


What You Make Possible



Troubleshooting Nexus 7000

BRKRST-3066



Session Goal

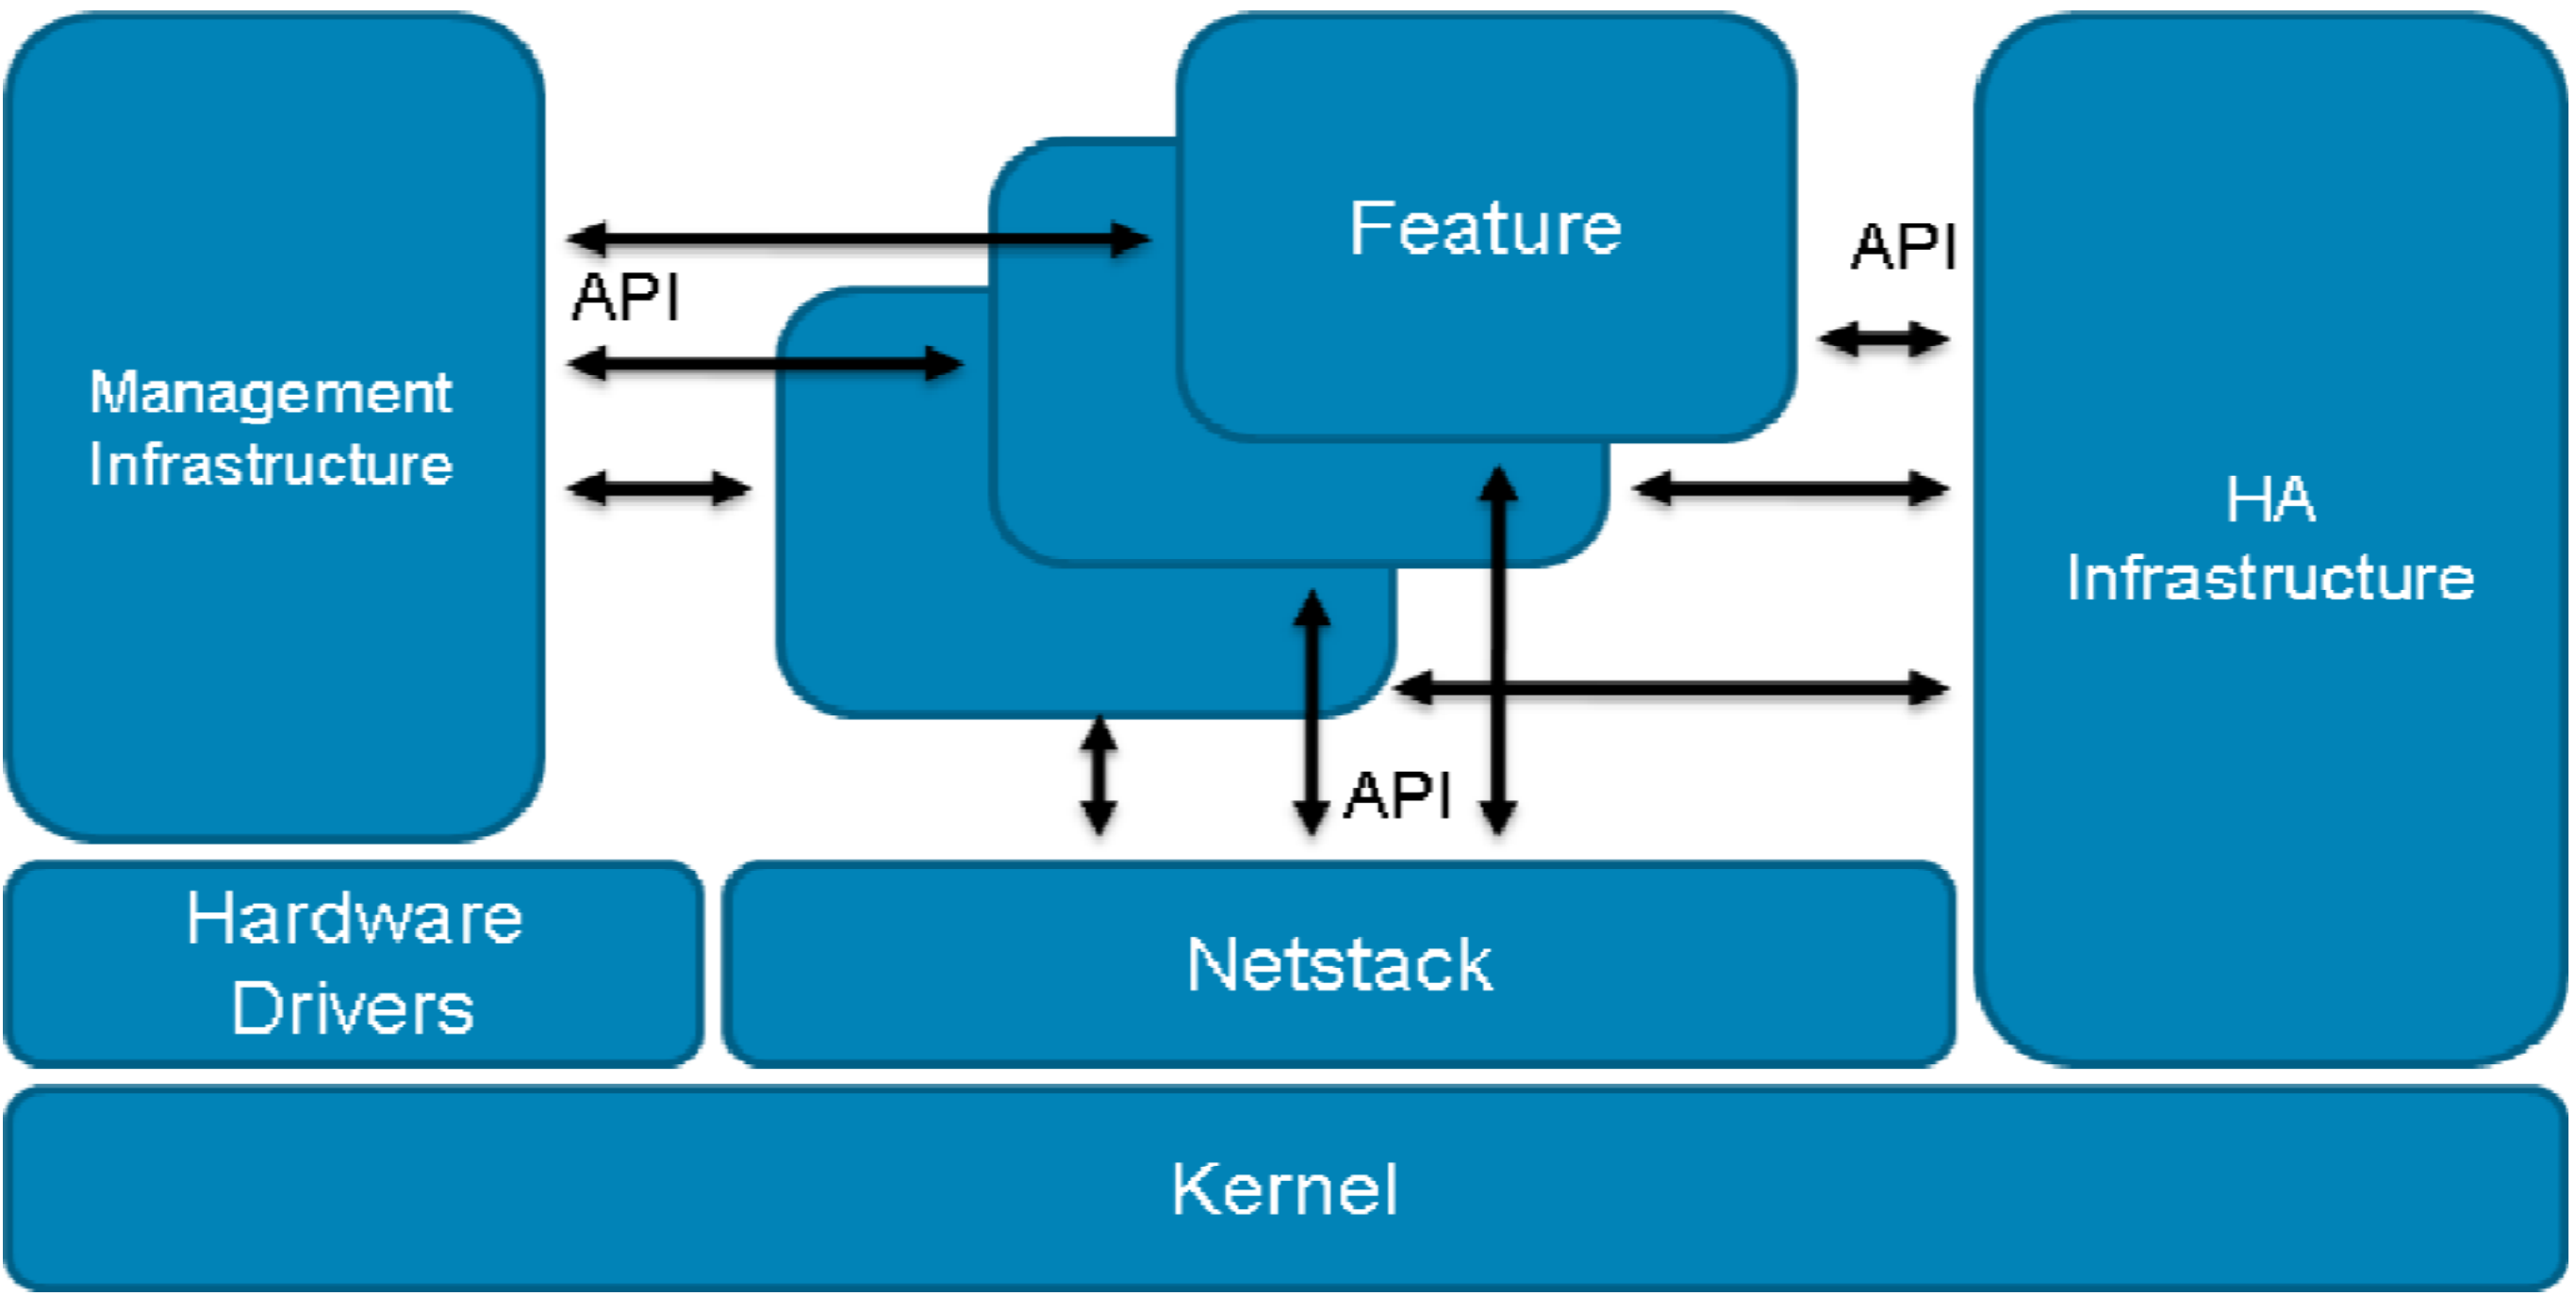
- To provide you with an understanding of the Cisco Nexus™ 7000 built-in troubleshooting tools and troubleshooting technique of NX-OS operating system
- Increase familiarity with Cisco Nexus™ 7000 logging and information capture
- Ensure you get the right information at the right time

Agenda

- **Introduce** NXOS software architecture and logging capability
- **Define** built-in troubleshooting tools that will assist with problem analysis and resolution
- **Demonstrate** troubleshooting tools used when investigating common network functions

NXOS Architecture Introduction

Traditional IOS Versus NX-OS – Software Modularity



NXOS Architecture Introduction

IOS vs NX-OS – Show Tech

- Logging output greatly increased; assists stateful information capture
- More data requires more effective filtering
- Capture feature related information for later analysis

```
Cat6K#show tech | redirect bootflash:6k.tech
```

```
Cat6K#dir bootflash:
```

```
Directory of bootflash:/
```

```
 9  -rwx      1194813  Feb 10 2012 06:42:30 +00:00  6k.tech
```

1.2MB file to review

```
N7K1# show tech >> bootflash:giant.tech
```

```
Show tech brief will take 4-6 minutes to complete. Please Wait ...
```

```
N7K1# dir bootflash:
```

```
120077767  Feb 10 16:18:40 2012  giant.tech
```

120MB file to review

Agenda

- NXOS Architecture Introduction
- Built-in Troubleshooting Tools
- Troubleshooting Nexus 7000
 - CPU
 - Control-Plane – CoPP
 - Hardware
 - vPC
 - Unicast Layer 2 and Layer 3 Forwarding and ARP
 - Multicast Layer 2 and Layer 3 Forwarding
 - QoS

Troubleshooting Tools

CLI Capture and Filter – eigrp, start-time, last

- Powerful CLI based on Linux Kernel
- Built-in CLI filter techniques increases speed and relevance
- grep, egrep, include, last, wc, start-time
- Use the filtering options in real-time to increase your effectiveness

```
N7K1-VPC# show logging logfile start-time 2012 Jan 11 16:00:00
2012 Jan 23 16:00:08 Campus_N7K1-VPC %ETHPORT-5-IF_DOWN_NONE: Interface port-channel40 is down (None)
```

```
N7K1-VPC# show logging logfile | egrep "MAC|VLAN|ETHPM"
2012 Jan 23 11:42:49 Campus_N7K1-VPC %ETHPORT-3-IF_ERROR_VLANS_SUSPENDED: VLANs 1,5-
9,11,100,200,203,1002-1005,1111 on Interface port-channel10 are being suspended.
```

```
N7K1-VPC# show ip route | egrep "local$"
*via 1.1.1.1, Lo0, [0/0], 04:30:29, local
```

Troubleshooting Tools

Debugging Log and Filter – debug logfile, debug-filter

- Debugging per feature logged to bootflash
- Set a filter file and log to a new file

```
N7K1-VPC# debug-filter ip ospf interface vlan 64
N7K1-VPC# debug logfile offending_traffic
N7K1-VPC# debug ip ospf packets
N7K1-VPC# undebug all
```

```
N7K1-VPC# show debug logfile offending_traffic
```

```
2012 Mar 26 23:33:25.992586 ospf: 6467 [3981] (default) rcvd: prty:7 ver:2 t:HELLO len:44 rid:0.0.0.0
  area:0.0.0.0 crc:0xfdd2 aut:0 aukid:0 from 192.253.64.254/Vlan64
2012 Mar 26 23:33:25.992780 ospf: 6467 [3981] Invalid src address 192.253.64.254, should not be seen
  on Vlan64
2012 Mar 26 23:33:25.992966 ospf: 6467 [3981] Invalid src address 192.253.64.254, should not be seen
  on Vlan64
```

