



Troubleshooting Cisco Catalyst 9000 Series Switches

Michel Peters - Technical Leader Engineering
BRKTRS-3090

CISCO *Live!*

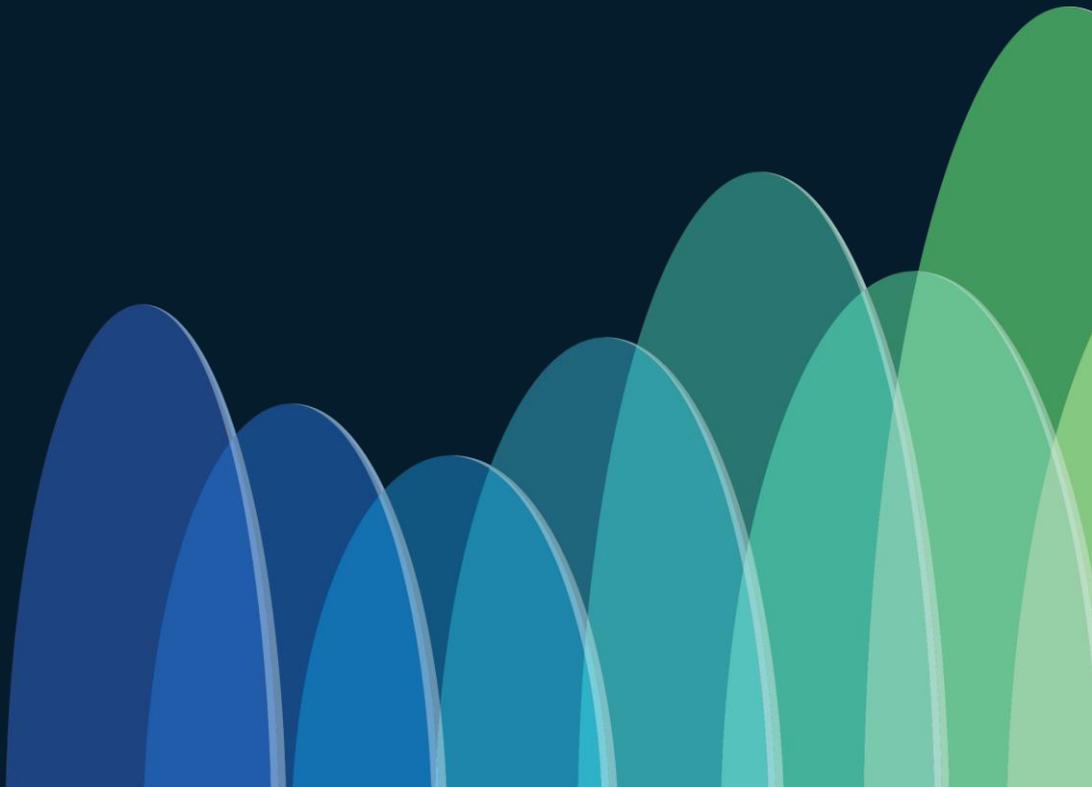


Agenda

- Introduction
- Memory and CPU
- Packet Drops
- Forwarding Verification



Introduction



Catalyst 9k family



	9200/9200L	9300/9300X	9400/9400X	9500/9500X	9600/9600X
Format	Stackable	Stackable	Chassis	Standalone	Chassis

9500X and 9600 Sup II based on Silicon One
 All others based on UADP Asics

Different models designed for
 different uses and performance



Catalyst IOS-XE Software release schedule

	Cupertino 17.9	Dublin 17.12	17.14	17.15	17.16
Next planned release	17.9.7	17.12.5	None	17.15.3	None
End of Software Maintenance	March 2025	March 2026	May 2025	-	-
End of Vulnerability Security Support	Sept 2026	Sept 2027	May 2025	-	-
Extended Maintenance Release	Yes	Yes	No	Yes	No

Recommended releases:

<https://www.cisco.com/c/en/us/support/docs/switches/catalyst-9300-series-switches/214814-recommended-releases-for-catalyst-9200-9.html>

Software Maintenance Upgrade(SMU)

- SMU's provide patches on top of IOS-XE
- - Hitless , No reload needed to apply on top of IOS-XE
- - Reboot, Reload is performed to apply SMU
- Apply with the command:

install add file <smu> activate commit

- Confirm with command :

show install summary

- Remove with :
install deactivate file <file>
- Do not just remove .bin from flash

File Information	Release Date	Size
Hitless/Recommended SMU, LISP Address-Resolution entry randomly disappears on 17.9.5 cat9k_iosxe.17.09.05.CSCwm47533.SPA.smu.bin Advisories 	16-Oct-2024	0.03 MB
Hitless/Recommended SMU, Catalyst 9K Switches: Unexpected reload in FED due to SIGFPE Arithmetic exception cat9k_iosxe.17.09.05.CSCwj55093.SPA.smu.bin Advisories 	26-Sep-2024	0.03 MB
Hitless/Recommended SMU, Hitless - September 2024 SMU PSIRT Bundle cat9k_iosxe.17.09.05.CSCwm01146.SPA.smu.bin Advisories 	25-Sep-2024	0.16 MB

