Cluster Configurator

Open Networking - Switch Cluster Deployments

Sharad Ahlawat

About PICA8

The Leader in Open Networking

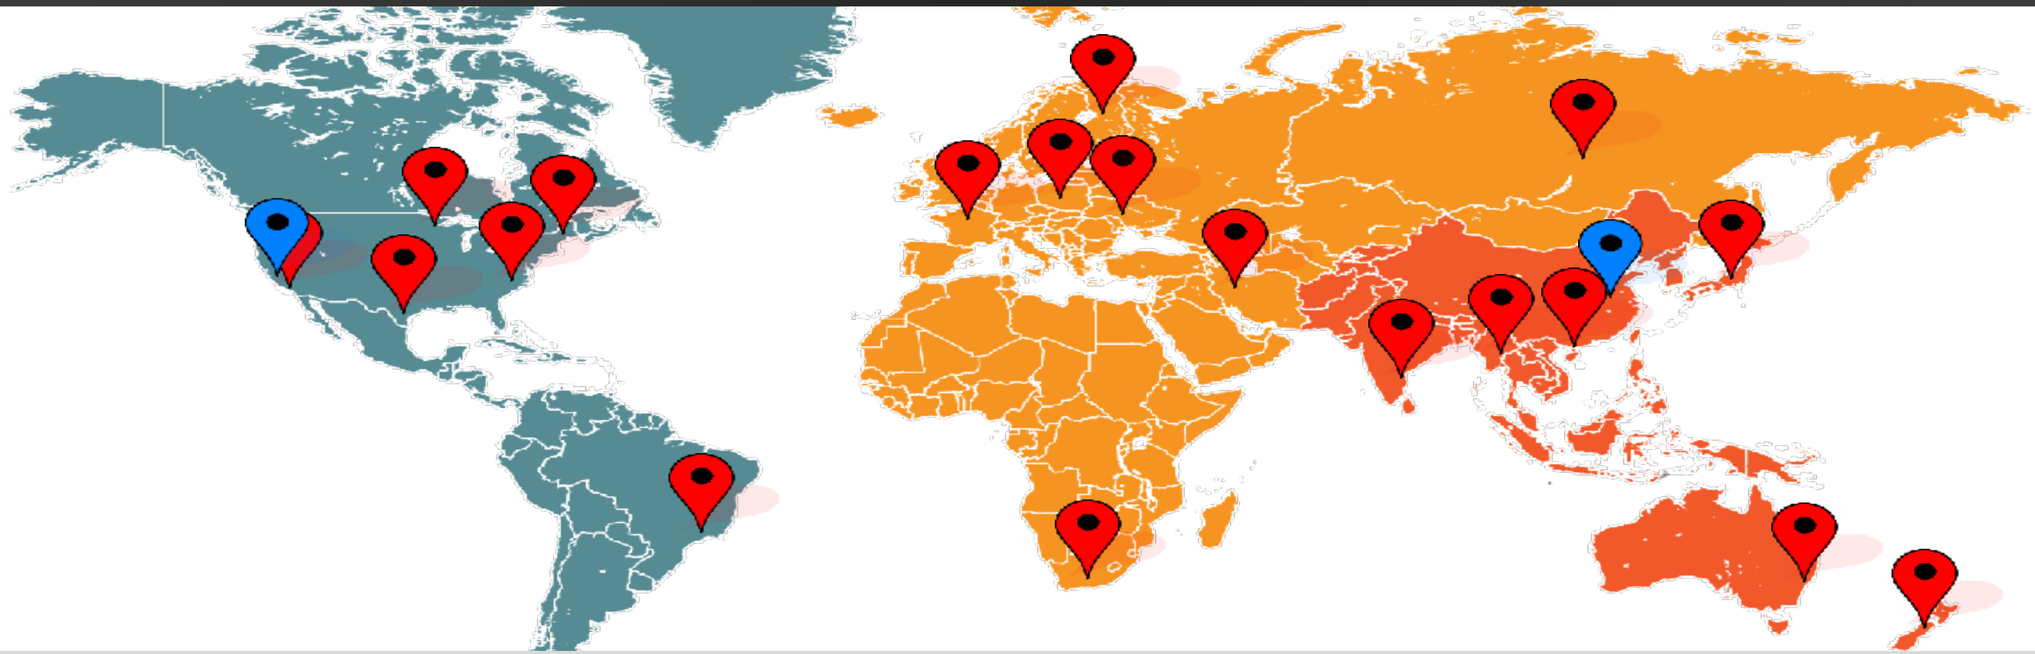

- Leading provider of SDN operating system software for white box switches
- Founded in 2009; first shipment in Q1, 2012
- Global footprint: 75+ employees, Silicon Valley HQ; Beijing R&D
- Over 500 customers and counting

PICA8 OPEN NETWORKING

Some of PICA8 Customers

Enterprise Networking Market

Challenges and Opportunities

- Customer premise equipment are due for a refresh cycle
- Shrinking IT budget and resource, need network infrastructure deployment and management efficiency
- Networking is getting complex and constantly facing security threats
- Current alternatives are vendor lock-in solutions
  - This includes services and applications (eg: stacking and chassis)
- Merchant silicon capable of addressing requirements
- Price, features per-port is software driven value
  - Traditionally these are premium hardware features
- Feature-set is well known and standard across customers.

The White Box Revolution

- White Box cost less, and provides freedom of choice
- Pioneered in Mega Data Centers, now gaining acceptance in Enterprise
- NOS separates hardware from application
- NOS enables different platforms
  - Different speed
  - Different bandwidth
  - Different ASIC features
  - Different use cases – WAN, LAN, Fabric
- NOS enables customized application
  - Customize the traffic distribution
  - Define the fail-over behaviors

PICA8's NOS - PICOS

- PICA8 NOS powers Access features on White Box
- PICOS runs on White Box switches from our partners
  - Provides White Box economics, and freedom of choice
  - Customers have saved up to 70% on CapEx and OpEx
- Efficient Provisioning and Operations
  - Automatically recognizes devices by type, then applies customized treatment for each device type (QoS, PoE, access controls)
  - Centrally managed security policies
  - Works with your existing infrastructure
- Industry Standard Configuration
- Easy customization by IT staff

CrossFlow Architecture

- **Additional** control plane based on CrossFlow (PICOS integrated OpenFlow)
- **Separate** control plane for security so you can respond to new threats quickly without affecting ACLs
- **Dynamic** fine-grain selective monitoring for performance analysis and network

Use Case: Retail Stores or Branch Offices

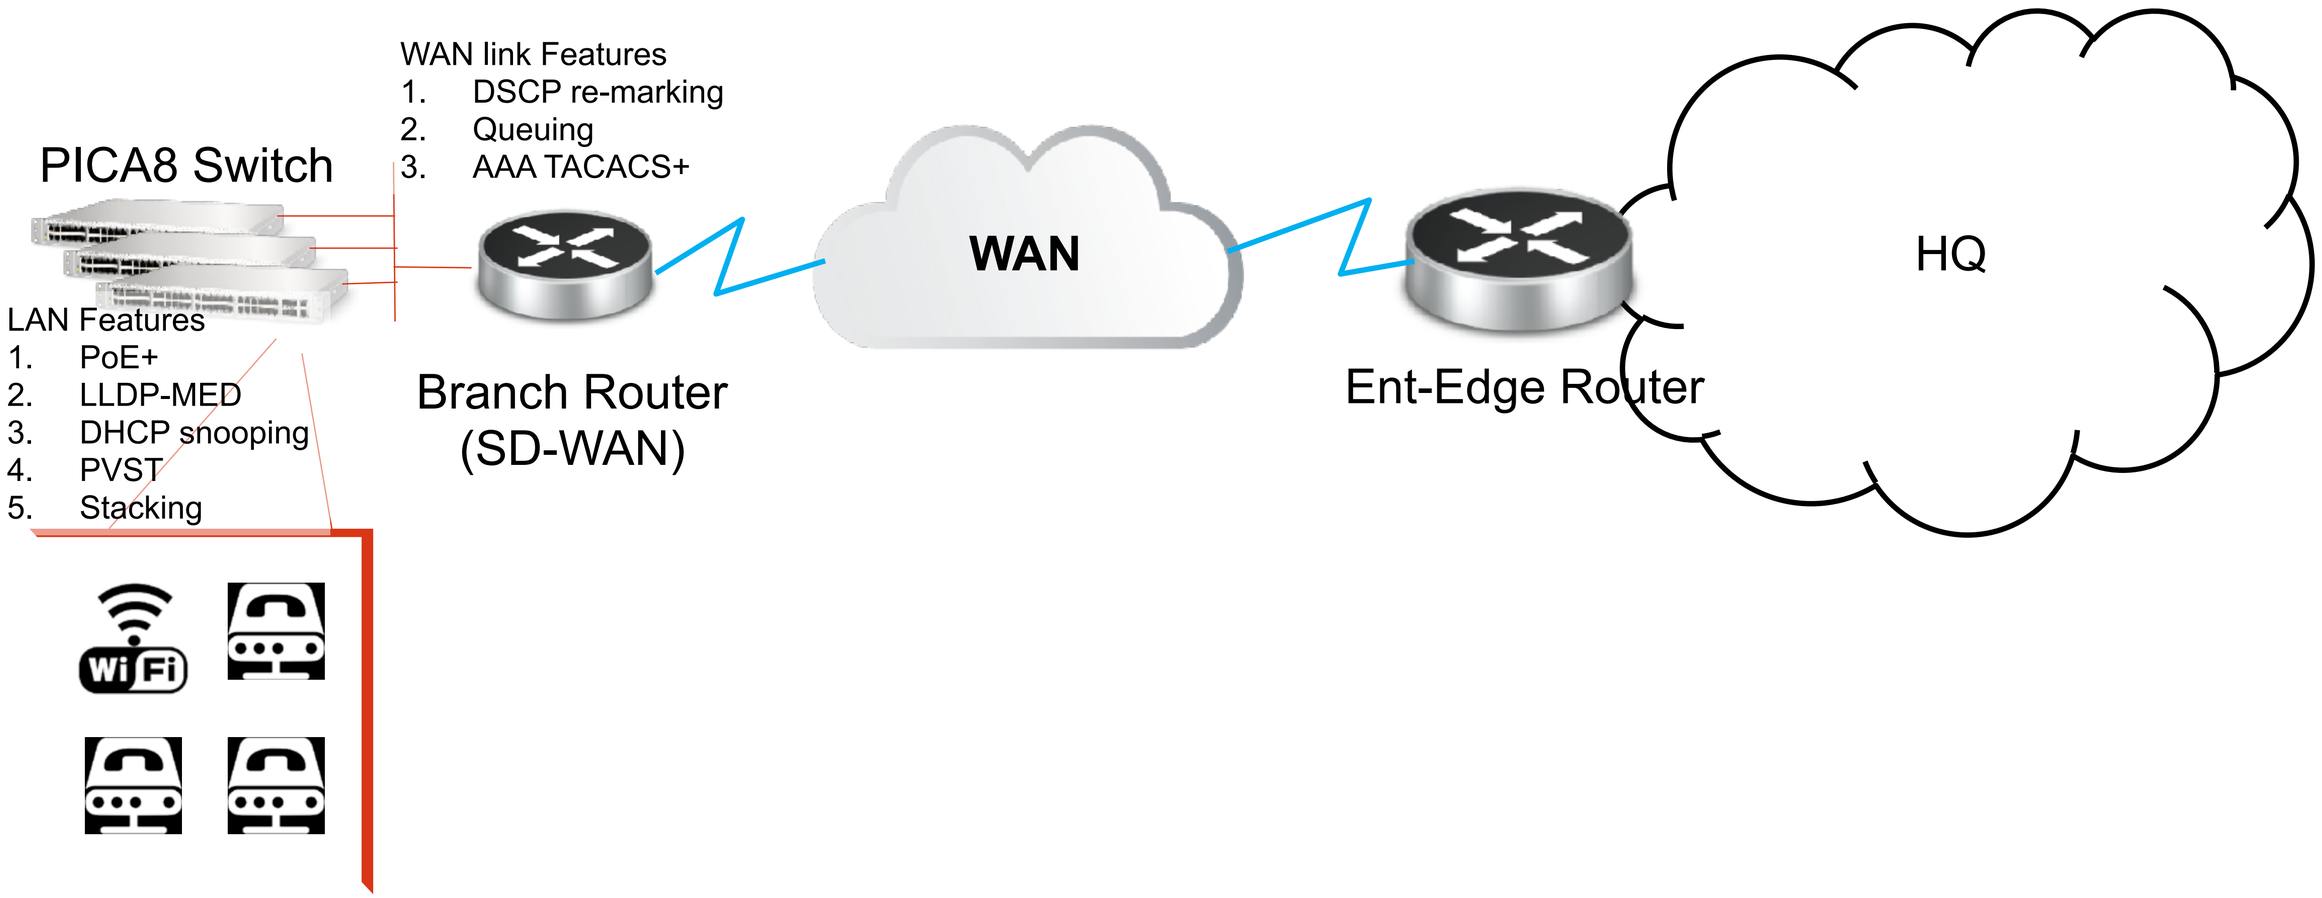

Use Case: Campus LANs

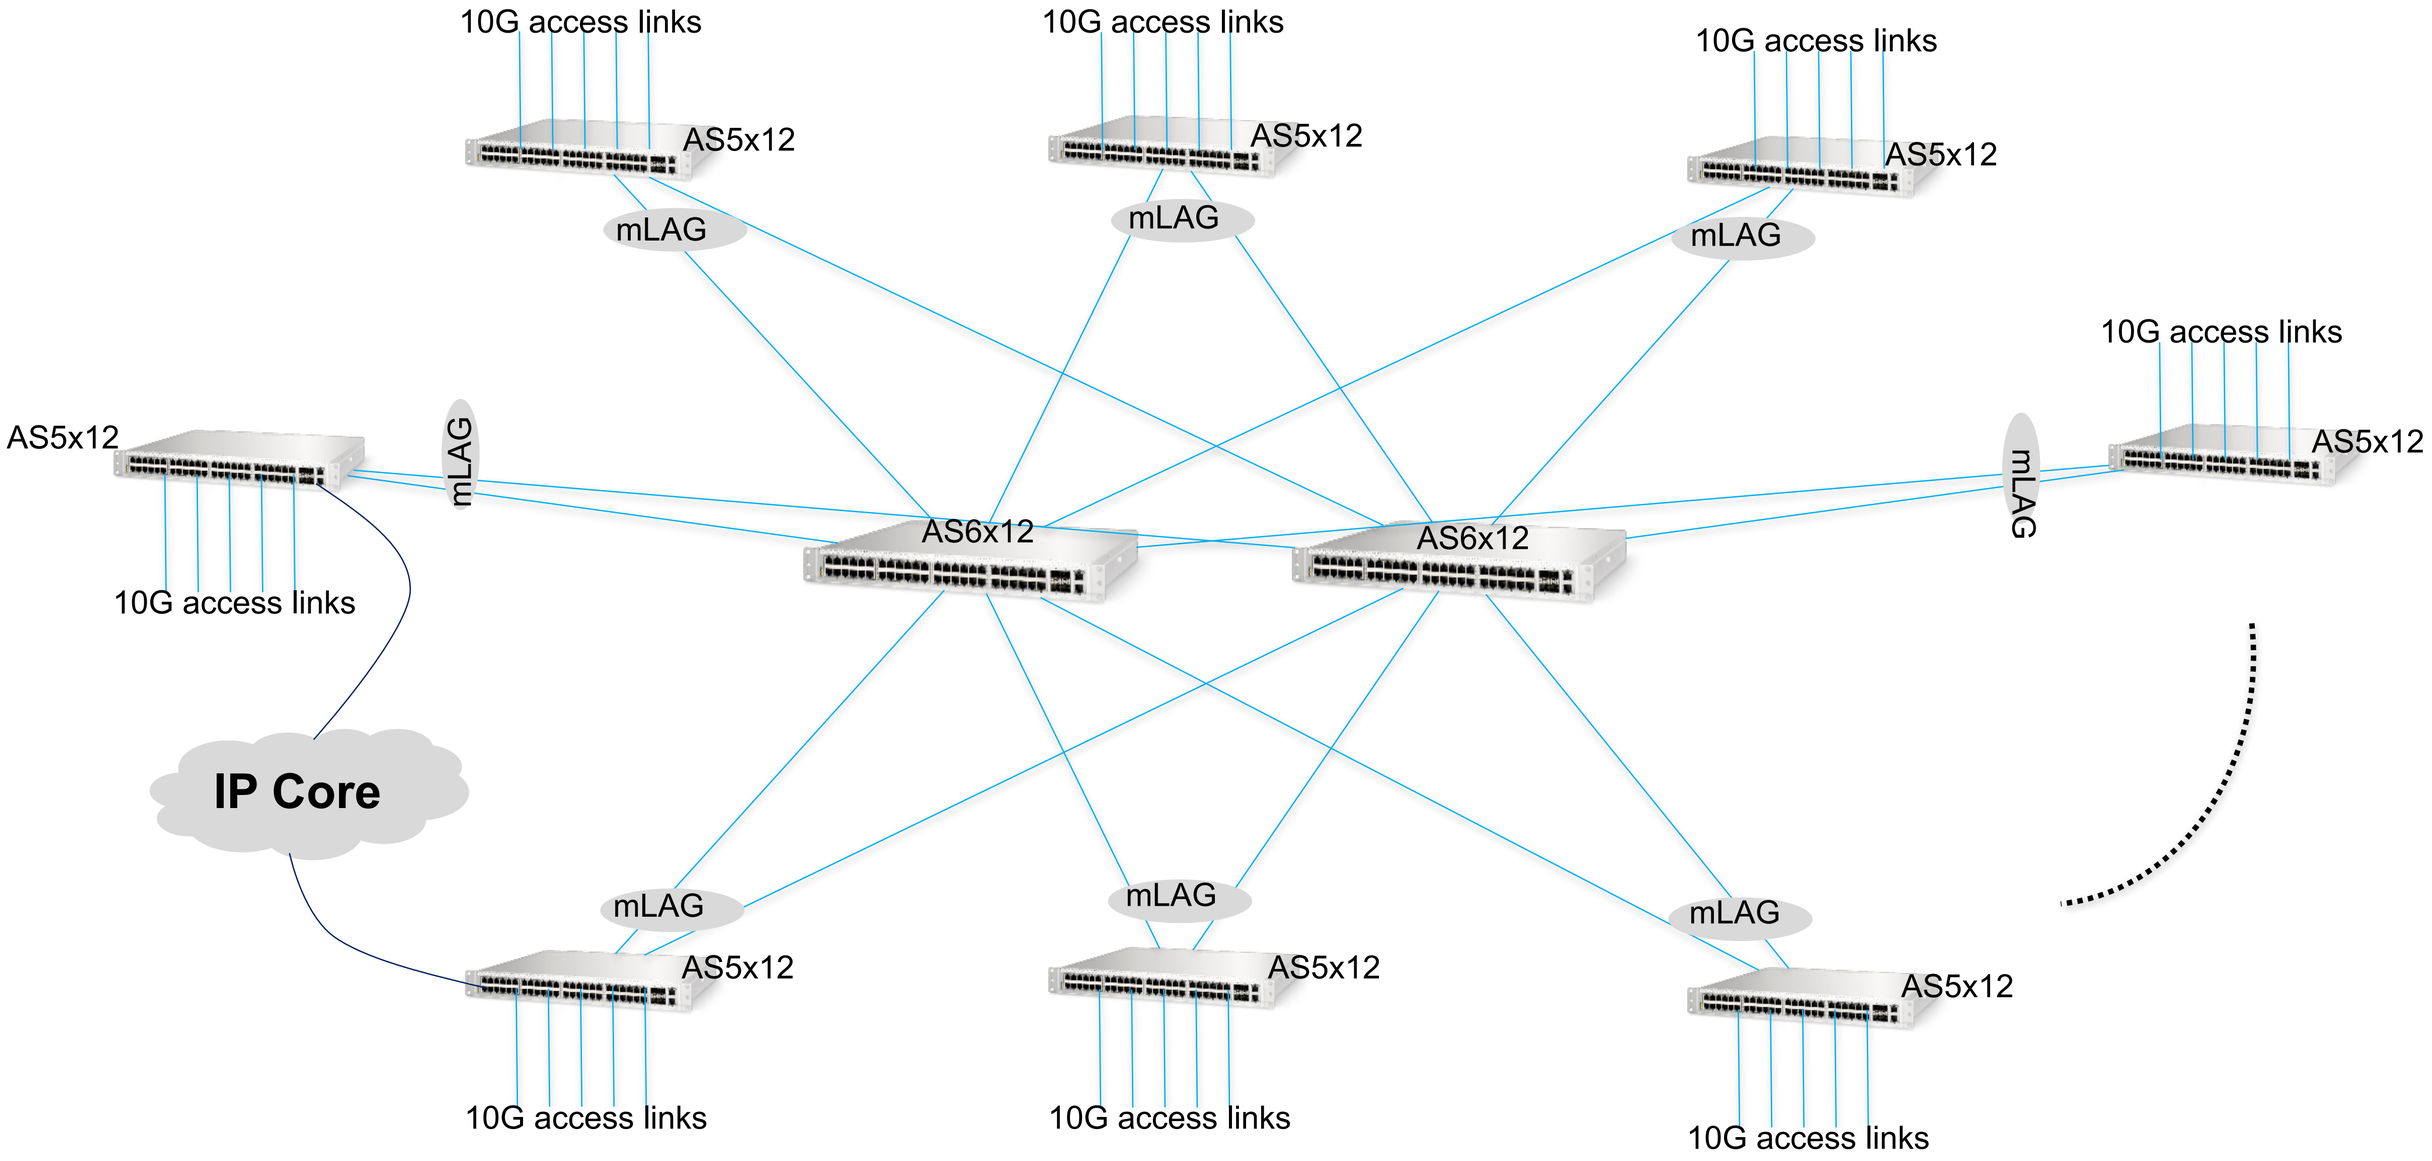

Features Required in Enterprise Switches

- Switching and Routing with RPVST+ and VRRP
- Sophisticated QoS
  - classification + re-marking + queuing to support voice, video and data
- PoE to support Voice, WiFi and other devices
- Recognize devices and provide the right network and power profile
  - LLDP-MED, CDP
- Secure remote management
  - AAA TACACS+, SNMPv3
- Switch protection
  - CoPP, BPDU guard
- Network Protection
  - DHCP snooping, IGMP snooping
- Stacking and Chassis configurations