Troubleshooting Tools

Logging and Feature History – show logging logfile, show <feature> internal

- Syslogs and feature event history per Virtual Device Context (VDC)
- Feature interaction tracked through 'event-history'
- Use syslogs with feature logging to compare feature behaviour

```
N7K1-VPC# show logging logfile
```

```
2012 Jan 23 17:20:12 Campus_N7K1-VPC %ETHPORT-3-IF_ERROR_VLANS_SUSPENDED: VLANs 50 on Interface port-channel20 are being suspended. (Reason: Vlan is not allowed on Peer-link)
```

```
N7K1-VPC# show system internal ethpm event-history interface port-channel 20
```

```
29) FSM:<port-channel20> Transition at 500252 usecs after Mon Jan 23 17:20:12 2012
```

```
Previous state: [ETH_PORT_FSM_ST_TRUNK_UP]
```

```
Triggered event: [ETH_PORT_FSM_EV_LOGICAL_CHG]
```

```
Next state: [ETH_PORT_FSM_ST_WAIT_LOGICAL_CHANGE_TRUNK]
```

Troubleshooting Tools

Logging and Feature History – Show Tech Collection

```
Campus_N7K1-VPC# show tech | grep "'show "  
'show license usage vdc-all LAN_ENTERPRISE_SERVICES_PKG'  
'show license usage vdc-all LAN_ADVANCED_SERVICES_PKG'  
'show license usage vdc-all SCALABLE_SERVICES_PKG'  
'show license usage vdc-all TRANSPORT_SERVICES_PKG'  
'show license usage vdc-all ENHANCED_LAYER2_PKG'  
'show system internal license event-history'  
'show system internal license mem-stats'  
'show system internal license mem-stats detail'  
'show system redundancy ha status'  
'show system redundancy status'  
'show system internal eltm errors`
```

```
N7K1-VPC# show tech-support stp
```

```
N7K1-VPC# show tech-support vtp
```

```
N7K1-VPC# show tech-support pixm
```

```
N7K1-VPC# show tech-support forwarding 12 unicast
```


Troubleshooting Tools

Data Plane Traffic Capture - SPAN

- Data Plane SPAN in hardware
- VACL, Local SPAN, ERSPAN
- ACL Log to flash

```
N7K1-VPC (config) # monitor session 1 type ?
  acl-capture          Create an acl-capture session
  erspan-destination  Create an erspan destination session
  erspan-source        Create an erspan source session
  local                Create a local session

N7K1-VPC (config) # monitor session 1 type local
N7K1-VPC (config-monitor) # source interface e2/1
N7K1-VPC (config-monitor) # destination interface e2/2
N7K1-VPC (config-monitor) # no shut
```

Troubleshooting Tools

Onboard Logging & Diagnostic Monitoring – show logging enabled, show diagnostic

- Persistent, per module logging
- Review event history for failure detection
- Always check diagnostics before troubleshooting!

```
N7K1# show logging onboard module 2 exception-log
```

```
-----  
Module: 2  
-----
```

```
Exception Log Record : Wed Nov 16 09:36:28 2011 (608385 us)
```

```
N7K1# show diagnostic events
```

```
1) Event:E_DEBUG, length:115, at 475956 usecs after Mon Jan 23 16:17:07 2012
```

```
  [104] Event_INFO: TestName->ASICRegisterCheck TestingType->health monitoring module->5 Result->pass  
  Reason->Success
```

Troubleshooting Tools

CPU Traffic Capture - Ethalyzer

- Ethalyzer built in Wireshark capture utility
- Filter capture based on granular 'tcpdump' syntax
- Display capture natively from NXOS console using Wireshark filters

```
N7K1# ethalyzer local interface inband capture-filter 'arp' limit-captured-frames  
100 write bootflash:arp.pcap
```

```
Capturing on inband  
100
```

```
Program exited with status 0.
```

```
N7K1# ethalyzer local read bootflash:arp.pcap limit-captured-frames 100
```

```
2012-01-26 10:13:18.697098 00:00:00:00:01:23 -> ff:ff:ff:ff:ff:ff ARP Gratuitous ARP for 0.0.0.0 (Request)  
2012-01-26 10:13:18.697182 00:00:00:00:01:23 -> ff:ff:ff:ff:ff:ff ARP Gratuitous ARP for 0.0.0.0 (Request)  
2012-01-26 10:13:18.697430 00:00:00:00:01:23 -> ff:ff:ff:ff:ff:ff ARP Gratuitous ARP for 0.0.0.0 (Request)  
2012-01-26 10:13:18.697680 00:00:00:00:01:23 -> ff:ff:ff:ff:ff:ff ARP Gratuitous ARP for 0.0.0.0 (Request)
```

Agenda

- NXOS Architecture Introduction
- Built-in Troubleshooting Tools
- Troubleshooting Nexus 7000
 - CPU
 - Control-Plane – CoPP
 - Hardware
 - vPC
 - Unicast Layer 2 and Layer 3 Forwarding and ARP
 - Multicast Layer 2 and Layer 3 Forwarding
 - QoS

Troubleshooting CPU

Identify Process and Timeline – show system resources, show proc cpu

- N7K Sup1 DualCore CPU with robust preemptive scheduling
- High CPU process is not always an issue
- Per VDC output; identify the process using CPU resources

```
N7K1# show system resources
Load average: 1 minute: 0.40 5 minutes: 0.36 15 minutes: 0.28
Processes : 1091 total, 1 running
CPU states : 0.5% user, 4.0% kernel, 95.5% idle
Memory usage: 8260888K total, 3164848K used, 5096040K free
```

```
N7K1# show proc cpu sort | grep -v 0.0
PID      Runtime(ms)   Invoked    uSecs   1Sec   Process
-----  -
3530     259509       119851    2165    1.9%   platform
```

Please note that only processes from the requested vdc are shown above

Troubleshooting CPU

Identify Process and Timeline – show process cpu history

```
N7K1# show proc cpu history
```

```
<snip>
```

```

                                     1
555554435555544354657844550
466546476420229956261360720
100                                *
90                                 *
80                                 *
70                                 **
60  ***      *      *  ****  *  *
50 *****  ***** * ***** ***
40 *****
30 *****
20 *****
10 #####
0...5...1...1...2...2...3...3...4...4...5...5...6...6...7.
   0    5    0    5    0    5    0    5    0    5    0    5    0
CPU% per hour (last 72 hours)
* = maximum CPU%   # = average CPU%
```

- History shows combined utilisation across VDC
- Compare problem **timeline** with CPU spikes and average
- Identify the start of utilisation issues
- Snapshot to bootflash file for **historical trending**

Boot up 26 hours earlier

Troubleshooting CPU

Module CPU Health Check – show system resources

- CPU per module; manages module processes, no network traffic

```
module-1# show system resources
```

```
Load average: 1 minute: 0.24 5 minutes: 0.40 15 minutes: 0.34
Processes : 69 total, 2 running
CPU states : 3.0% user, 0.0% kernel, 97.0% idle
Memory usage: 1035776K total, 443632K used, 592144K free
```

```
module-1# show system internal processes cpu
```

```
top - 14:06:36 up 21 days, 15:34, 1 user, load average: 0.27, 0.42, 0.34
Tasks: 71 total, 3 running, 68 sleeping, 0 stopped, 0 zombie
Cpu(s): 9.8%us, 7.1%sy, 2.6%ni, 78.8%id, 0.0%wa, 0.1%hi, 1.6%si, 0.0%st
Mem: 1035776k total, 443784k used, 591992k free, 0k buffers
Swap: 0k total, 0k used, 0k free, 93352k cached
```

| PID | USER | PR | NI | VIRT | RES | SHR | S | %CPU | %MEM | TIME+ | COMMAND |
|------|------|----|----|-------|------|------|---|------|------|-----------|--------------|
| 1935 | root | 25 | 5 | 39304 | 8352 | 3628 | R | 81.1 | 0.8 | 3240:42 | stats_client |
| 1921 | root | -2 | 0 | 37296 | 9476 | 4116 | S | 15.8 | 0.9 | 1534:01 | naxos |
| 1917 | root | -2 | 0 | 35656 | 6600 | 4524 | S | 2.0 | 0.6 | 377:50.48 | eureka |

Troubleshooting CPU

Receiving Traffic Sent to CPU – show proc cpu sort

```
N7K1# show proc cpu | ex 0.0
```

| PID | Runtime (ms) | Invoked | uSecs | 1Sec | Process |
|------|--------------|---------|-------|------|----------|
| 3904 | 204 | 74 | 2764 | 0.9% | netstack |
| 3939 | 477316 | 717760 | 665 | 0.9% | R2D2_usd |

```
CPU util : 12.0% user, 4.0% kernel, 84.0% idle
```

Please note that **only processes from the requested vdc are shown above**

Per VDC instance of netstack;
each process utilization
separate from other

```
N7K1# switchto vdc VPC
```

```
N7K1-VPC# show proc cpu sort | ex 0.0
```

| PID | Runtime (ms) | Invoked | uSecs | 1Sec | Process |
|------|--------------|---------|-------|-------|----------|
| 5861 | 360 | 126 | 2861 | 11.3% | netstack |
| 5840 | 162 | 38 | 4273 | 4.7% | arp |

```
CPU util : 5.0% user, 21.5% kernel, 73.5% idle
```

Please note that **only processes from the requested vdc are shown above**

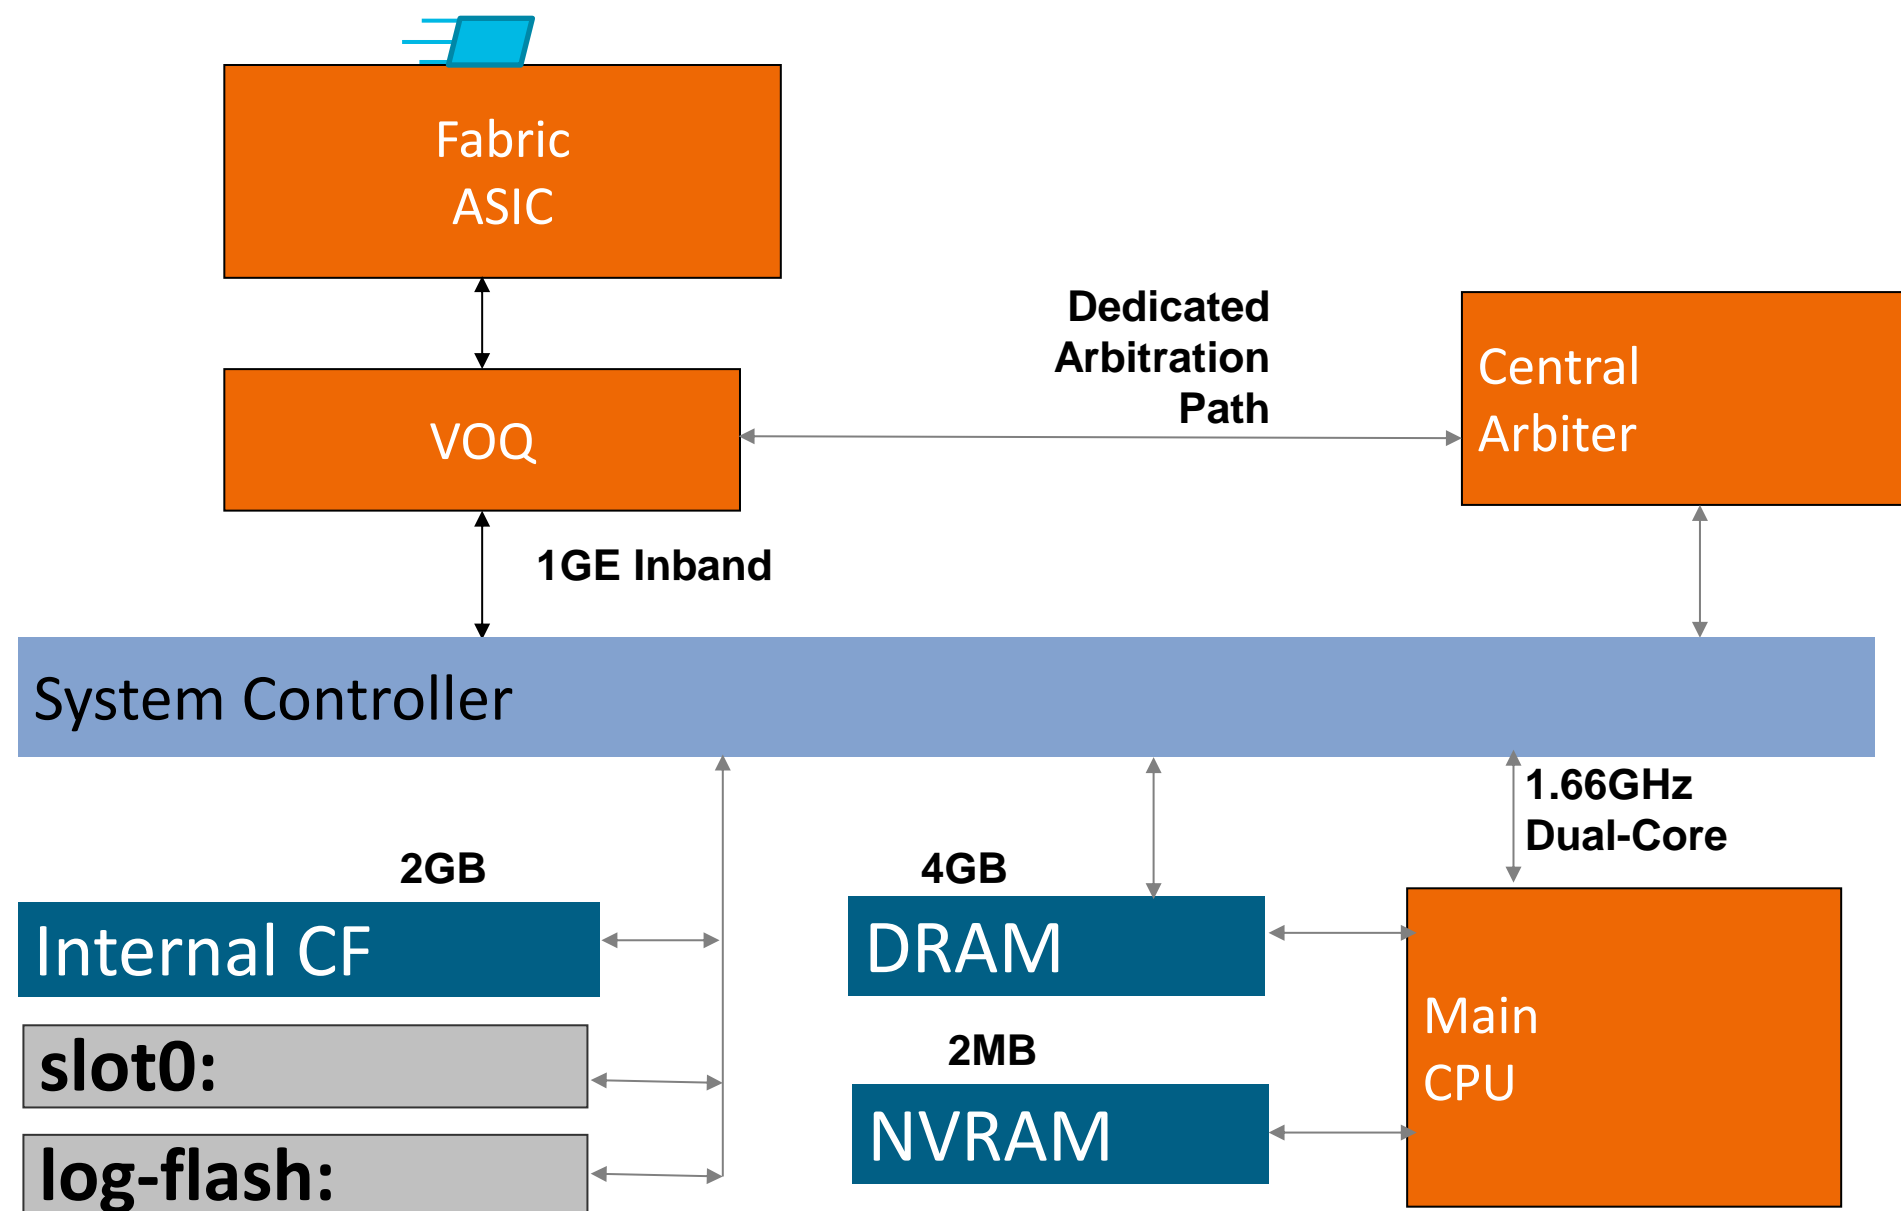
Netstack and ARP are running
higher than usual; we have
traffic hitting CPU