Install Summary

```
9400_21#sh install summary
```

```
[ R0 R1 ] Installed Package(s) Information:
```

```
State (St): I - Inactive, U - Activated & Uncommitted,
```

```
          C - Activated & Committed, D - Deactivated & Uncommitted
```

```
-----  
Type  St   Filename/Version  
-----
```

```
IMG   C    17.12.04.0.4826
```

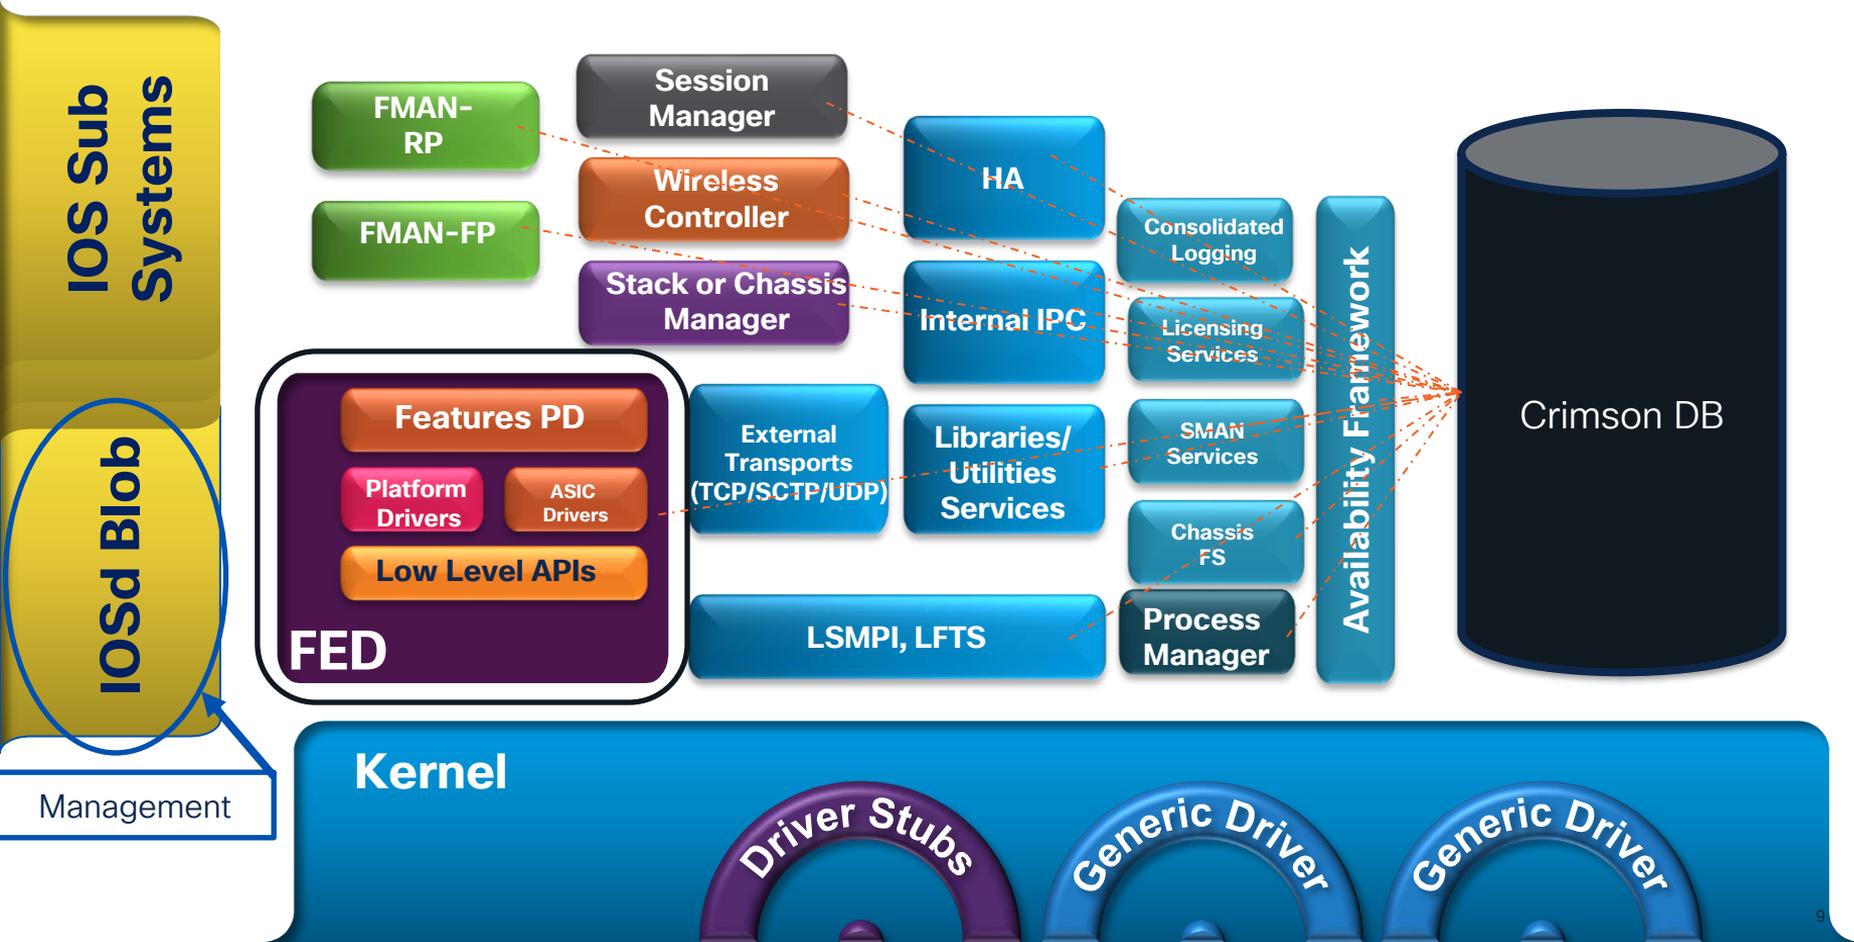
```
PKG   C    bootflash:C9800-SW-iosxe-wlc.17.09.05.SPA.bin
```

```
SMU   C    bootflash:cat9k_iosxe.17.12.04.CSCwk22887.SPA.smu.bin
```

```
SMU   D    bootflash:cat9k_iosxe.2025-01-29_05.38.0.CSCwn37241.SSA.smu.bin
```

