

IP fabrics - reloaded

Joerg Ammon <jammon@extremenetworks.com>
Senior Principal Systems Engineer
2017-11-09

Extreme Networks Acquisition update

- Oct 30, 2017: Extreme Networks Completes Acquisition of Brocade's Data Center Networking Business
- Extreme Networks expects to generate over \$1 billion in annual revenues
- Extreme acquires customer relationships, personnel and technology assets from Brocade
- continued support for the acquired existing data center networking product lines



Agenda

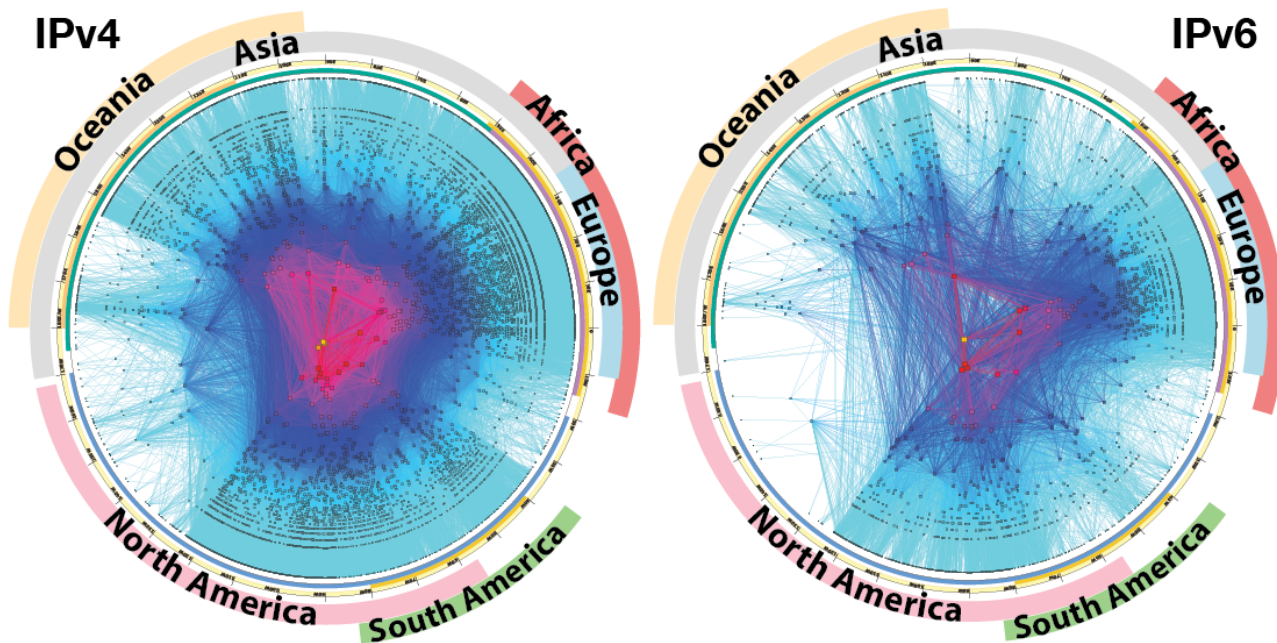
- Why eBGP
- BGP made as simple as possible
- Some EVPN address-family internals
- Data Center Interconnect
- Unified bridging
- VXLAN vs MPLS
- Automation
- Visibility



Why BGP?