Agenda

- NXOS Architecture Introduction
- Built-in Troubleshooting Tools
- Troubleshooting Nexus 7000
 - CPU
 - Control-Plane – CoPP
 - Hardware
 - vPC
 - Unicast Layer 2 and Layer 3 Forwarding and ARP
 - Multicast Layer 2 and Layer 3 Forwarding
 - QoS

Troubleshooting CoPP

Control Plane Policing – Protect CPU from your network



Troubleshooting CoPP

Control Plane Policing – Protect CPU from your network

- Protect CPU from network traffic
 - arp, nd (ipv6)
 - DHCP traffic
 - Glean traffic (no arp or nd)
- Two stage policing
 - Rate-limit in hardware per-module before sending to CPU
 - Policy based policing traffic that have passed rate-limiters
- May require **tweaking** to match your network conditions

Troubleshooting CoPP

Control Plane Policing – show policy-map interface control-plane

```
N7K1# show policy-map int control
Control Plane
  service-policy input: copp-system-policy

class-map copp-system-class-critical (match-any)
  match access-group name copp-system-acl-bgp

class-map copp-system-class-exception (match-any)
  match exception ip option
  match exception ip icmp unreachable
  match exception ipv6 option
  match exception ipv6 icmp unreachable
  set cos 1
  police cir 360 kbps, bc 250 ms
  module 2 :
    conformed 0 bytes; action: transmit
    violated 0 bytes; action: drop
  module 3 :
    conformed 0 bytes; action: transmit
    violated 0 bytes; action: drop
  module 10 :
    conformed 273344 bytes; action: transmit
    violated 0 bytes; action: drop
```

Granular class based matching;
IP and MAC ACL in one class

Exception traffic; CPU issues no
longer created by ICMP exception
generation

Implemented per-module; check
for violated traffic

Troubleshooting CoPP

Control Plane Policing – show policy-map interface control-plane

- Filter output to confirm affected classes
- Check class-map definitions if violated traffic does not match your expectation

```
N7K1# show policy-map int control | egrep "class-map|violated"
```

```
class-map copp-system-class-critical (match-any)
```

```
violated 0 bytes; action: drop
```

```
violated 0 bytes; action: drop
```

```
violated 0 bytes; action: drop
```

```
class-map copp-system-class-management (match-any)
```

```
violated 0 bytes; action: drop
```

```
violated 0 bytes; action: drop
```

```
violated 0 bytes; action: drop
```

```
class-map class-default (match-any)
```

```
violated 0 bytes; action: drop
```

```
violated 0 bytes; action: drop
```

```
violated 274930048 bytes; action: drop
```

Routing protocols class; no drops

TFTP; SNMP; FTP protocols

Low rate policer on default class;
100Kbps default dropped traffic

Troubleshooting CoPP

Control Plane Policing – show hardware rate-limit

- Per module rate-limiter in addition to control-plane policy

```
N7K1# show hardware rate-limiter
```

```
Units for Config: packets per second
```

```
Rate Limiter Class
```

```
Parameters
```

```
-----  
layer-3 control
```

```
Config      : 10000  
Allowed     : 0  
Dropped     : 0  
Total       : 0
```

Review layer-3 control if control-plane instability

```
<snip>
```

```
copy
```

Copy to sup used for Exceptions; limited before CoPP

```
Config      : 30000  
Allowed     : 3934270  
Dropped     : 829800  
Total       : 4764070
```

```
receive
```

Traffic destined to Sup

```
Config      : 30000  
Allowed     : 6875754  
Dropped     : 0  
Total       : 6875754
```

- Hardware Rate-Limiter performed by forwarding engine hardware
- Global system wide hardware feature
- Limit based on packet definitions

Troubleshooting CoPP

Review Traffic Sent to CPU – show hardware internal inband cpu-mac stats

N7K-1# show hardware internal cpu-mac inband stats

```

RMON counters
total packets          Rx          Tx
good packets          779905245  1421785114
total octets (hi)      779905245  1421650279
total octets (low)    0          0
total octets (low)    172302724342  192974265660
Error counters
-----+-----
Rx no buffers ..... 1203243
Throttle statistics
Throttle interval ..... 2 * 100ms
Packet rate limit ..... 32000 pps
Tick counter ..... 12414130
Rx packet rate (current/max) 4993 / 20296 pps
Tx packet rate (current/max) 60 / 3474 pps
MAC counters
MAC0 (R2D2)
Rx          Tx          MAC1 (CPU)
Rx          Tx          Rx          Tx
XOFF packets auto-generated          5447
XOFF packets          7590855  6731953
XON packets          0          18561642
    
```

Total number of frames received and sent by CPU

Traffic tail dropped; after CoPP and RL

CPU bound traffic current pps /maximum pps reached

Troubleshooting CoPP

Review Traffic Sent to CPU – ethalyzer local interface

- Start by capturing all traffic; define filters based on first capture
- Capture to text file or .pcap
- Local review or export to wireshark

Troubleshooting CoPP

Review Traffic Sent to CPU – ethanalyzer local interface

```
N7K1# ethanalyzer local int inband limit-captured-frame 100 write  
bootflash:cpu.pcap
```

```
Capturing on inband  
100  
Program exited with status 0.
```

```
N7K1# ethanalyzer local int inband decode-internal limit-captured-frames 100 >  
cpu.txt
```

Creates pcap file which can later be analysed by GUI wireshark

Provides information from internal system headers, not allowed with 'write' keyword

Filter syntax: http://www.wireshark.org/docs/wsug_html_chunked/ChCapCaptureFilterSection.html

Troubleshooting CoPP

Review Traffic Sent to CPU – ethanalyzer local read

```
N7K1# ethanalyzer local read bootflash:cpu.pcap
```

```
2012-01-26 10:02:20.387538 172.16.60.123 -> 172.16.50.1  UDP Source port: 63  Destination port: 63
2012-01-26 10:02:20.387691 172.16.60.123 -> 172.16.50.1  UDP Source port: 63  Destination port: 63
2012-01-26 10:02:20.388066 172.16.60.123 -> 172.16.50.1  UDP Source port: 63  Destination port: 63
2012-01-26 10:02:20.388566 172.16.60.123 -> 172.16.50.1  UDP Source port: 63  Destination port: 63
2012-01-26 10:02:20.388940 172.16.60.123 -> 172.16.50.1  UDP Source port: 63  Destination port: 63
```

Local analysis completed from
CLI prompt

Troubleshooting CoPP

Review Traffic Sent to CPU – ethanalyzer local read

```
N7K1# ethanalyzer local read bootflash:cpu.pcap detail
```

```
Frame 1 (92 bytes on wire, 60 bytes captured)
```

```
Arrival Time: Jan 26, 2012 10:02:20.387538000
```

```
[Time delta from previous captured frame: 0.000000000 seconds]
```

```
[Time delta from previous displayed frame: 0.000000000 seconds]
```

```
[Time since reference or first frame: 0.000000000 seconds]
```

```
Frame Number: 1
```

```
Frame Length: 92 bytes
```

```
Capture Length: 60 bytes
```

```
[Frame is marked: False]
```

```
[Protocols in frame: eth:ip:udp:data]
```

```
Ethernet II, Src: 00:24:97:36:81:3f (00:24:97:36:81:3f), Dst: 04:c5:a4:e9:ac:44  
(04:c5:a4:e9:ac:44)
```

```
Destination: 04:c5:a4:e9:ac:44 (04:c5:a4:e9:ac:44)
```

```
Address: 04:c5:a4:e9:ac:44 (04:c5:a4:e9:ac:44)
```

```
.... 0 .... = IG bit: Individual address (unicast)
```

```
.... 0. .... = LG bit: Globally unique address (factory default)
```

Wireshark output; display filters help to find relevant packets

Use filter in next capture to increase the relevance

Troubleshooting CoPP

Ethalyzer Examples

Capture ICMP traffic from specific source to destination

```
ethalyzer local inter inband capture-filter 'icmp && src host 192.168.0.2 && dst host 192.168.0.1'
```

Capture ALL traffic from specific source to destination using IP

```
ethalyzer local inter inband capture-filter 'src host 192.168.0.2 && dst host 192.168.0.1'
```

Capture ALL traffic from specific source to destination using MAC

```
ethalyzer local inter inband capture-filter 'ether src 00:1b:0d:ef:5a:c0 && ether dst 00:11:92:fa:b2:c0'
```

Capture ARP traffic destined to an interface or SVI having IP 192.168.88.3

```
ethalyzer local inter inband capture-filter 'arp && src host 192.168.88.3'
```

Capture ALL SNMP traffic

```
ethalyzer local inter inband capture-filter 'udp port 161 || tcp port 161' limit-captured-frames 1000
```

Capture ALL HSRP v1 Control Packets

```
ethalyzer local interface inband capture-filter "dst net 224.0.0.2 and udp dst port 1985"
```

Agenda

- NXOS Architecture Introduction
- Built-in Troubleshooting Tools
- **Troubleshooting Nexus 7000**
 - CPU
 - Control-Plane – CoPP
 - **Hardware**
 - vPC
 - Unicast Layer 2 and Layer 3 Forwarding and ARP
 - Multicast Layer 2 and Layer 3 Forwarding
 - QoS

Troubleshooting Hardware

Module – show module, reload module <x>

```
N7K1# show module
Mod  Ports  Module-Type                Model                Status
---  -
2    32     10 Gbps Ethernet XL Module N7K-M132XP-12L      ok
3    48     1000 Mbps Optical Ethernet N7K-M148GS-11      testing
5    0      Supervisor module-1X      N7K-SUP1           active *
6    0      Supervisor module-1X      N7K-SUP1           ha-standby
10   48     10/100/1000 Mbps Ethernet N7K-M148GT-11      ok

<snip>
Mod  Online Diag Status
---  -
2    Pass
3    Untested
5    Pass
6    Pass
10   Pass

N7K1# reload module 3
This command will reload module 3. Proceed[y/n]? [n] y
reloading module 3 ...
```

Module booting during troubleshooting

Reload from CLI denies module access to power shelf; only reset will resolve physical issue

Troubleshooting Hardware

Online Diagnostics Review – show diagnostic result, show diagnostic event

```
N7K1# show diagnostic result mod 3 | grep -v ^$
```

```
Current bootup diagnostic level: complete
Module 3: 1000 Mbps Optical Ethernet Module
Test results: (. = Pass, F = Fail, I = Incomplete,
U = Untested, A = Abort, E = Error disabled)
1) ASICRegisterCheck-----> .
2) PrimaryBootROM-----> .
3) SecondaryBootROM-----> .
4) EOBCPortLoopback-----> .
5) OBFL-----> .
6) PortLoopback:
Port   1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16
-----
      .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .
Port  17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32
-----
      .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .
Port  33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48
-----
      .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .
```

Eliminates empty lines

Per port diagnostic test. Initiated from module CPU. F could be transient or survive a module reset

```
N7K1# show diagnostic events error
```

```
1) Event:E_DEBUG, length:217, at 362751 usecs after Wed Jan 25 03:26:21 2012
   [103] Event_ERROR: TestName->SpineControlBus TestingType->health monitoring module->5 Result->fail
   Reason-> XBar
```


Troubleshooting Hardware

Fabric – show module, show hardware capacity

```
N7K1# show module xbar
```

| Xbar | Ports | Module-Type | Model | Status |
|------|-------|-----------------|-----------------|--------|
| 1 | 0 | Fabric Module 1 | N7K-C7010-FAB-1 | ok |
| 2 | 0 | Fabric Module 1 | N7K-C7010-FAB-1 | ok |
| 3 | 0 | Fabric Module 1 | N7K-C7010-FAB-1 | ok |

```
N7K1# show hardware capacity fabric-utilization
```

Fabric Planes:

A -- Unicast fabric interface

B -- Multicast/Multidestination fabric interface

```
-----PEAK FABRIC UTILIZATION-----
```

| I/O | FABRIC | | | Ingress | | Egress | |
|------|--------|------|-------|---------|----------------|--------|----------------|
| Slot | Mod | Inst | Plane | Util | Time | Util | Time |
| 2 | 1 | 1 | A | 0% | 01-23@16:41:46 | 0% | 01-23@16:41:46 |
| 2 | 1 | 1 | B | 0% | 01-23@16:41:46 | 0% | 01-23@16:41:46 |

