

# Что нового в EOS

Александр Беспалов, системный инженер ([ab@arista.com](mailto:ab@arista.com))

Российская команда инженеров Arista ([ru-se@arista.com](mailto:ru-se@arista.com))

# Where to get information?

- Documentation

- Arista.com => Support => Product Documentation

- <https://www.arista.com/en/support/product-documentation>

- Release Notes

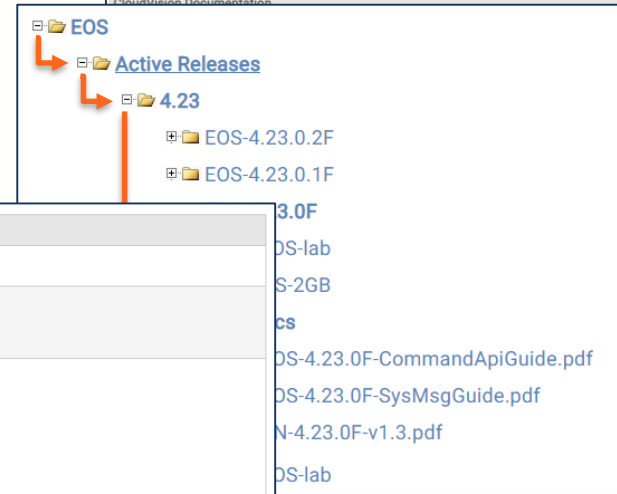
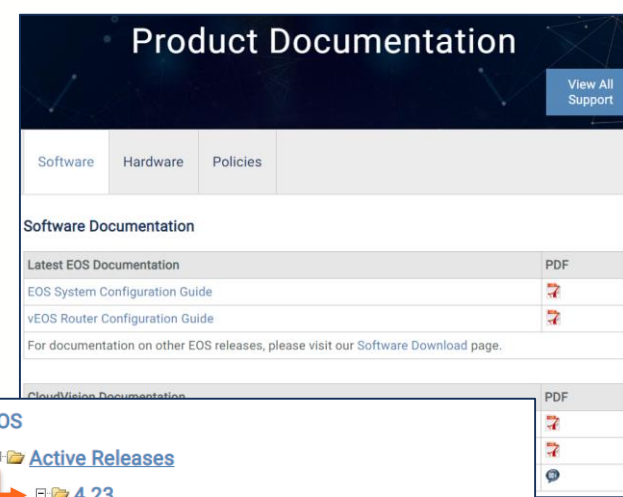
- Arista.com => Support => Software Download

- <https://www.arista.com/en/support/software-download>

- Transfer of Information (TOI)

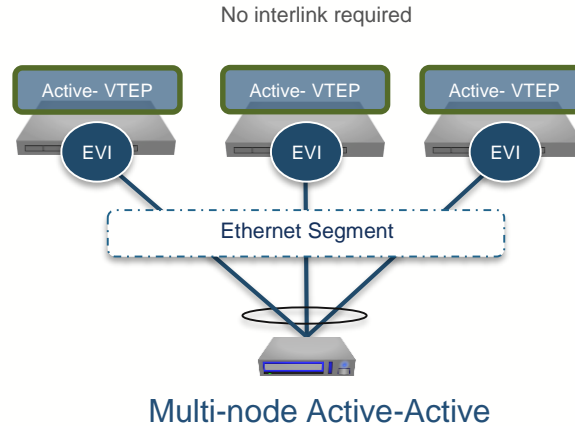
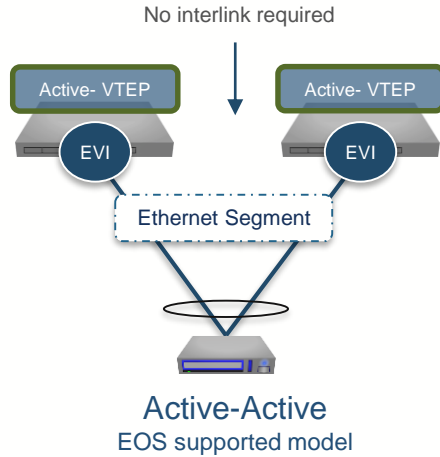
- Arista.com => EOS Central => TOI

