Arista 7130

The Latest Innovations in Low-latency Networks

Paul Davy, Q2 2022
pauld@arista.com

ARISTA

SwitchApp

- Specialty low-latency is great for some…
  - But sometimes it's too weird.
  - There is a need for a familiar, low-latency, L2/3 switch

- So Arista built a switching pipeline in FPGA...
  - FPGAs are merchant silicon.
  - EOS uses the FPGA as a just-another-switch-chip
  - Arista forwarding logic, not Arista silicon.
  - Different configuration profiles give new logic

SwitchApp - Now the **fastest* Layer 3** switch too

- **Full-featured** 1/10G Layer 2/3 switch -- *implemented in FPGA*
  - Ultra-low latency packet forwarding averaging 89 ns (L2) or 130 ns (L3)*
  - Full cut-through architecture
  - 48x 1/10G ports
  - 10k Uni- or multicast MAC addresses
  - 30k unicast, 22k multicast routes
- Fully **integrated** with EOS, running on 7130LB devices
  - Standard EOS CLI and protocols: STP, LLDP, IGMP, LAG
  - Standard EOS L3 stack: BGP, PIM, OSPF
  - Standard Management, telemetry, protocols, forwarding plane, CloudVision
- Layer 3 features are now generally available from EOS 4.28.0F
  - Download at: https://mako.arista.com/dyn/softwareportal/releases/#switchapp