Xbar utilization max, check previous history for trending info

```
N7K1 (config) # poweroff xbar 1
```

Power off xbar before trying hardware swap

Troubleshooting Hardware

Fabric – show module, show hardware capacity

```
Campus_N7K1-VPC# sh int e2/2
Ethernet2/2 is up
admin state is up, Dedicated Interface
  Belongs to Po10
  Hardware: 10000 Ethernet, address: c89c.1d82.07f9 (bia c89c.1d82.07f9)
  Description: 40G Trunk btwn 7ks
  MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec
  Last link flapped 4d11h
  Last clearing of "show interface" counters never
  1 interface resets
  30 seconds input rate 11192 bits/sec, 14 packets/sec
  30 seconds output rate 11096 bits/sec, 14 packets/sec
  Load-Interval #2: 5 minute (300 seconds)
    input rate 10.89 Kbps, 19 pps; output rate 10.74 Kbps, 19 pps
  RX
    698 unicast packets  7616146 multicast packets  28542 broadcast packets
    7645386 input packets  553579374 bytes
    0 jumbo packets  0 storm suppression packets
    0 runts  0 giants  0 CRC  0 no buffer
    0 input error  0 short frame  0 overrun  0 underrun  0 ignored
    0 watchdog  0 bad etype drop  0 bad proto drop  0 if down drop
    0 input with dribble  0 input discard
    0 Rx pause
```

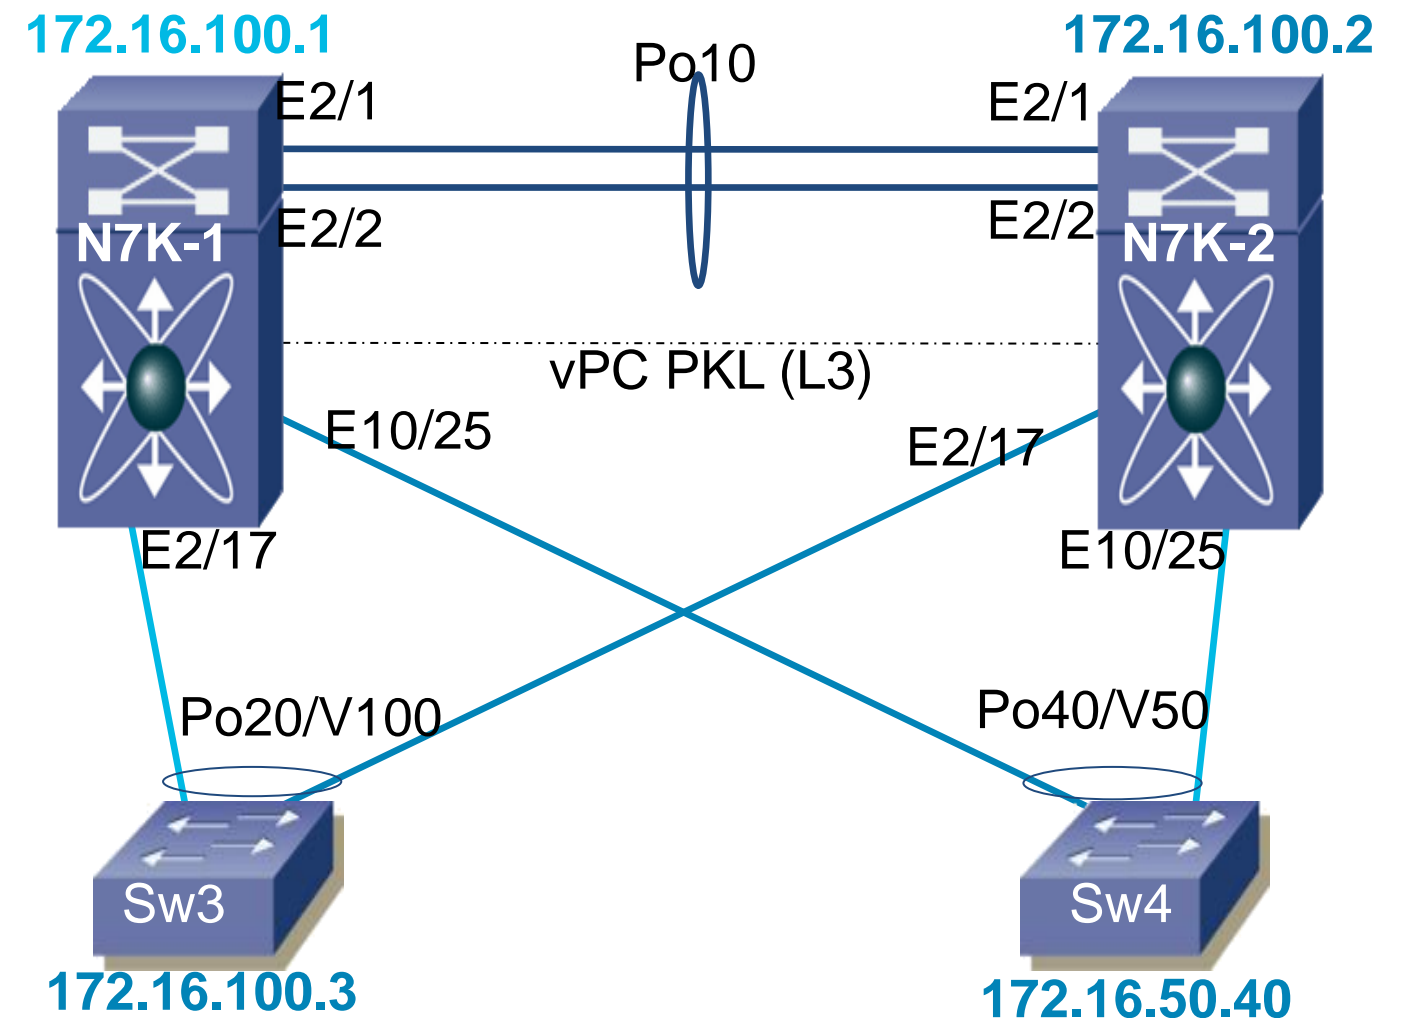
Agenda

- NXOS Architecture Introduction
- Built-in Troubleshooting Tools
- Troubleshooting Nexus 7000
 - CPU
 - Control-Plane – CoPP
 - Hardware
 - vPC
 - Unicast Layer 2 and Layer 3 Forwarding and ARP
 - Multicast Layer 2 and Layer 3 Forwarding
 - QoS

Troubleshooting vPC

Distributed Layer 2 Forwarding

- Designed as an alternative to STP
- Appear as one L2 device to the network
- Peers must maintain consistent network view
- Recovery mechanisms protect forwarding



Troubleshooting vPC

Global Consistency Check – show vpc brief

```
N7K1-VPC# show vpc brief
```

```
vPC domain id          : 1
Peer status             : peer adjacency formed ok
vPC keep-alive status  : peer is alive
Configuration consistency status : success
Per-vlan consistency status : success
Type-2 consistency status : failed
Type-2 inconsistency reason : SVI type-2 configuration incompatible
vPC role                : secondary
Number of vPCs configured : 2
Peer Gateway            : Enabled
```

Secondary role, vPCs shut if peer-link is down

Peer gateway enabled; routing for peer

```
<snip>
```

```
vPC Peer-link status
```

| id | Port | Status | Active vlans |
|----|------|--------|---------------------------|
| 1 | Po10 | up | 1,5-9,11,100,200,203,1111 |

Type-2 consistency check failed, but vlans still passing on vPC and peer-link

```
vPC status
```

| id | Port | Status | Consistency | Reason | Active vlans |
|----|------|--------|-------------|---------|---------------------------|
| 20 | Po20 | up | success | success | 1,5-9,11,203 |
| 40 | Po40 | up | success | success | 1,5-9,11,100,200,203,1111 |

Troubleshooting vPC

Global Consistency Check – show vpc consistency parameters global

```
N7K1-VPC# show vpc consistency-parameters global
```

```
Type 1 : vPC will be suspended in case of mismatch
```

| Name | Type | Local Value | Peer Value |
|---|------|---|---|
| STP Mode | 1 | Rapid-PVST | Rapid-PVST |
| STP Disabled | 1 | None | None |
| <snip> | | | |
| STP Port Type, Edge BPDUFilter, Edge BPDUGuard | 1 | Normal, Disabled, Disabled | Normal, Disabled, Disabled |
| STP MST Simulate PVST | 1 | Enabled | Enabled |
| VTP domain | 2 | TEST | - |
| VTP version | 2 | 1 | - |
| VTP mode | 2 | Server | - |
| VTP password | 2 | | - |
| VTP pruning status | 2 | Disabled | - |
| Interface-vlan admin up | 2 | 5-6,100 | 5-6,100,200 |
| Interface-vlan routing capability | 2 | 1,5-6,100 | 1,5-6,100,200 |
| Allowed VLANs | - | 1,5-9,11,100,200,203,1 002-1005,1111 | 1-3,5-12,16,20-23,32,1 00,156,200-201,203,220 -222,230,508,555,921,9 93-994,1000,1111,1221, 1999,2211-2213,2233-22 34,2901 |
| Local suspended VLANs | - | 1002-1005 | - |

VTP enabled locally,
disabled on remote

SVI 200 not
configured locally

Mismatched allowed
VLANs

Troubleshooting vPC

MAC Address Synchronised – show mac address-table

- MAC addresses synchronized via CFS
- First Hop Redundancy Protocol addresses forwarded by both peers
- Peer-gateway allows routing for peer SVI MAC

```
N7K1-VPC# show mac add vlan 100
```

```
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since last seen,+ - primary entry using vPC Peer-Link
VLAN    MAC Address      Type      age      Secure NTFY Ports/SWID.SSID.LID
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
G 100    0000.0c07.ac01   static    -        F      F      vPC Peer-Link (R)
G 100    04c5.a4e9.aac4   static    -        F      F      vPC Peer-Link (R)
G 100    04c5.a4e9.ac44   static    -        F      F      sup-eth1 (R)
* 100    0021.d87c.2740   dynamic   0        F      F      Po20
```

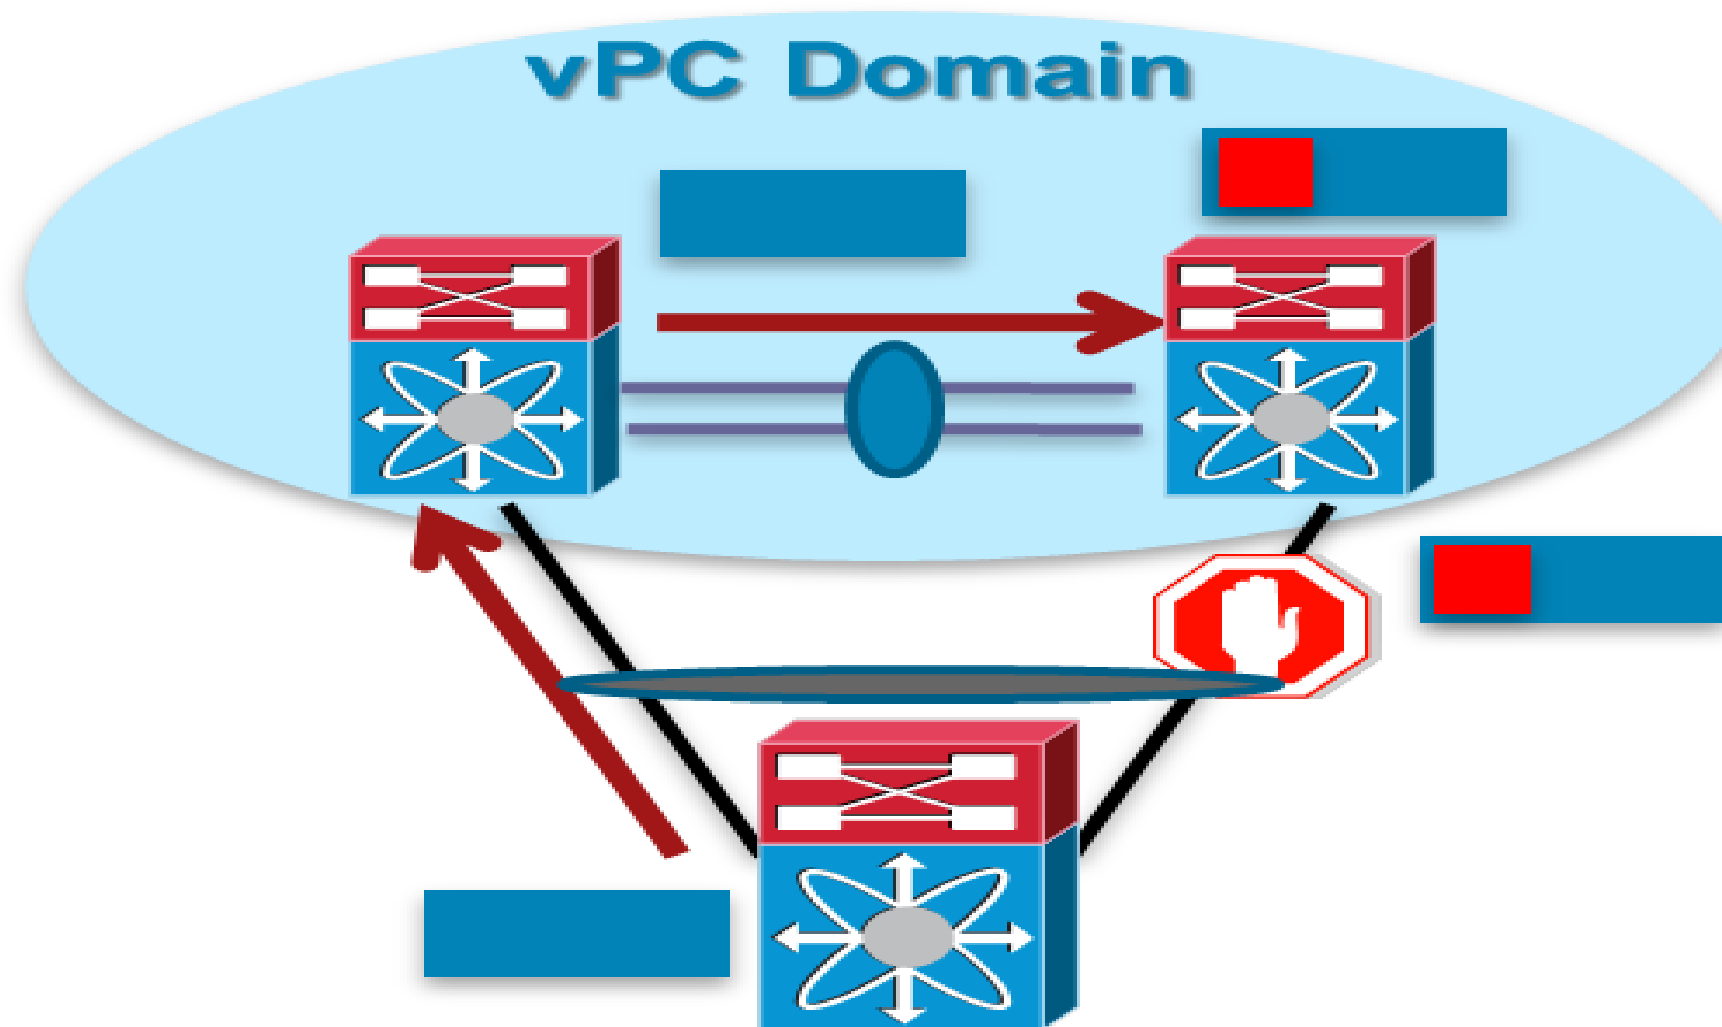
```
N7K1-VPC# show run vpc | egrep "peer-gateway"
```

```
peer-gateway
```