- <https://eos.arista.com/toi/>



EOS Release	TOIs
4.23	<ul style="list-style-type: none"><li>EOS-4.23.0F</li></ul>
4.22	<ul style="list-style-type: none"><li>EOS-4.22.1F</li><li>EOS-4.22.0F</li></ul>
4.21	<ul style="list-style-type: none"><li>EOS-4.21.6F</li><li>EOS-4.21.5F</li><li>EOS-4.21.3F</li><li>EOS-4.21.1F</li><li>EOS-4.21.0F</li></ul>
4.20	<ul style="list-style-type: none"><li>EOS-4.20.6F</li><li>EOS-4.20.5F</li><li>EOS-4.20.1F</li></ul>

# EVPN-VXLAN All-Active Support



- **EOS 4.22.0** (7500R/R2, 7280R/R2), **EOS 4.22.1** (7300X3, 7050X3)
- Support for multi-homing CE nodes to multiple PE nodes
  - Standard LAG CE-PE; same shared Ethernet segment (ESI)
  - Multi-vendor interop on same ES
  - no peer link, no limit in the number VTEPs in the ES
- Multihoming uses EVPN additional routes:
  - Type-1: mass withdrawal, aliasing
  - Type-4: Designated Forwarder election


# EVPN-VXLAN All-Active Support

## Configuration:

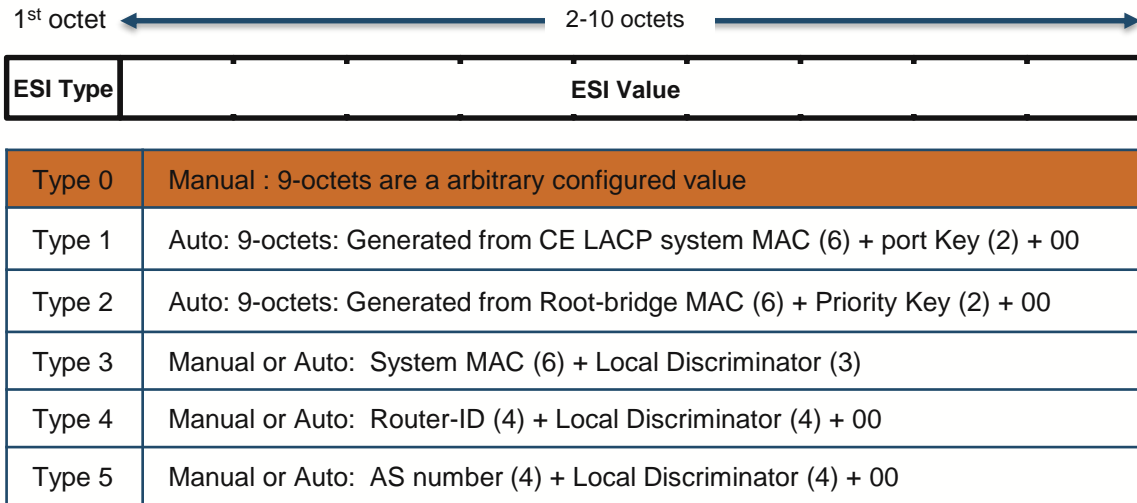
```
!
interface Port-Channel100
  switchport trunk allowed vlan 10-13,20
  switchport mode trunk
!
  evpn ethernet-segment
    identifier 0000:0000:0001:0411:1111
    route-target import 00:00:00:00:01:04
    lacp system-id 0000.0000.0104
!
```

## Route Target Community Options:

- Manually configured (implemented in EOS)
- Auto-derived from ESI:

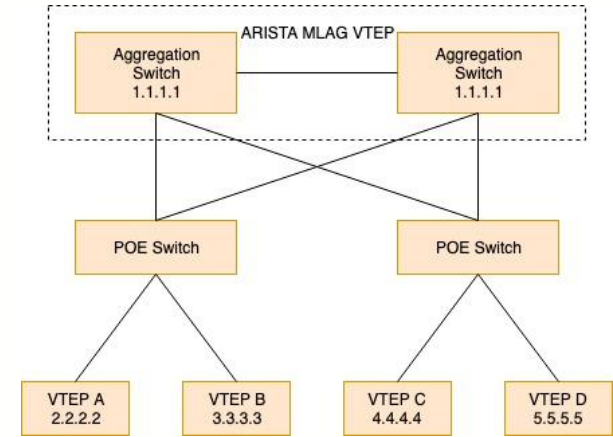


## ESI Structure:



# VXLAN VTEP to VTEP Bridging

- Some “special” topologies cannot support VXLAN
- Ethernet Split-Horizon
  - prevents sending what was received (on the same port)
  - what a VTEP receives from VXLAN cannot be sent back into VXLAN
- Example of suboptimal (but realistic) scenarios:
  - Topologies: Hub-and-Spokes, star, or daisy-chains
  - Use cases: Campus, WAN, Access-Points distribution
- Supported on:
  - WiFi Access Points
  - 720XP
  - 7050X3
  - 7300X3
- **EOS 4.22.1**



```
!  
interface Vxlan1  
    vxlan source-interface Loopback1  
    vxlan udp-port 4789  
    vxlan bridging vtep-to-vtep source-vtep tx disabled  
    vxlan vlan 10 vni 110  
    vxlan flood vtep 21.0.0.6  
!
```

# VXLAN Flood-lists Learning over the Data-plane

- Normally VXLAN VTEPs are learned via BGP-EVPN or static configuration
- 100s of VXLAN endpoints w/o EVPN support (modern WiFi APs)
- Scenario:
  - WiFi Access-Points distribution
- Solution:
  - VXLAN VTEP learned on first packet arrival
  - VTEPs subnet can be restricted to IP ranges
  - VTEP removed on all MACs age out
- Supported on:
  - Platform independent
- **EOS 4.23.0**

```
!  
interface Vxlan1  
    vxlan source-interface Loopback1  
    vxlan udp-port 4789  
    vxlan flood vtep learned data-plane  
    vxlan learn-restrict vtep 1.1.1.1/24  
    vxlan vlan 10 vni 110  
    vxlan flood vtep 21.0.0.6  
!
```

# Support for GRE tunnel interface

- Implementation:

- Up to 256 interfaces
- Hardware based forwarding
- Pv4 over IPv4 GRE tunnel interfaces
- Dynamic routing protocols over GRE:
  - » BGP
  - » ISIS
- Overlay VRF support

- Platforms:

- **EOS 4.22.1** (720XP, 7050SX3)
- **EOS 4.22.2** (7050X, DCS-7050X2, 7050X3, 7060X, 7060X2, 7250X, 7260X, 7260X3 and 7300)
- **EOS 4.23.0** (7010T)
- **EOS 4.23.1** (7368X4)

```
!  
interface Tunnel 10  
    tunnel mode gre  
    ip address 192.168.1.1/24  
    tunnel source 10.1.1.1  
    tunnel destination 10.1.1.2  
    tunnel path-mtu-discovery  
    tunnel tos 10  
    tunnel ttl 10  
    vrf overlayVrf  
!
```

# LANZ Mirroring

- During congestion **some** packets experience higher forwarding latency due to buffering. How do I know which traffic exactly experienced higher latency?
- Solution:
  - Start mirror traffic if configured buffer high-threshold is reached
  - Some platforms can mirror to GRE
- Platforms:
  - **EOS 4.22.1** (7010T, 7150S, 7050X, 7050X2, 7060X, 7300X)
  - **EOS 4.23.0** (720XP, 7050SX3, 7368X4)