- SMU's should be Activated and Committed
- Show install summary also shows packages (ewlc) and base image (IMG)

IOS-XE graphical overview



Show and Debug Commands

- Many of the features run primarily on IOSd
- Those features provide the regular debugs and show commands

```
Switch#show ip route  
Switch#show cdp neig  
Switch#debug ip packet 101
```

- Note: Debugs/show commands shows IOSd's view of the system
ex, debug ip packet shows ip packets IOSd has visibility on.
- Processes out of IOSd do communicate with IOSd,
ex, show access-session gets access-session information from
Session Manager Progress

IOS-XE Specific Show Commands

- Many processes outside IOSd are non-platform specific
Ex: Forwarding manager (fman) , WebUI (nginx), Wireless (wncd)
- Debug/show commands similar on all IOS-XE Devices

```
Switch#show platform software ip switch active R0 cef  
ASR_1k#show platform software ip rp active cef  
CSRv#show platform resources
```

- Specify location for command execution:
 - Active, Standby or switch number
 - RP/R0, Route Processor,
 - FP/F0 , Forwarding Processor

Catalyst 9k Platform Specific Show Commands

- Forwarding Engine Driver(FED) is the Catalyst 9000 platform specific layer
- FED layer commands under software fed or software fed:
show/set/debug platform hardware|software fed

```
9300#show platform software fed switch active ifm mappings
9300#show platform software fed switch 5 ifm mappings
9400#show platform software fed active ifm mappings
```

- Each Supervisor module and each switch in a stack runs its own Instance of FED.

Debugging outside IOSd process

- To facilitate debugging/logging trace logs are available for all processes using always on tracing
- Tracing levels set with granularity (default notice). Always on tracing
- Common processes: - smd, session manager
 - fed, forwarding engine driver
 - wncd, wireless process

```
Switch#set platform software trace smd switch active R0 dot1x-all verbose
Switch#show platform software trace level smd switch active R0 | inc dot1x
dot1x                               Notice
dot1x-all                           Verbose
dot1x-redun                          Notice
Switch#set platform software trace all notice ← set all back to default
```

Displaying trace logs

- Tracelog files are stored in `crashinfo:/logs` in binary format.
- Analyzing traces: `show logging process <process>`
- Bundle creation of binary traces: “request platform software trace archive”
- Appending command with internal could provide more detail

```
Edge_1#sh logging process smd | inc RADIUS
2022/06/06 23:24:03.268912 {smd_R0-0}{1}: [radius] [24732]: (info): RADIUS: Send Accounting-Request to
10.48.91.222:1813 id 1813/184, len 850
2022/06/06 23:24:03.268937 {smd_R0-0}{1}: [radius] [24732]: (info): RADIUS: authenticator e5 d1 b7 4d 8b e9 d5
06 - 14 b9 8d b6 8c 29 93 94
2022/06/06 23:24:03.268945 {smd_R0-0}{1}: [radius] [24732]: (info): RADIUS: Vendor, Cisco [26] 211
2022/06/06 23:24:03.268954 {smd_R0-0}{1}: [radius] [24732]: (info): RADIUS: Cisco
AVpair [1] 205 "cts-pac-opaque="
2022/06/06 23:24:03.268960 {smd_R0-0}{1}: [radius] [24732]: (info): RADIUS: Vendor, Cisco [26] 36
2022/06/06 23:24:03.268966 {smd_R0-0}{1}: [radius] [24732]: (info): RADIUS: Cisco
AVpair [1] 30 "dc-profile-name=Cisco-Device"
```

Make sure to enable traces in the right places. Debugs on IOSd might not be aware of full picture