Peer MAC address with G-bit set, route traffic destined to peer

Troubleshooting vPC

Loop Avoidance – Peer Gateway



Troubleshooting vPC

Data Collection – show tech vpc, show tech stp, show tech pixm

```
N7K1-VPC# show tech-support vpc | grep "`show`"
```

```
`show version`  
`show module`  
`show vpc brief`  
`show vpc role`  
`show running-config vpc`  
`show system internal vpcm event-history global`  
`show system internal vpcm event-history errors`  
`show system internal vpcm event-history msgs`  
`show system internal vpcm event-history interactions`  
`show system internal vpcm mem-stats detail`  
`show system internal vpcm info all`  
`show system internal vpcm info global`  
`show cfs internal ethernet-peer database`  
`show spanning-tree`
```

```
N7K1-VPC# show tech-support stp
```

```
N7K1-VPC# show tech-support vtp
```

```
N7K1-VPC# show tech-support pixm
```

```
N7K1-VPC# show tech-support forwarding 12 unicast
```

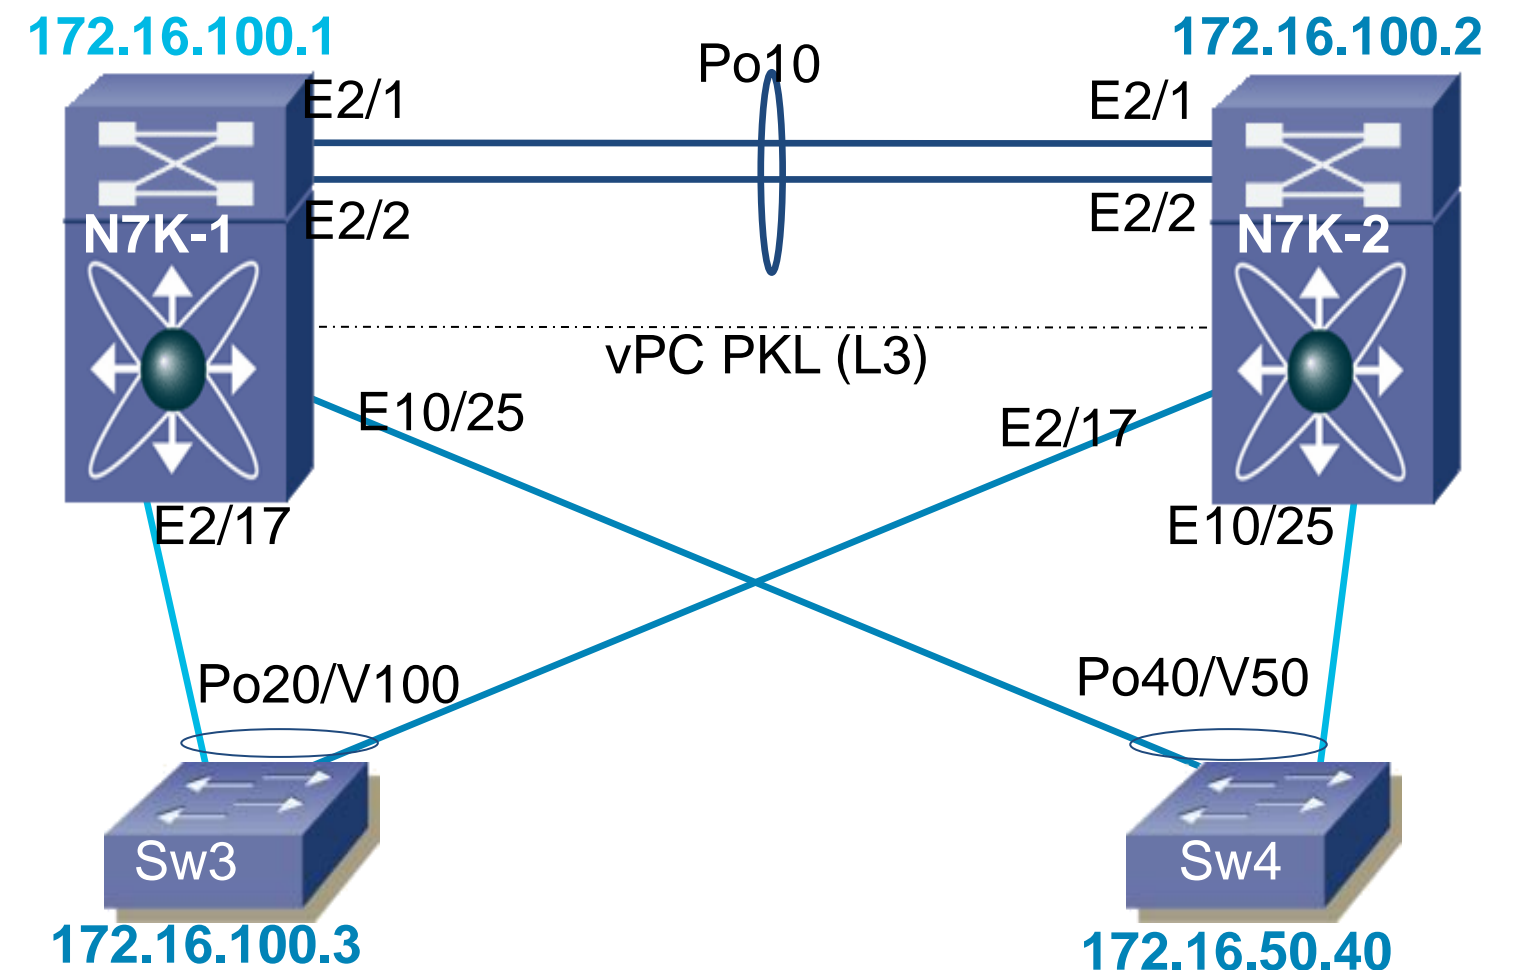
Agenda

- NXOS Architecture Introduction
- Built-in Troubleshooting Tools
- Troubleshooting Nexus 7000
 - CPU
 - Control-Plane – CoPP
 - Hardware
 - vPC
 - Unicast Layer 2 and Layer 3 Forwarding and ARP
 - Multicast Layer 2 and Layer 3 Forwarding
 - QoS