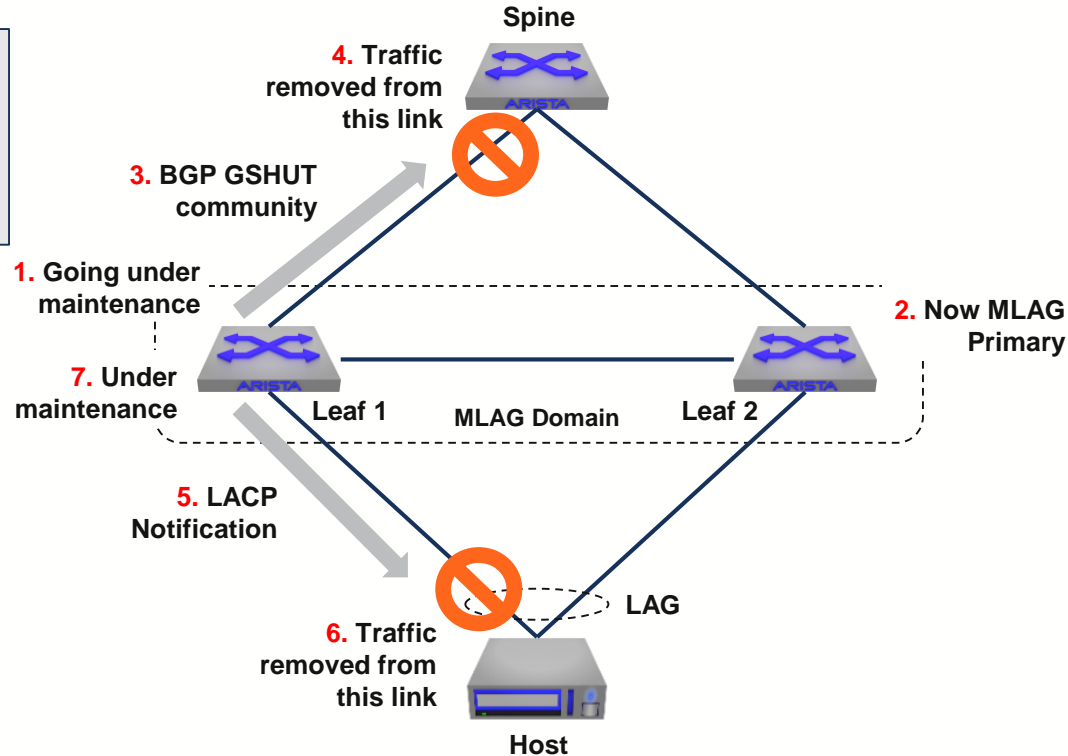
```
!  
queue-monitor length mirror  
queue-monitor length mirror destination Ethernet <port>  
queue-monitor length mirror destination tunnel mode gre source <src_ip> destination <dst_ip>  
!
```

# Maintenance Mode: Leaf (MLAG+BGP)

- No upgrade allowed because of the convergence outage
- How to get new features?
- Solution:
  - Maintenance Mode for MLAG+BGP
  - Race to absolute 0 outage
  - You can now upgrade an L2/L3 leaf with an absolute minimum outage
  - Complement the ASU feature (whose use case is single-homed services)
- **EOS 4.22.0**

# Maintenance Mode: Leaf (MLAG+BGP)

```
!  
maintenance  
  unit System  
    quiesce  
!
```




# Accelerated Software Upgrade

```
switch# reload fast-boot
```


- How to minimize Leaf upgrade downtime with single homed hosts?
- Solution:
  - Accelerated Software Upgrade (ASU)
  - Exists for several years in Arista
  - Headless traffic forwarding on system reboot/upgrade
  - Requires graceful restart
  - Different approaches on boot up (depends on agent and software release)
    - » Approach #1: New image reconstructs state from hardware
    - » Approach #2: New image regenerates state from network peers, and then bulk update hardware state
  - No interface flaps
  - Overall downtime from milliseconds to several seconds
- **EOS 4.22.1**: support for L2 EVPN with VXLAN transport

# Partial BGP Multipath-Relax Based on First $n$ ASNs


Default behaviour  
multipath-relax



EOS 4.15  
no multipath-relax



EOS 4.22.1  
multipath-relax *match 2*



```
[no | default] bgp bestpath as-path multipath relax [match <n>]
```