SwitchApp looks, feels and tastes like any EOS switch

*NOT STAC BENCHMARKS

SwitchApp - Now the **fastest* Layer 3** switch too

133 ns*

Average Latency

SwitchApp looks, feels and tastes like any EOS switch

*NOT STAC BENCHMARKS

Full cut-through

- RFC2544 tests show full cut-through for all tested packet lengths

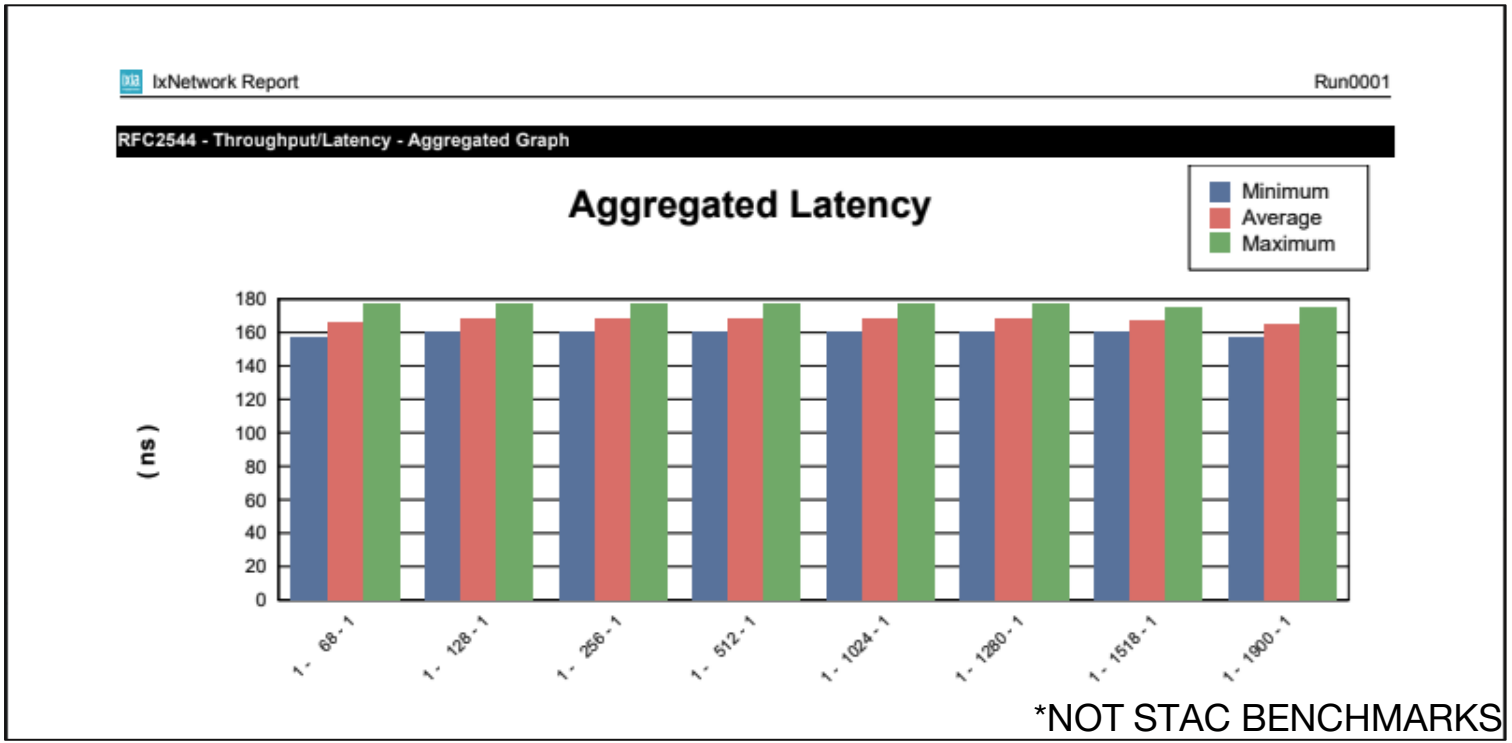

*NOT STAC BENCHMARKS

Key use case: Exchange Facing Switch

A set of clients need to talk to a set of services

- Exchange facing switches have:
  - modest east-west traffic
    - Trading engines don’t communicate with each other (much).
  - asymmetry
    - Order traffic is low bandwidth
    - Market data comes from services
  - a well defined feature set
    - Features required to peer with exchanges: BGP, PIM
    - orders are unicast, market data is multicast.
    - Converged PTP for sync.
    - Specific ACL and NAT features
  - low latency focus
    - Nanoseconds matter