A Switch Upgrade Opportunity - Cisco Catalyst 3850

- Legacy Product
  - Cisco Catalyst 3850 48-port 1GE POE+ switch
  - TCO (3 yrs): $10,000

- Replacement with Open Networking
  - Edge-Core AS 4610-54P 48-port 1GE POE+switch
  - PICOS NOS Software
  - TCO (3 yrs): $3,000

- Savings
  - Deployment at 3000 sites: 5000 Switches
  - $35,000,000

Stacking and Chassis - traditional switch clusters

- Both Stacking and Chassis are port aggregation solutions with stacking being the poor man’s solution.

- Stacking
  - String multiple switches together.
  - Software and Hardware tied together.
  - Proprietary, Lock-in, Identical HW for stack, Special Stack Ports.
  - Switch failure in a stack affects performance of all switches.
  - One control plane split across two switches.
  - Replacements and scaling up/down workflows require the entire stack to be taken offline.

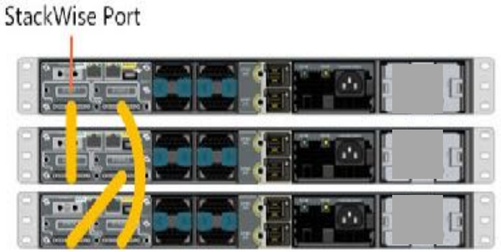

- Chassis
  - Offers better reliability and resiliency compare to stacking.
  - Software and Hardware tied together.
  - Proprietary, Expensive, Lock-in, Fixed Slots.

CC - enabling White Box switch clustering

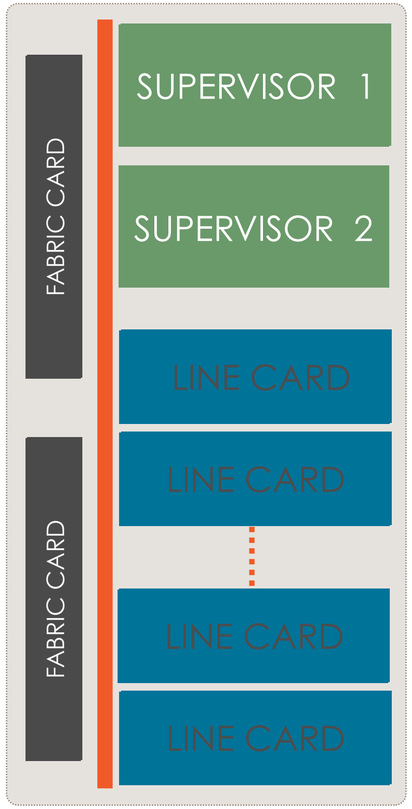

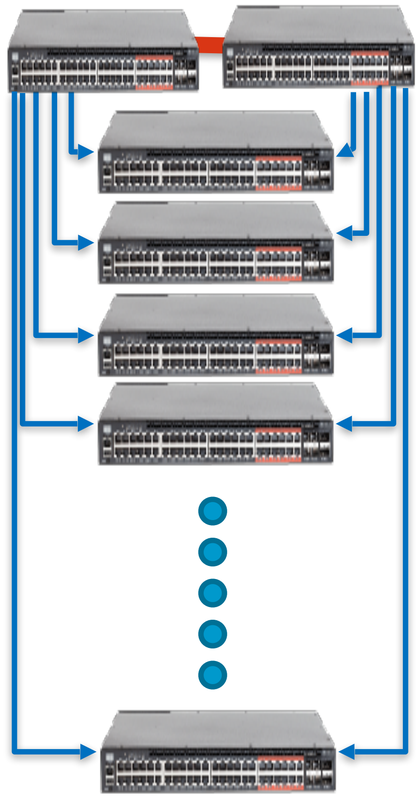

- The CC fabric is a MLAG mesh between the switches
- Leaf's are attached to each of two Spines
- Network traffic is load balanced across the LAG links
- STP is disabled to provide Active/Active connectivity and no blocked links
- CC internal management plane uses:
  - Isolated network subnet
  - Isolated VLAN trunked across all the fabric LAG links
  - SSH protocol for secure communication between the spines and leafs in the cluster
- CC commands to configure and manage cluster operations

A White Box Chassis solution with Edgecore Switches (WBC-E)

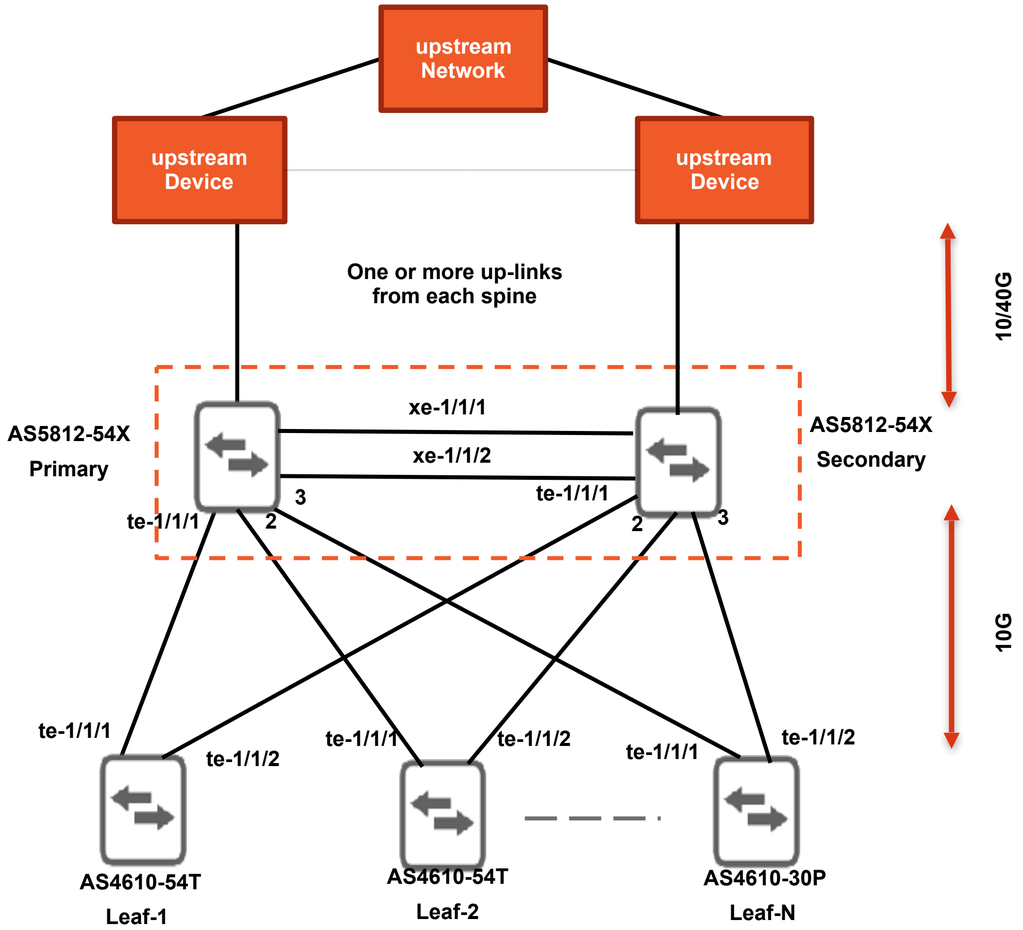

- Spine images for AS5812-54X
  - 48 x 10G
  - 6 x 40G
- Leaf images for AS4610-54T
  - 48 x 1G or 24 x 1G (48 and 24 port models)
  - 6 x 10G
- up to to 16 Leaf switches or more distributed across racks or even the campus.
- Max user ports
  - 768 x 1G
  - 96 x 10G
  - 8 x 40G
- **CC Enabled**

WBC-9E - A convenient Chassis Frame solution to upgrade aging 6509s

- Full feature **PICOS with CC**.
- 9 White Box Edgecore switches.

  2 x AS5812-54X

  7 x AS4610-54P

  or a mix of AS4610-54T, AS4610-30P and AS4610-30T.
- Custom Switch **Chassis Frame** that mounts onto the 19"rack.

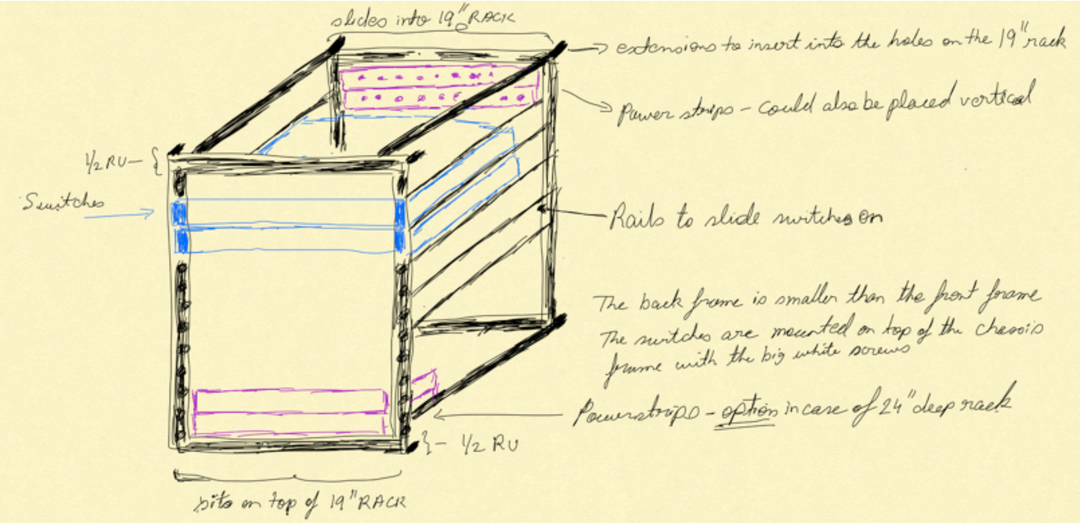

- The Switches and PDUs are mounted on the frame and the switches themselves are held in place in the chassis frame using thumb screws for easy insertion and removal.
- Short size fitted DAC cables, neatly bundled and labeled to help quickly and correctly connect the switches onsite.
- Power requirements:
  - 4610-54p – w/o PoE is 100W – w/both PS providing full PoE power ~1800W
  - 5812-54x – no PoE is 350W
  - So for the WBC-9E, all switches would be at max 1400W and then extra for PoE
- PDUs:
  - 30A PDUs would deliver 120V/3600W or 240V/7200W of power.
  - Not all 4610 ports will be configured for max PoE at any given time as a single 4610 power supply is only capable of ~900W, also most VOIP phones only consume 7.5W or 15W, so two of these PDUs would be sufficient for the WBC-9E.
  - For 36" deep racks the PDUs can be mounted on the chassis frame back top corner but for 24" deep racks the frame would be extended by 2 RU and the switches would be mounted in the front, below the switches.
  - The switches have short size fitted power cords connecting the switches to the PDUs, ensuring PDU/circuit redundancy.
  - The PDUs save on power outlets, as all outlets may already be used and finding another 16 additional outlets could be an issue.
  - If the user wants to use their own RACK PDU's they can opt to not have these PDU's preinstalled in their WBC and prewired up to the switches.
  - The PDUs could be upgraded to switched PDUs and managed by the CC application if desired.

A White Box Stack solution with Edgecore Switches (WBS-E)

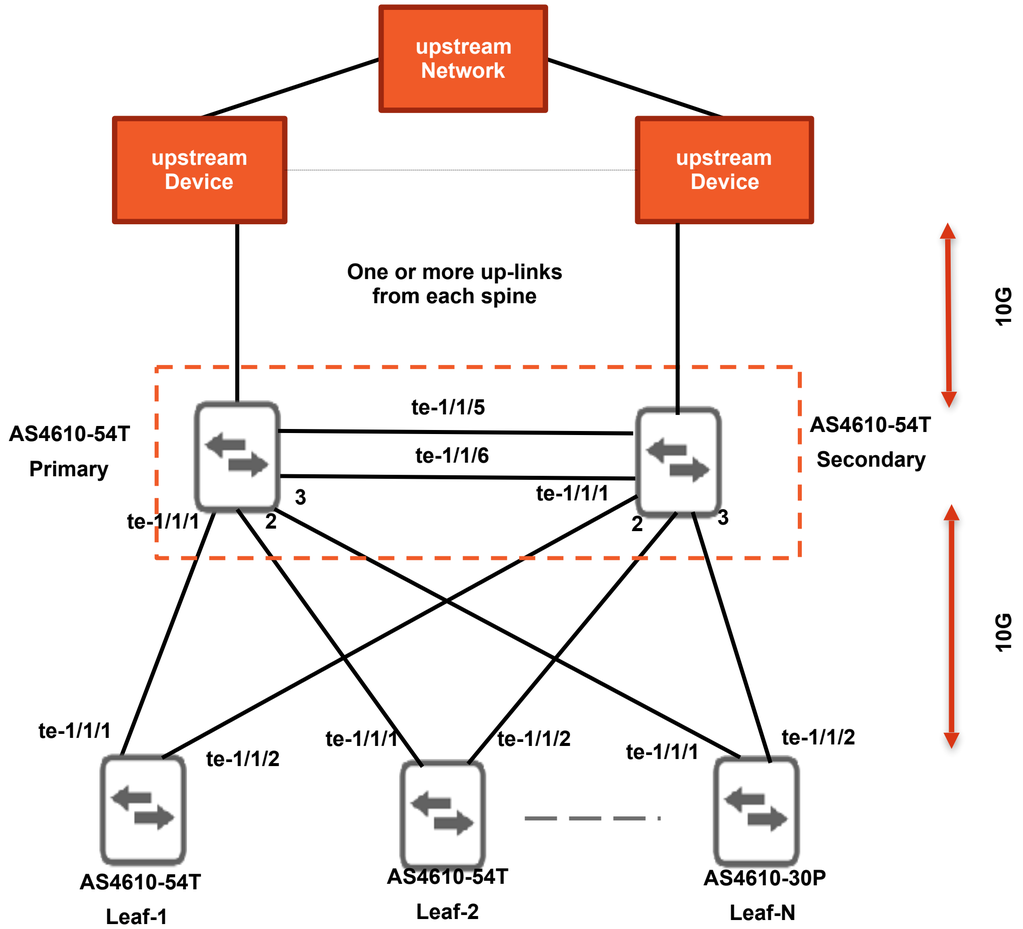

- Spine images for AS4610-54T
  - 48 x 1G or 24 x 1G (48 and 24 port models)
  - 6 x 10G
- Leaf images for AS4610-54T
  - 48 x 1G or 24 x 1G (48 and 24 port models)
  - 6 x 10G
- up to to 4 Leaf switches.
- Max user ports
  - 96 to 288 x 1G
  - 4 to 16 x 10G
- **CC Enabled**

A Chassis Upgrade Opportunity - Cisco Catalyst 6500

- Modular chassis switch. 3,4,6,9,13 slots (6503, 6504, 6506, 6509, 6513).

- Shipping since 1999. Thousands of 6500 are EOL now. No support, no warranty. Urgently need a replacement plan.

- “Catalyst 6500 has generated more than $45B in cumulative revenue for Cisco, thanks to its large footprint of close to 800,000 systems, 110 million ports shipped thus far to some 45,000 customers” - source: Lippis Report, 2012

- Cisco plans to retain its hold onto this market segment with the Cat 6500E and the Cat 9600 https://www.eweek.com/networking/end-of-an-era-cisco-bids-adieu-to-catalyst-6000-hello-to-catalyst-9600

- Customers embracing Open Networking are looking for more flexible and efficient solutions to replace their 6500 deployments with an alternative that offers the features of a chassis with the modularity and price of stacking.

