

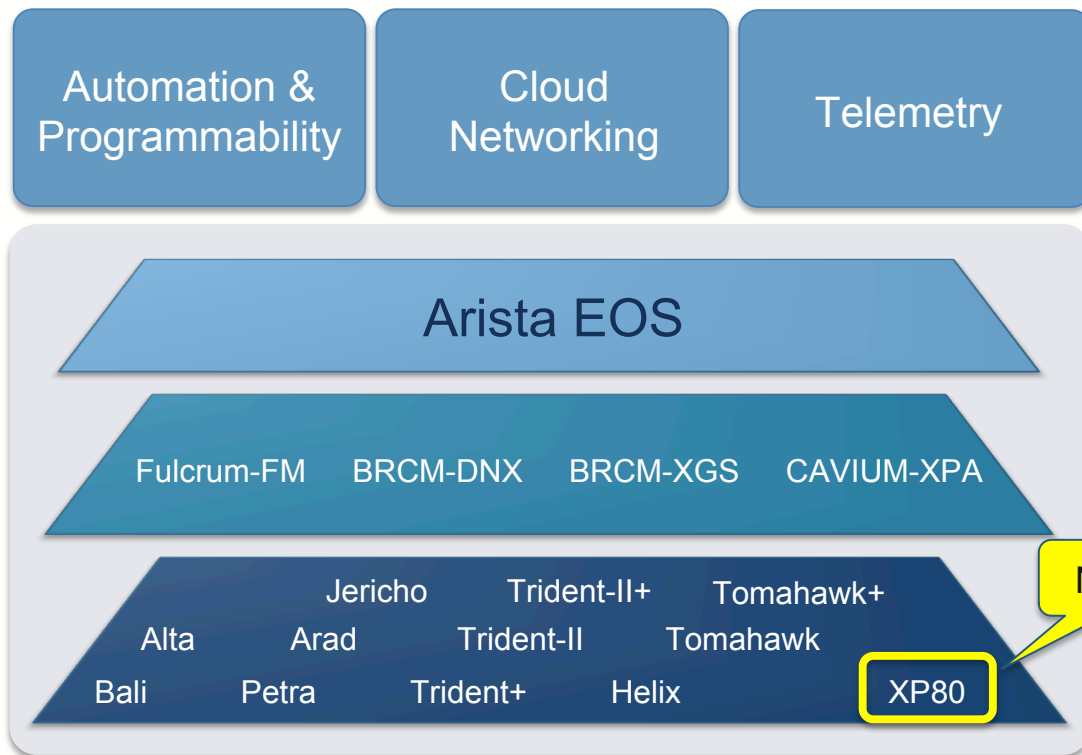
Arista 7160 Series and AlgoMatch™

Under Embargo Until 6am EDT December 5th 2016

Confidential. Copyright © Arista 2016. All rights reserved.

ARISTA

Proven Arista EOS Architecture Flexibility



Consistent and Open

One Single EOS

4 Architectures

12 Silicon Families

Introducing Arista 7160 Series



7160-32CQ – 32 x QSFP100

Q1'17



7160-48YC6 – 48 x 25G & 6 x 100G

Q1'17



7160-48TC6 – 48 x 10G-T & 6 x 100G

Q1'17

- AlgoMatch – 2-6X scale and efficiency
- Future Proof – Programmable packet Pipeline
- Dynamic Buffering – 24MB Fully shared
- Open Standards – IEEE 25GbE & Consortium
- Flexible Profiles – Common L2 & L3 resource

- Wire speed L2 & L3 with VXLAN
- 128-way ECMP, IPv4 & IPv6 Routing
- Wire speed 6.4Tbps, 1.2Bpps
- 48K ACLs – flexible and efficient
- VXLAN Bridging and Routing
- Latency 2-3usec for demanding applications
- Single Image Common Arista EOS
- Comprehensive Provisioning and Monitoring