Memory and CPU

Looking at the Kernel CPU information

- Kernel CPU utilization can be seen per switch/supervisor module

```
9300_54#show proc cpu platform sorted location switch active R0
CPU utilization for five seconds: 12%, one minute: 9%, five minutes: 5%
Core 0: CPU utilization for five seconds: 14%, one minute: 10%, five minutes: 5%
Core 7: CPU utilization for five seconds: 14%, one minute: 10%, five minutes: 6%
  Pid      PPid      5Sec      1Min      5Min  Status      Size  Name
-----
 14854    13498    51%      38%      22%  S           951436 linux_iosd-imag
 15322    15160    42%      30%      15%  S           313236 fed main event
   2960         2      8%       5%       2%  R            0 lsmpi-xmit
   2961         2      5%       4%       2%  S            0 lsmpi-rx
```

- linux_iosd-image process as seen on kernel is IOSd process
- Lsmpi processes are part of Linux Shared Memory Punt Interface
- Not all processes run on all switches.
- Changing the location will display cpu for standby or member switches

IOSd CPU utilization

```
9300_54#show proc cpu sort
```

```
CPU utilization for five seconds: 25%/1%; one minute: 22%; five minutes: 11%
```

PID	Runtime (ms)	Invoked	uSecs	5Sec	1Min	5Min	TTY	Process
40	9474304	54714487	173	22.95%	16.12%	8.15%	0	ARP Input
24	95502	12431773	7	0.07%	0.00%	0.00%	0	IPC Mcast Pendl
41	152951	13280277	11	0.07%	0.00%	0.00%	0	ARP Background
53	102409	12729984	8	0.07%	0.00%	0.00%	0	Dynamic ARP Insp
83	4759850	48620416	97	0.07%	0.03%	0.02%	0	IOSD ipc task
121	5602021	50896336	110	0.07%	0.04%	0.05%	0	Crimson config p
136	5112131	792052381	6	0.07%	0.03%	0.02%	0	L2 LISP Punt Pro

- Inside IOSd process many functions and processes run
Examples: BGP, RIP, CEF, ARP, UDLD, CDP, SSH, SNMP, Telnet
- CPU not involved in data plane forwarding
- Control Plane Policing enabled by default to throttle traffic

CPU utilization in IOSd

- CPU history gives an overview of CPU history for last minute, hour and day

```
9300_54#show processor cpu history
          444      444      443      444      444      444      444      444      44
1189111 18711 1188911 18821111883111 7841111188711 19871111189
100
 90
 50  **      **      **      **      **      **      ***      ***
 40  *#*     *#*     *#*     *#*     *#*     *#*     *#*     *#*
 30  ##*     ##*     ##*     ##*     ##*     ##*     ##*     ##*
 20  ##*     ##*     ##*     ##*     ##*     ##*     ##*     ##*
 10  ###     ###     ###     ###     ###     ###     ###     ###
    0.....5.....1.....1.....2.....2.....3.....3.....4.....4.....5.....5.....6
          0      5      0      5      0      5      0      5      0      5      0
          CPU% per minute (last 60 minutes)
          * = maximum CPU%   # = average
```

- Pattern of Interval and length of CPU spikes can point to trigger.
- Average high CPU more cause for concern then short spikes

Catching CPU spikes in IOSd

- CPU spikes often occur when not actively watching the switch.
- IOS allows CPU monitoring :
process cpu threshold type total rising <%> interval <s> falling <%> interval <s>
- When threshold are exceeded syslog is generated
- Syslog can be used by Embedded Event Manager to automate collection of data during time of failure

```
%SYS-1-CPURISINGTHRESHOLD: Threshold: Total CPU Utilization(Total/Intr): 48%/15%, Top 3 processes(Pid/Util): 40/32%, 94/0%, 105/0%
```

Top 3 processes,
ARP input is highest

```
9300_54#show proc cpu | inc 40 | 94 | 105 | PID
```

PID	Runtime(ms)	Invoked	uSecs	5Sec	1Min	5Min	TTY	Process
40	19640520	69264320	283	22.48%	20.61%	10.14%	0	ARP Input
94	652410	6428606	101	0.07%	0.00%	0.00%	0	PuntInject Keepa
105	3477489	6127722	567	0.00%	0.02%	0.02%	0	Crimson flush tr

Determining Cause of Inband traffic

```
9300_54#sh platform software fed switch active punt rates interfaces
```

```
Packets per second averaged over 10 seconds, 1 min and 5 mins
```

Active interfaces
sending to cpu

```
=====
```

Interface Name	IF_ID	Recv 10s	Recv 1min	Recv 5min	Drop 10s	Drop 1min	Drop 5min
GigabitEthernet1/0/11	0x00000013	5999	1707	2074	0	0	0
Vlan192	0x00000036	5999	1707	2074	0	0	0

```
=====
```

```
9300_54#sh platform software fed switch active punt cpuq rates
```

```
Packets per second averaged over 10 seconds, 1 min and 5 mins
```