CAIDA's IPv4 vs IPv6 AS Core AS-level Internet Graph

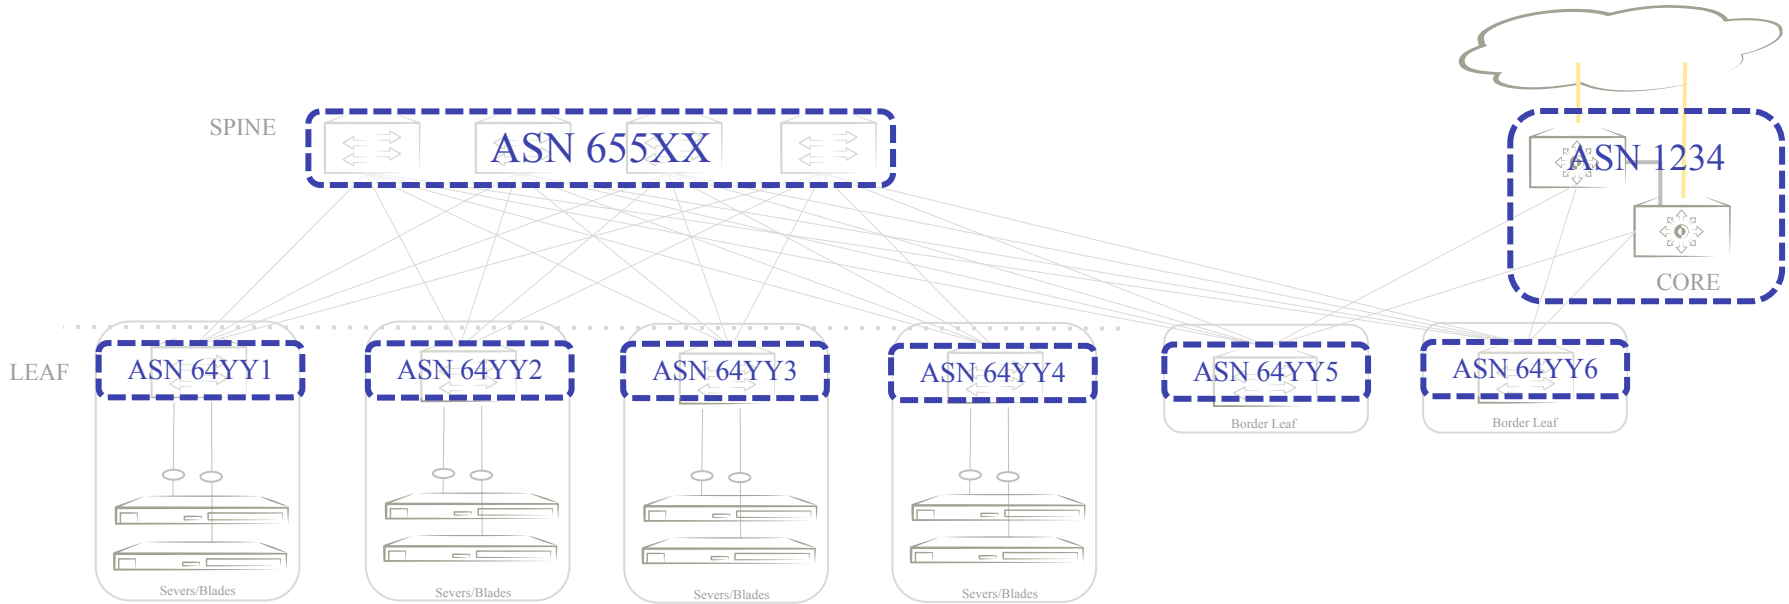
Archipelago July 2015



Copyright © 2015 UC Regents. All rights reserved.

Source: https://www.caida.org/research/topology/as_core_network/pics/2015/ascore-2015-jan-ipv4v6-standalone-1200x710.png

Making your life with an IP fabric easy...



- One Private - ASN per Rack
- One Private - ASN for SPINE
- One Private ASN for Super Spine
- 4 Byte ASN for Scale
- Public ASN at Core/Edge
- Unnumbered IP on physical links
- LLDP to communicate loopback address
- Congruent topology for IP and EVPN