WBC-9E vs Cisco Catalyst 6509E

	WBC-9E	Cisco 6509E
Enabling technology	CC on PICOS w/ standard MLAG	Proprietary Fabric
CapEx of 236GE ports	$54,800	$173,500
5-year OpEx of 236GE	$10,800	$42,300
Minimum setup cost	$16,500	$74,500
Maximum GE ports	768	236
Non-blocking bandwidth	3.84 Tbps	0.72 Tbps
Hardware choice	Multiple white boxes with multiple configurations	Only Cisco 6500 cards
Ease of management	Single IP management	Single IP management

PICOS and Open Networking

Programmable **I**nternetworking & **C**ommunication **O**perating **S**ystem (**PICOS**)
form
Programmable **I**nternetworking & **C**ommunication **A**rchitecture / Infinite(8) possibilities (**PICA8**)

PICOS Architecture

Modular

- ✓ Unmodified Debian kernel
- ✓ Modular architecture
- ✓ Open source technologies
- ✓ ASIC abstraction

Scalable

- ✓ Separation assures extensibility and performance of each plane

Open

- ✓ Standard Debian packages
- ✓ All scripting frameworks and languages available
- ✓ Open APIs

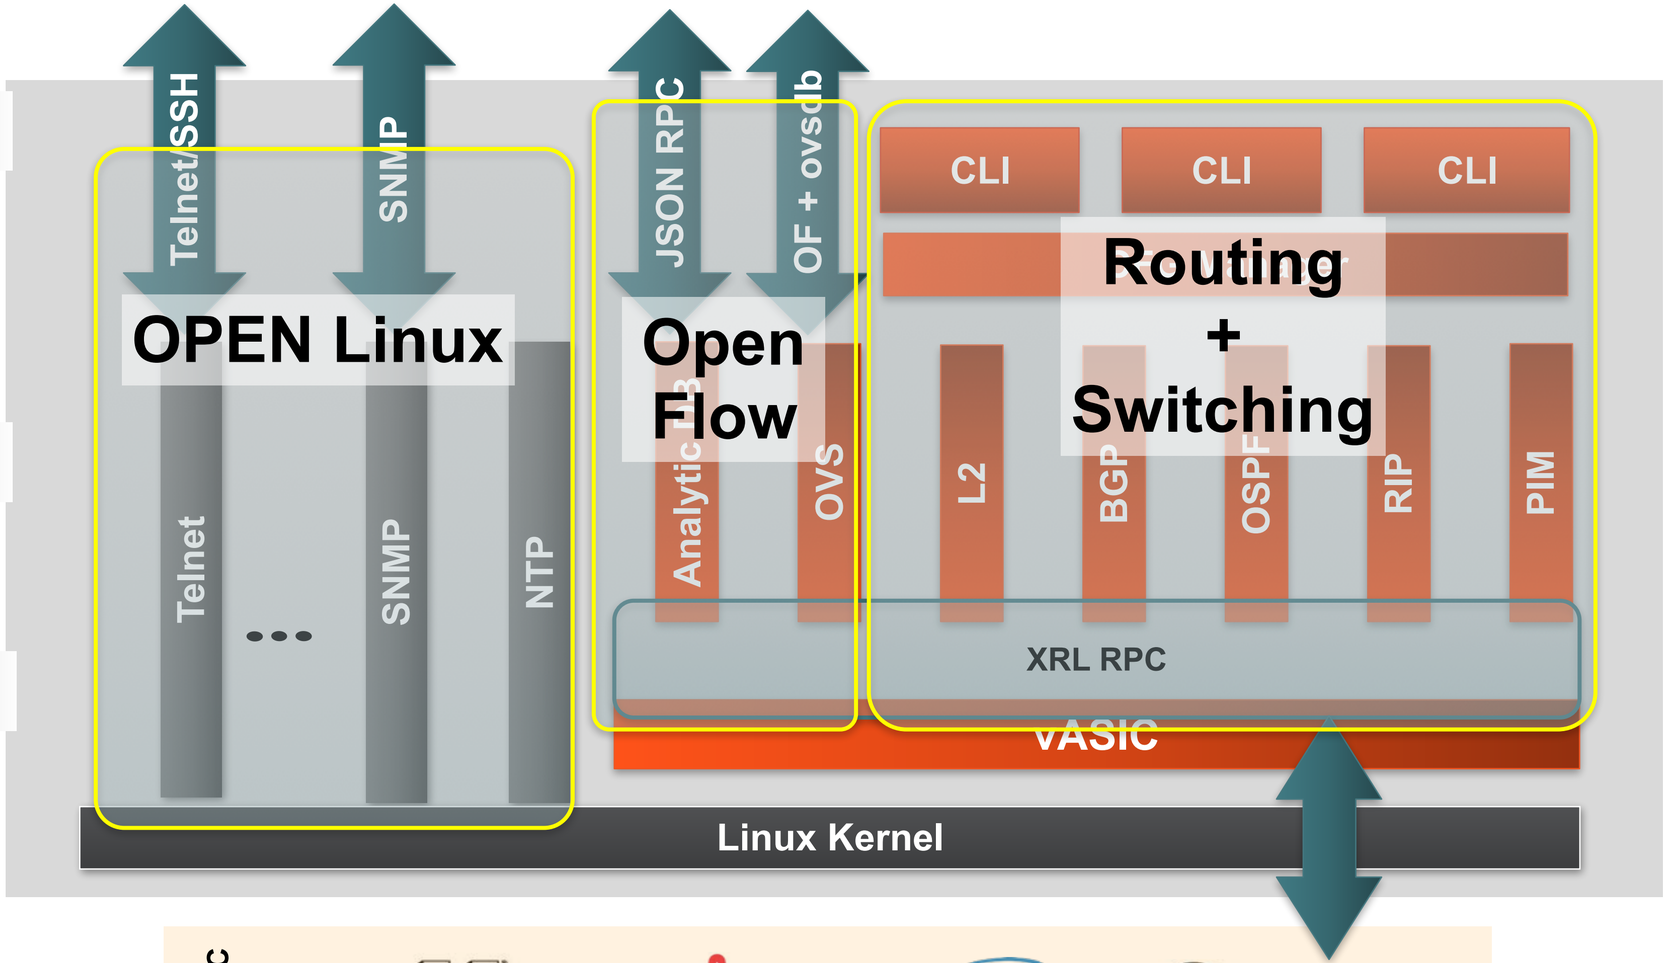

PICA8 OPEN NETWORKING

PICOS Full Feature QoS Support

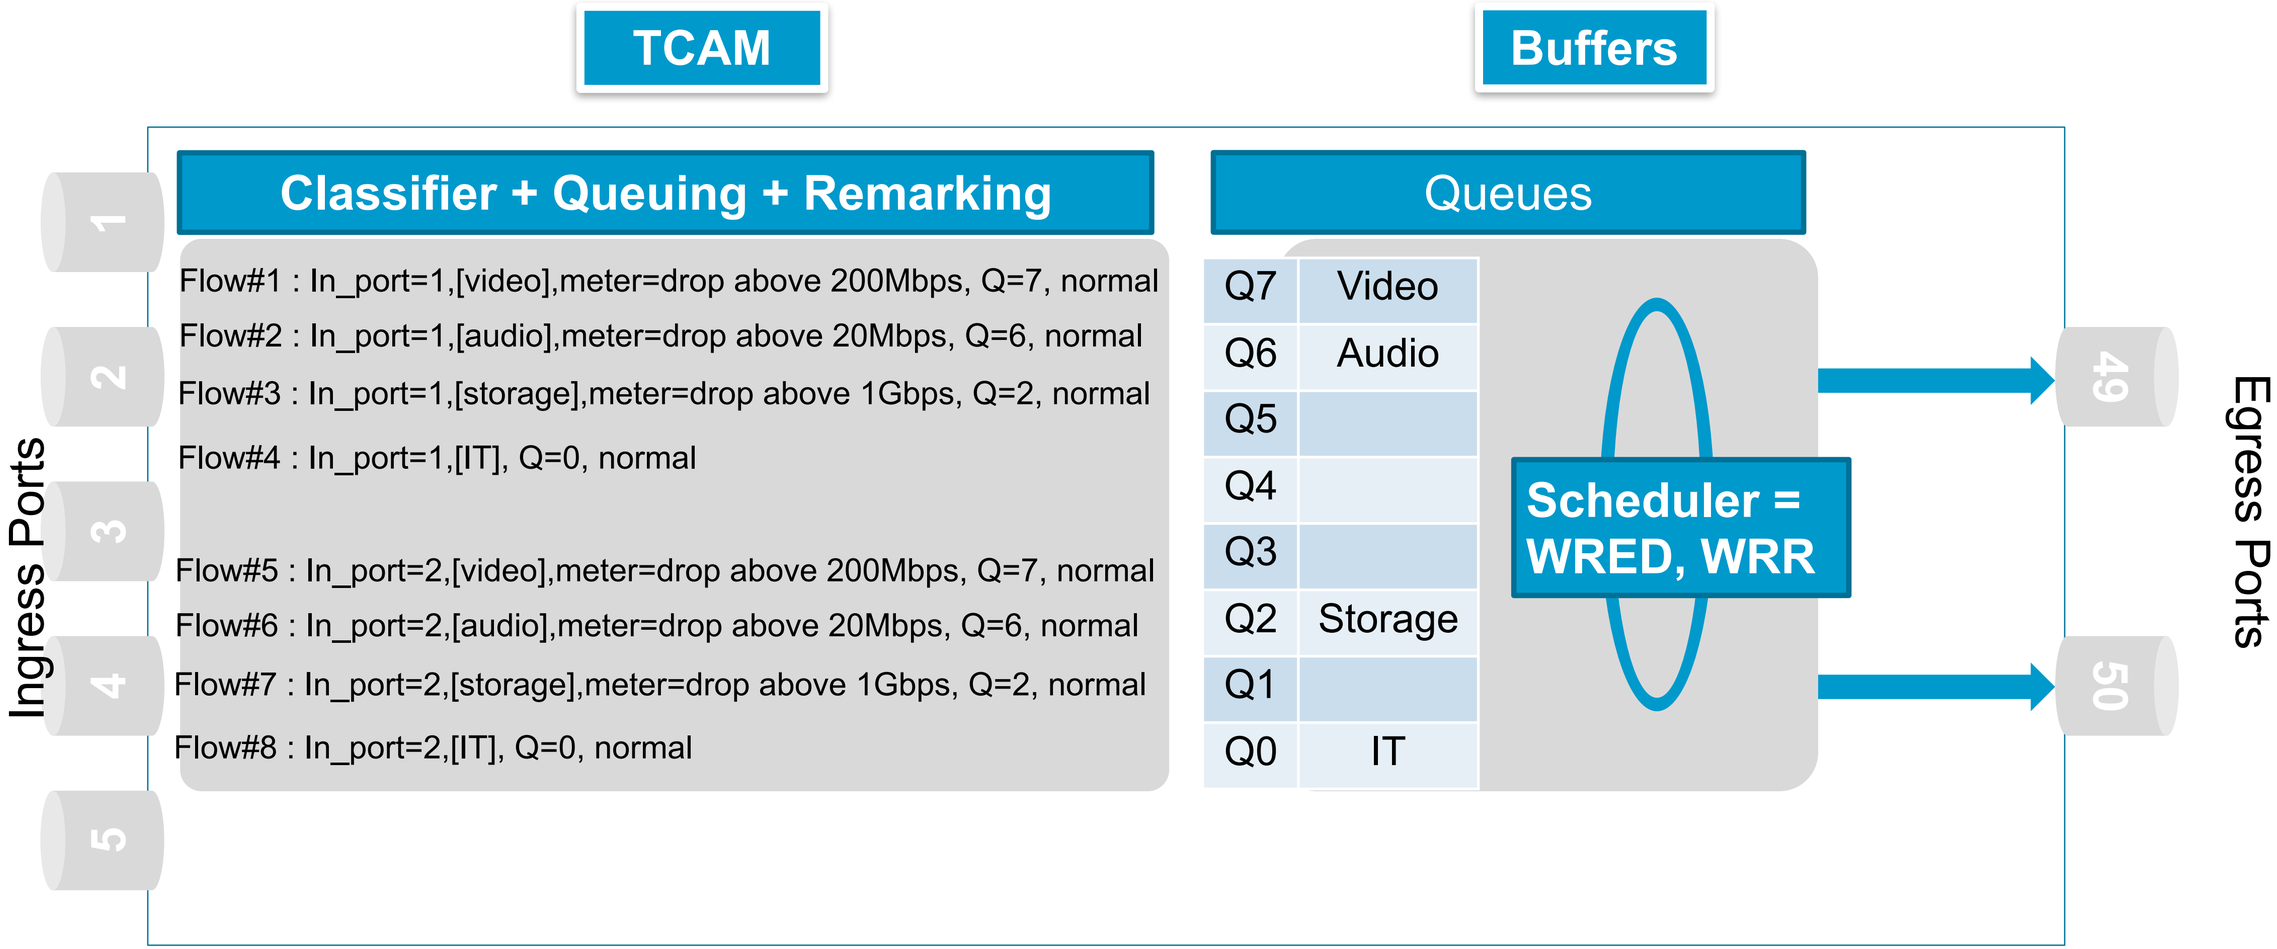

Current OpenFlow implementations in the market

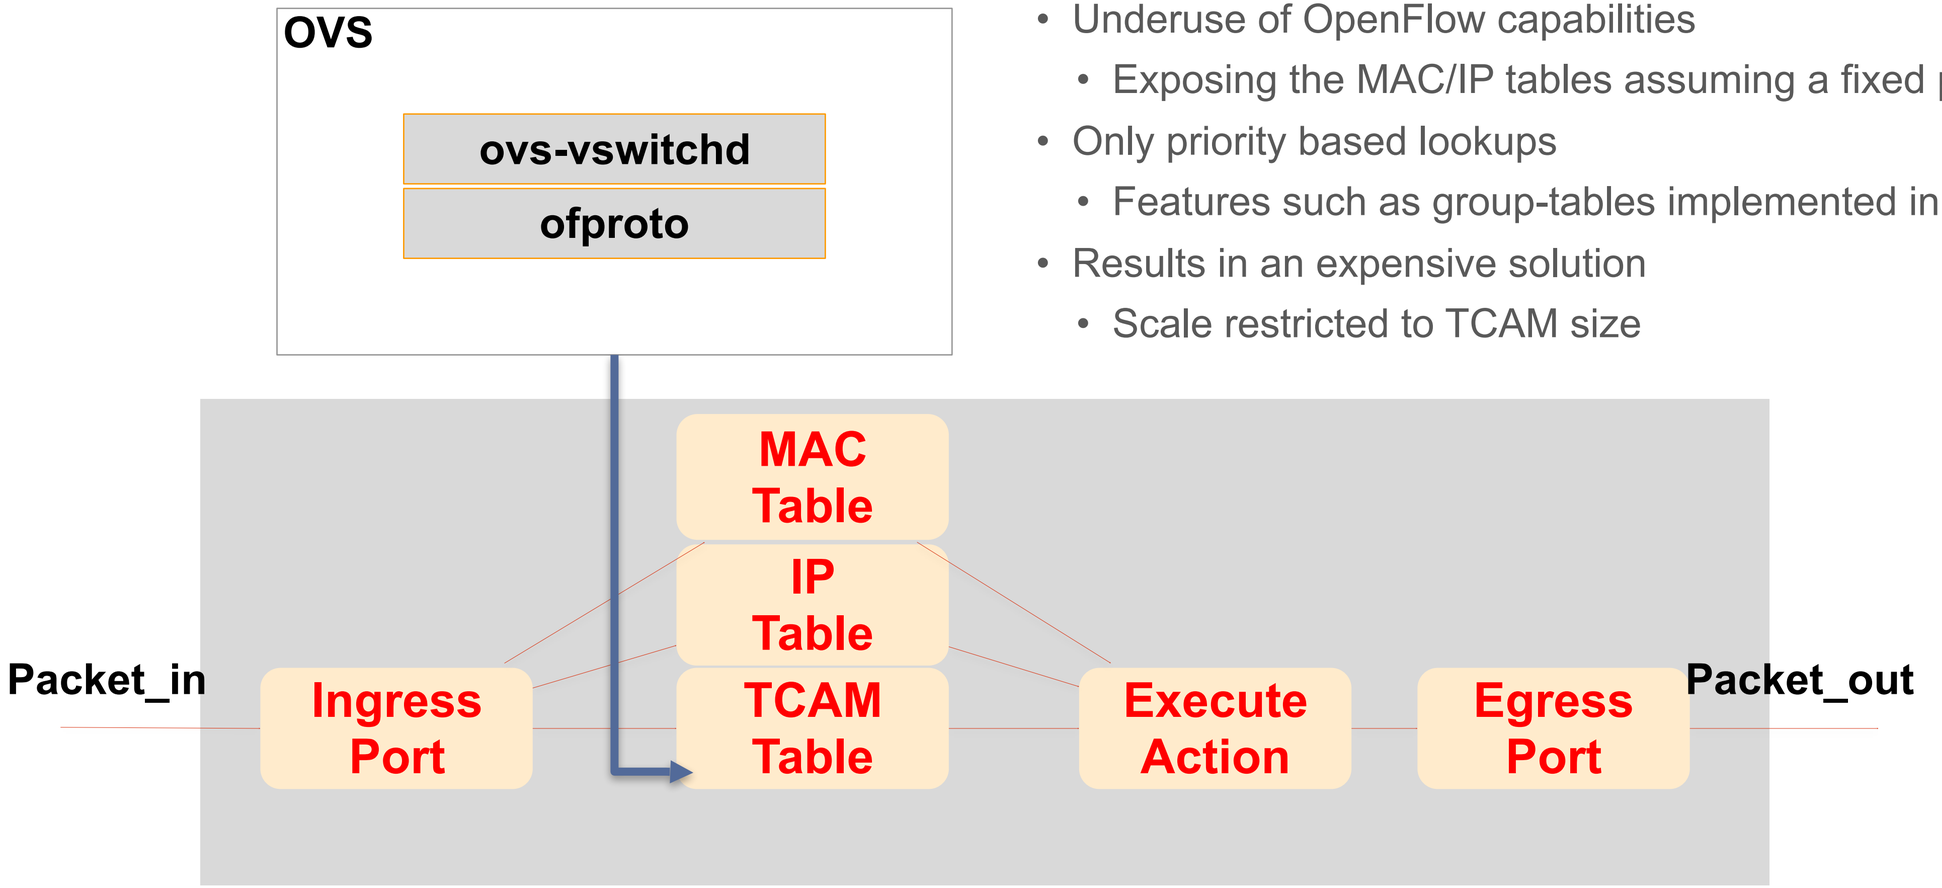

- Under utilization of ASIC capabilities
  - OVS programs the TCAM tables
- Underuse of OpenFlow capabilities
  - Exposing the MAC/IP tables assuming a fixed pipeline
- Only priority based lookups
  - Features such as group-tables implemented in TCAM
- Results in an expensive solution
  - Scale restricted to TCAM size