Per Queue Statistics

```
=====
```

Q no	Queue Name	Rx 10s	Rx 1min	Rx 5min	Drop 10s	Drop 1min	Drop 5min
0	CPU_Q_DOT1X_AUTH	0	0	0	0	0	0
1	CPU_Q_L2_CONTROL	0	0	0	0	0	0
5	CPU_Q_FORUS_ADDR_RESOLUTION	6000	1874	2414	0	0	0

```
=====
```

Queue
information

Drops typically seen on Port
Asic not in FED punt path

Using Embedded Packet Capture

- EPC provides packet capture on interface or control plane level
- Data capture in hardware, traffic copied to EPC process for capturing
- Packet capture rate limited to protect CPU , good for visibility, not for performance captures

```
9300_54#monitor capture CL interface GigabitEthernet 1/0/11 in
9300_54#monitor capture CL match any
9300_54#monitor capture CL start
9300_54#monitor capture CL stop
9300_54#monitor capture CL export location flash:cl.pcap
9300_54#sh monitor capture file flash:cl.pcap display-filter arp
Starting the packet display ..... Press Ctrl + Shift + 6 to exit
 352  57.9 00:00:ca:fe:ca:fe ff:ff:ff:ff:ff:ff ARP 60 Who has 192.168.5.1? Tell 192.168.2.8
 353  57.9 00:00:ca:fe:ca:fe ff:ff:ff:ff:ff:ff ARP 60 Who has 192.168.5.2? Tell 192.168.2.8
 354  57.9 00:00:ca:fe:ca:fe ff:ff:ff:ff:ff:ff ARP 60 Who has 192.168.5.3? Tell 192.168.2.8
 355  57.9 00:00:ca:fe:ca:fe ff:ff:ff:ff:ff:ff ARP 60 Who has 192.168.5.4? Tell 192.168.2.8
 356  57.9 00:00:ca:fe:ca:fe ff:ff:ff:ff:ff:ff ARP 60 Who has 192.168.5.5? Tell 192.168.2.8
```

Where and what to capture

Export capture to file

Displays capture from flash or buffer

Determine packets hitting a specific CPU queue

- FED allows limited packet captures done in punt and inject path
- Allows filter to be setup to collect specific data.
- Use filter “fed.queue==<queue>” for collecting specific queue
- Detail and brief packet display. Detail contains internal forwarding information

```
9300_54#debug plat soft fed switch active punt packet-capture set-filter "fed.queue == 2"
9300_54#debug platform software fed switch active punt packet-capture [start|stop]
9300_54#show platform software fed switch active punt packet-capture brief
Punt packet capturing: disabled. Buffer wrapping: disabled
Total captured so far: 4096 packets. Capture capacity : 4096 packets
Capture filter : "fed.queue == 5"
----- Punt Packet Number: 3, Timestamp: 2023/05/31 13:14:45.888 -----
 interface : physical: GigabitEthernet1/0/11[if-id: 0x00000013], pal: Vlan192 [if-id: 0x00000036]
 metadata  : cause: 7 [ARP request or response], sub-cause: 1, q-no: 5, linktype: MCP_LINK_TYPE_IP
 ether hdr : dest mac: ffff.ffff.ffff, src mac: 0000.cafe.cafe
 ether hdr : ethertype: 0x0806 (ARP)
```

Metadata indicates
why it was punted

Control Plane Policing HW stats

- Control plane policing drops frames on port ASIC's to protect CPU
- Do not modify defaults without clear understanding of the impact

```

9300_1#show plat hardware fed switch active qos queue stats internal cpu policer
          CPU Queue Statistics
          (default) (set)
QId PlcIdx Queue Name           Enabled Rate Rate Queue Queue
          Drop (Bytes) Drop (Frames)
-----
 4     2   Routing Control             Yes  5400  5400    0      0
 5    14   Forus Address resolution       Yes  4000  65000 2212622868 345698644
          CPU Queue Policer Statistics
Policer Policer Accept Policer Accept Policer Drop Policer Drop
Index   Bytes         Frames         Bytes         Frames
-----
12      0             0             0             0
13     9562218      122704        453808601     7079824
          CPP Classes to queue map
PlcIdx  CPP Class                               : Queues
-----
 0      system-cpp-police-data                 : ICMP GEN/ BROADCAST/ ICMP Redirect/
14      system-cpp-police-forus                : Forus Address resolution/ Forus traffic/
    
```

Control Plane policing drops

Rate changed from Default rate

Platform Memory

```
9600_SVL#sh processes memory platform sorted location switch 1 R0
System memory: 15993924K total, 6935004K used, 9058920K free,
Lowest: 9025140K
```

Pid	Text	Data	Stack	Dynamic	RSS	Name
5373	256860	1488024	136	488	1488024	linux_iosd-imag
31175	203	885528	136	132984	885528	fed main event
26748	7236	295848	136	3808	295848	fman_rp
15847	8158	263112	136	4692	263112	fman_fp_image
27259	450	235088	136	4720	235088	dbm
28641	2583	228456	164	101856	228456	confd
30179	87	208392	136	7832	208392	pubd
25903	336	193308	136	388	193308	ndbmand
4272	790	177200	136	33732	177200	smand
18087	178	176024	136	13104	176024	sessmgrd