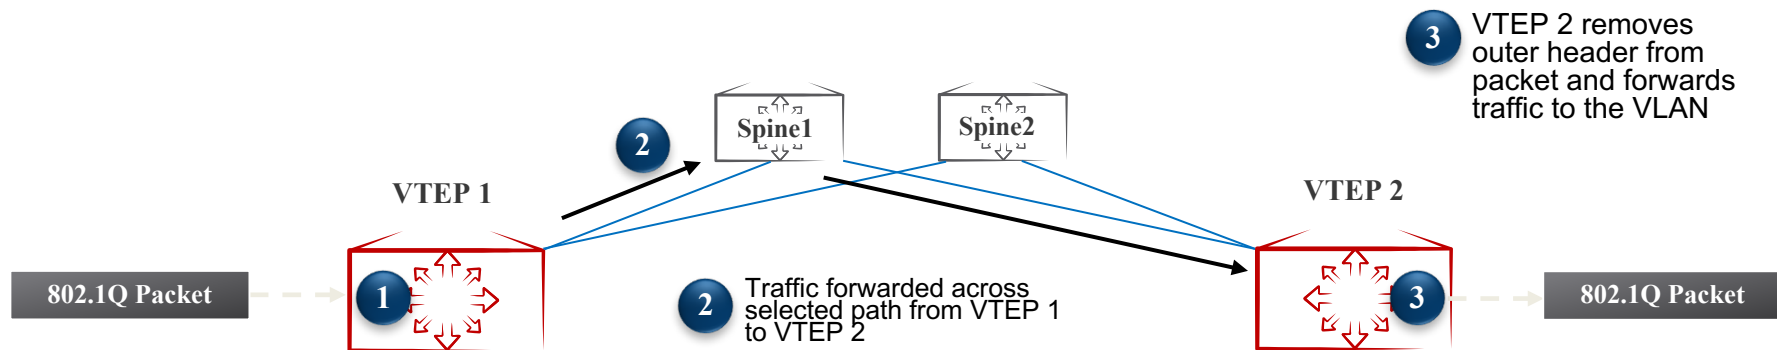
Details of an EVPN BGP update – Inclusive Multicast Route

```
Leaf-1# show bgp evpn route type inclusive-multicast detail
```

```
Route Distinguisher: 300:301
```