PICOS OpenFlow implementation with Table Type Patterns

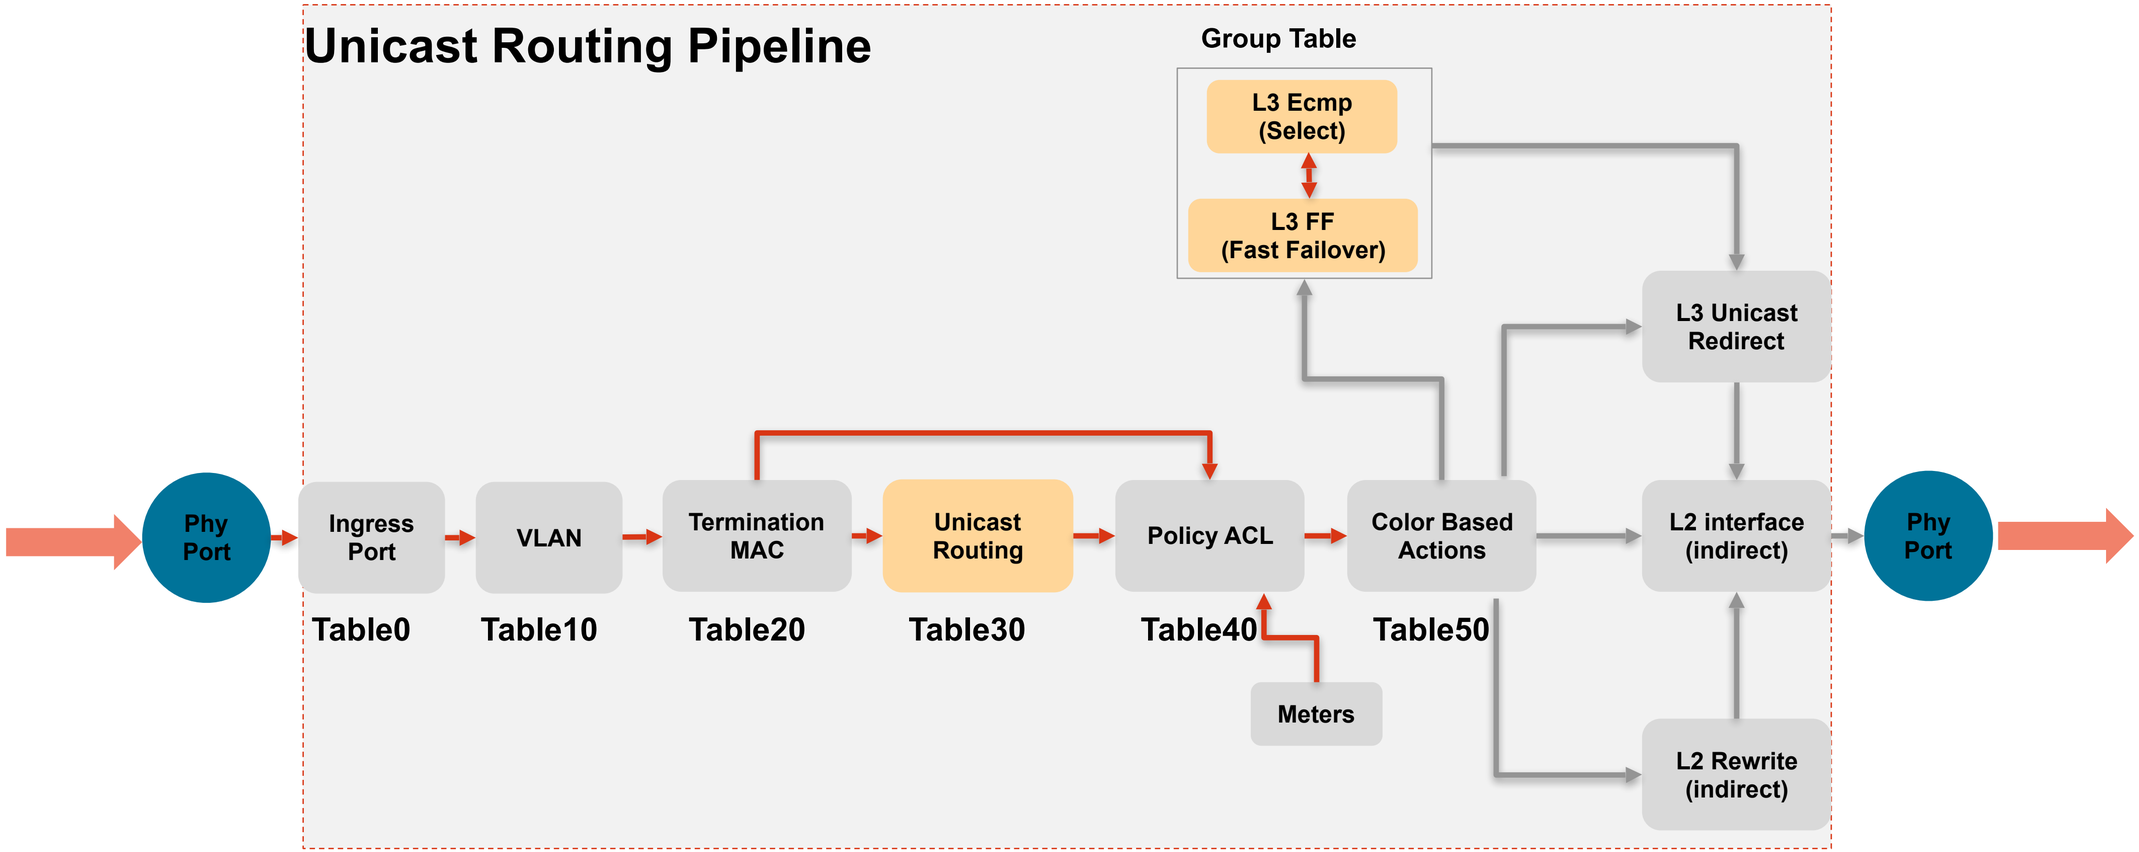

- Utilize all ASIC tables (VLAN+MAC+IP+TCAM) via OpenFlow
- Enables flexible pipelines
  - Utilize priority or LPM algorithms for lookups
- Scale comparable to incumbent routers/switches

PICOS OpenFlow Provisioning

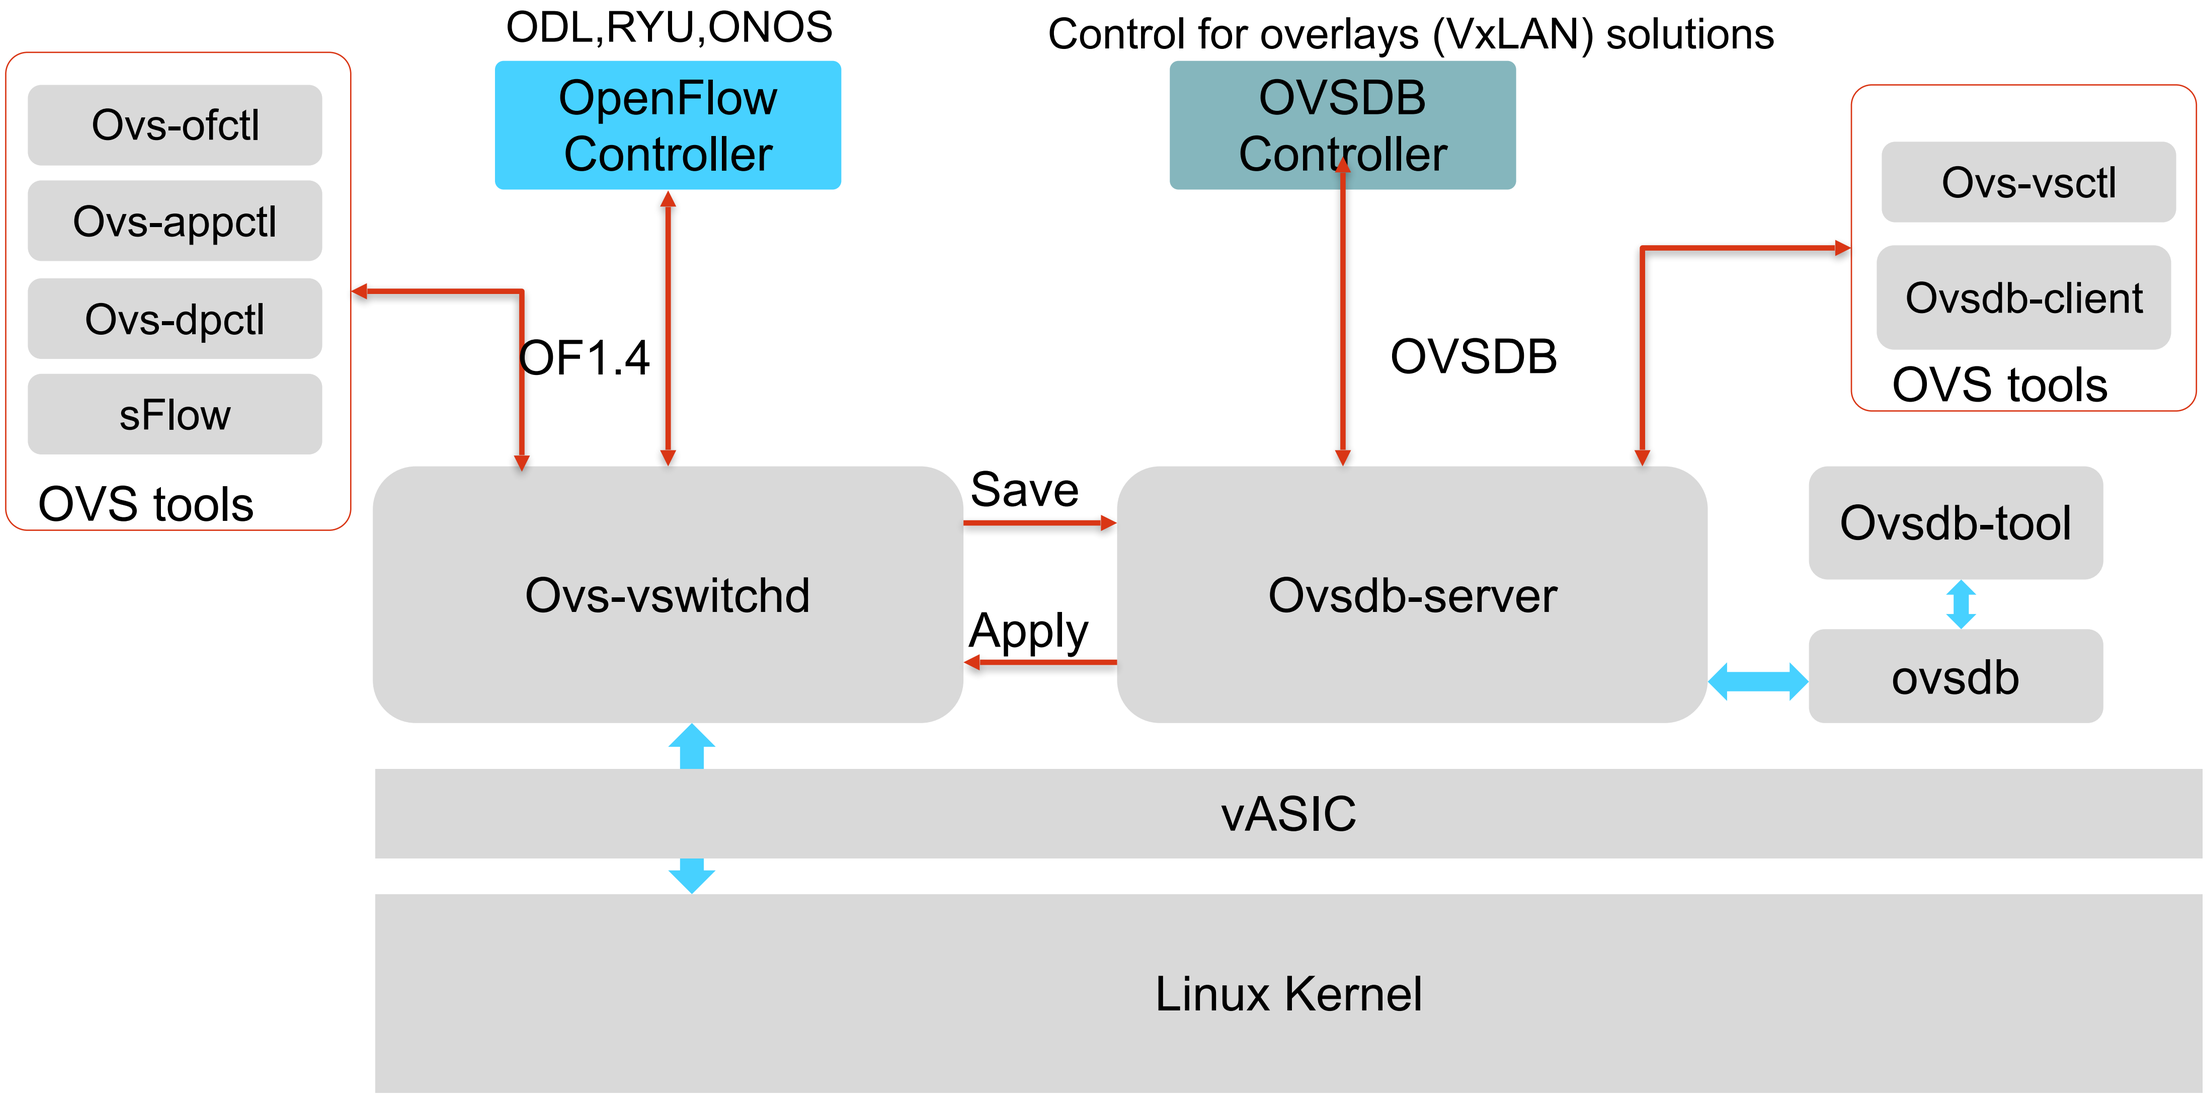

CrossFlow Ports

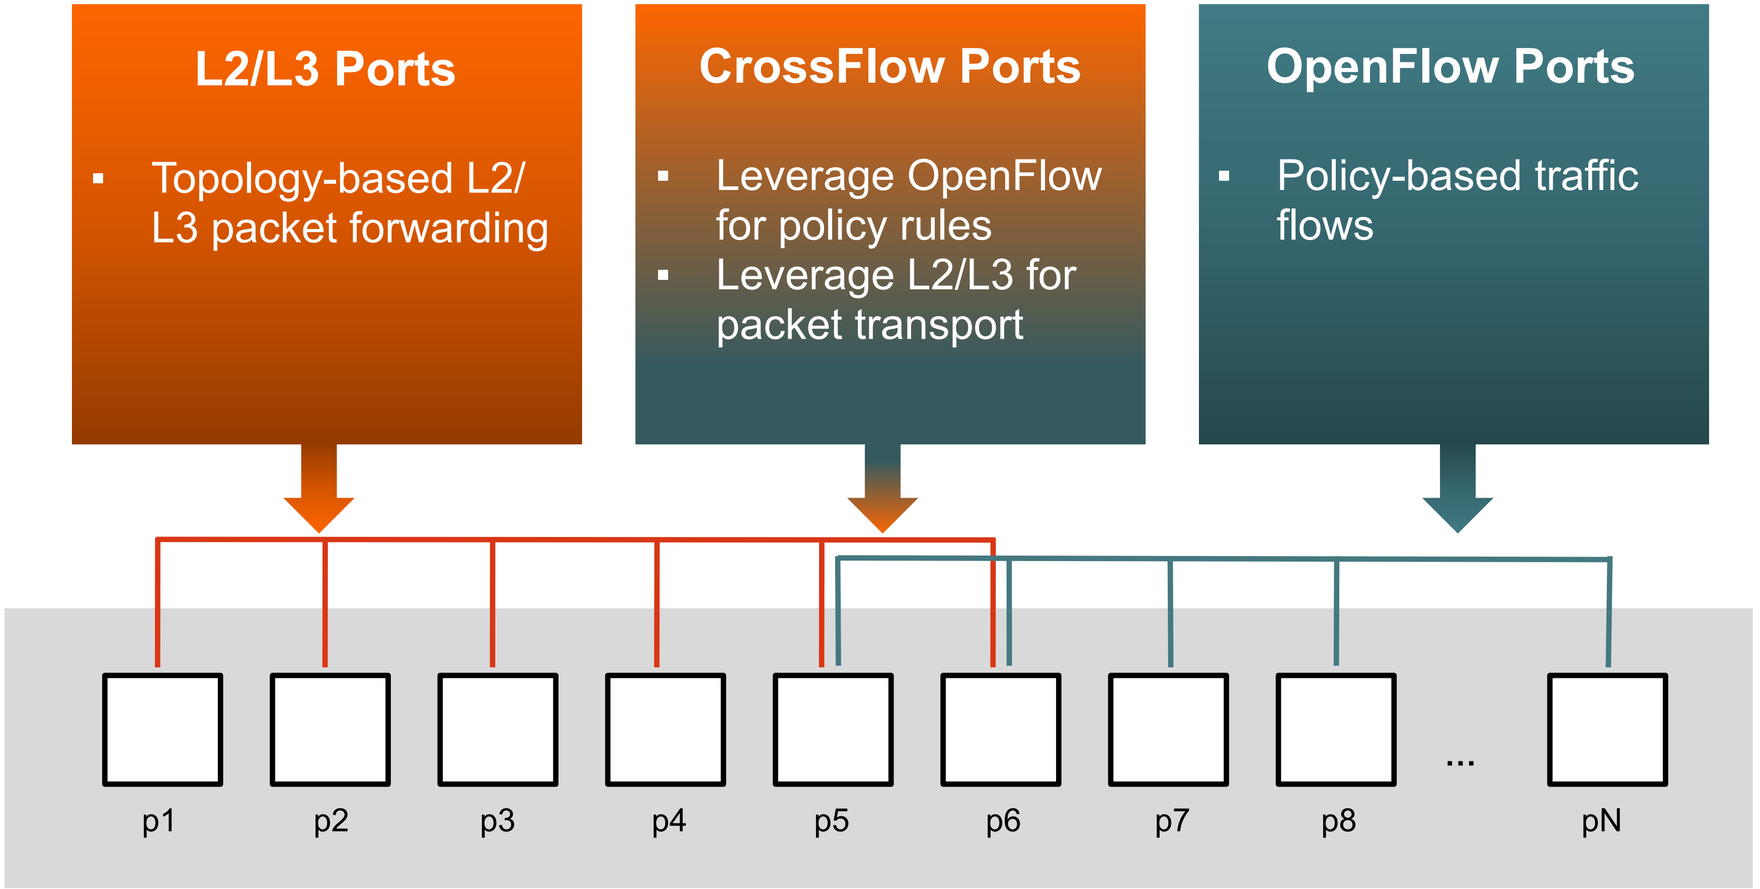

PICOS loaded onto bare metal switches

L2/L3 and OVS share all hardware resource.

Allocate resources in Switching CLI and the remaining are used for OVS.

- Ingress ACL entries count
- Egress ACL entries count
- L2 MAC entries count
- L3 route entries count

set interface stm …

CrossFlow Packet Flow

Packets on L2/L3 ports	Packets on CrossFlow ports	Packets on OpenFlow ports
• L2 • L3 • ACL, Filter • Mac learning • Can flood to L2/L3 ports and crossflow ports	• L2 • L3 • ACL, Filter • OpenFlow Flow • Mac learning • Can flood to L2/L3 ports and crossflow ports	• Flow-based. l2/l3 flow, TCAM flow(ingress/egress) • Default drop flow • No mac learning • Cannot flood

CrossFlow - Accelerate True SDN Adoption

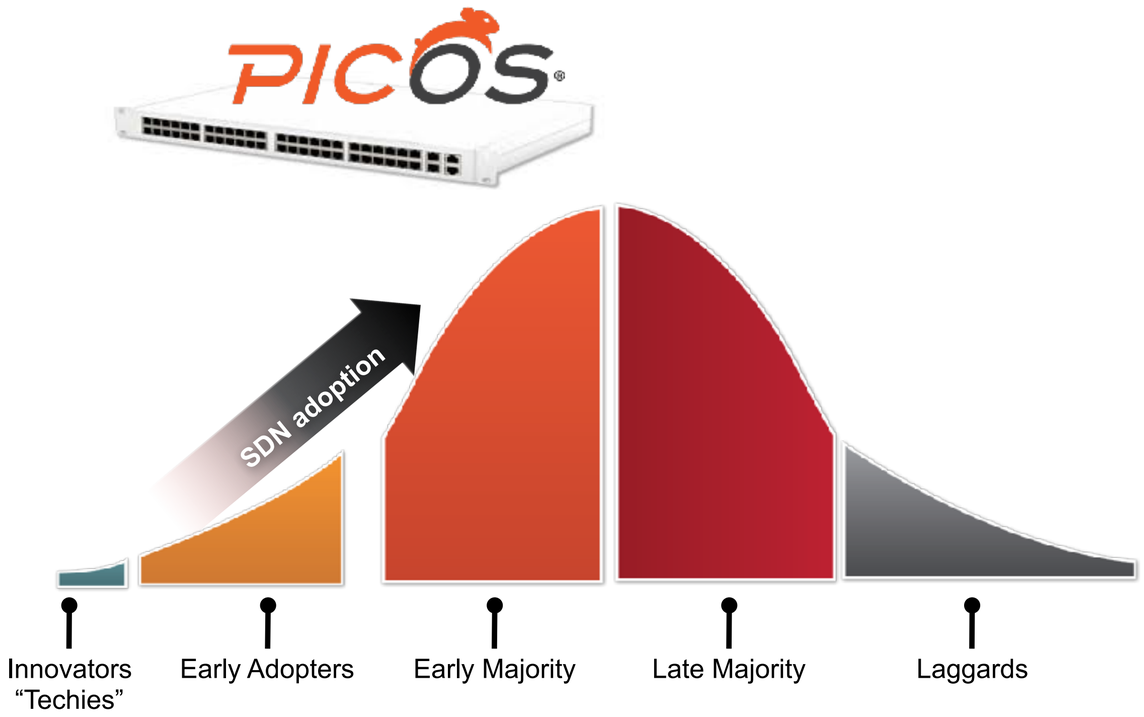

CrossFlow Networking enables:

- Policy driven – network driven by business logic and operational needs
- Flexibility to use OpenFlow with full integration to legacy equipment and L2/L3 protocols
- Reduced CapEx and Operational complexity

The CC Application

Cluster Configurator (CC)

What should a Switch Cluster management solution offer

- Simplify initial bring up
- Streamline Installation and Operations
- Repeatability of deployment and configuration
- Flexibility to mix and match different switches based on speed and desired connectivity requirements
- Redundant Control/Data-Plane
- Campus-wide distributed cluster with LR fiber links
- Open Standards based and non-proprietary solution
- Deployment in existing L2 or L3 customer topologies

Introducing the PICOS Cluster Configurator application

- Single Management IP
- Single CLI shell for all switches
- No OOB connectivity required for leafs
- Commands to configure and manage cluster operations
- Aggregation and relay of Syslog and SNMP
- Centrally managed SSH, NTP etc. configurations
- Single point of failure resiliency
  - Redundant Configuration-Plane
  - Redundant Control/Data-Plane
- Auto/ZTP provisioning for scaling and replacements with no downtime
- Simplified Licensing
- Web Interface

CC - The Open Switch Cluster solution

- The CC application delivers a flexible and scalable approach to address customer networking port aggregation requirements
- CC simplifies customer deployment and operational workflows
- MLAG Spine & Leaf Topology, reduces dependency on STP and increases link utilization
- Standardized template for MLAG deployment making for easy deployment and troubleshooting, key for modular repeatability in an enterprise scale deployment
- L2 Rapid PVST or L3 OSPF/BGP network boundary on the cluster spine switches connecting to the upstream aggregation/core network

Cluster Configurator (CC) Benefit Summary

Greater Availability: Enabling MLAG technology increases network availability and redundancy. Reduces network downtime and enables undisrupted business operations.