Diagram illustrating the mapping of process names to their respective components:

- linux_iosd-imag → IOSd
- fed main event → FED
- fman_rp → Forwarding manager
- fman_fp_image → Forwarding manager
- sessmgrd → SMD

- Kernel memory utilization is available per switch
- linux_iosd-image process is IOSd
- Resident Set Size(RSS), memory occupied by each Process

Per process memory allocation

```
9300_54#show platform software memory smd switch active R0 brief
```

module	allocated	requested	allocs	frees
-----	-----	-----	-----	-----
Summary	10651060	10392132	461338	445155
OBJ_SMTEST_EVENTLOG	6528072	6496040	2002	0
AAA_CHUNK_ATTR_SUBLIST	1085842	1085762	5	0
OBJ_DC	625435	429323	12260	3
chunk	343897	342425	97	5
OBJ_RCLSRV_SEND_BUF_CH	263045	262981	4	0
smd	213359	213295	4	0
AAA_ATTR_LIST_HANDLE_I	133193	133145	3	0
process	133172	133124	3	0
eventutil	120685	119757	2074	2016
OBJ_EPM_CACHE_HASH_ELE	97688	97128	44	9
AAA_CHUNK_ATTR_HEADER	78225	78145	5	0
OBJ_EAP_ALLOC_HDL	65552	65536	1	0
AAA_MLIST_ID	65552	65536	1	0

Detailed memory usage per process , details vary per process

IOSd Memory

```
9600_SVL#sh processes memory sorted holding
```

```
Processor Pool Total: 2959189564 Used: 687714160 Free: 2271475404
reserve P Pool Total: 102404 Used: 88 Free: 102316
lsmpi_io Pool Total: 6295128 Used: 6294296 Free: 832
```

PID	TTY	Allocated	Freed	Holding	Getbufs	Retbufs	Process
0	0	488208064	77059624	381785696	0	0	*Init*
238	0	137113496	1017784	136266568	0	0	IP ARP Adjacency
82	0	50195904	2966104	27577736	0	0	IOSD ipc task
565	0	720962728	696333696	24623088	0	0	DHCPD Receive
4	0	42371936	17822928	22358600	0	0	RF Slave Main Th

- IOSd runs as a process , but does still provides some kernel features like memory management for all processes running inside IOSd
- Processor Pool: Pool for Processes on IOSd
- lsmpi_io : Linux Shared Memory Punt Interface memory, IO buffers

IOSd Memory Allocator

```
9600_SVL#sh processes memory 565
Tracekey : 1#01327ced3d92c0f78c7959f154fa0114
Process ID: 565
Process Name: DHCPD Receive
Total Memory Held: 24623088 bytes
```

```
Processor memory Holding = 24623088 bytes
size = 10549344, count = 321, pc = :5644D0D8C000+934D623
size = 10450752, count = 318, pc = :5644D0D8C000+548EF85
size = 1478880, count = 45, pc = :5644D0D8C000+861423D
size = 1032040, count = 120, pc = :5644D0D8C000+8615EB3
size = 985920, count = 30, pc = :5644D0D8C000+8608243
```

- Allocator PC indicates what line of code allocated the memory
- Not all increased in memory are leaks, memory might get released again or functions might just need more memory
- Monitor over time and multiple boxes to determine possible pattern

Packet Drops



Checking interface status

Err-disabled => check
show interface status errdisable
or show log for cause

Vlan column shows routed(L3) ,
trunk or base vlan

Half Duplex 1Gb/s, possible
duplex mismatch?

```
9500_1#show interfaces status | ex notc
Port          Name          Status      Vlan      Duplex  Speed  Type
Twe1/0/7      Twe1/0/7      connected   routed    a-full  a-1000 10/100/1000BaseTX SFP
Twe1/0/12     Twe1/0/12     err-disabled 1          full    10G    SFP-10GBase-CU1M
Twe1/0/15     Twe1/0/15     connected   routed    full    10G    SFP-10GBase-CU3M
Twe1/0/20     Twe1/0/20     connected   4094     full    10G    SFP-10GBase-CU1M
Twe1/0/21     Twe1/0/21     connected   routed    a-half  a-1000 10/100/1000BaseTX SFP
Twe1/0/22     Twe1/0/22     disabled    1          auto    auto    10/100/1000BaseTX SFP
Hu1/0/26      Hu1/0/26      connected   trunk     full    40G    QSFP 40G CU3M
Hu1/0/27      Hu1/0/27      connected   4094     full    100G   QSFP 100G CU2M
```