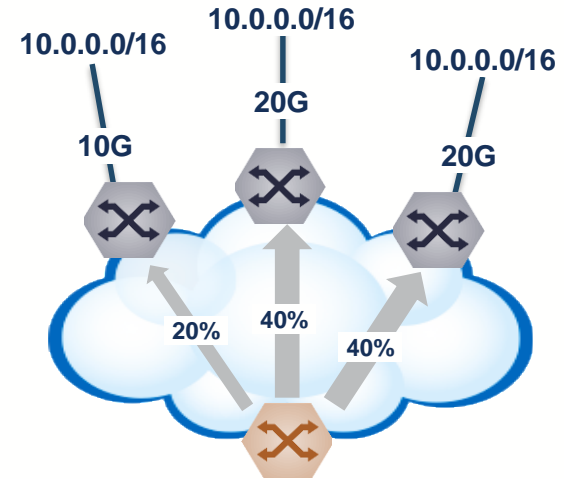
# Selective BGP AS-Path Multipath-Relax

- How to select which routes will be used for multipath?
- Solution:
  - New “set bgp bestpath as-path weight <>” knob under route-map CLI
  - Matched paths also tagged as multipath-relaxed
  - During route selection consider the AS-path weights before comparing the AS-path contents
  - Default weight is 0
  - Path with higher as-path weight is preferred
- **EOS 4.23.0**

```
!  
route-map <r-map-name> <permit/deny> <seq>  
    match <conditions>  
    set bgp bestpath as-path weight <1..65535>  
!
```

# BGP UCMP Support in ArBGP

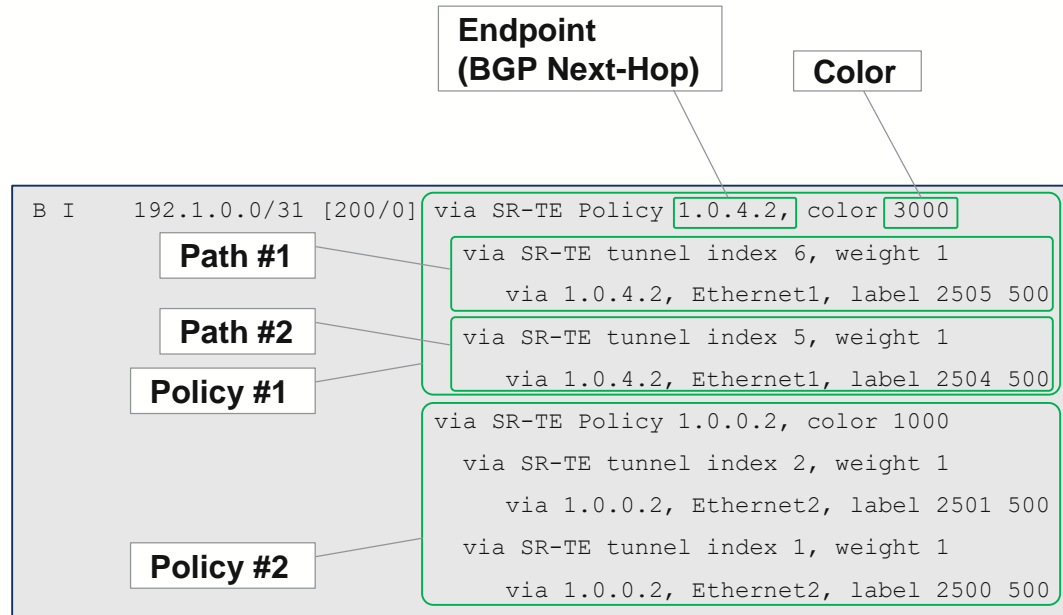
- BGP UCMP is using link-bandwidth communities to calculate path weights
- Supported since EOS 4.15.4 using Gated (ribd)
- BGP UCMP is now supported ArBGP:
  - **EOS 4.22.1:**
    - » Adjust UCMP
    - » Receive link-bandwidth community
    - » Set link-bandwidth community
  - **EOS 4.23.0:**
    - » Adjust Auto feature added. Receiving link bandwidth will be added to link-bandwidth community received from a BGP peer





# Segment Routing Traffic Engineering Policy

- SR-TE Policy consists of:
  - SID lists (paths) with weights for UCMP (up to 16)
  - Preference for path selection
  - Binding SID. Mandatory in EOS
- SR policy identification:
  - Endpoint (IPv4 or IPv6 destination)
  - Color. To select particular paths
- Policy sources:
  - BGP
  - Static configuration
- **EOS 4.23.0**