Better Performance: The ability to span access and aggregation network tiers and interconnect n-number of switches in deployments with the ability to do easy sparing of leaf switches.

Scalability and Flexibility: Pay-as-you-grow scalability on fixed configuration switches -- from 1GbE to 100GbE allows flexible growth as network requirements follow the ASICs performance curve.

Dual Control Plane implementation: Leverage PICA8's CrossFlow functionality that enables OpenFlow/OVS to exert control on active L2/L3 ports without impacting network traffic.

Large Product Portfolio: Choice of different deployment configurations, offering different switch platforms with varying bandwidth and port density options. No vendor lock-in, no vendor-imposed scale limitations.

Orchestration and Management: Auto detection and provisioning for new switches allowing configuration updates across the stack through a single operation and centralized CLI using a single IP address.

Resiliency and Redundancy: Automatic switch failover in a fully resilient deployment. No dependency on any protocol requiring re-convergence. Uses SSH for secure connection between the nodes and LLDP for neighbor discovery.

Backup & Restore: Ability to perform up to 3 backups and restore in addition to factory default rollback behavior.

Log Aggregation: Provides aggregated Syslog and SNMP capability which can be also relayed to external servers.

Reduced OpEx/CapEx: Network OpEx and CapEx plummets, even compared to heavily discounted legacy alternatives.

Open Networking: Simple, Economical, Vendor agnostic, Flexible, Modular solution that leverages no proprietary protocols or applications and delivers a fully programmable and extensible networking platform.

CC for the - IT Network Director

- Performance
  - Broadcom ASIC – same as used by Top vendors
- Reliability
  - Broadcom ASIC – same for every switch branded or White Box
- No vendor Lock-in
  - No proprietary protocols
  - Trusted and industry standardized protocols and applications
- Equipment Cost savings
- Support Cost savings
  - Software Upgrades
  - Hardware Upgrades

CC for the - IT Network Operations Manager

- Installation
  - Get Switch with CC image from System Integrator
  - Cable up the switches
  - Power up the switches
  - You have an operational switch cluster
- Integrate with different existing deployed network topologies
- Scale Cluster
  - Cable up Leaf Switch with CC image
  - Single command Copy existing Leaf configuration to New Leaf
- Replacement
  - Remove Leaf
  - Insert Leaf Switch with CC image
  - Single command Restore previous Leaf configuration to New Leaf
- Leverage PICOS's feature rich protocol support and multi-device support to streamline connectivity
- Leverage PICOS CrossFlow architecture to deploy solutions for previously unsolvable requirements
- Integrate new ASIC capabilities quicker, as soon as they are available from the ODMs
- Support
  - Direct support from PICA8 or from System Integrator with SLAs

CC for the - IT Network Administrator

- Management/Orchestration Application on a switch cluster running PICOS
- CC application runs on the Spine switches extending the XORP CLI
  - Mgmt IP addresses are assigned based on LLDP port information
- Design: MLAG Spine & Leaf Topology
  - CC auto creates 2 high speed port LAG for connectivity between spines
  - CC dynamically creates 1-16 static MLAG ports for connections to Leafs
  - Automatic configuration of IP address and Hostnames on Leafs are based on the Port numbers connected on Spines
  - Switches communicate via SSH using cluster unique shared keys
- Services enabled on Spines for aggregation north-bound
  - Syslog
  - SNMP
  - NTP
- Upgrades by copying new package files to the leafs and installing them.
- Direct apt-get support through proxy on Spines.
- Enjoy the power and flexibility of Open Networking – Deploy your own applications on PICOS

CC GUI

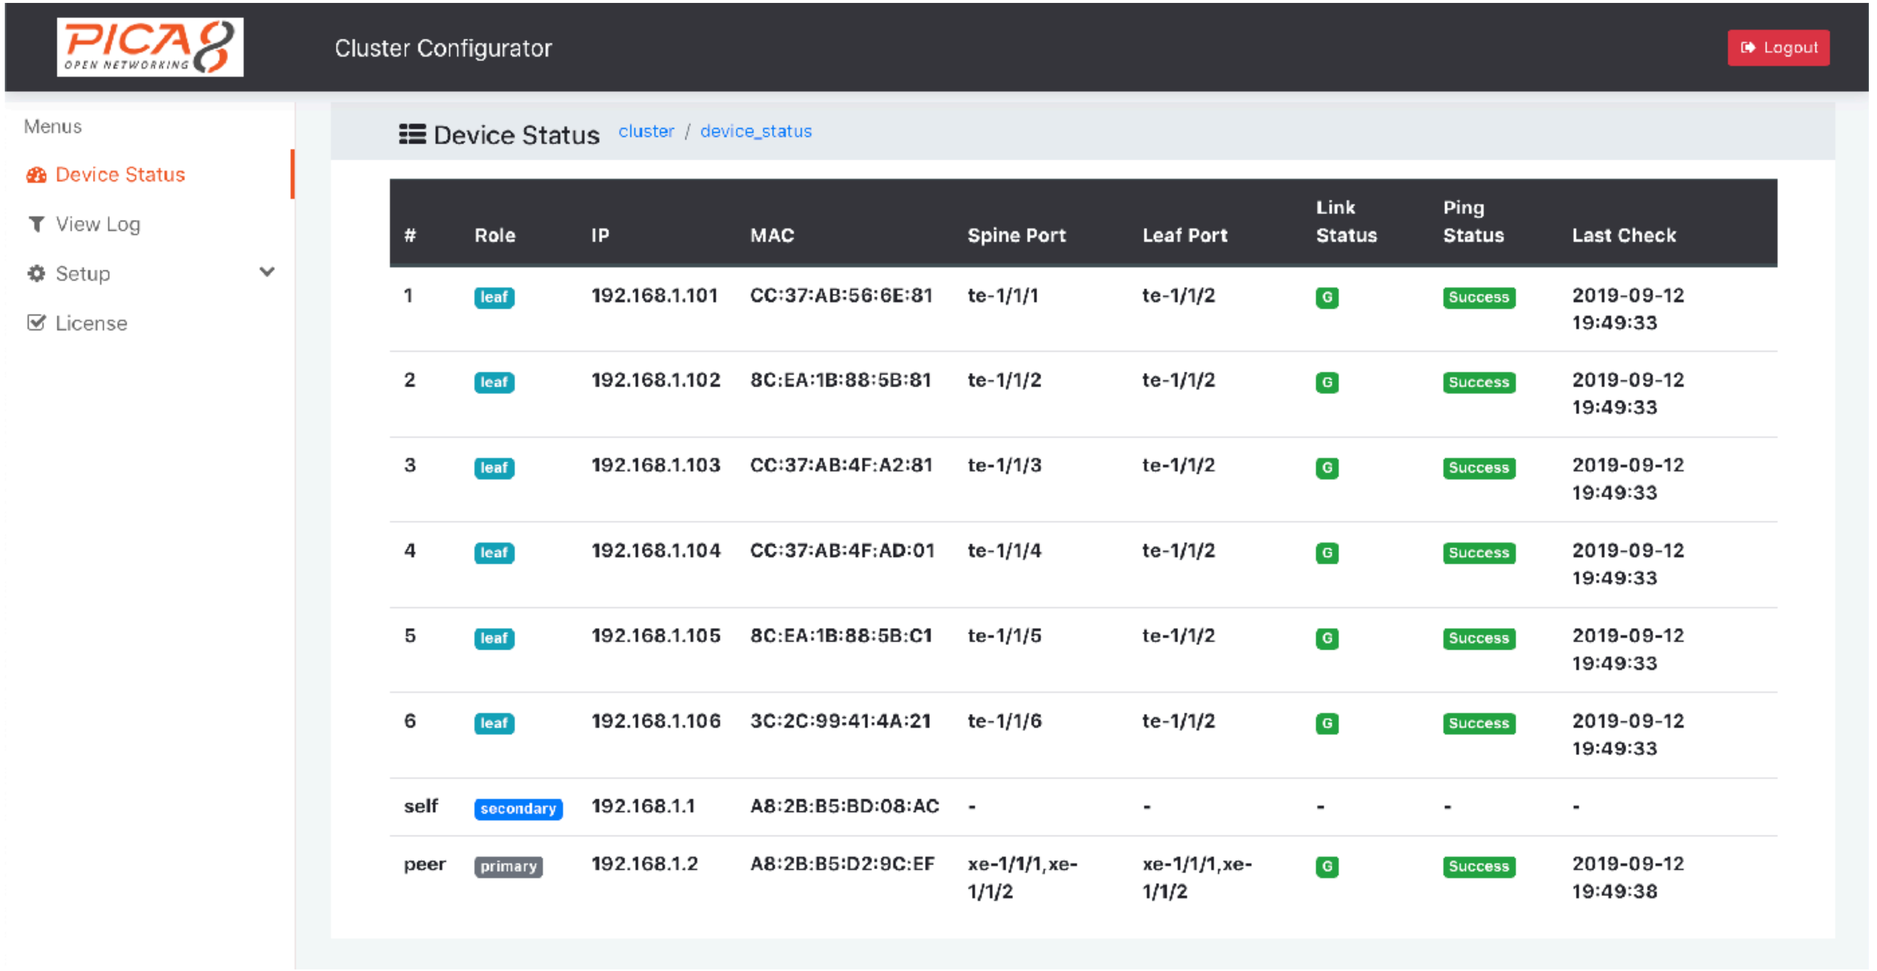

#	Role	IP	MAC	Spine Port	Leaf Port	Link Status	Ping Status	Last Check
1	leaf	192.168.1.101	CC:37:AB:56:6E:81	te-1/1/1	te-1/1/2	G	Success	2019-09-12 19:49:33
2	leaf	192.168.1.102	8C:EA:1B:88:5B:81	te-1/1/2	te-1/1/2	G	Success	2019-09-12 19:49:33
3	leaf	192.168.1.103	CC:37:AB:4F:A2:81	te-1/1/3	te-1/1/2	G	Success	2019-09-12 19:49:33
4	leaf	192.168.1.104	CC:37:AB:4F:AD:01	te-1/1/4	te-1/1/2	G	Success	2019-09-12 19:49:33
5	leaf	192.168.1.105	8C:EA:1B:88:5B:C1	te-1/1/5	te-1/1/2	G	Success	2019-09-12 19:49:33
6	leaf	192.168.1.106	3C:2C:99:41:4A:21	te-1/1/6	te-1/1/2	G	Success	2019-09-12 19:49:33
self	secondary	192.168.1.1	A8:2B:B5:BD:08:AC	-	-	-	-	-
peer	primary	192.168.1.2	A8:2B:B5:D2:9C:EF	xe-1/1/1,xe-1/1/2	xe-1/1/1,xe-1/1/2	G	Success	2019-09-12 19:49:38

Status Screen

Provides a quick overview of the switches in the cluster. Depending on where the VRRP VIP is pinned the "self" switch could be the "primary" or "secondary" switch of the cluster. The IP address are from the internal CC network and not the IPs of the ethernet management ports. The Spine and Leaf ports are from the perspective of the "self" switch. Link state is (Y)ellow if a LAG link is down or (R)ed if both links are down. The last check is timestamp for the last connectivity check run by the CC application.

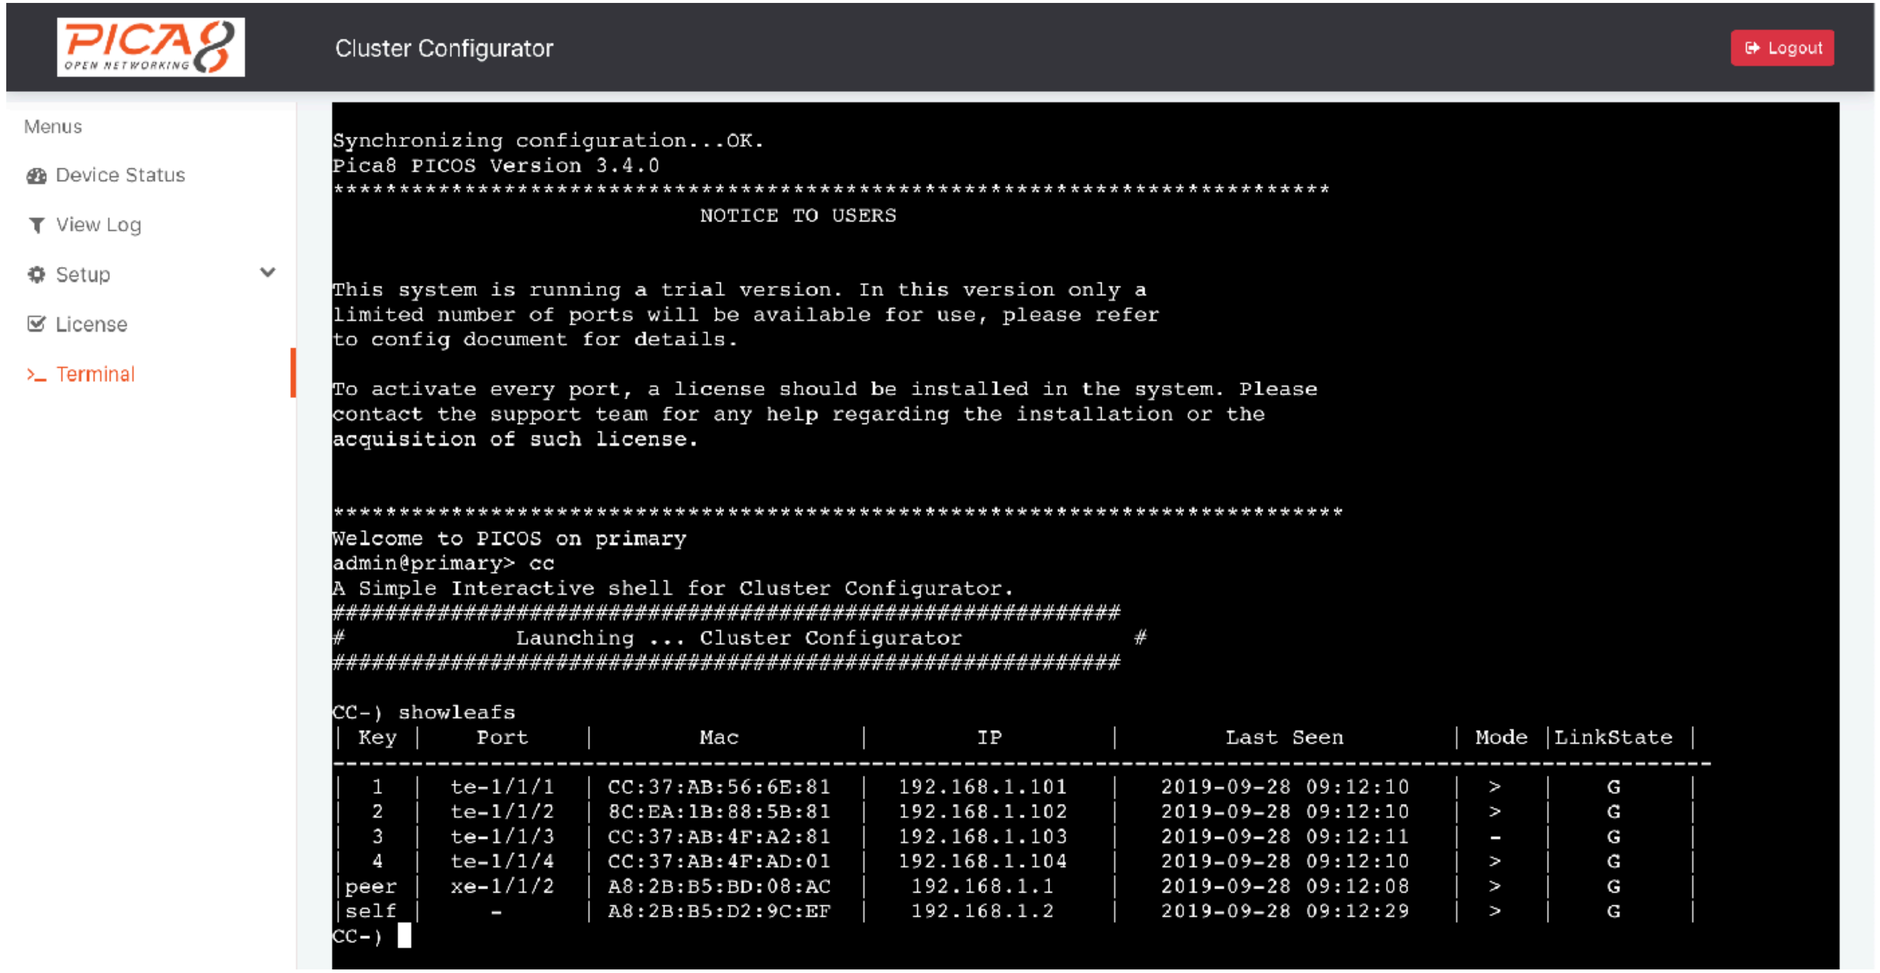

The GUI Integrated CLI Shell

A single CLI shell for the entire switch cluster. Auto logs in to the switch using credentials entered during login. After launching the CC application you can configure and manage any switch in the cluster. The shell supports copy and paste and persistently maintains its connection state even if you switch away to another screen.