Show interface status command gives a quick overview of interface status

Ethernet Statistics

```
Switch#show controllers ethernet-controller gi 5/0/48
Transmit GigabitEthernet5/0/48 Receive
1562496684 Total bytes 2968958225 Total bytes
 5032561 Unicast frames 6004241 Unicast frames
700808558 Unicast bytes 1807110661 Unicast bytes
1269484 Multicast frames 2789759 Multicast frames
861688062 Multicast bytes 1161847500 Multicast bytes
 1 Broadcast frames 1 Broadcast frames
 0 Cos 0 Pause frames 0 Cos 0 Pause frames
1236978 Minimum size frames 871517 Minimum size frames
1892419 65 to 127 byte frames 2181611 65 to 127 byte frames
1941967 128 to 255 byte frames 2712229 128 to 255 byte frames
685594 256 to 511 byte frames 1260418 256 to 511 byte frames
20261 512 to 1023 byte frames 900135 512 to 1023 byte frames
524827 1024 to 1518 byte frames 868091 1024 to 1518 byte frames
 0 8192 to 16383 byte frames 0 8192 to 16383 byte frames
 0 16384 to 32767 byte frame 0 16384 to 32767 byte frame
 0 > 32768 byte frames 0 > 32768 byte frames
 0 Late collision frames 0 SymbolErr frames
 0 Excess Defer frames 0 Collision fragments
 0 Good (1 coll) frames 0 ValidUnderSize frames
 0 Good (>1 coll) frames 0 InvalidOverSize frames
 0 Deferred frames 0 ValidOverSize frames

LAST UPDATE 361 msec AGO
```

Ethernet controller statistics give more detailed port statistics

Frame size distribution

Error statistics

Verifying link utilization

Input/output rates show average over 5 minutes (default). Traffic might be bursty in nature

```
Switch#show interfaces | inc line|rate
```

```
Vlan1 is up, line protocol is up , Autostate Enabled
 5 minute input rate 0 bits/sec, 0 packets/sec
 5 minute output rate 0 bits/sec, 0 packets/sec
GigabitEthernet0/0 is administratively down, line protocol is down
 5 minute input rate 0 bits/sec, 0 packets/sec
 5 minute output rate 0 bits/sec, 0 packets/sec
GigabitEthernet1/0/11 is up, line protocol is up (connected)
 30 seconds input rate 575000 bits/sec, 958 packets/sec
 30 seconds output rate 126975000 bits/sec, 10473 packets/sec
GigabitEthernet1/0/12 is down, line protocol is down (notconnect)
 5 minute input rate 0 bits/sec, 0 packets/sec
 5 minute output rate 0 bits/sec, 0 packets/sec
GigabitEthernet1/0/13 is down, line protocol is down (notconnect)
 5 minute input rate 0 bits/sec, 0 packets/sec
 5 minute output rate 0 bits/sec, 0 packets/sec
```

load-interval has a range of 30-600 seconds

```
Switch#show controllers utilization
```

Port	Receive Utilization	Transmit Utilization
Gi1/0/1	0	0
Gi1/0/11	0	92
.		
Gi1/0/24	0	0
Tel1/1/1	0	0
Tel1/1/4	9	0

Total Ports : 33

Bandwidth in %
Current load

```
Total Ports Receive Bandwidth Percentage Utilization : 0
Total Ports Transmit Bandwidth Percentage Utilization : 0
Average Switch Percentage Utilization : 0
```

Tail Drops

```
9300_54#show interfaces gigabitEthernet 1/0/11 | inc output drops
  Input queue: 0/2000/0/0 (size/max/drops/flushes); Total output drops: 1277
9300_54#show controllers ethernet-controller gig 1/0/11 | inc Excess Def
      1277 Excess Defer frames                0 Collision fragments
SNMP:
SNMPv2-SMI::enterprises.9.2.2.1.1.27.8 = INTEGER: 1277
```

ith.

- Buffer allocation per class can be changed inside service-policy
- Global multiplier to increase buffers up to 1200%
qos softmax-queue-multiplier <percentage>
Only adjust when needed, drops are not always a problem
- Ensure proper classification of traffic and apply correct policies,
not all traffic is equal, file transfers don't mind buffering, voice does

QoS Hardware configuration

```
9300_45#sh plat hard fed switch active qos queue config interface gi 1/0/11
Asic:0 Core:1 DATA Port:11 GPN:1 LinkSpeed:0x1
```

DTS	Hardmax	Softmax	PortSMin	GlblSMin	PortStEnd			
0 1 5	200 12 3200	5 500 0 0 6 9600						
1 1 4	0 13 4800	5 750 2 300 6 9600						
Priority	Shaped/shared	weight	shaping_step	sharpedWeight				
0 0	Shared	50	0	0				
1 0	Shared	75	0	0				
Port	Port	Port	Port					
Priority	Shaped/shared	weight	shaping_step					
2	Shaped	254	255					
Weight0	Max_Th0	Min_Th0	Weigth1	Max_Th1	Min_Th1	Weight2	Max_Th2	Min_Th2
0 0	2709	0	0	3028	0	0	3400	0
1 0	3825	0	0	4275	0	0	4800	0

Hardmax : Reserved
Softmax : Global pool

Queue mode:
shaped or shared
Queue limit for shaping
Step/weight * speed

Drop thresholds for
queue/threshold
In Buffers (256 byte)

- QoS configured using service policies on interfaces
- Applied service-policy translated into Hardware settings that match HW capabilities

QoS hardware statistics

```
9300_45#sh platform hardware fed switch active qos queue stats interface gigabitEthernet 1/0/11
```

AQM Global counters

```
GlobalHardLimit: 7976 | GlobalHardBufCount: 0
GlobalSoftLimit: 11872 | GlobalSoftBufCount: 0
```