Arista 7160 Series Innovations

Dynamic Buffers

- High performance fully shared 24MB of packet buffer
- Extensive counters for latency and congestion analysis
- Intelligent Active Queue Management

Programmable Pipeline

- Add new protocols and encapsulations
- Quicker deployment of new solutions
- No forklift required – simple software updates

Flexible Profiles

- Adaptable forwarding tables with custom profiles
- Improved resource flexibility with smaller segments
- Maximum scale with no wasted capacity

IEEE 25GbE 802.3by

- Available with 25G SFP and 100G QSFP ports
- Full IEEE Standard and 25G consortium support
- Investment protection for 10G to 25G migration

AlgoMatch

- Flow matching for access control, policy and visibility
- 2-6X the scale of traditional methods with lower power
- Richer policy support

Modern Cloud Network Policy Requirements



Traditional / Legacy Approach

Apply Policy at Core and Aggregation

Provision Static Traffic Matching

Policy is IPv4 Optimized

Follows Rigid Architecture

Opaque resources, prone to overloads

Enforces silo application policies



Cloud Requirements

Flexible Spine and Leaf

Highly Dynamic Policies

Scale for both IPv4 & IPv6 equally

Reusable and Programmable Engines

Highly Deterministic Resources

Built to scale with applications

The Cloud has driven new approaches....

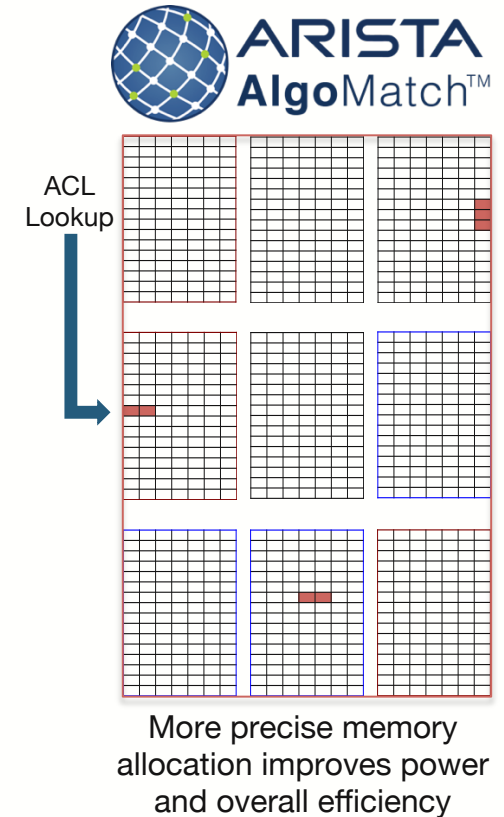

Arista AlgoMatch

Both hardware AND software innovation:

- EOS innovation for in-memory lookup using smarter matching algorithm at wire speed
- Replaces 20 years of traditional TCAM technology with more flexible, programmable engine
- A better way of programming filters that expands the scale, performance and efficiency

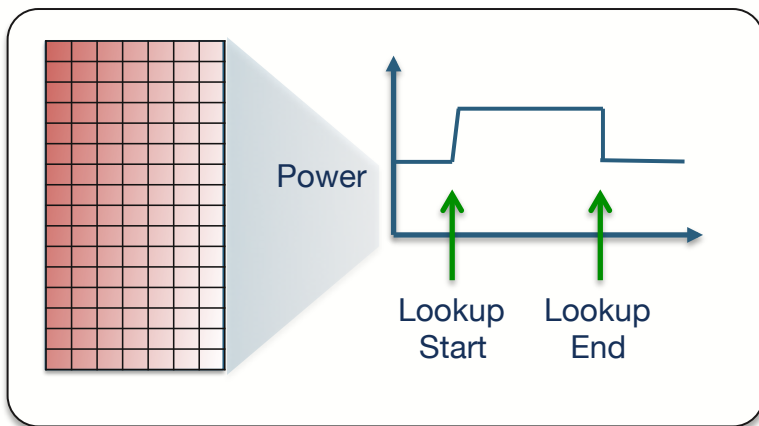
Benefits:

- Improved power efficiency: 50% lower
- Increased Scale: 2x IPv4, 4x IPv6
- Efficient L4 rules: allow enhanced security policy
- Counters: provide better visibility and flow filtering
- Streaming state to CloudVision for granular monitoring and analytics

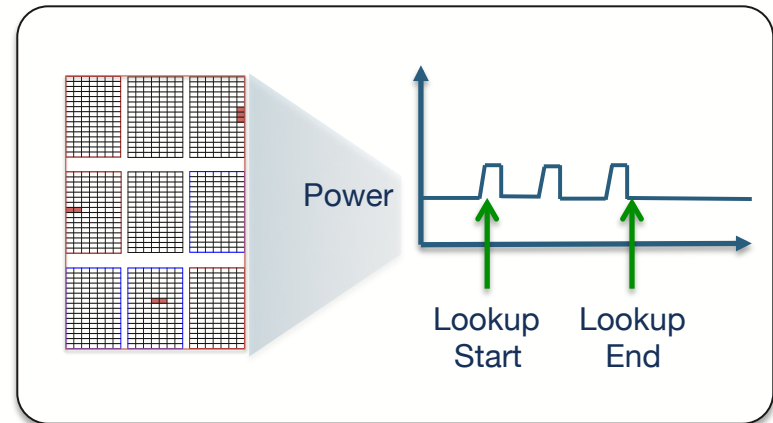


Improved efficiency and scale → Improved security and visibility

AlgoMatch - Power Efficient Searches



- Lookups checks all TCAM locations
- All lines in the TCAM are charged at the start for all lookups
- TCAM match and “don’t care” are high
- TCAM miss discharges
- TCAM power increases inline with capacity




- AlgoMatch power efficient searching
- Only check locations needed
- Only charge the banks being actively searched
- AlgoMatch power scales efficiently
- Lowers power by as much as 50%

7160 Series Programmable Pipeline

- Add new protocols and encapsulations through software
 - Avoid need for a chip replacement
 - New logic to change lookup, traffic management, packet modification, visibility
 - Recent encapsulations: VXLAN, NVGRE
- Benefits
 - Test and deploy without waiting for new silicon cycles
 - Extend the lifecycle of leaf and spine
 - No forklift required – simple software updates extend the lifecycle
 - Consistent architecture



Typical 24 month gap




Faster release via S/W Update

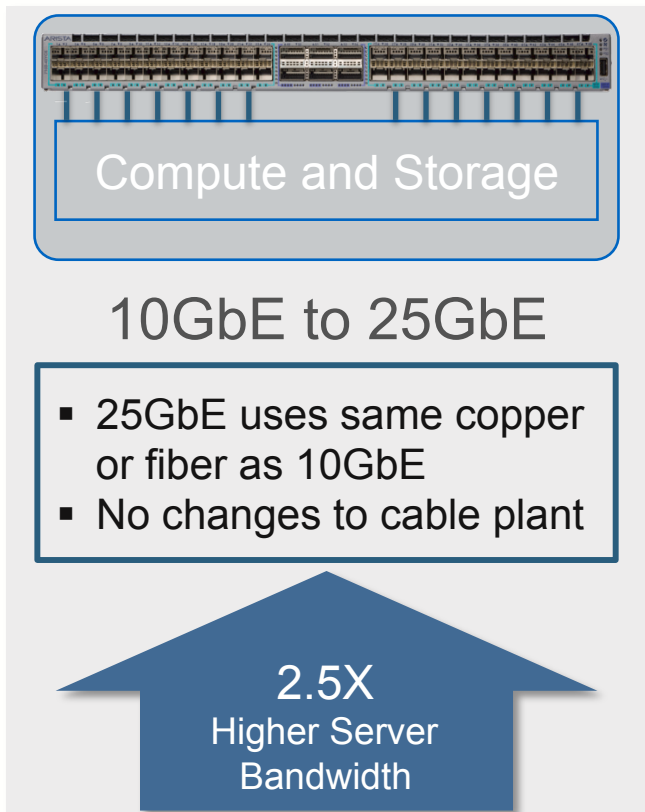
7160 Series Flexible Forwarding Profiles