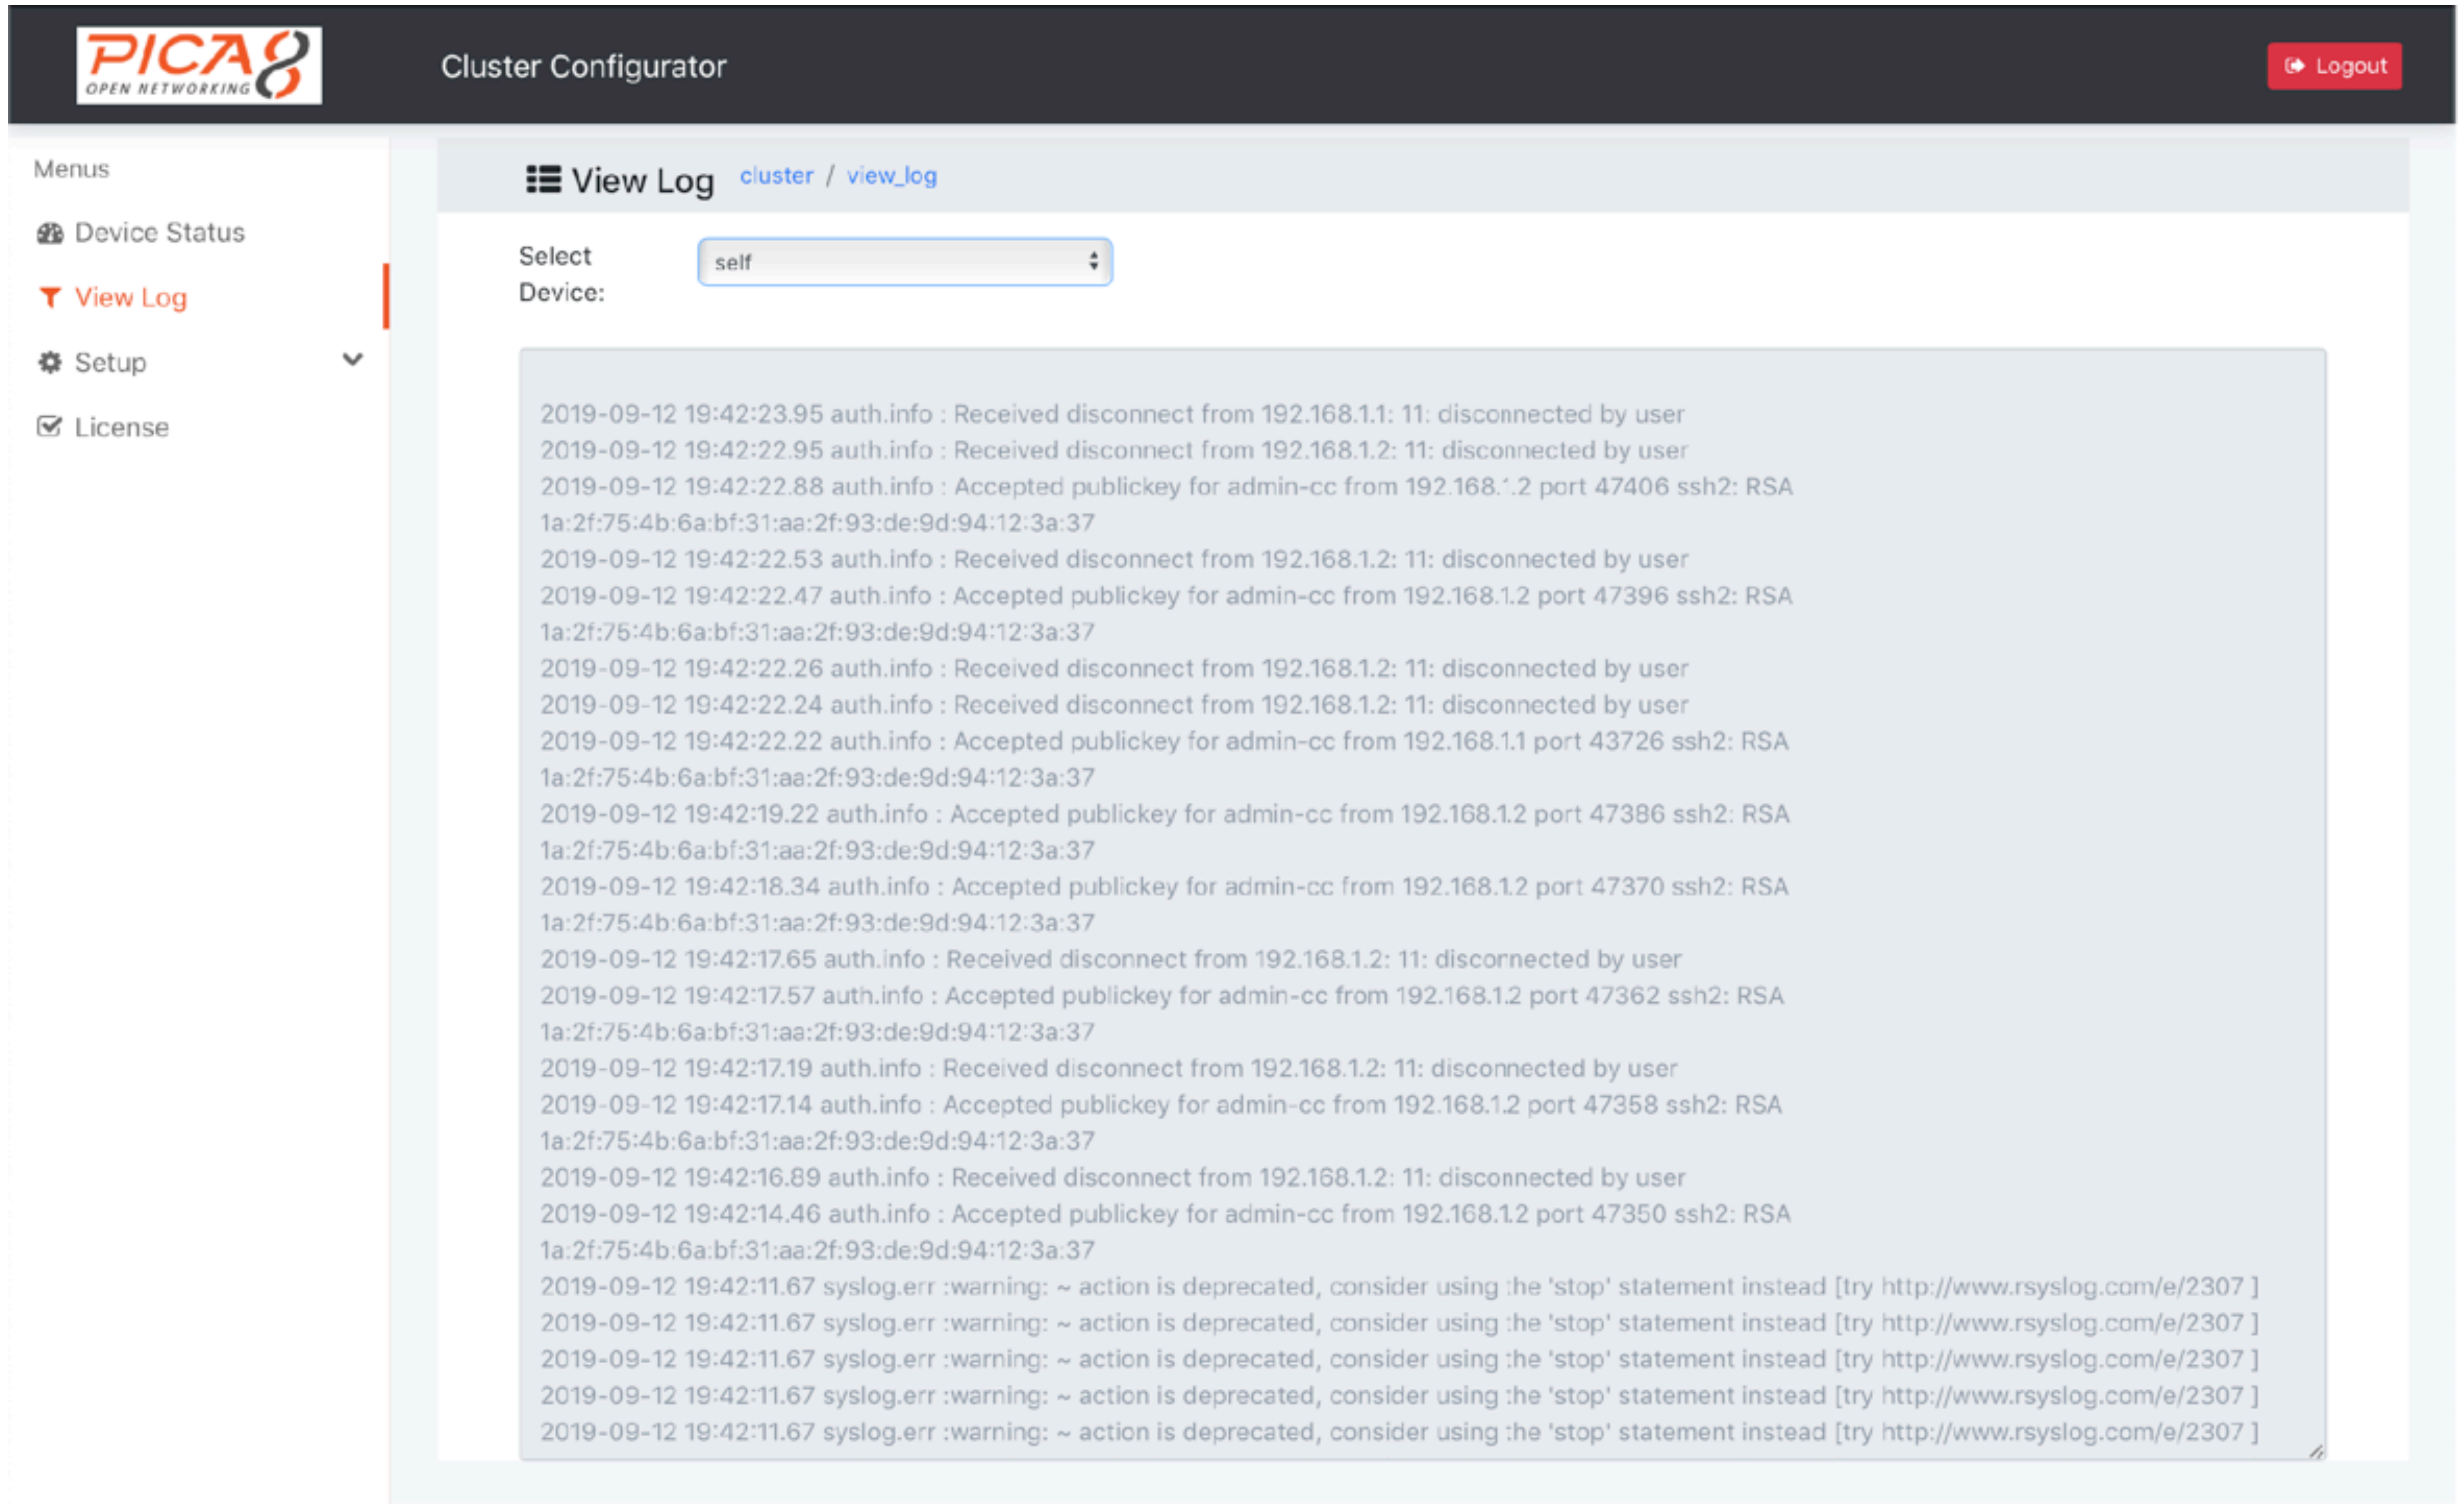

Switch Logs

Check syslog messages from any switch of the cluster. Messages are time descending order sorted. These syslog messages are also relayed to the user configured remote syslog server.

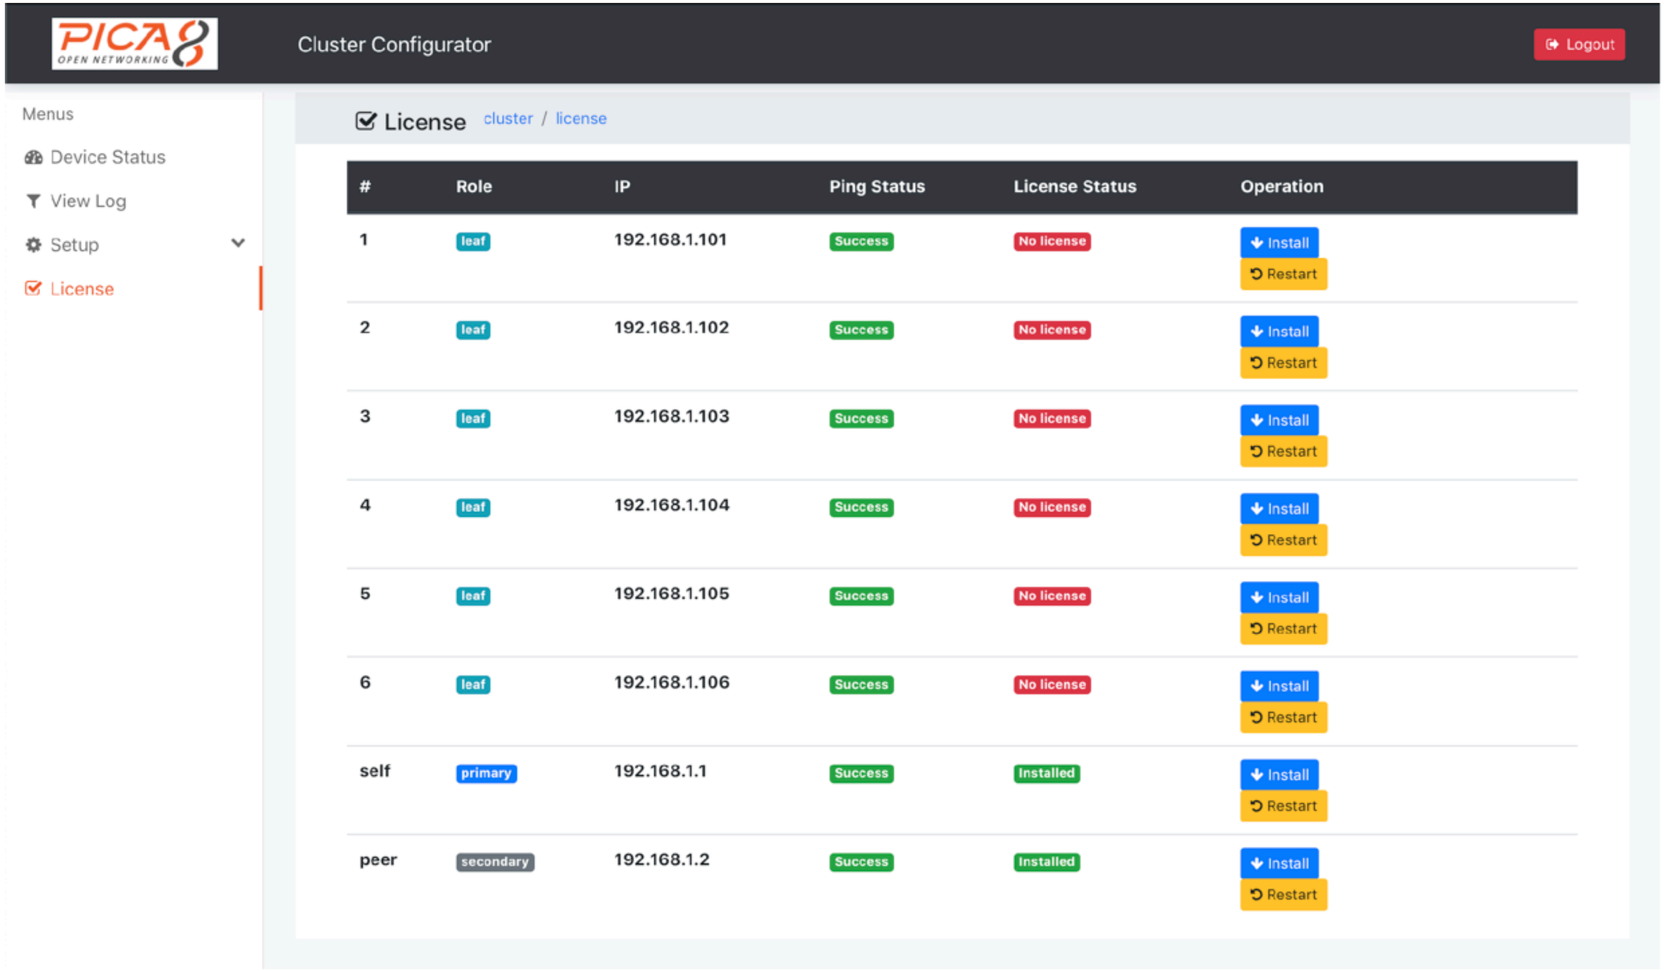

Configure Licenses

A single screen to check and manage licenses for all the switches in the cluster. After installing the license you need to restart the PICOS service on the switch. The leafs support simultaneous restart, for the spines you are limited to restarting the PICOS service on one spine at a time to maintain service continuity for the cluster. The Primary and Secondary roles on the spines may switch depending on how you restart the service on the spines.

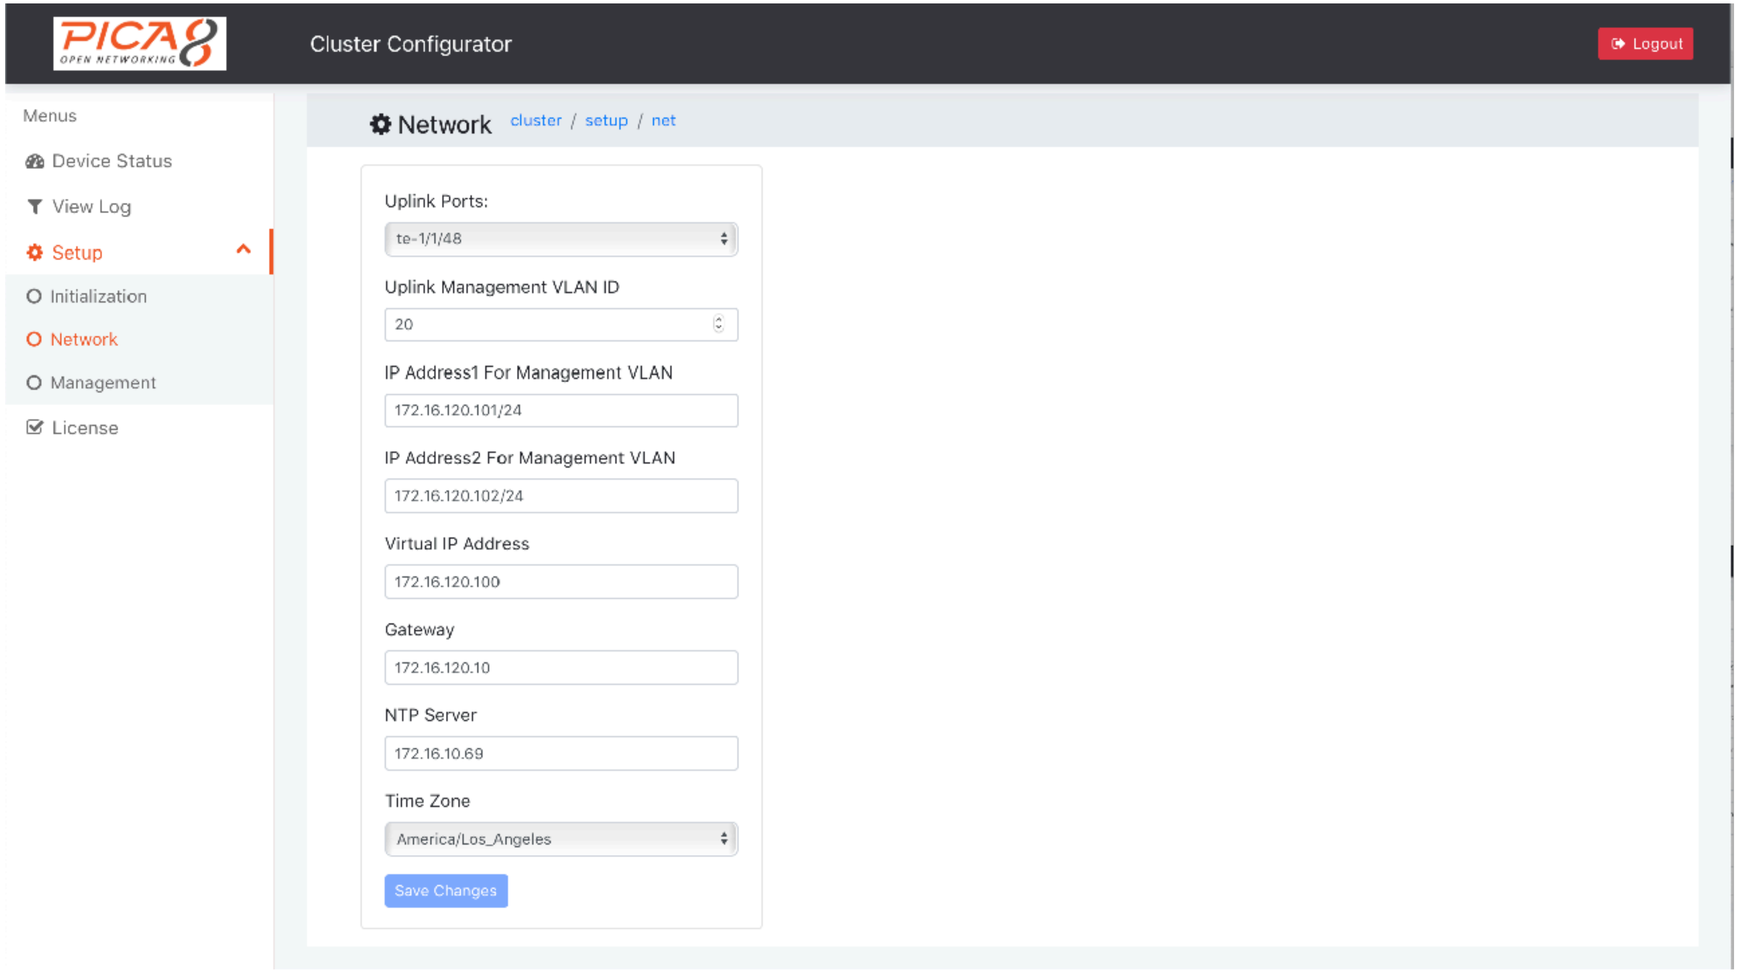

Network Connectivity settings

Configure the most common networking configurations required on the spines to connect to the upstream network of a switch cluster. The NTP and timezone settings are automatically synced down to all the leafs.