Troubleshooting Unicast Forwarding

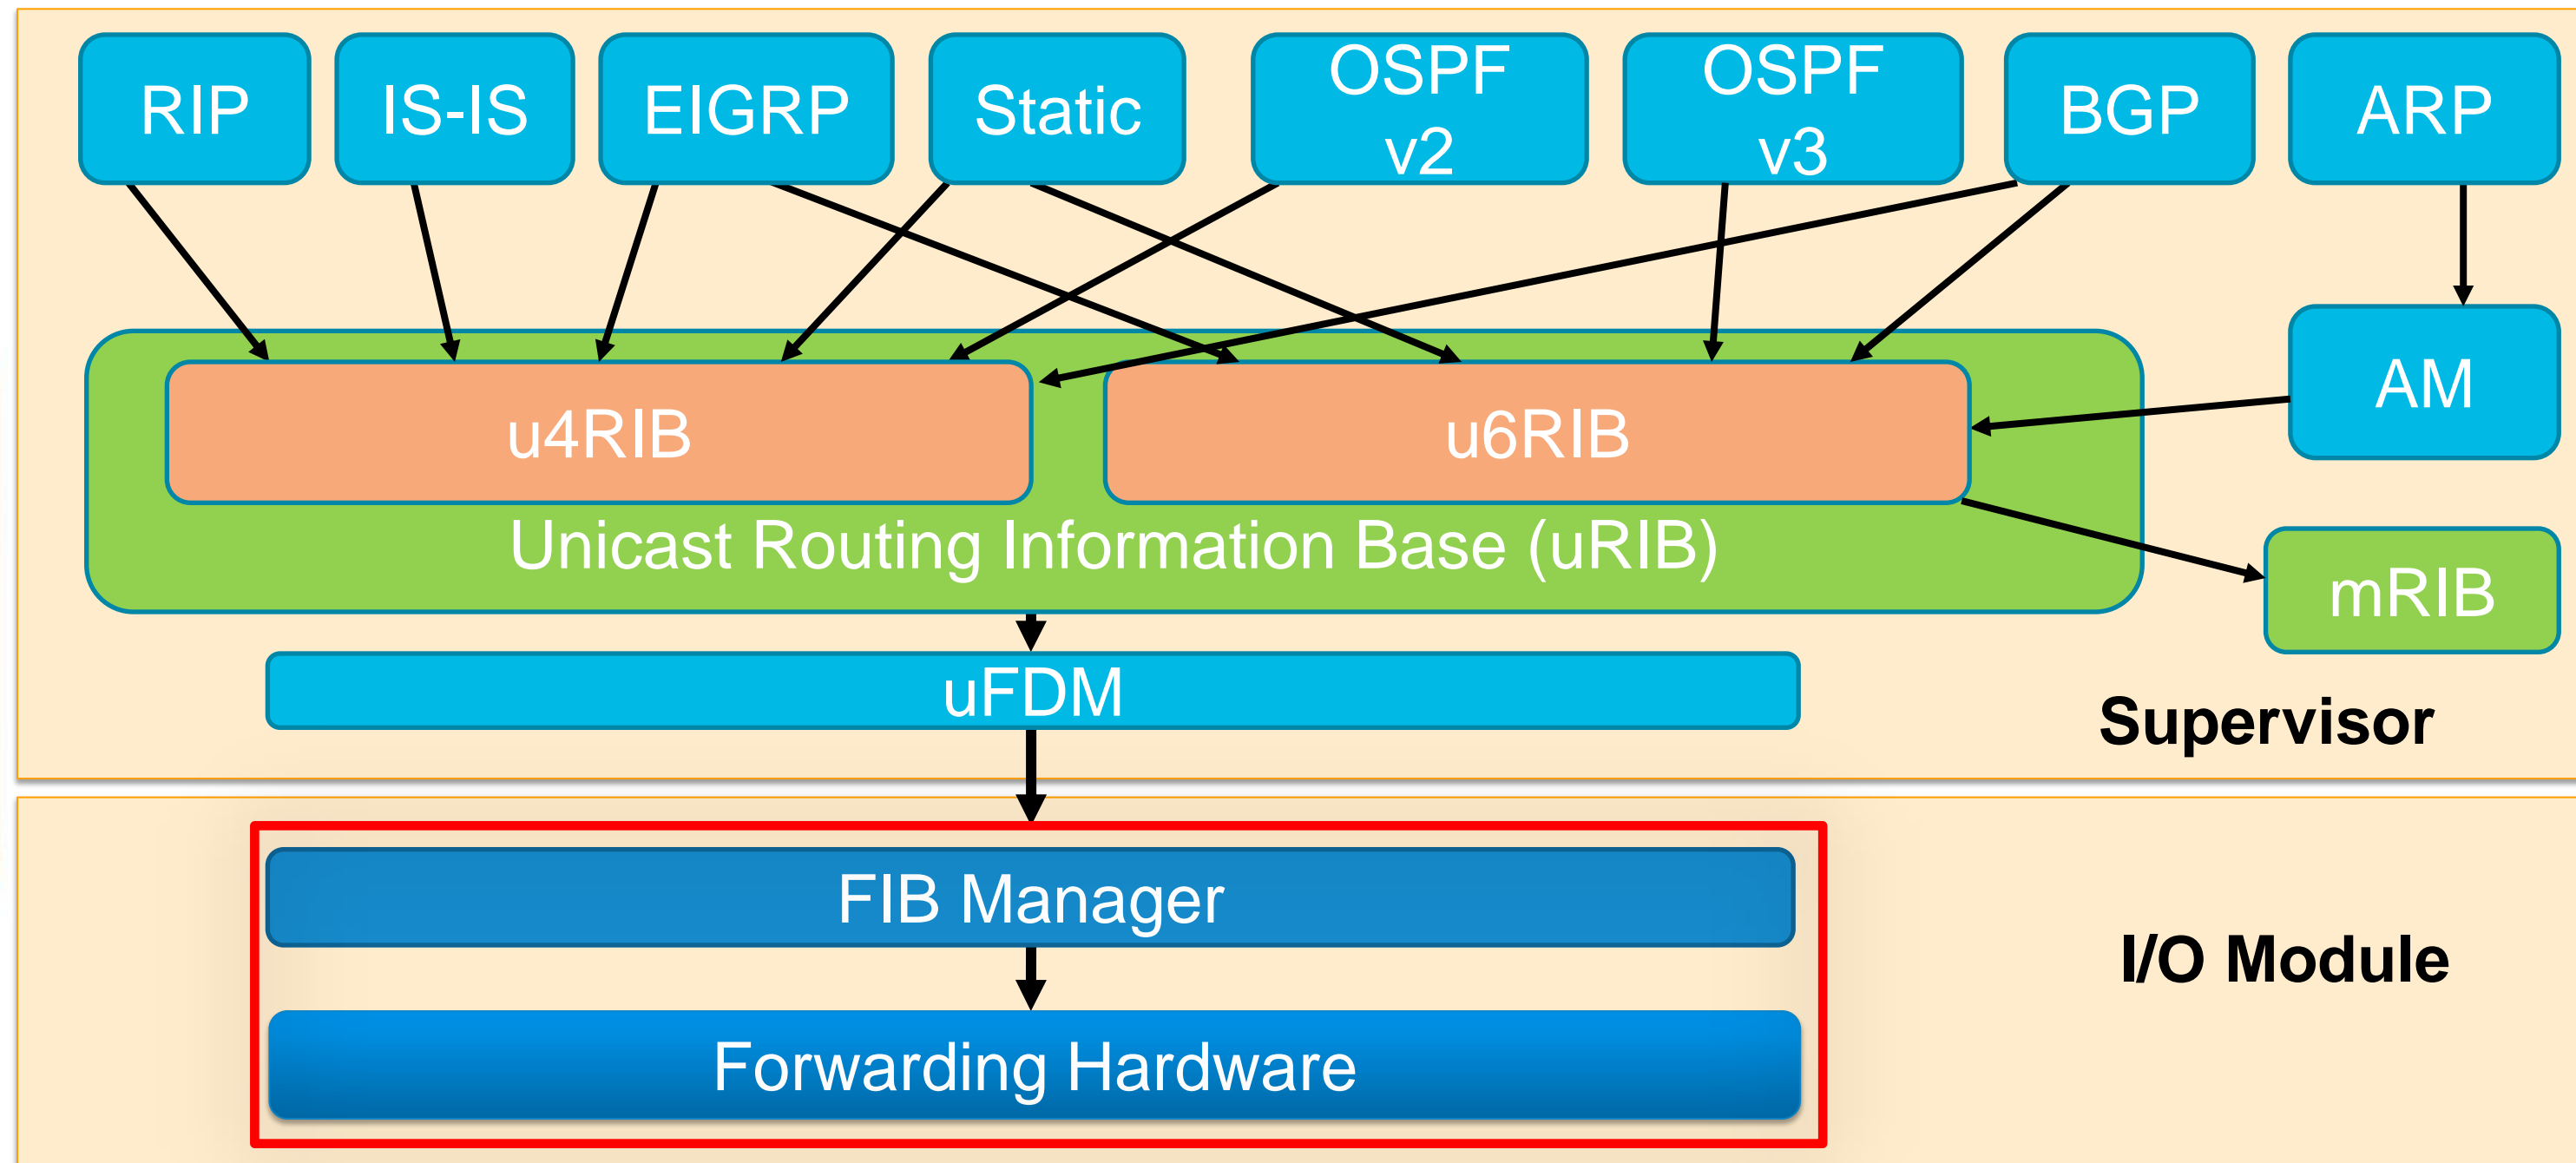
Forwarding L2 and L3 in Hardware

- L2 learning in hardware by forwarding engine per module
- L3 entries learned by CPU and pushed to hardware engine
- Verify software entries against **topology** diagrams
- Verify hardware entries against software entries



Troubleshooting Unicast Forwarding

Forwarding L2 and L3 in Hardware



Troubleshooting Unicast Forwarding

Software Tables – show ip route, show ip adjacency, show mac address-table

```
N7K1-VPC# show ip route 172.16.50.40
IP Route Table for VRF "default"
'*' denotes best ucast next-hop

172.16.50.40/32, ubest/mbest: 1/0, attached
   *via 172.16.50.40, Vlan50, [250/0], 00:02:09, am
```

```
N7K1-VPC# show ip adjacency 172.16.50.40
Flags: # - Adjacencies Throttled for Glean
IP Adjacency Table for VRF default
Total number of entries: 1
Address          MAC Address      Pref Source      Interface
172.16.50.40    0024.9736.813f  50   arp             Vlan50
```

```
Campus_N7K1-VPC# show mac address-table address 0024.9736.813f
Legend:
   * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
   age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN      MAC Address      Type      age      Secure NTFY Ports/SWID.SSID.LID
-----+-----+-----+-----+-----+-----+-----
* 50         0024.9736.813f  dynamic  150      F    F    Po40
```

Troubleshooting Unicast Forwarding

L2 Hardware Entries – show hardware mac-add, show system internal pixm

```
N7K1-VPC# show hardware mac add 2 vlan 50
```

| FE | Valid | PI | BD | MAC | Index | Stat | SW | Modi | Age | Tmr | GM | Sec | TR | NT | RM | RMA | Cap | Fld | Always |
|----|-------|----|----|----------------|---------|------|-------|------|------|-----|----|-----|----|----|----|------|-----|-----|--------|
| | | | | | | ic | | fied | Byte | Sel | | ure | AP | FY | | TURE | | | Learn |
| 0 | 1 | 0 | 59 | 0024.9736.813f | 0x00a48 | 0 | 0x003 | 0 | 9 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 1 | 0 | 59 | 0000.0000.0123 | 0x00a48 | 0 | 0x003 | 0 | 156 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 1 | 1 | 59 | 0000.0c07.ac01 | 0x00400 | 1 | 0x000 | 0 | 156 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 1 | 1 | 59 | 04c5.a4e9.ac44 | 0x00400 | 1 | 0x000 | 0 | 130 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 1 | 0 | 59 | 4180.0000.0127 | 0x07ff1 | 1 | 0x001 | 0 | 170 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 |
| 0 | 1 | 0 | 59 | 4180.0000.0128 | 0x07ff1 | 1 | 0x001 | 0 | 29 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 |

```
N7K1-VPC# show vlan internal bd-info vlan-to-bd 50
```

| VDC Id | Vlan Id | BD Id |
|--------|---------|-------|
| 4 | 50 | 59 |

Allocated bridge-domain matches in hardware table

DMAC sent to LTL index for PO40

```
N7K1-VPC# show system internal pixm info ltl 0xa48
```

| PC_TYPE | PORT | LTL | RES_ID | LTL_FLAG | CB_FLAG | MEMB_CNT |
|---------|------|--------|------------|------------|------------|----------|
| Normal | Po40 | 0x0a48 | 0x16000027 | 0x00000000 | 0x00000002 | 1 |

Troubleshooting Unicast Forwarding

L3 Hardware Entries – show ip fib route <dest>, show system internal forward

```
N7K1-VPC# show ip fib route 172.16.50.40 module 2
```

```
IPv4 routes for table default/base
```

| Prefix | Next-hop | Interface | Labels |
|-----------------|--------------|-----------|--------|
| 172.16.50.40/32 | 172.16.50.40 | Vlan50 | |

```
N7K1-VPC# show system internal forwarding ip route 172.16.50.40 module 2
```

| Dev | Prefix | PfxIndex | AdjIndex | LIFB | LIF |
|-----|-----------------|----------|----------|------|------|
| 1 | 172.16.50.40/32 | 0x2033 | 0x43015 | 0 | 0x3b |

```
Campus_N7K1-VPC# show system internal forwarding adjacency mod 2 entry 0x43015 det
```

```
Device: 1   Index: 0x43015   DMAC: 0024.9736.813f   SMAC: 00c5.a4e9.ac44  
LIF: 0x3b (Vlan50)   DI: 0x0   ccc: 4   L2_FWD: NO   RDT: NO  
packets: 0   bytes: 0   zone enforce: 0
```

Troubleshooting Unicast Forwarding

Scale and Utilisation – show ip route summary, show hardware internal

```
N7K1-VPC# show ip route summary
```

```
IP Route Table for VRF "default"
```

```
Total number of routes: 45
```

```
Total number of paths: 47
```

```
Best paths per protocol:
```

```
am : 4
```

```
local : 9
```

```
Backup paths per protocol:
```

```
None
```

```
<snip>
```

```
Number of routes per mask-length:
```

```
/8 : 1 /24: 6 /30: 1 /32: 37
```

```
N7K1-VPC# show hardware internal forwarding table utilization module 2
```

```
Note: Utilization may not reach the maximum.
```

```
Module 2 usage:
```

| Route Type | Used (Log/Phys) | %Used | Free (Log/Phys) | %Free | Total (Log/Phys) |
|---------------|--------------------|-------|--------------------|-------|---------------------|
| IPv4 Unicast: | 67/67 | 0 | | | |
| L2VPN Peer: | 0/0 | 0 | | | |
| MPLS: | 0/0 | 0 | | | |
| | | | 904967/904967 | 99 | 905120/905120 |

Troubleshooting Unicast Forwarding

Data Collection – show tech forwarding, show tech netstack, show tech arp

```
N7K1-VPC# show tech-support forwarding L3 unicast | grep "`show`"
```

```
`show forwarding route summary vrf all`  
`show forwarding route max-display-count 100000 vrf all`  
`show forwarding vrf all adjacency`  
`show forwarding ipv6 route summary vrf all`  
`show forwarding ipv6 route max-display-count 100000 vrf all`  
`show forwarding vrf all ipv6 adjacency`  
`show forwarding trace`  
`show forwarding internal errors`  
`show forwarding internal error counts`  
`show forwarding internal unicast counts vdc all`  
`show forwarding internal message counts`
```

```
N7K1-VPC# show tech-support netstack | grep "`show`" | wc -l  
212
```

```
N7K1-VPC# show tech-support arp | grep "`show`"
```

```
`show running-config arp`  
`show ip arp internal event-history cli`  
`show ip arp vrf all`  
`show ip arp static vrf all`  
`show ip arp summary vrf all`  
`show ip arp tunnel-statistics`
```

