Fabric Modules

Mixing second-generation linecard modules with first-generation fabric modules will limit the second-generation linecard modules to some capabilities of the first-generation fabric.

Some of these capability limits are:

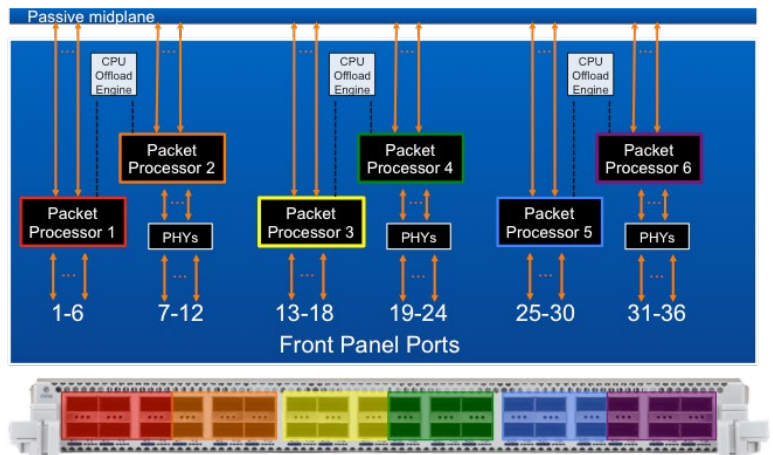
1. Fabric links operate at 6.25 GHz rather than 11.5 GHz (approximately 45% lower fabric capacity)
2. The encoding is 8b/9 rather than the more efficient 64b/66 (~approximately 9% less efficient)
3. Fabric maximum cell size is limited to 128 bytes rather than 256 bytes (~5% less efficient)
4. At most 144 active fabric links to a linecard rather than up to 192 active fabric links on the highest density modules (Up to 25% less performance)
5. Multicast expansion in the fabric is limited to 16K groups rather than 64K groups

Points (1) through (3) above mean that the crossbar fabric bandwidth capacity to second-generation modules is approximately half what it would otherwise be with second-generation fabric modules. Point (4) limits the maximum performance only on the highest density second-generation linecards.

How this converts to real performance depends on the individual modules themselves and the grouping of ports across packet processors and the number of operational fabric links:

DCS-7500E-36Q-LC (36 x QSFP+):

- Port groups of 6 x QSFP+ ports (6x40G) per forwarding engine.
- Linecard module performance with 6 x first-generation fabrics:
 - o Any 2 x QSFP+ (80 Gbps) can be used in each port group with N+1 fabric redundancy
 - o Up to 100 Gbps total bandwidth in each port group with no fabric redundancy (reduced performance on a fabric failure)



DCS-7500E-72S-LC (48 x SFP+ and 2 x MXP):

- Three port groups covering:
 - o SFP+ ports 1-20
 - o SFP+ ports 21-34 and MXP port 49/1-12
 - o SFP+ ports 35-48 and MXP port 50/1-12
- Linecard module performance with 6 x first-generation fabrics:
 - o Up to 120 Gbps in each port group with N+1 fabric redundancy
 - o Up to 150 Gbps in each port group with no fabric redundancy (reduced performance on a fabric failure)

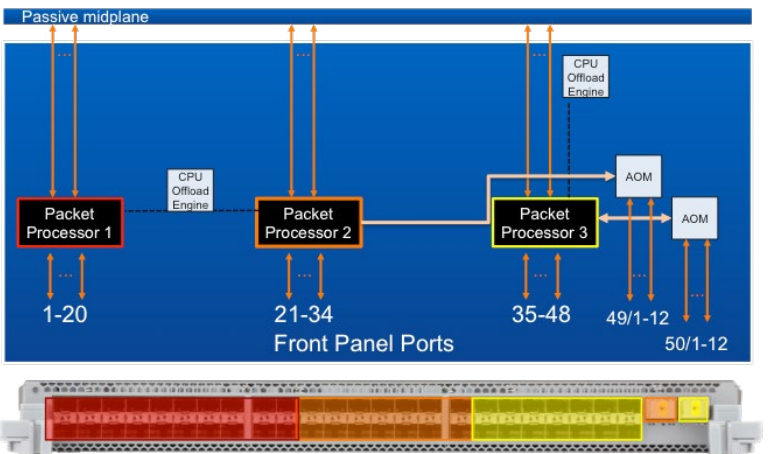


Figure 8: Arista DCS-7500E-72S-LC 48 x SFP+/SFP and 2 x MXP Linecard port groups

DCS-7500E-48S-LC (48 x SFP+):