- Forwarding table profiles
 - Allows efficient use of common resources
 - Higher maximum scale
 - Small segment size for more combinations
 - Up to 448K logical resources available
- Benefits
 - Same system used in multiple scenarios
 - Achieve higher scale with no wasted capacity
 - Avoid scale limits associated with fixed systems
 - Efficient resource usage lowers power

- Example Profiles include
 - Balanced - general Leaf and Spine
 - Large Hosts - L2 Overlay VXLAN
 - Large Routing - DC Leaf-Spine
 - Large ARP / Host scale - Education & Hospitality



The Need for 25GbE




- Unrelenting traffic growth and capacity
- 2.5X performance with the same infrastructure
- Increase server throughput
- Storage devices increased performance
- Consistent design and cabling as 10G
- Long term investment and migration using 10/25G



- 7160 Series 25GbE with SFP25 & QSFP100
- Full IEEE Standard and 25G consortium support
- Investment protection for 10G to 25G migrations

Design for Capacity Growth and Server Transitions

- Increasing East-West traffic is driving a need for additional uplink capacity
- Containers are creating new traffic profiles that requires more bandwidth
- Server workloads are more intensive 10G -> 25G and 40G -> 100G
- New undersubscribed design options allow for efficient flow balancing



Expanded Options for Leaf Spine Solutions


- Fixed Systems with 100G Uplinks

48 x 25G SFP and 10G-T with 6 40/100GbE Uplink

- Leaf switch for flexible ECMP designs
- Flexible Uplinks: Choice of 6 x 40GbE or 6 x 100GbE
- Each QSFP100 port: 1x40G or 100G, 4x10G or 25G
- Flexible profiles, programmable pipeline
- Broad features with scalable resources
- 24MB Buffer for speed changes and microburst