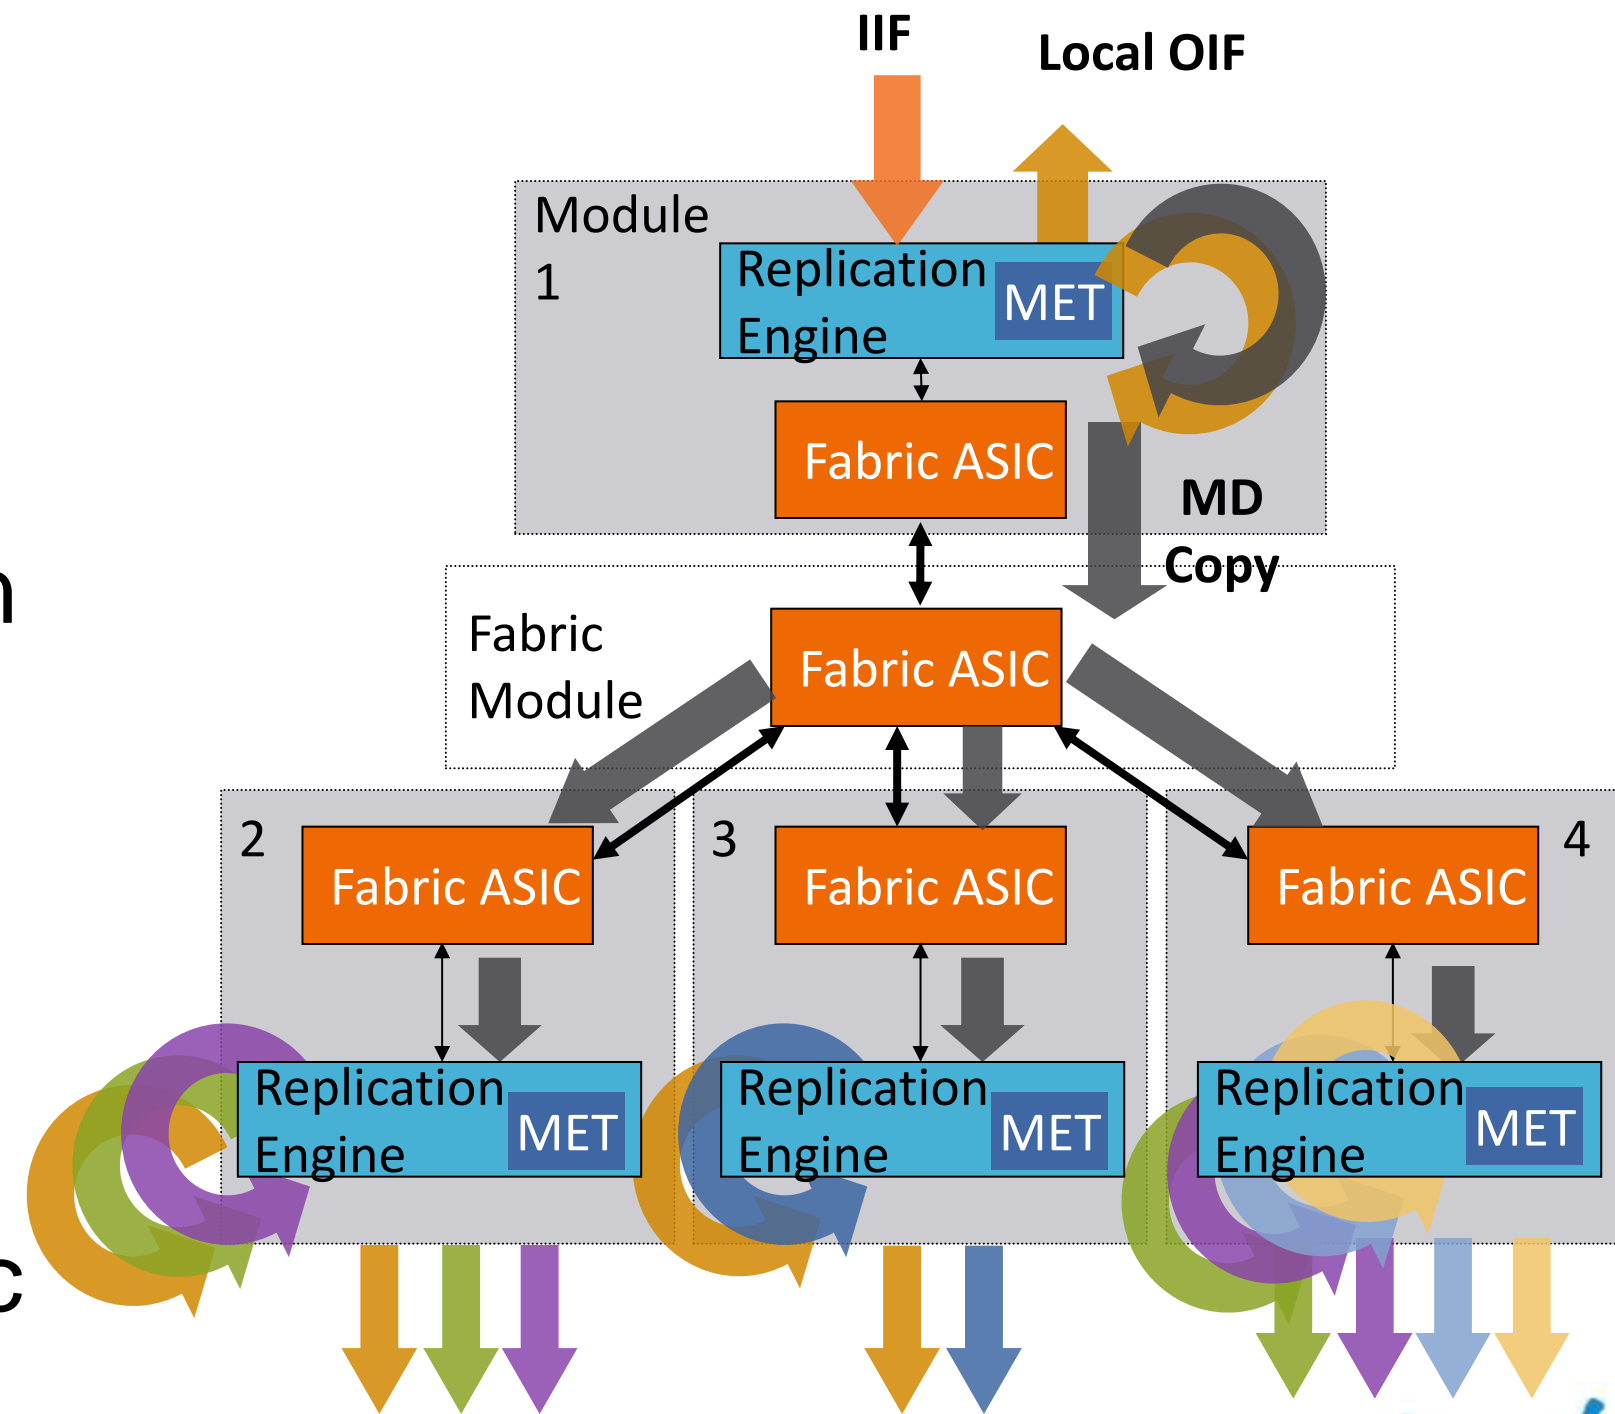

Agenda

- NXOS Architecture Introduction
- Built-in Troubleshooting Tools
- Troubleshooting Nexus 7000
 - CPU
 - Control-Plane – CoPP
 - Hardware
 - vPC
 - Unicast Layer 2 and Layer 3 Forwarding and ARP
 - Multicast Layer 2 and Layer 3 Forwarding
 - QoS

Troubleshooting Multicast Forwarding

Distributing Streams in Hardware

- Software learning and hardware entries mirror IPv4 distribution model
- Conservation of L2 replication based on L3 address
- Egress replication forwarding conserves hardware resources
- Single copy sent across fabric conserves bandwidth



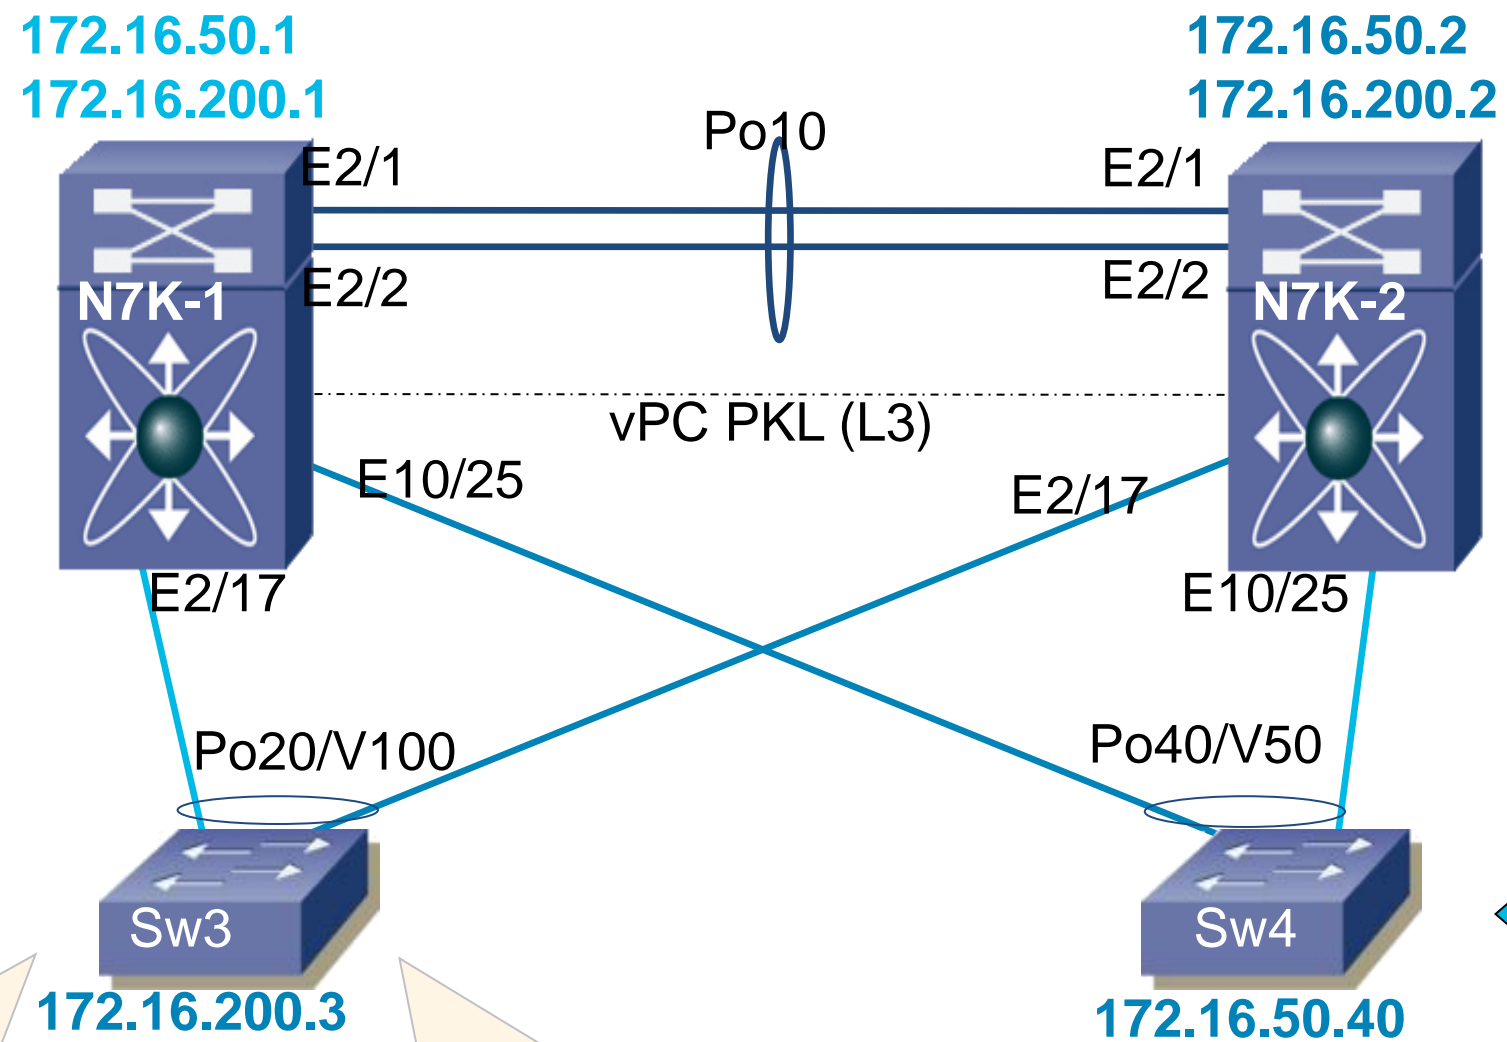
Troubleshooting Multicast Forwarding

Layer 3

- Forwarding information propagates from Supervisor to LC.
- Below process modules are involved.
 - **MRIB**
 - **MFDM PI**
 - **MFDM PD**
 - **MFIB**
 - **Hardware** (FIB TCAM, FIB DRAM, Adjacency Table)

Troubleshooting Multicast Forwarding

Example Topology




L2 Stream:
igmp join-group 239.0.0.200

L3 Stream:
igmp join-group 239.0.0.70

L2 Stream from VLAN200

L3 Stream from VLAN50



Troubleshooting Multicast Forwarding

L2 Software Entries– show ip igmp snooping

- Querier must be present; typically points towards mrouter

```
Campus_N7K1-VPC# show ip igmp snooping group 239.0.0.200 vlan 200
```

```
Type: S - Static, D - Dynamic, R - Router port, F - Fabricpath core port
```

| Vlan | Group Address | Ver | Type | Port list |
|------|---------------|-----|------|-----------|
| 200 | 239.0.0.200 | v2 | D | Po20 |

```
Campus_N7K1-VPC# show ip igmp snooping mrouter vlan 200
```

```
Type: S - Static, D - Dynamic, V - vPC Peer Link
```

```
I - Internal, F - Fabricpath core port
```

```
U - User Configured
```

| Vlan | Router-port | Type | Uptime | Expires |
|------|-------------|------|--------|----------|
| 200 | Po20 | D | 6d03h | 00:04:59 |
| 200 | Vlan200 | ID | 6d03h | 00:04:20 |

```
Campus_N7K1-VPC# show ip igmp snooping querier vlan 200
```

| Vlan | IP Address | Version | Expires | Port |
|------|--------------|---------|----------|--------------------|
| 200 | 172.16.200.1 | v2 | 00:03:33 | Vlan200 (internal) |

Troubleshooting Multicast Forwarding

L2 MFDM Software Entries– show forwarding distribution multicast

```
Campus_N7K1-VPC# show forwarding distribution ip igmp snooping vlan 200 group  
239.0.0.200
```

```
Vlan: 200, Group: 239.0.0.200, Source: 0.0.0.0
```

```
Outgoing Interface List Index: 5
```

```
Reference Count: 3
```

```
Platform Index: 0x7fed
```

```
Vpc peer link exclude flag set
```

```
Number of Outgoing Interfaces: 2
```

```
port-channel10
```

```
port-channel20
```

```
Campus_N7K1-VPC# show forwarding distribution multicast outgoing-interface-list L2 5
```

```
Outgoing Interface List Index: 5
```

```
Reference Count: 3
```

```
Platform Index: 0x7fed
```

```
Vpc peer link exclude flag set
```

```
Number of Outgoing Interfaces: 2
```

```
port-channel10
```

```
port-channel20
```


Troubleshooting Multicast Forwarding

L2 Hardware Entries– show system internal ip igmp snooping

```
Campus_N7K1-VPC# show system internal ip igmp snooping vlan 200 group 239.0.0.200
module 2
```

```
VDC: 4
```

| Vlan | Group | Source | Epoch | RID | DTL | hwptr |
|------|-------------|--------|-------|-----|--------|---------|
| 200 | 239.0.0.200 | | 1 | 5 | 0x7fed | 0x1e60f |

```
Campus_N7K1-VPC# show system internal ip igmp snooping vlan 200 group 239.0.0.200
module 10
```

```
VDC: 4
```

| Vlan | Group | Source | Epoch | RID | DTL | hwptr |
|------|-------------|--------|-------|-----|--------|---------|
| 200 | 239.0.0.200 | | 1 | 12 | 0x7fed | 0x1e60f |

```
Campus_N7K1-VPC# show system internal pixm info ltl 0x7fed
```

```
MCAST LTLs allocated for VDC:4
```

```
=====
```

| LTL | IFIDX | LTL_FLAG | CB_FLAG | MI[0] |
|--------|------------|----------|---------|-------|
| 0x7fed | 0x0000000c | 0x00 | 0x0002 | 0x002 |

```
-----
IFIDX          LTL
```

```
-----
```

| | |
|------|--------|
| Po20 | 0x0a42 |
| Po10 | 0x0a40 |

Troubleshooting Multicast Forwarding

L3 Software Entries – show ip pim, show ip mroute

```
Campus_N7K1-VPC# show ip mroute
```

```
IP Multicast Routing Table for VRF "default"
```

```
(* , 239.0.0.70/32) , uptime: 5d21h, pim ip  
Incoming interface: loopback0, RPF nbr: 1.1.1.1  
Outgoing interface list: (count: 1)  
Vlan200, uptime: 5d21h, pim
```

```
(172.16.50.123/32, 239.0.0.70/32) , uptime: 5d21h, ip pim mrib  
Incoming interface: Vlan50, RPF nbr: 172.16.50.123, internal  
Outgoing interface list: (count: 1)  
Vlan200, uptime: 5d21h, pim
```

```
Campus_N7K1-VPC# show ip pim interface brief
```

```
PIM Interface Status for VRF "default"
```

| Interface | IP Address | PIM DR Address | Neighbor Count | Border Interface |
|-----------|--------------|----------------|----------------|------------------|
| Vlan200 | 172.16.200.1 | 172.16.200.1 | 2 | no |
| Vlan50 | 172.16.50.1 | 172.16.50.1 | 2 | no |
| loopback0 | 1.1.1.1 | 1.1.1.1 | 0 | no |

Troubleshooting Multicast Forwarding

L3 Hardware Entries – show forwarding multicast route

```
Campus_N7K1-VPC# show forwarding multicast route group 239.0.0.70 source
172.16.50.123
slot 2
=====
(172.16.50.123/32, 239.0.0.70/32), RPF Interface: Vlan50, flags:
  Received Packets: 0 Bytes: 0
  Number of Outgoing Interfaces: 1
  Outgoing Interface List Index: 4
    Vlan200 Outgoing Packets:1100739 Bytes:870975864
slot 10
=====
(172.16.50.123/32, 239.0.0.70/32), RPF Interface: Vlan50, flags:
  Received Packets: 422488 Bytes: 334427892
  Number of Outgoing Interfaces: 1
  Outgoing Interface List Index: 4
    Vlan200 Outgoing Packets:370782 Bytes:293426795
```

```
Campus_N7K1-VPC# show forwarding multicast outgoing-interface-list 4
slot 2
Outgoing Interface List Index: 4
  Reference Count: 4
    Vlan200
```