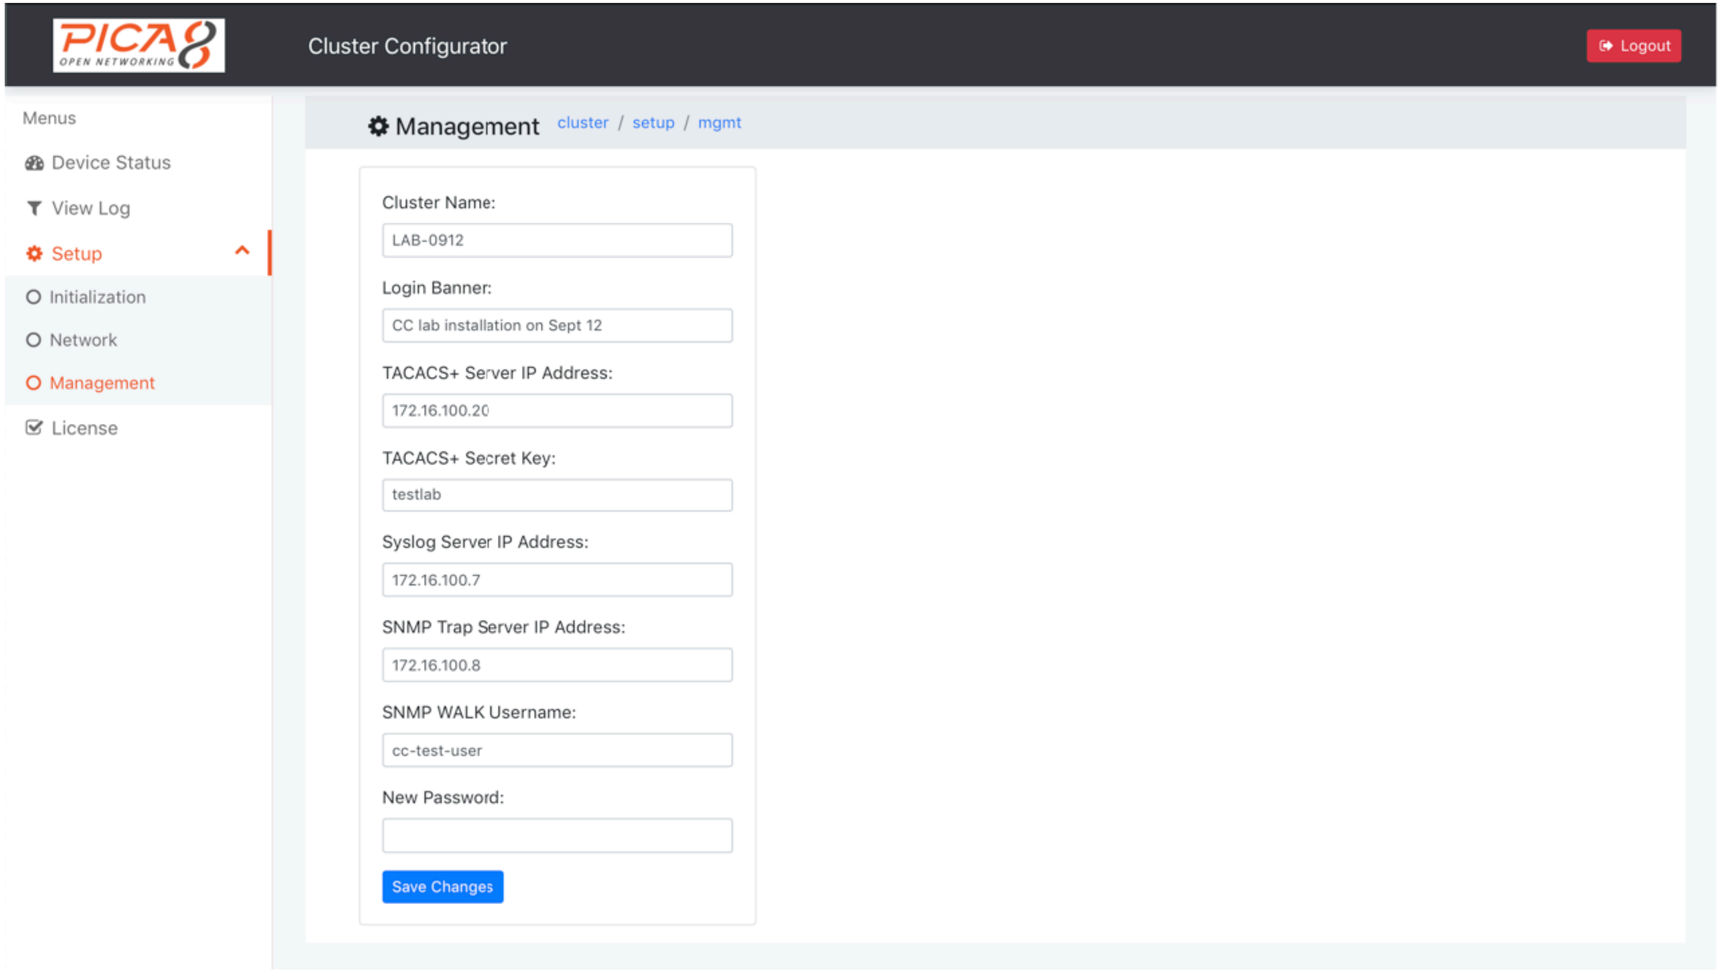

Network Management settings

Configure the most common network management configurations required to operationalize a switch cluster. The Cluster Name, Login Banner are automatically synced down to all the leafs. TACACS is used for AAA control on the spine switches; in a CC deployment the leafs do not have any external network management connectivity, everything is configured via the spines. The Syslog/SNMP server is the target for all Syslog/SNMP messages from all the switches in the cluster relayed via the spine rsyslogd/snmp-relay service. The SNMP Walk user name is the password for walking the spine or leaf switch MIBs remotely. The New Password if configured changes the default admin password for all switches in the cluster.

```
CC-) showleafs
| Key |    Port    |        Mac         |        IP        |        Last Seen        | Mode |LinkState |
-------------------------------------------------------------------------------------------------------------
|  1  |  te-1/1/1  | CC:37:AB:56:6E:81  |  192.168.1.101   |  2019-09-12 19:53:37    |  >   |    G     |
|  2  |  te-1/1/2  | 8C:EA:1B:88:5B:81  |  192.168.1.102   |  2019-09-12 19:53:37    |  >   |    G     |
|  3  |  te-1/1/3  | CC:37:AB:4F:A2:81  |  192.168.1.103   |  2019-09-12 19:53:37    |  >   |    G     |
|  4  |  te-1/1/4  | CC:37:AB:4F:AD:01  |  192.168.1.104   |  2019-09-12 19:53:37    |  >   |    G     |
|  5  |  te-1/1/5  | 8C:EA:1B:88:5B:C1  |  192.168.1.105   |  2019-09-12 19:53:37    |  >   |    G     |
|  6  |  te-1/1/6  | 3C:2C:99:41:4A:21  |  192.168.1.106   |  2019-09-12 19:53:37    |  >   |    G     |
|peer |  xe-1/1/2  | A8:2B:B5:BD:08:AC  |   192.168.1.1    |  2019-09-12 19:53:44    |  >   |    G     |
|self |     -      | A8:2B:B5:D2:9C:EF  |   192.168.1.2    |  2019-09-12 19:53:52    |  >   |    G     |
CC-)
.last_tb            file_del            setup
_                   file_get            showleafs
checktopo           file_list           showspines
config_copy         file_push           ssh
config_purge        help                switch_erase
config_restore      manage_license      switch_reinit
config_restore_ovs  push_autorun        switch_rekey
config_save         rem                 viewlog
config_save_ovs     runscript
exit                send
CC-)
```

The CC CLI

You can SSH to the VRRP VIP or the individual IP address of the spines switches and after logging you can launch the "cc" application to gain configuration and management access to all the switches in the cluster. Multiple CC commands have been made available to help you configure and manage the cluster. "showleafs" output includes the "Mode" which shows the current configuration prompt on the switches.

CC - Commands 1/2

Status	Initialization	Configuration

• showleafs
• showspines
• checktopo
This command checks the spine and leaf cabling connections and alerts you if it detects an inconsistency.
• viewlog
	• setup
	  • init
	  • net
	  • mgmt
	• manage_license
		• send
		Relays config commands to a set or all switches
		Supports %n-m% numerical expansion to minimize configuration commands; only a single expansion is supported per config line
		• runscript
		Pushes a Linux Shell script to the switch and runs it on the switch
		Generally used for Linux and OVS/OpenFlow configurations.
		• ssh
		• _
		Used for local switch CLI command execution
		• rem
		config remark / comment

PICOS documentation – https://docs.pica8.com

CC - Commands 2/2

Operational

- switch_reinit

  Restart PICOS service on the switch
- switch_erase

  Erase all configuration on the switch
- switch_rekey

  Reinitialize the SSH key used between the spines and leafs

File Management

- file_get

  get a file onto the spine switch from an upstream location
- file_list
- file_del
- file_push

  push a file to the leaf switch from the spine
- push_autorun

  push a file to the switch and install it as the PICOS pre-start or post-start execute autorun script for customization of switch boot process

Config Management

- config_save

  Spines/Leaf replacements are automatically restored to previous config - if a previous config is not there you end up with the starter config - so remember to config_save for spines after major changes
- config_restore
- config_save_ovs
- config_restore_ovs
- config_purge

  Deletes all configuration backups for a switch slot. This followed up with a switch_erase and switch_reinit will completely reset a switch
- config_copy

  copy configuration from one switch to another

  be careful to not use copy once Switch-IP specific configurations like MLAG are configured or if switch software/hardware configs are different in any other way

CC Topologies

Showcasing 6 different cases

Upstream Topology Interoperability 1/6

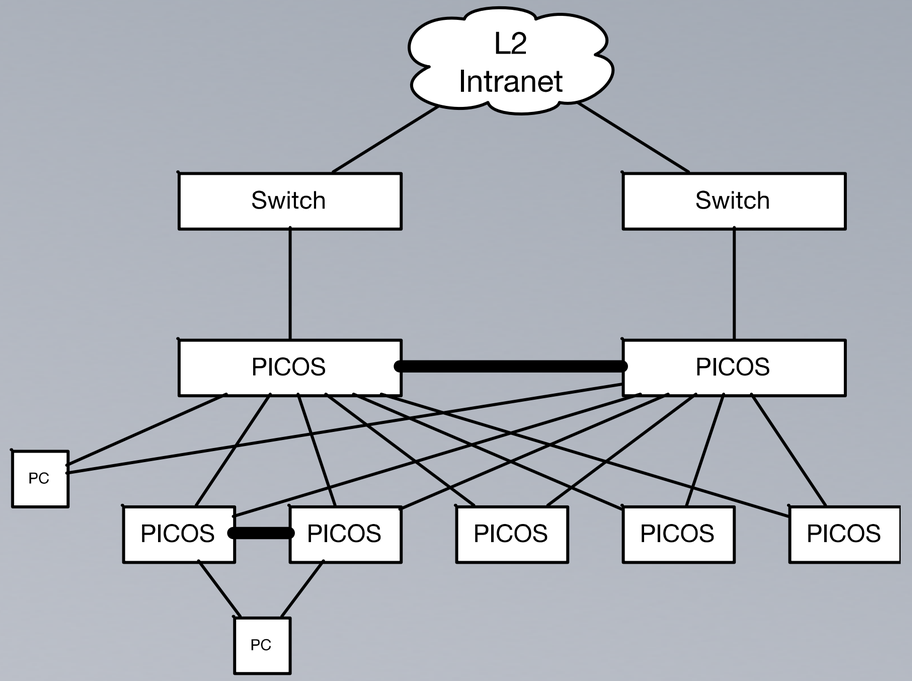

External only VLANS

VLAN 10 - mgmt

External and Cluster VLANS

VLAN 11 - data

VLAN 12 - voice

Upstream: multiple topologies, some create topology loops and let STP protocols handle them.

redundant STP managed upstream network connection

L2 to L2 Topo #1

Upstream Topology Interoperability 2/6

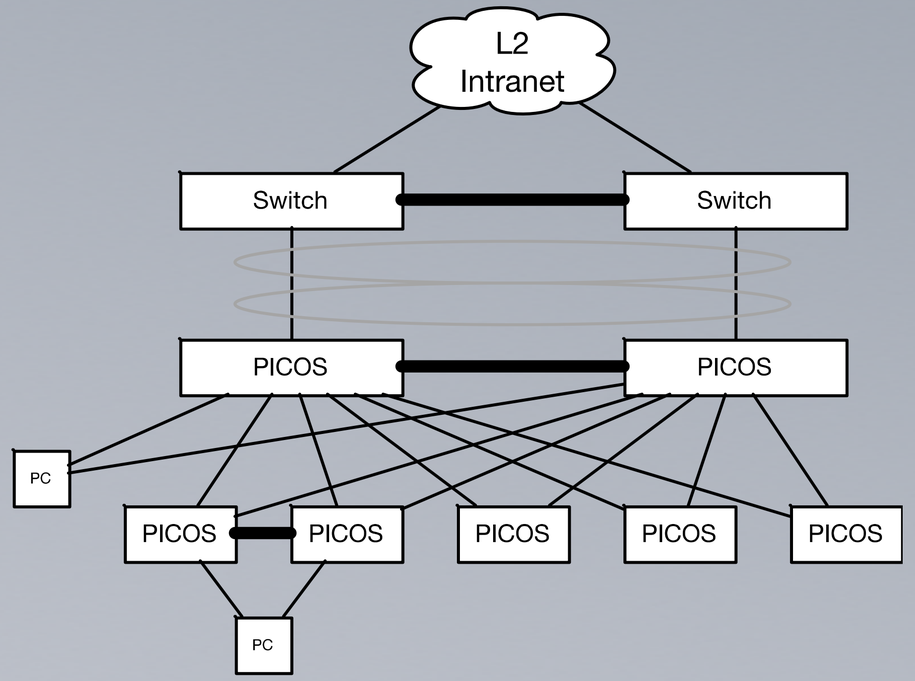

External only VLANS

VLAN 10 - mgmt

External and Cluster VLANS

VLAN 11 - data

VLAN 12 - voice

Upstream: connection to upstream switches is using a LAG that eliminates the requirement for STP and utilizes all uplinks distributing traffic flows across them.

LAG connectivity to upstream network

L2 to L2 Topo #2

Upstream Topology Interoperability 3/6

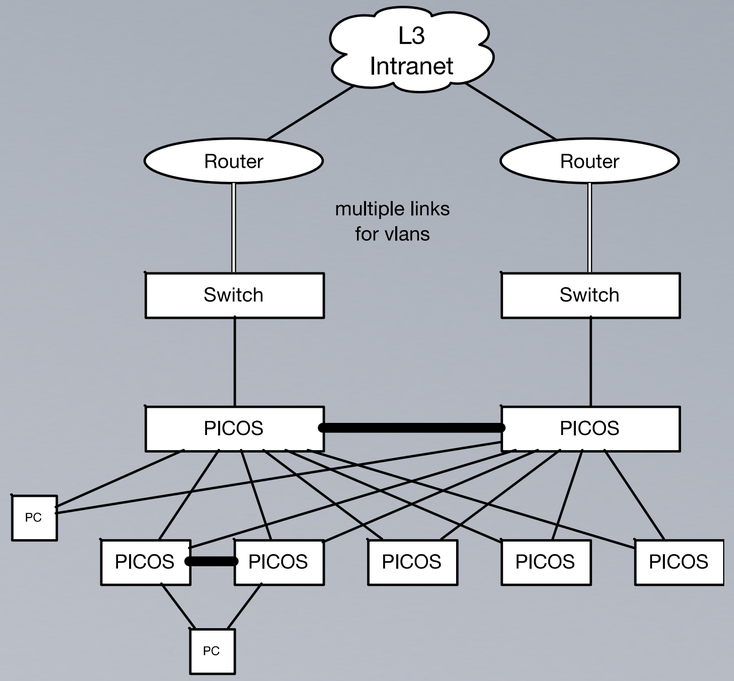

External only VLANS

VLAN 20 - mgmt

External and Cluster VLANS

VLAN 21 - data

VLAN 22 - voice

Upstream: connection to upstream network is over trunks connected to aggregation switches which connect to core routers. There are no L2 loops in the topology and the routers run their VRRP/HSRP sessions for VLAN 20 over the LAG link between the spines.

upstream network via switch trunks

L2 to L3 Topo #1

Upstream Topology Interoperability 4/6

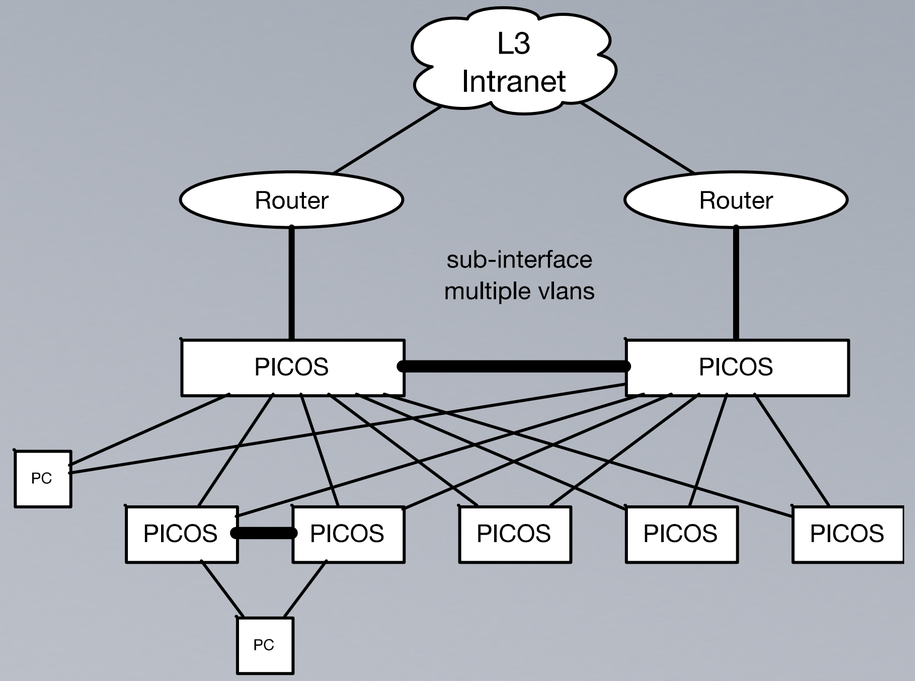

External only VLANS

VLAN 30 - mgmt

External and Cluster VLANS

VLAN 31 - data

VLAN 32 - voice

Upstream: connection to upstream network is over trunks connected to sub-interfaces on the routers. There are no L2 loops in the topology and the routers run their VRRP/HSRP sessions for VLAN 30 over the LAG link between the spines.

upstream network via router sub-interfaces

L2 to L3 Topo #2

Upstream Topology Interoperability 5/6

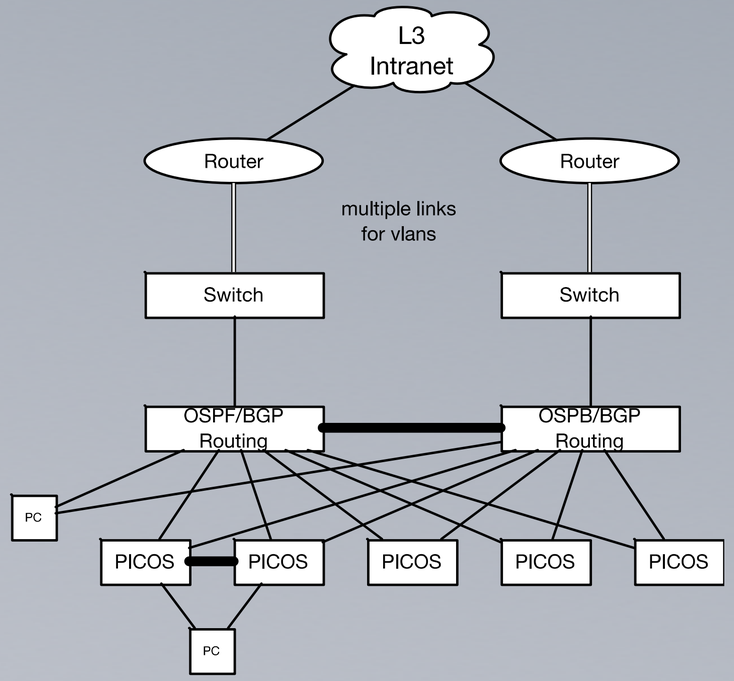

External only VLANS

- VLAN 20 - mgmt
- VLAN 21 - data
- VLAN 22 - voice

Cluster only VLANS with VRRP VIPs and DHCP relay to upstream DHCP server

- VLANS 2101, 2102 - data
- VLAN 2201, 2202 - voice

OSPF routing on spines

Upstream: connection to upstream network is over trunks connected to aggregation switches which connect to core routers. There are no L2 loops in the topology and the routers run their VRRP/HSRP sessions for VLAN 20, 21 and 22 over the LAG link between the spines.