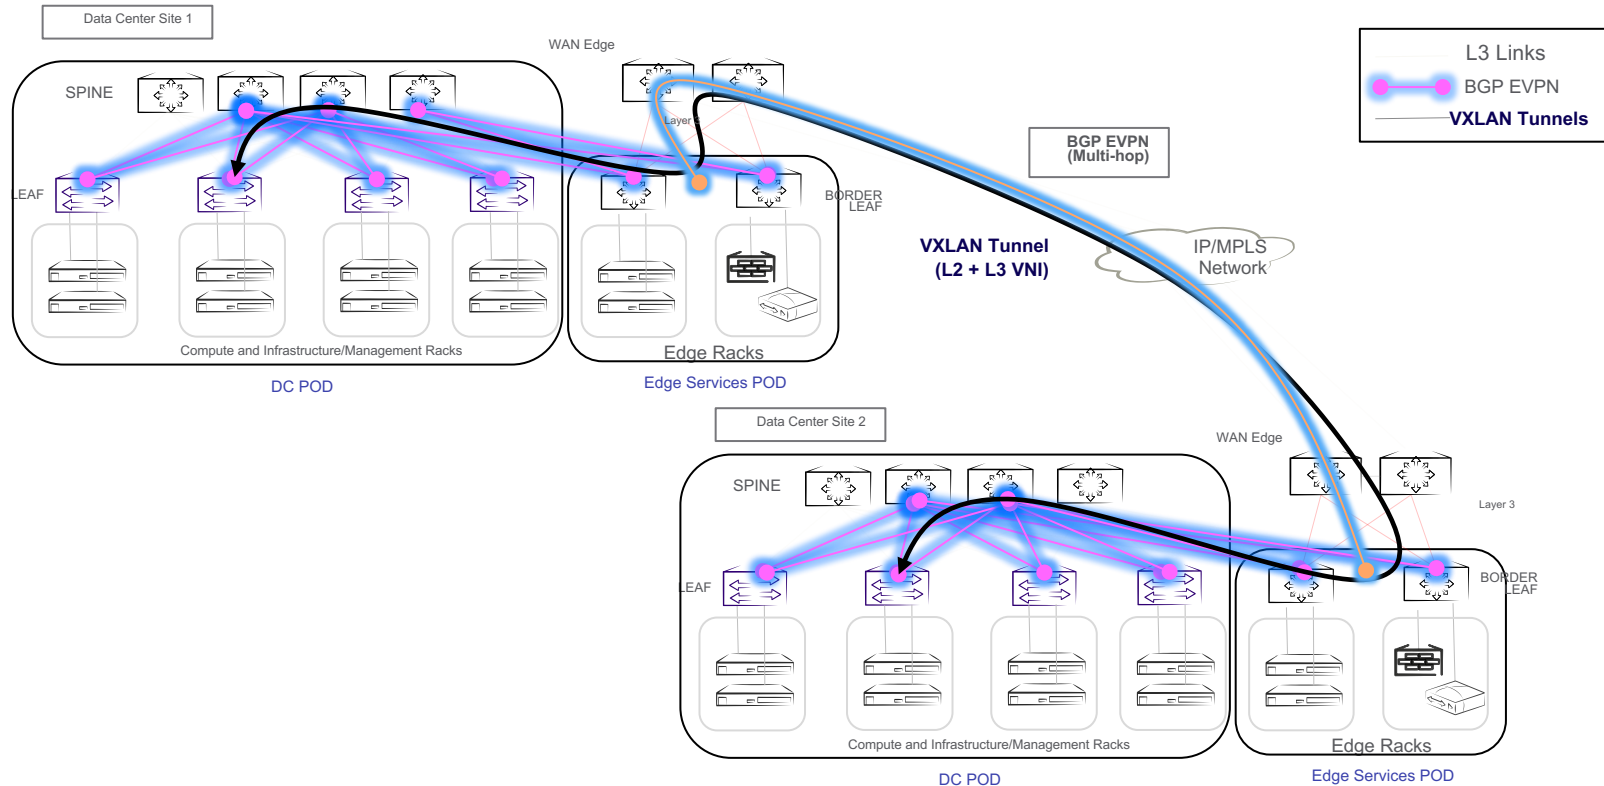
```
2      Prefix: IMR:[0][IPv4:172.32.254.203],   Status: BE,   Age: 3d23h51m16s
      NEXT_HOP: 172.32.254.203, Learned from Peer: 172.32.254.1 (64512)
      LOCAL_PREF: 100, MED: none, ORIGIN: incomplete, Weight: 0
      AS_PATH: 64512 65000
      Extended Community: ExtCom:03:0c:00:00:00:00:00:08 RT 300:300
      PMSI Attribute Flags: 0x00000000 Label-Stack: 0x00001c85
      Tunnel-Type: 0x00000006 Tunnel-IP: 172.32.254.203
      Extended Community: ExtCom: Tunnel Encapsulation (Type Vxlan)
      Adj_RIB_out count: 1, Admin distance 20
      L2_vni: 7301
      RD: 300:301
```