Troubleshooting Multicast Forwarding

Data Collection – show tech ip pim, show tech forwarding

```
N7K1-VPC# show tech-support ip pim | grep "`show`"
```

```
`show running-config pim`  
`show ip pim group-range vrf all`  
`show ip pim interface vrf all`  
`show ip pim neighbor vrf all`  
`show ip pim route vrf all`  
[snip]
```

```
N7K1-VPC# show tech-support forwarding l3 multicast | grep "`show`"
```

```
`show forwarding multicast outgoing-interface-list`  
`show forwarding ip multicast route summary vrf all`  
`show system internal forwarding ip multicast route summary`  
`show forwarding ipv6 multicast route summary vrf all`  
`show system internal forwarding adjacency multicast`  
[snip]
```

```
N7K1-VPC# show tech-support ip multicast | grep "`show`"
```

```
`show tech-support ip igmp`  
`show running-config igmp`  
`show ip igmp route vrf all`  
[snip]  
`show tech-support ip msdp`
```

Agenda

- NXOS Architecture Introduction
- Built-in Troubleshooting Tools
- Troubleshooting Nexus 7000
 - CPU
 - Control-Plane – CoPP
 - Hardware
 - vPC
 - Unicast Layer 2 and Layer 3 Forwarding and ARP
 - Multicast Layer 2 and Layer 3 Forwarding
 - QoS

Troubleshooting QoS

Utilising Hardware Resources Effectively

- Nexus 7000 QoS is on by default and cannot be disabled
- System defined classification consistent across all VDCs
- Trust on all ports, marking can be modified with policy
- Similar queuing to prior platforms; added VoQ Fabric and MQC implementation



Troubleshooting QoS

Types of QoS

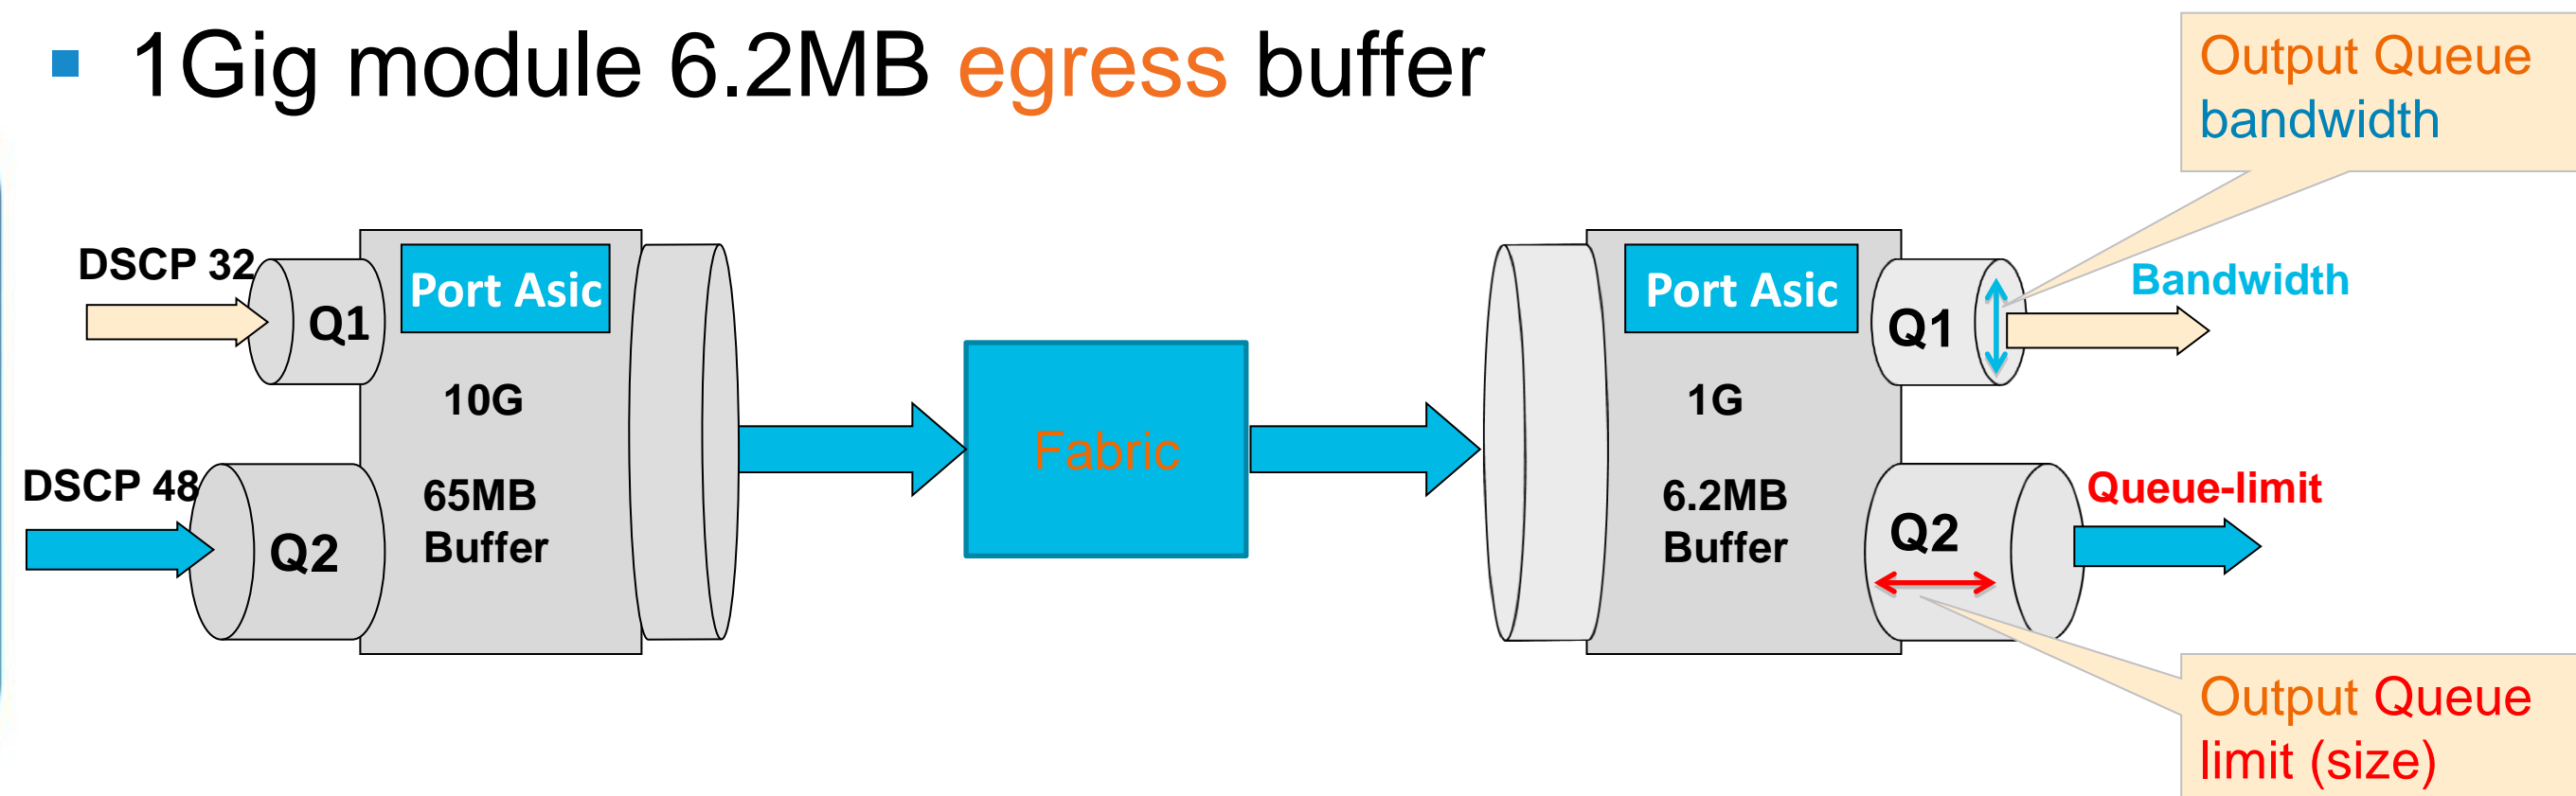
Types of QoS with Nexus 7K

- **Queuing**
 - Applied on and processed by Port ASIC
- **QOS**
 - Applied on forwarding engine ASIC
 - Processed during lookup by forwarding engine ASIC
- **Network**
 - Applied system wide, has significance within N7000 system
 - Used to influence the way traffic is handled within the system.

Troubleshooting QoS

How Drops Occur – Big Pipe to little pipe

- 10 Gig Module 65MB **ingress** buffer
- 1Gig module 6.2MB **egress** buffer



- Hardware drops due to buffer oversubscription
- Must ask: how much traffic is appropriate for **this** network?

Troubleshooting QoS

System wide Class-map configuration

- Queuing class-map names are static, based on port-type and queue
- Configurable only in default VDC
 - Changes apply to ALL ports of specified type in ALL VDCs
 - Changes are traffic disruptive for ports of specified type

Troubleshooting QoS

System wide Class-map configuration

N7K1(Config)#class-map type queuing match-any ?

1G Egress Port Type

1p3q4t-out-pq1
1p3q4t-out-q-default
1p3q4t-out-q2
1p3q4t-out-q3

1G Ingress Port Type

2q4t-in-q-default
2q4t-in-q1

10G Ingress Port Type

8q2t-in-q-default
8q2t-in-q1
8q2t-in-q2
8q2t-in-q3
8q2t-in-q4
8q2t-in-q5
8q2t-in-q6
8q2t-in-q7

10G Egress Port Type

1p7q4t-out-pq1
1p7q4t-out-q-default
1p7q4t-out-q2
1p7q4t-out-q3
1p7q4t-out-q4
1p7q4t-out-q5
1p7q4t-out-q6
1p7q4t-out-q7

Troubleshooting QoS

Classification and Queuing – show policy-mpa interface

```
N7K1-VPC# show policy-map int e2/17
```

```
Service-policy (queuing) output: default-out-policy
```

```
SNMP Policy Index: 301990010
```

```
Class-map (queuing): out-pq1 (match-any)
```

```
priority level 1
```

```
queue-limit percent 16
```

```
queue dropped pkts : 0
```

```
Class-map (queuing): out-q2 (match-any)
```

```
queue-limit percent 1
```

```
queue dropped pkts : 0
```

```
Class-map (queuing): out-q3 (match-any)
```

```
queue-limit percent 1
```

```
queue dropped pkts : 0
```

```
Class-map (queuing): out-q-default (match-any)
```

```
queue-limit percent 82
```

```
bandwidth remaining percent 25
```

```
queue dropped pkts : 0
```

Priority queueing on by default

Output drops will increment if queue limit exceeded

Troubleshooting QoS

Classification and Queuing – show queuing interface

```
N7K1-VPC# show queuing int e2/17
```

```
Interface Ethernet2/17 TX Queuing strategy: Weighted Round-Robin
```

```
Port QoS is enabled
```

```
Queuing Mode in TX direction: mode-cos
```

```
Transmit queues [type = 1p7q4t]
```

| Queue Id | Scheduling | Num of thresholds |
|----------------------|------------|-------------------|
| 1p7q4t-out-q-default | WRR | 04 |
| 1p7q4t-out-q2 | WRR | 04 |
| 1p7q4t-out-q3 | WRR | 04 |
| 1p7q4t-out-q4 | WRR | 04 |
| 1p7q4t-out-q5 | WRR | 04 |
| 1p7q4t-out-q6 | WRR | 04 |
| 1p7q4t-out-q7 | WRR | 04 |
| 1p7q4t-out-pq1 | Priority | 04 |

Scheduling and queue-limit read from hardware should match your configuration

```
WRR configuration read from HW
```

```
WRR bandwidth ratios: 25[1p7q4t-out-q-default] 15[1p7q4t-out-q2] 11[1p7q4t-out-q3] 11[1p7q4t-out-q4] 11[1p7q4t-out-q5] 11[1p7q4t-out-q6] 11[1p7q4t-out-q7]
```

```
queue-limit ratios configuration read from HW
```

```
queue-limit ratios: 78[1p7q4t-out-q-default] 1[1p7q4t-out-q2] 1[1p7q4t-out-q3] *1[1p7q4t-out-q4] *1[1p7q4t-out-q5] *1[1p7q4t-out-q6] *1[1p7q4t-out-q7] 16[1p7q4t-out-pq1]
```

```
* means unused queue with mandatory minimum queue-limit
```


Troubleshooting QoS

Hardware Drop – show hardware internal statistics

```
N7K1-VPC# show hardware internal error mod 2
```

```
|-----|  
| Device:R2D2                               Role:MAC                               Mod: 2 |  
| Last cleared @ Mon Jan 23 11:42:12 2012 |  
|-----|
```

Tail drop only typically applicable

```
Instance:5
```

| ID | Name | Value | Ports |
|-------|------------------------------|-------------------|---------------|
| 37920 | r2d2_tx_taildrop_drop_ctr_q2 | 0000000000000030 | 17,19,21,23 - |
| 37936 | r2d2_tx_taildrop_drop_ctr_q3 | 00000000000010650 | 17,19,21,23 - |

```
|-----|  
| Device:Naxos                               Role:MAC SECURITY                               Mod: 2 |  
| Last cleared @ Mon Jan 23 11:42:12 2012 |  
|-----|
```

Per port output read from hardware device "Naxos"

```
Instance:10
```

| ID | Name | Value | Ports |
|------|-------------------------|------------------|-------|
| 8194 | mac_egress_taildrop_ctr | 0000000000010680 | 21 - |
| 8195 | mac_egress_taildrop_ctr | 0000000000010680 | 23 - |

Summary

- NXOS offers a huge increase in available information over traditional data centre platforms
- Familiarising yourself with the tools available gives you the best chance to understand your own network
- Start capturing the right information at the right time



References

- Nexus 7000 Troubleshooting Guide

http://docwiki.cisco.com/wiki/Cisco_Nexus_7000_Series_NX-OS_Troubleshooting_Guide

- NXOS YouTube Intro

<http://www.youtube.com/user/nxs7000>

- NXOS vs IOS Comparison Guide

http://docwiki.cisco.com/wiki/Cisco_Nexus_7000_NX-OS/IOS_Comparison_Tech_Notes

- Ethalyzer Capture Filters

http://www.wireshark.org/docs/wsug_html_chunked/ChCapCaptureFilterSection.html

Q & A



Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2013 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 8 March 12:00pm-2:00pm



Cisco *live!* 365

Don't forget to activate your Cisco Live 365 account for access to all session material,

communities, and on-demand and live activities throughout the year. Log into your Cisco Live portal and click the "Enter Cisco Live 365" button.

www.ciscoliveaustralia.com/portal/login.wv

Cisco *live!*