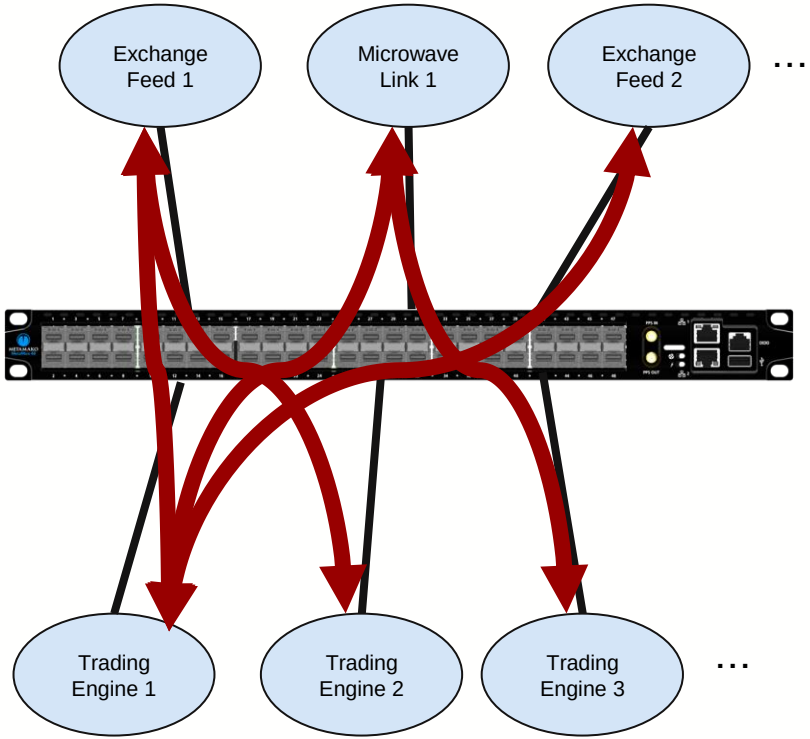

Arista SwitchApp Profiles

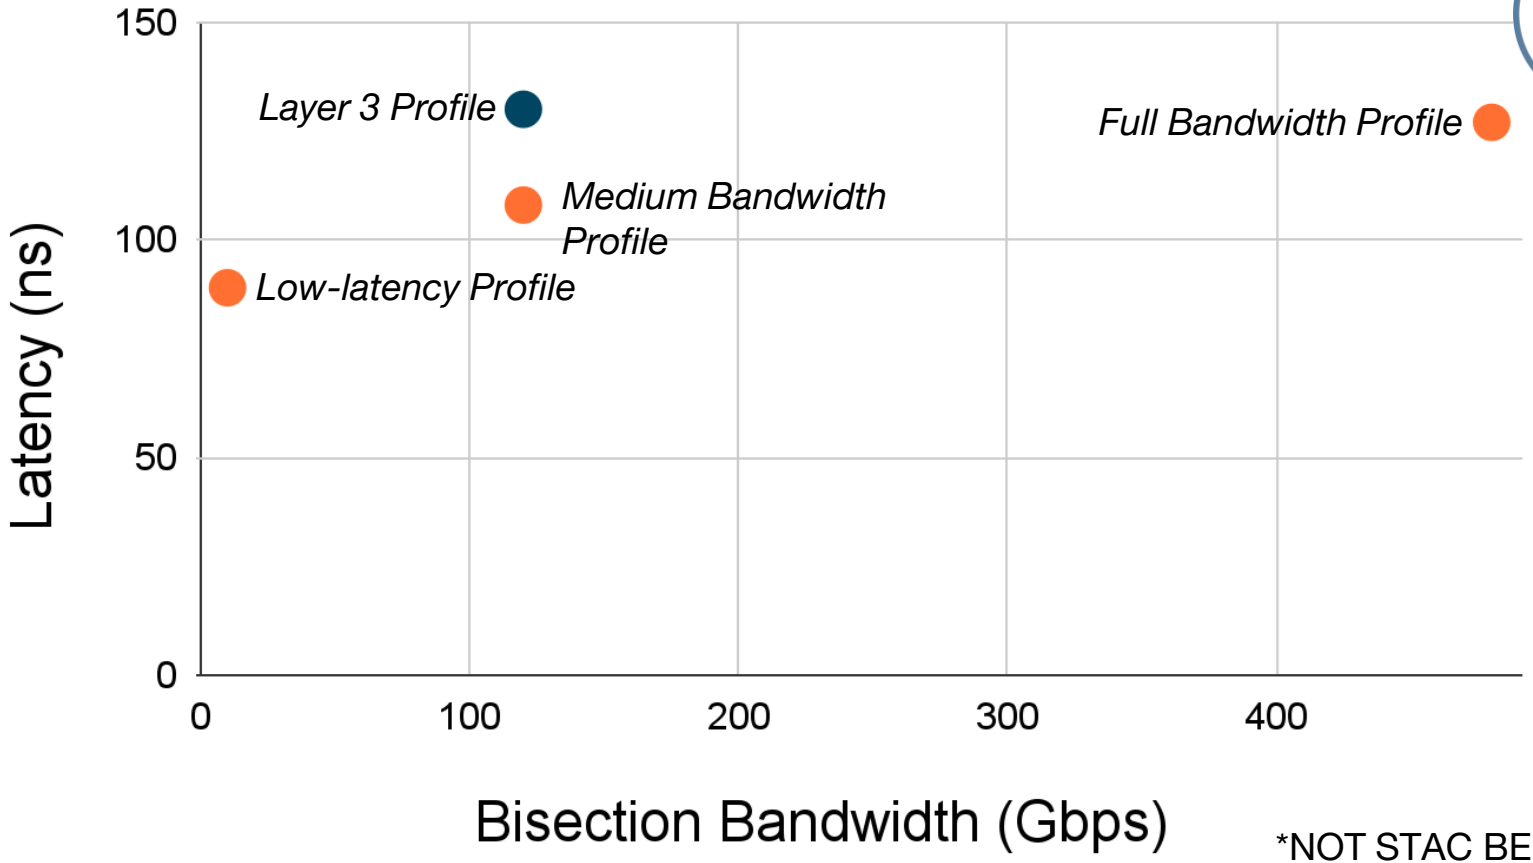

*NOT STAC BENCHMARKS

ARISTA

What is bisection bandwidth?