VXLAN Overview and Traffic Distribution

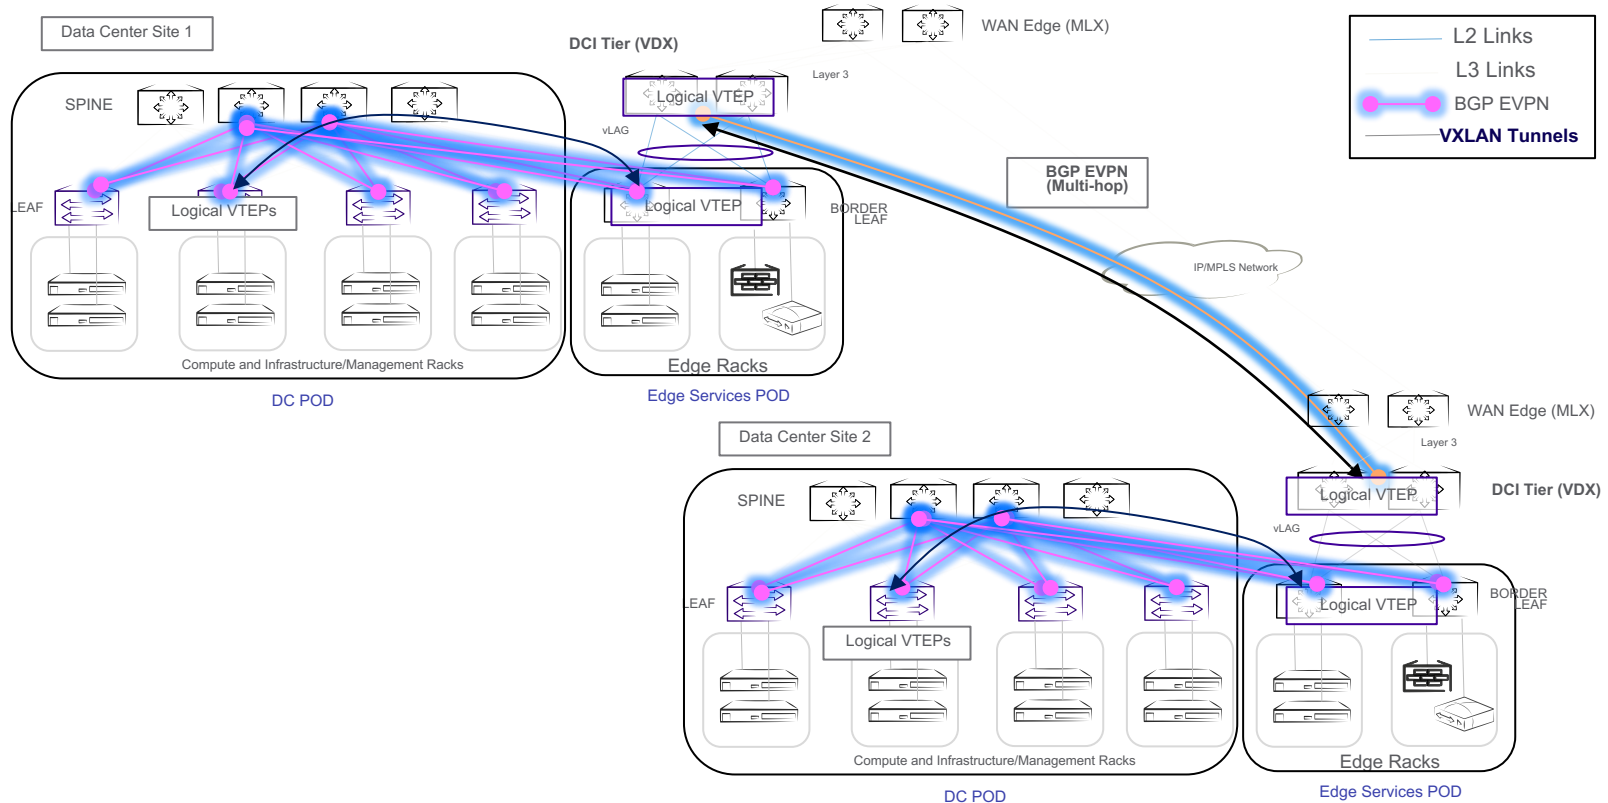


- A single VXLAN tunnel exists between VTEP 1 and VTEP 2
 - Multiple physical paths available for traffic
 - VXLAN traffic is distributed across these paths via ECMP
 - Source UDP port varies to ensure traffic distribution