The upstream routers learn of the 2101, 2102, 2201, 2202 networks via OSPF updates.

routed network via switch trunks

L3 to L3 Topo #1

Upstream Topology Interoperability 6/6

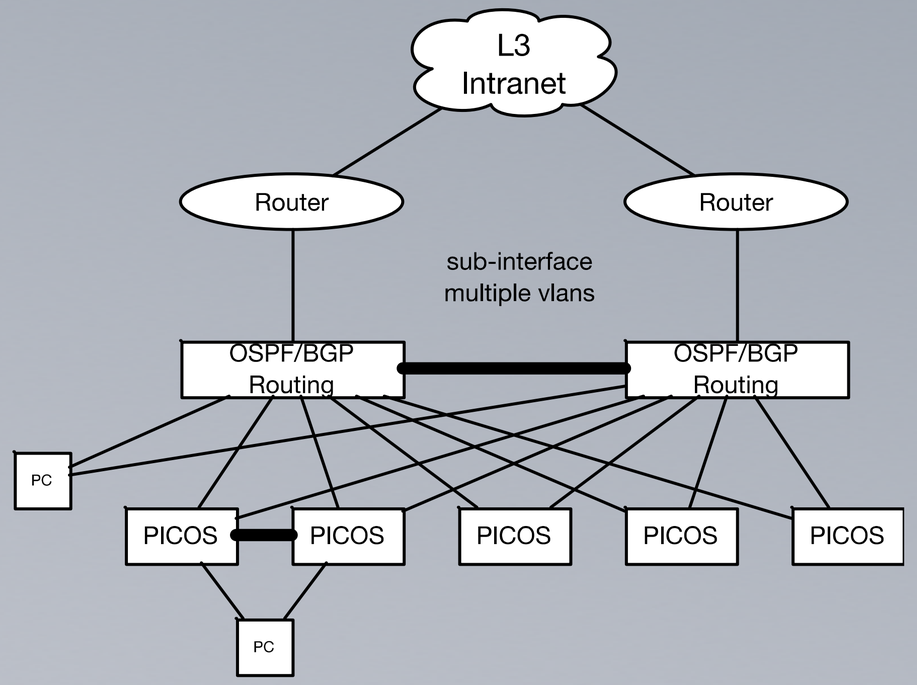

External only VLANS

- VLAN 30 - mgmt
- VLAN 31 - data
- VLAN 32 - voice

Cluster only VLANS with VRRP VIPs and DHCP relay to upstream DHCP server

- VLANS 2101, 2102 - data
- VLAN 2201, 2202 - voice

OSPF routing on spines

Upstream: connection to upstream network is over trunks connected to sub-interfaces on the routers. There are no L2 loops in the topology and the routers run their VRRP/HSRP sessions for VLAN 30, 31 and 32 over the LAG link between the spines.

The upstream routers learn of the 2101, 2102, 2201, 2202 networks via OSPF updates.

routed network via sub-interfaces

L3 to L3 Topo #2

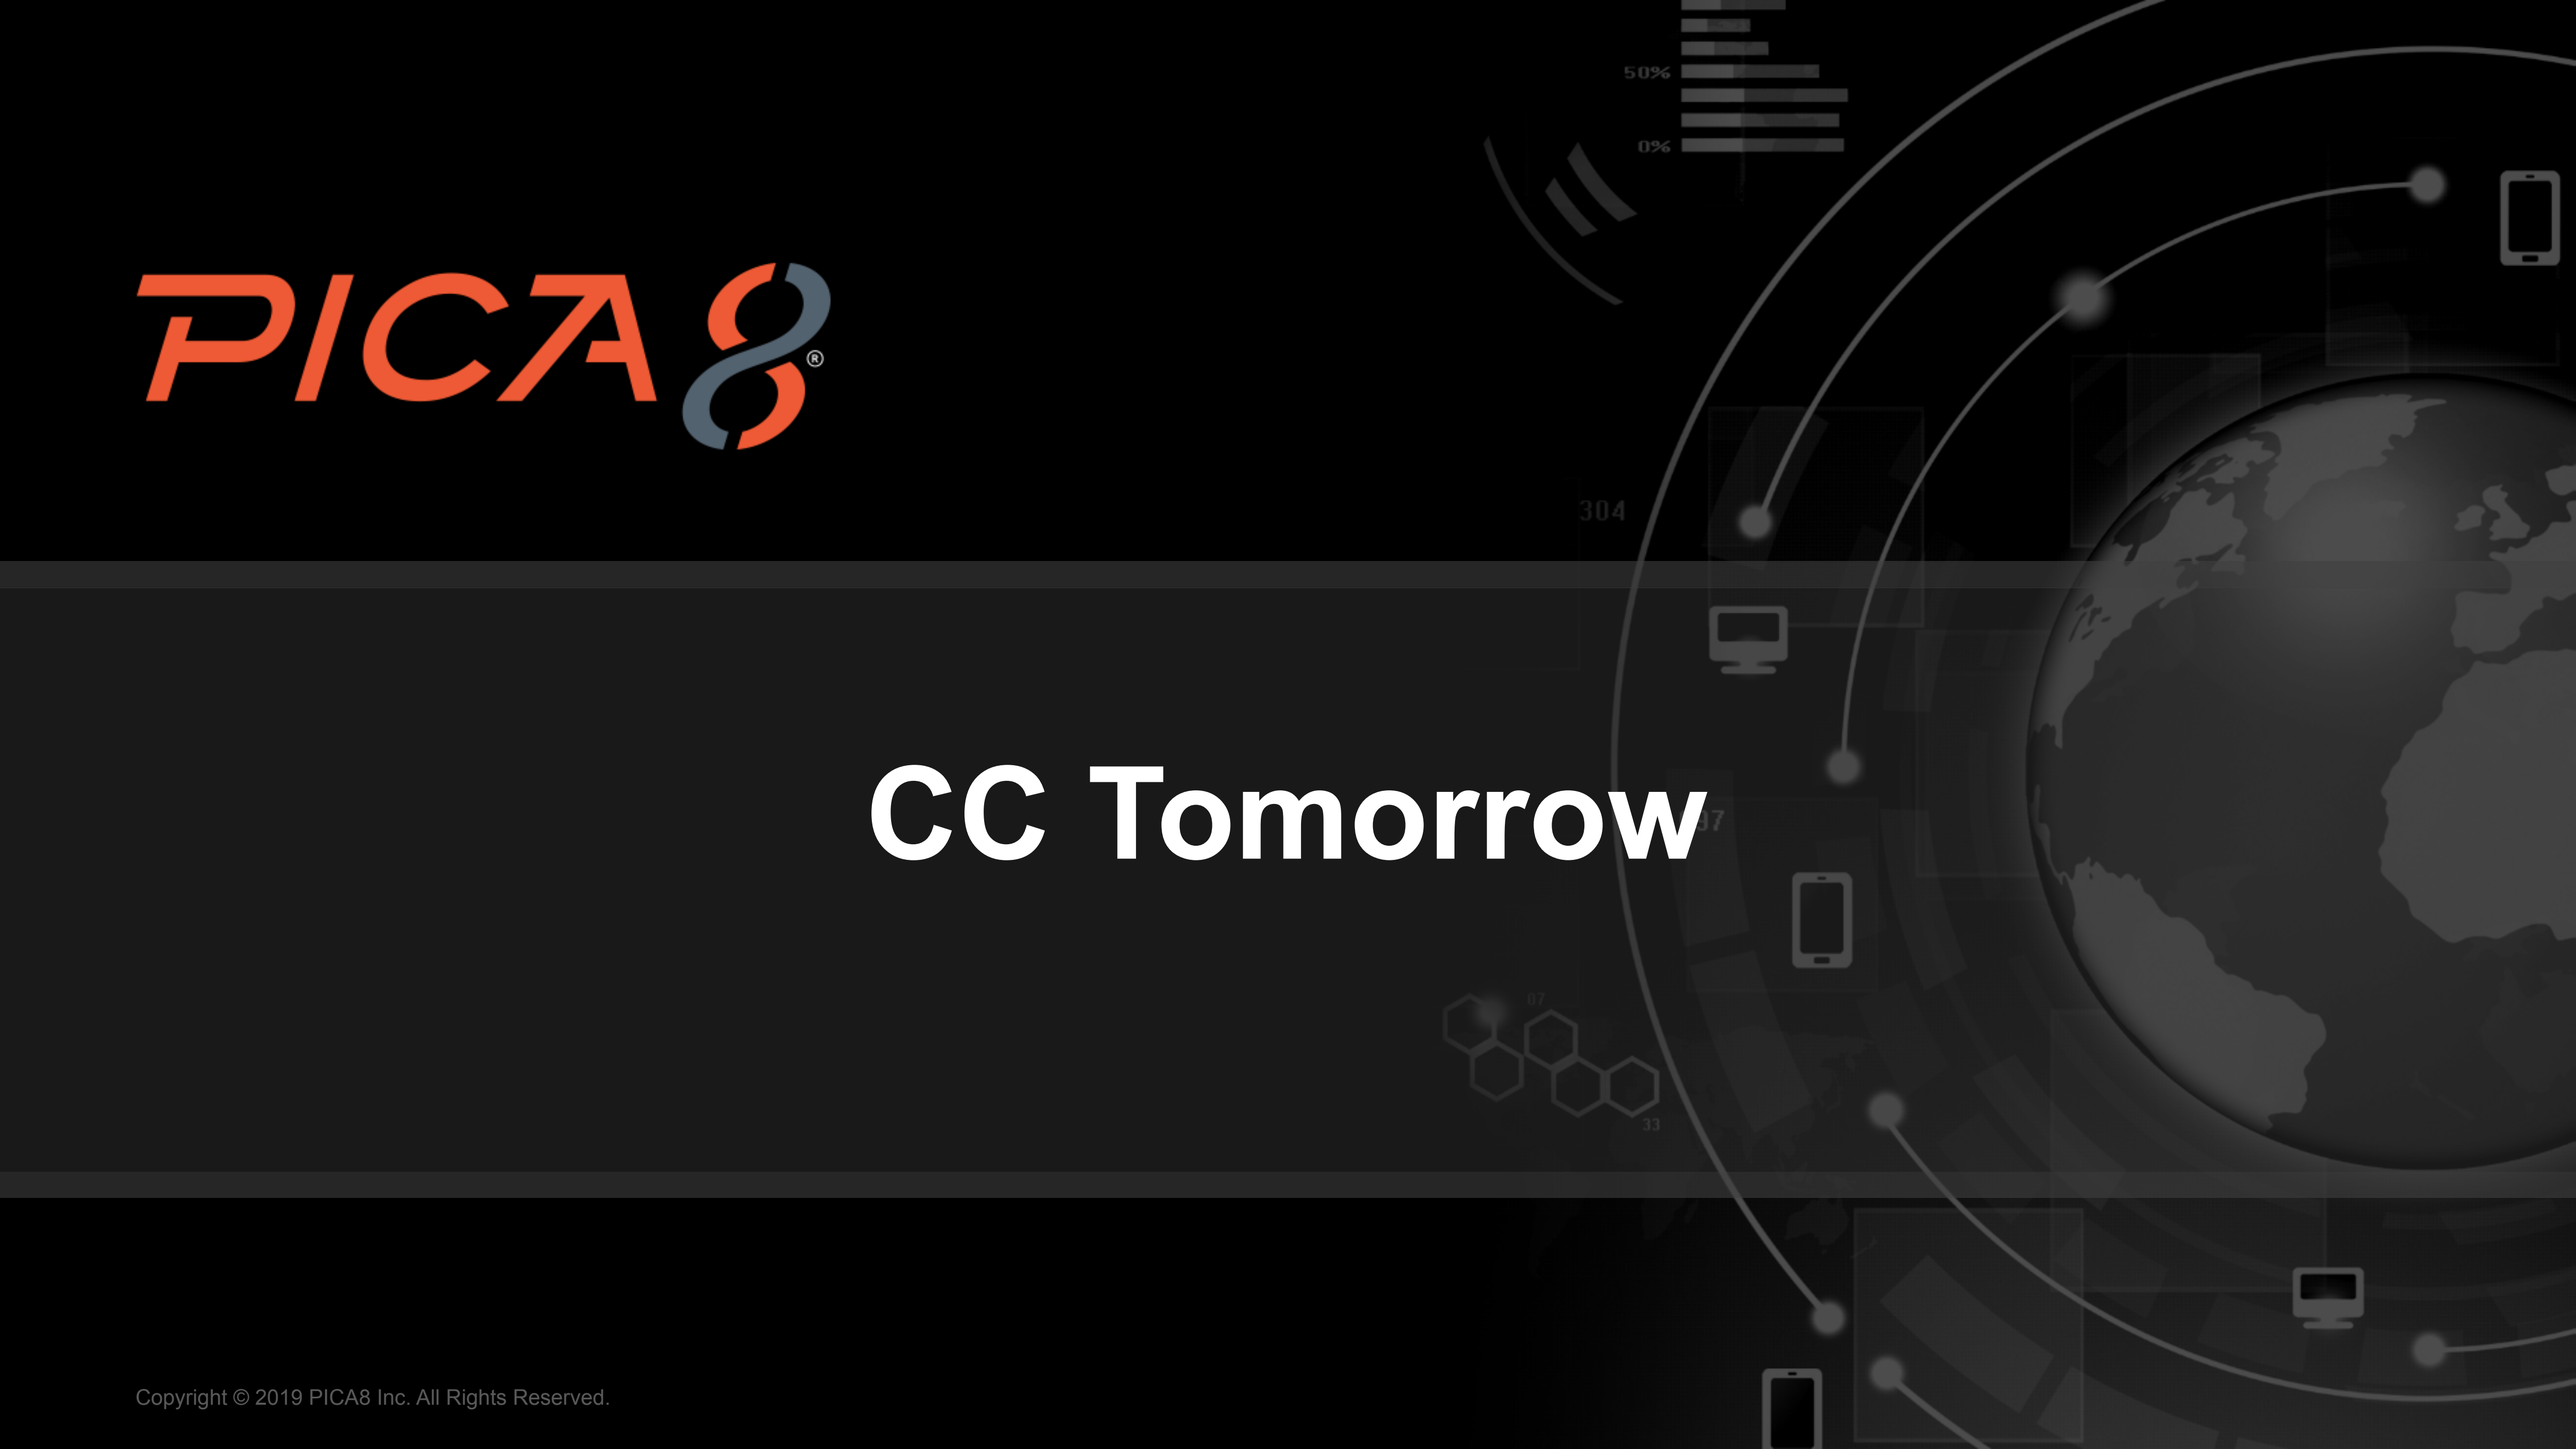

CC Tomorrow

CC & PICOS Feature Pipeline

- GraphQL API
- CC package upgrade
- CC snips
  - Abstraction of user deployment building blocks
- CrossFlow snips

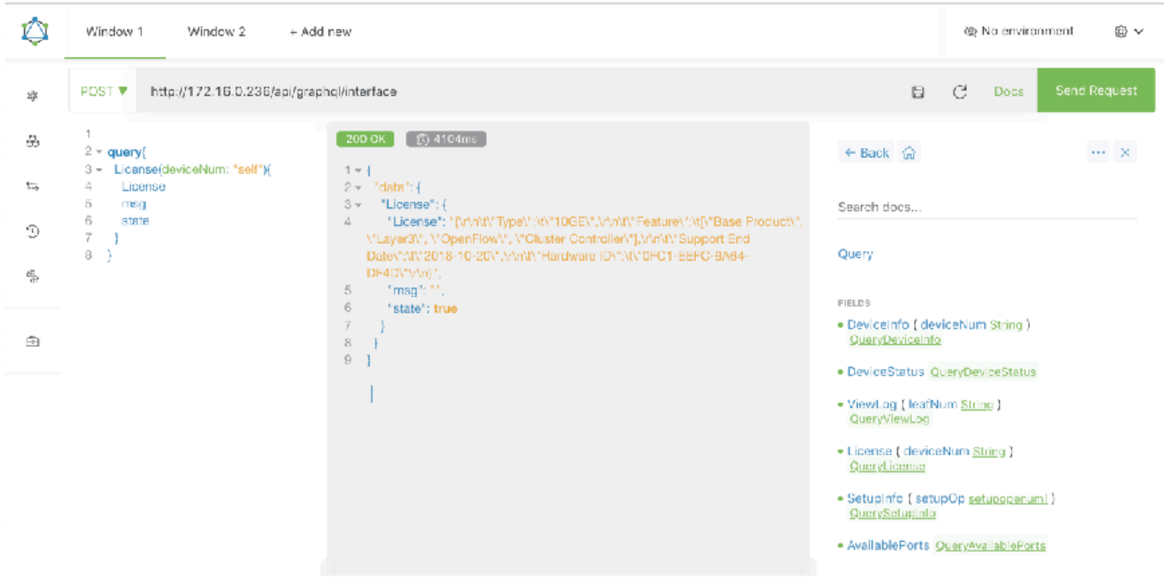

- VRF enhancements
  - VRF DHCP Relay Support
  - VRF SSH Support
- CC NAC integration
- Next-Gen MLAG / STP model
  - Eliminates all PICOS deployment topology restrictions

Questions

sharad.ahlawat@pica8.com