- The traffic that can pass through the pipeline
- E.g. for unicast traffic, 4x 10G endpoints = 40G

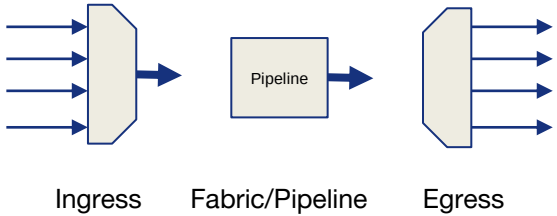

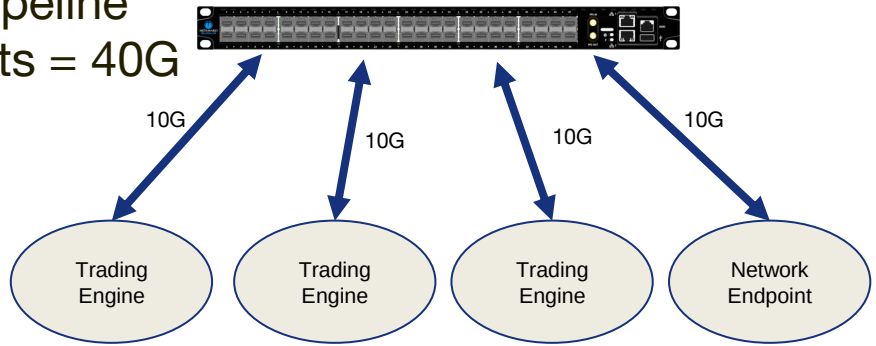

4x 10G endpoints = 40G
possible bisection bandwidth

So I can only have 12 nodes at L3 on SwitchApp!?!?

- Unicast traffic is usually limited bandwidth in finance
- Multicast traffic is replicated at egress

Important Note: **Multicast replication on egress**

- Multicast replication is on egress
  - Replicated multicast does not consume extra bisection bandwidth

- e.g. multicast market data from an exchange feed is max 10G of our bisection bandwidth.

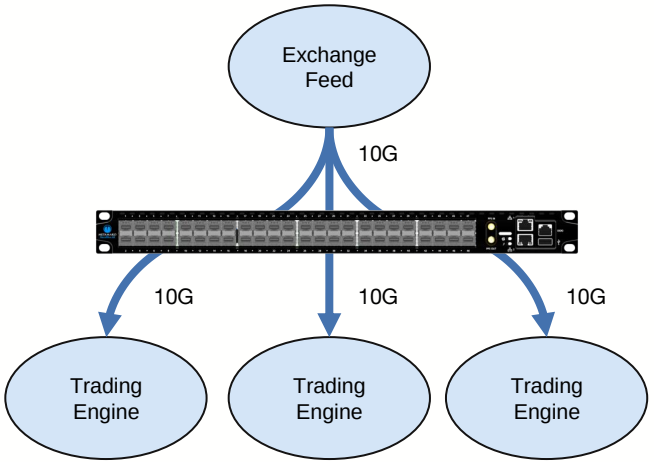

10G Bisection Bandwidth

SwitchApp **GA Features**

Features released in EOS 4.28.0

- Layer 2
  - RSTP
  - PVST
  - SVIs
  - Lane bonding
  - 1G
- LLDP
- Port counters
- Configurable profiles
  - L2: Low latency, mid-bandwidth, full-bandwidth
  - L3: mid-bandwidth
- Layer 3
  - BGP
  - PIM
  - Static and dynamic multicast routing
  - OSPF

Tested using Arista's standard, high quality, infrastructure.

SwitchApp **Roadmap Features**

Coming soon:

- PTP Boundary Clock
- LANZ
- MLAG
- NAT
- ACLs
  - RACL
  - ACL resources shared with L3 Unicast resources
- VRF Support
- Multicast boundary join control, ACLs.
- BFD

Released as software updates

It’s time to try **SwitchApp**

- Compatible with any DCS-7130-48LB device

Come and talk to us!