BGP EVPN based DCI: L3 Handoff




BGP EVPN based DCI: L2 Handoff



Bridge domain and unified bridging

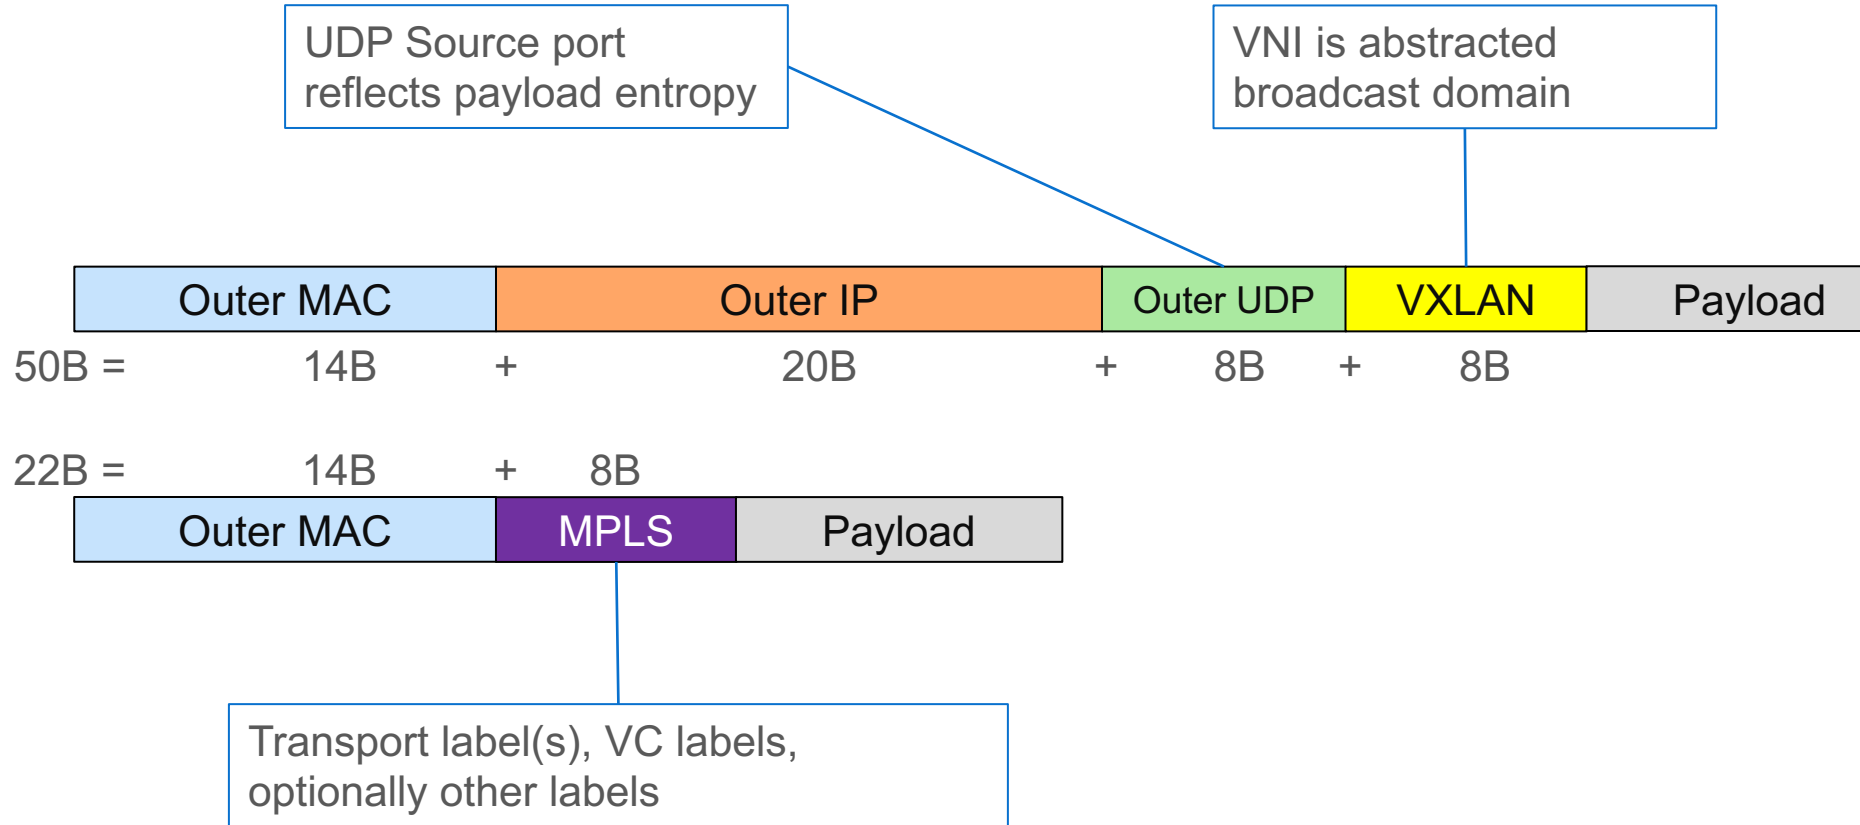


Redundant connections into EVPN - simplified



- Configure / auto derive an ESI
- Autodiscover PEs attached to ES
- Elect designated forwarder (per VLAN)
- Split-horizon (impacts data plane)
 - MPLS: ESI-label
 - VXLAN: Track source IP-addresses and filter
- Aliasing: Associate remote MAC with all PEs announcing the same ES