48x 10G-T and 6x 100GbE / 40GbE




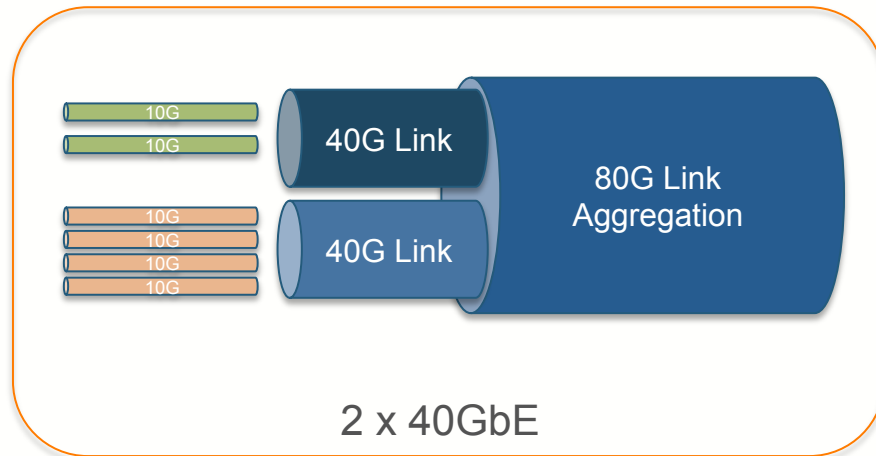
48x SFP25 and 6 x 100GbE / 40GbE

100G Uplink Efficiency

Multiple ports in LAG is good but not a replacement for higher speed links

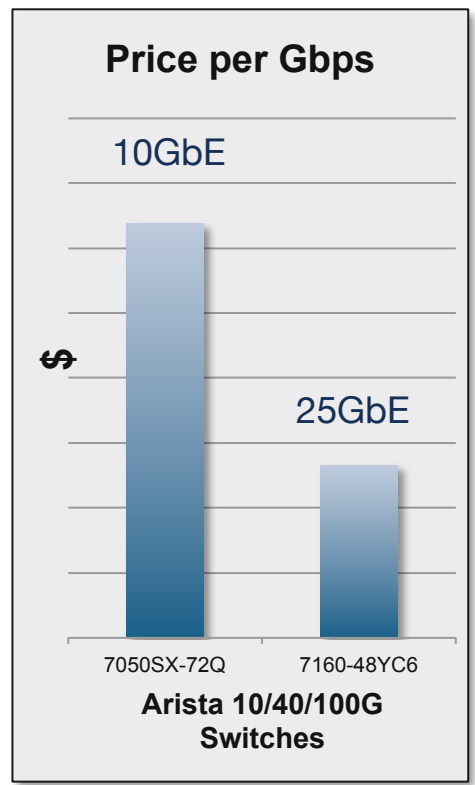
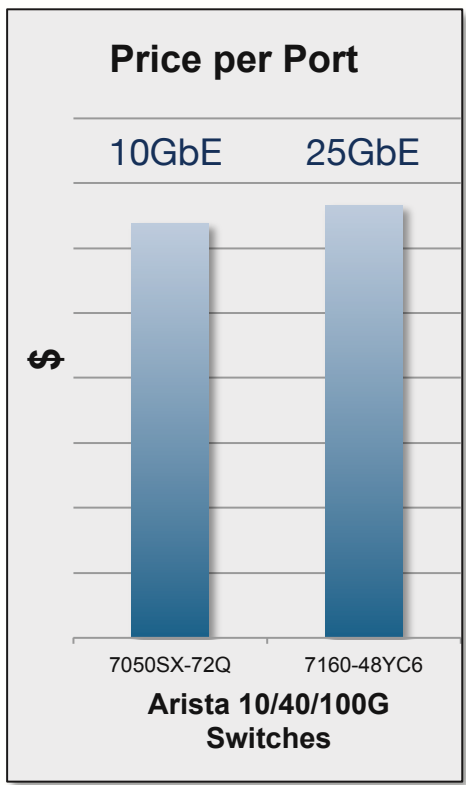
Challenges with Link Aggregation:

- Uneven flow distribution can lead to unbalanced traffic patterns
- Increased probability of hash collisions with servers sustaining 10GbE
- “2x40GbE” total bandwidth is 20% lower than 100GbE
- Effective bandwidth of 80G significantly lower after hashing efficiency



25GbE - Higher Bandwidth with Lower Cost Per Port

- Hyper-scale server bandwidth translating to switches
- 25GbE offers investment protection
- Expect 10GbE to migrate to 25GbE quickly due to both pricing and compatibility



Expanded 7060X and 7280R Fixed Systems



Q1'17

7060CX2-32S
32 x 40/100GbE / 128 x 10/25GbE
Expanded 22MB Buffer
6.4Tbps / 3.3Bpps
450ns latency

- 40% Larger Packet Buffer
- IEEE 25GbE Support
- 10G to 100G Flexible Configurations
- Wire speed Performance
- Low Latency, Cloud Features & Scale
- Leaf and Spine Deployment Options



Q4'16

7280QR-C72
56 QSFP+ / 16 QSFP100
6.4Tbps / 2.88Bpps
10/25/40/50/100GbE
Ultra Deep 16GB Buffer

- Consistent 7280R Series Features
- Flexible 10/25/40/50/100G Options
- Future Proof 25GbE and 50GbE
- Enabler for Next Gen Leaf and Spine
- Wire speed L2 & L3 with VXLAN
- Broad deployment options


Pricing and Availability

AlgoMatch will ship for new Fixed systems in Q1'17

Fully supported as part of EOS, with no separate licensing

7160 Series Systems are shipping in Q1'17

Pricing is comparable to 7060X Series at less than \$1000 per 100G port



7060CX2-32S is shipping in January 2017

7280QR-C72 is orderable and generally available now



* All specifications subject to change



Thank You

www.arista.com