

SwitchApp

Low latency Layer 2 switching in as low as 92 ns



92 ns
Latency



Deterministic



Layer 2
Switching

Arista SwitchApp is a full-featured 1/10/40G Layer 2 switching app implemented in FPGA, optimised for low latency and integrated with EOS.

SwitchApp runs on Arista's 7130LB devices, implementing a full Layer 2 cross connect for up to 48 switch ports. Rather than requiring a one-size-fits-all switch architecture, as would be necessary in fixed silicon, Arista's logic engineers have optimised several pipelines, each tuned for the intended use case. The flexibility and programmability of FPGA fabric allows customers to obtain the best possible performance.

For example, a simple L2 switch may achieve a lower latency than one that has more features. A switch with a low bi-section bandwidth achieves a lower latency than one with non-blocking bandwidth.

SwitchApp leverages the world's best network operating system – EOS – to provide access to a rich, and growing, set of features. As new features are added, the hardware logic can be updated. Whole new switch architectures are deployed via software update.

FEATURES*	BENEFITS
Low latency	Forward packets at an ultra low latency of 92-130 ns* from any port to any other port. SwitchApp has been architected with latency as a first class design goal.
Bandwidth	Use the non-blocking bandwidth profiles to provide up to 480 Gbps between ports.
Full L2 Switch Pipeline	Implement a full-featured L2 switch pipeline, allowing for VLAN, Multicast, ACLs, port aggregation (and MLAG). In use cases that don't require advanced L2 features, take advantage of a "Cut down L2" mode to achieve the absolute lowest latency. Full latency measurements are available for each build and each feature profile.
EOS Integration	Take advantage of the familiar EOS configuration experience and feature set. Protocols such as RSTP, LLDP and IGMP work as per other traditional L2 switches. CLI, telemetry, counters, configuration automation and more are equally available.
Deterministic	Know and rely upon your system's latency for fairness, or implement an ideal execution environment for your orders. Full cut-through architecture in all profiles results in no latency variation with different length packets.
Scale	Use 10,000 Unicast MAC addresses or Multicast groups and leverage up to 20,000 ACLs.
FPGA-based	Leverage the re-programmability of FPGAs to optimise switching architectures to support different use cases. New feature requests, normally reserved for hardware updates, can be implemented via software update. Future improvements which will be delivered via software update include Layer 3 routing capabilities (including BGP, PIM, OSPF, VRRP, and more).
L1-Integration	Keep the benefits of Layer 1 switching, alongside SwitchApp. 7130 devices which support both SwitchApp and L1 switching, can do both simultaneously. Use L1 switching to implement tapping, low latency distribution, and media conversion, all converged in the same device as the L2 switch.
Agile Ports	Adapt from 1G, to 10G and 40G without costly upgrades.

Optimized for

- Arista 7130LB devices with embedded Xilinx Ultrascale+ FPGA(s).

Use cases

- Multi-layer MLAG-based leaf-spine fabric, including redundant connections.
- Exchange-facing switch with multiple exchange connections, L1 tapping, and redundancy.
- L2 Multicast pub/sub.

Arista EOS

SwitchApp running on an Arista 7130 LB device uses the same EOS software as all Arista products, simplifying network administration. Arista EOS is a modular switch operating system with a unique state sharing architecture that cleanly separates switch state from protocol processing and application logic. Built on top of a standard Linux kernel, all EOS processes run in their own protected memory space and exchange state through an in-memory database. This multi-process state sharing architecture provides the foundation for in-service-software updates and self-healing resiliency. With Arista EOS, advanced monitoring and automation capabilities such as Zero Touch Provisioning, VMTracer and Linux based tools can be run natively on the switch with the powerful multi-core x86 CPU subsystem.

Latency Profiles

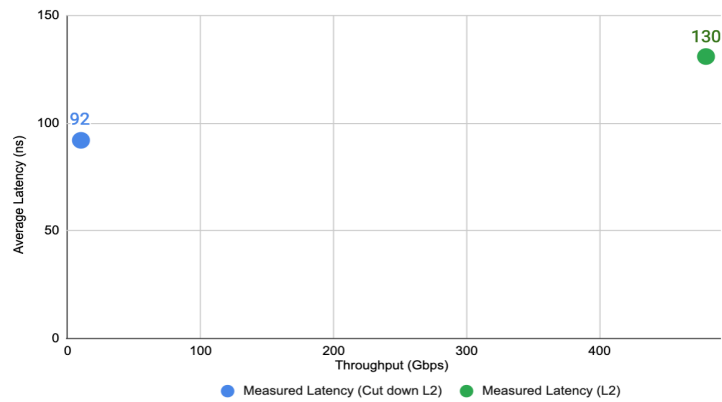


Table 1.1 SwitchApp measured latency for 10G and 480G profiles.

* Average latency quoted

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 © 2020 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. 09/20