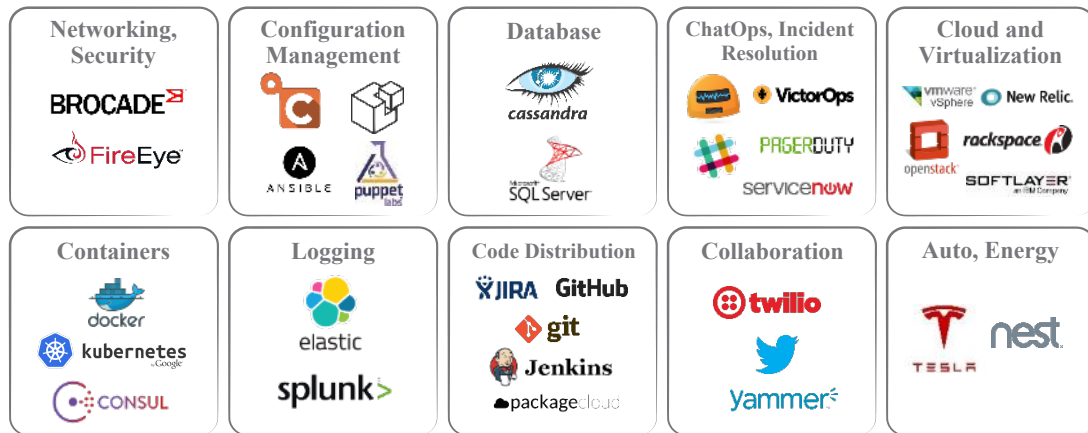

Header comparison: VXLAN vs MPLS



Extreme Workflow Composer – powered by



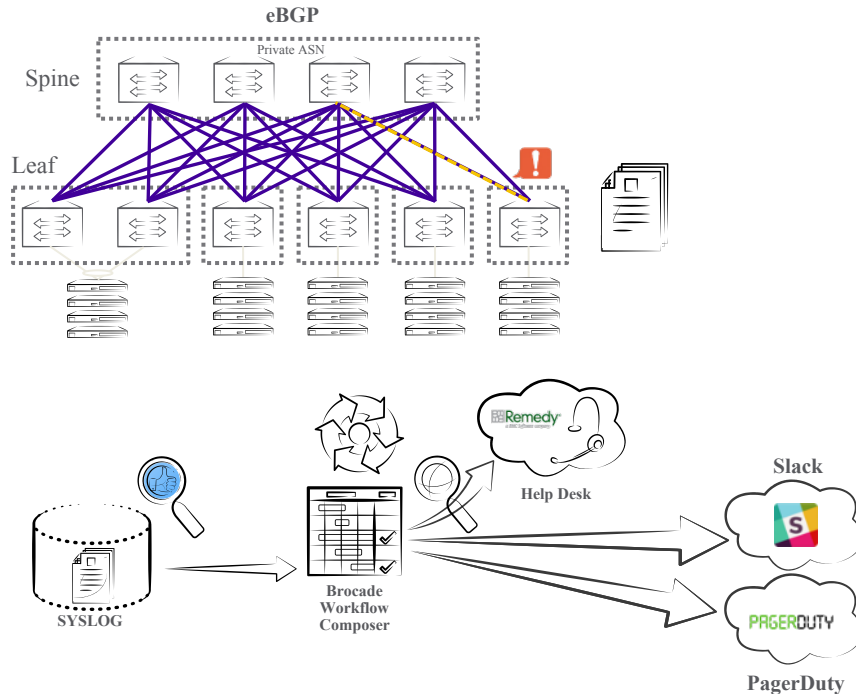
Workflow-centric, cross-domain, DevOps-inspired automation