# Arista 7170 Software Packages

- NAT
  - Static NAT (14k entries)
  - Dynamic NAT (135k entries)
  - Multicast NAT
  - NAT64 translation (MAP-T)
- Stateless loadbalancer
  - Up to 1000 VIPs
  - Up to 64k servers
  - Special kernel module required for servers

**BAREFOOT**  
NETWORKS



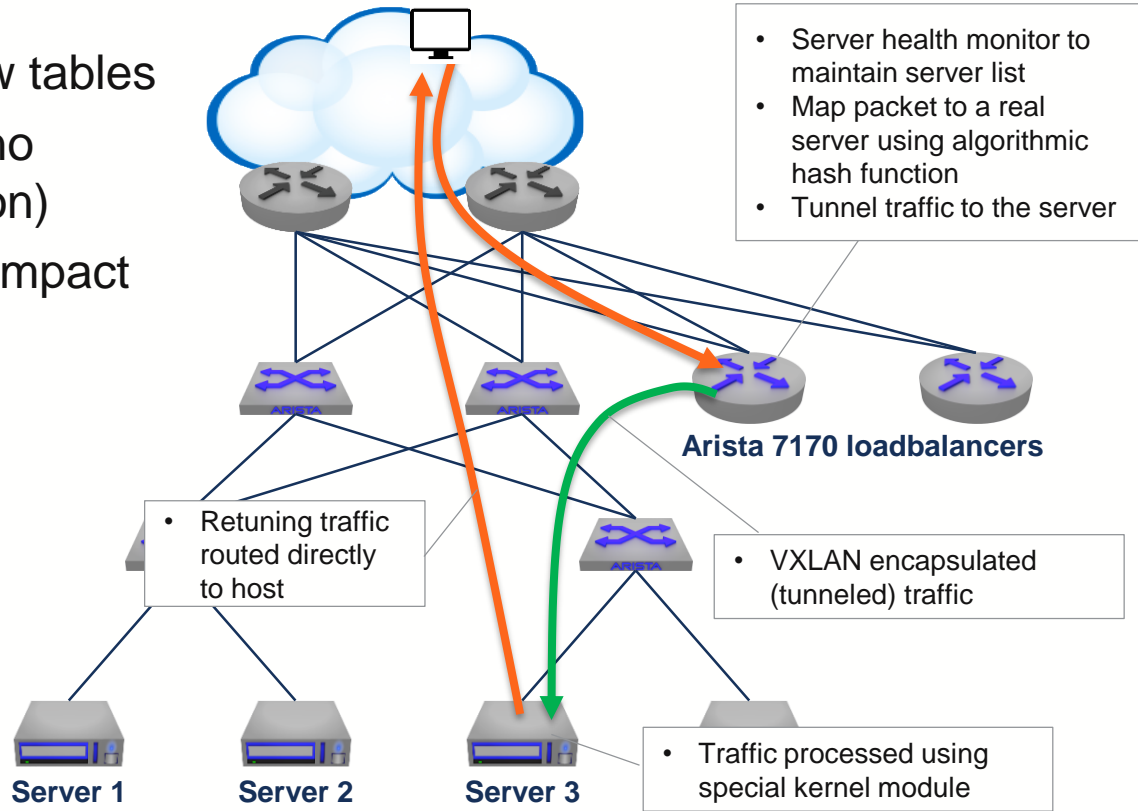
64 Ports 100GbE – 12.8Tbps



32 Ports 100GbE – 6.4Tbps

# Arista 7170 Stateless Loadbalancer

- Not needed to maintain flow tables
- Graceful server withdraw (no impact on any active session)
- Hitless server addition (no impact on any active session)
- Scaling:
  - Up to 1000 VIPs
  - Up to 64k servers



# Linux Upgrade



- Migration from Fedora to CentOS 7.5
- Wider set of tools and 3<sup>rd</sup> party software
- Easier to deploy 3<sup>rd</sup> party extensions
- **EOS 4.23.0**



# Thank You

[www.arista.com](http://www.arista.com)