- Three port groups covering:
 - o SFP+ ports 1-16
 - o SFP+ ports 17-32
 - o SFP+ ports 33-48
- Linecard module performance with 6 x first-generation fabrics:
 - o Up to 120 Gbps in each port group with N+1 fabric redundancy
 - o Up to 150 Gbps in each port group with no fabric redundancy (reduced performance on a fabric failure)

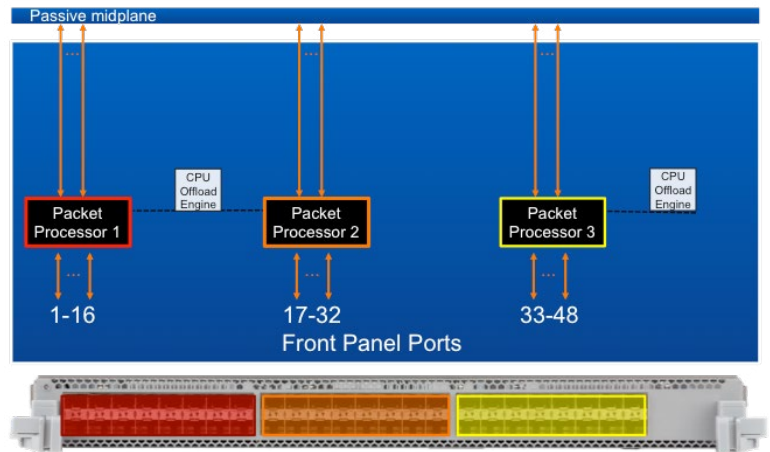


Figure 9: Arista DCS-7500E-48S-LC 48 x SFP+/SFP Linecard port groups

DCS-7500E-12CM-LC (12 x MXP):

- Port groups of 2 x MXP ports (1x100G/3x40G/12x10G) per forwarding engine.
- Linecard module performance with 6 x first-generation fabrics:
 - o Up to 80 Gbps can be used in each port group with N+1 fabric redundancy
 - o Up to 100 Gbps can be used in each port group with no fabric redundancy (reduced performance on a fabric failure)

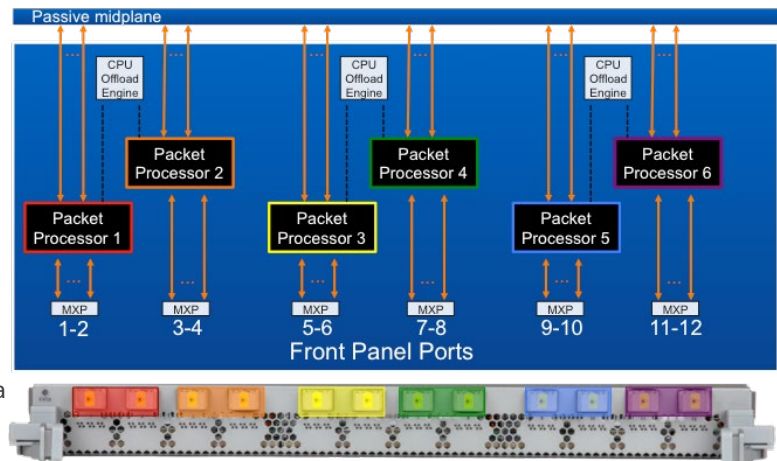


Figure 10: Arista DCS-7500E-12CM-LC 12 x MXP Linecard port groups

Conclusion

Where many vendors continue to promise investment protection and don't deliver, the Arista 7500 Series delivers complete forward and backward interoperability and investment protection across all components that make up an Arista 7500 Series switch.

This investment protection provides a flexible migration and upgrade path for existing Arista 7500 switches and fulfills Arista's promise of investment protection, expanding the performance, without expensive forklift upgrades.

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. 09/15