Troubleshooting and Remediation



Example: Turning 2:00 a.m. calls into 10:00 a.m. follow-ups



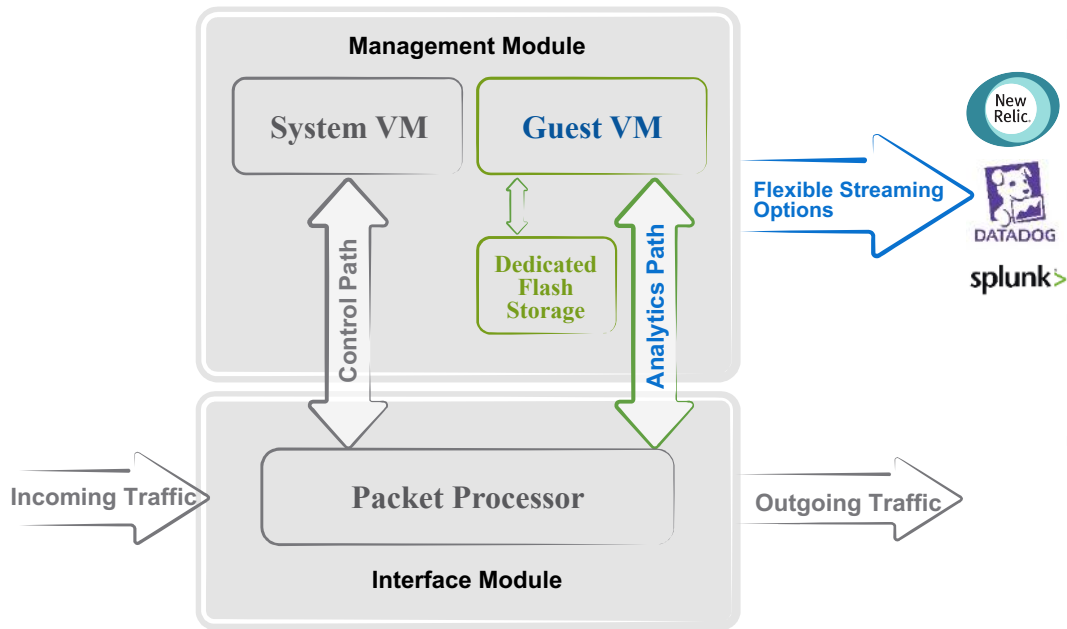
Switch Failure (BGP)

Elapsed Time	Event
	Switch link goes down.
Seconds	Switch sends SYSLOG message.
Seconds	WC SYSLOG sensor matches error message (e.g., REGEX).
Seconds	WC sensor triggers BGP to troubleshoot workflow.
Seconds	WC action extracts information: <ul style="list-style-type: none">• Switch IP, switch peer IP address, egress interface
<1 Minute	WC action logs into switch, performs the following actions: <ul style="list-style-type: none">• Execute workflow to determine service and interface state<ul style="list-style-type: none">• If interface state is down, attempt to restart interface
<1 Minute	WC action creates help desk ticket with outputs from above.
Seconds	WC initiates an alert workflow if interface cannot be restarted: <ul style="list-style-type: none">• Include help desk ID and URL• Post message to Slack for network ops team visibility• Launch incident to PagerDuty to notify operator

Total elapsed time: <5 minutes

Pervasive visibility

- Enhancing operational efficiencies with reduced MTTR



- Guest VM runs in KVM environment on every Brocade SLX switch and Extreme SLX router enabling visibility for network-wide insight
- Dedicated analytics path ensures high performance without disruption to network traffic
- Flexible, high-performance data streaming for off device processing, reporting and archiving
- Dedicated analytics storage provides real-time capture for easy and fast access

Embedded architecture on every Extreme SLX device provides pervasive visibility for improved SLAs



Thank You

WWW.EXTREMENETWORKS.COM