Highest number of buffers in use by interface since last issuing command

```
High Watermark Soft Buffers: 429 <--- clear on read
```

Asic:0 Core:1 DATA Port:10 Hardware Enqueue Counters

Q	Buffers (Count)	Enqueue-TH0 (Bytes)	Enqueue-TH1 (Bytes)	Enqueue-TH2 (Bytes)	Qpolicer (Bytes)
0	200	0	385820	46085690	0
1	0	0	0	0	0

Current buffer count per queue

Asic:0 Core:1 DATA Port:0 Hardware Drop Counters

Q	Drop-TH0 (Bytes)	Drop-TH1 (Bytes)	Drop-TH2 (Bytes)	SBufDrop (Bytes)	QebDrop (Bytes)	QpolicerDrop (Bytes)
0	0	0	412312	0	0	0
1	0	0	0	0	0	0

- At Asic level there are 8 Queues/3 Thresholds
- Enqueue/Drop Counters available per queue/per threshold
- Buffers (count) show currently assigned buffers to Queue (256 bytes)
- To enable high water mark counter monitoring.
set platform hardware fed switch active qos port-monitor interface <if>

CISCO Live!

Doppler ASIC packet forwarding drop counters

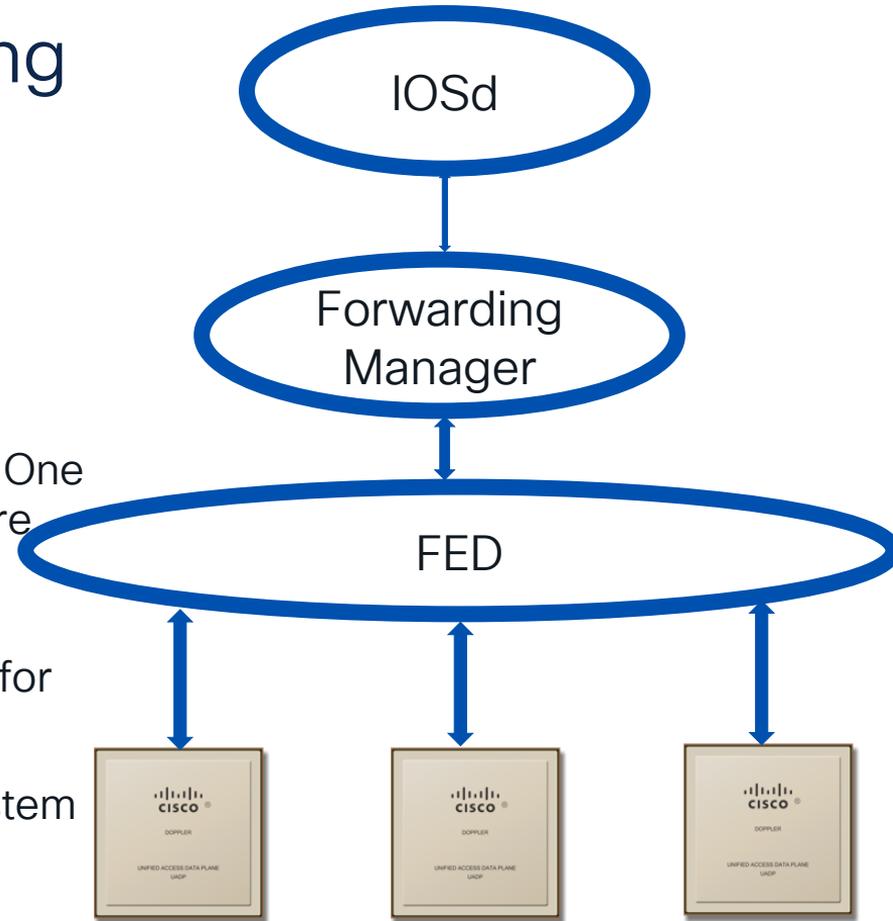
```
9300_45#sh platform hardware fed switch active fwd-asic drops exceptions
***EXCEPTION STATS ASIC INSTANCE 0 (asic/core 0/0)***
=====
Asic/core |          NAME          |   prev   |  current  |  delta
=====
0  0  NO_EXCEPTION           35364016  35364108  92
0  0  IPV4_CHECKSUM_ERROR    0         0         0
0  0  ROUTED_AND_IP_OPTIONS_EXCEPTION 2         2         0
0  0  CTS_FILTERED_EXCEPTION 0         0         0
0  0  AUTH_DRIVEN_DROP      0         0         0
0  0  PKT_DROP_COUNT        0         3732      3732
0  0  ALLOW_DOT1Q_EXCEPTION_COUNT 0         0         0
0  0  ALLOW_PRIORITY_TAGGED_EXCEPTION_COUNT 0         0         0
0  0  IGR_EXCEPTION_L5_ERROR 0         363       363
0  0  IP_UNICAST_TTL_REACHED_ZERO 0         0         0
0  0  MISC_FATAL_ERROR      0         0         0
```

- Every packet passing through Port Asic gets parsed by the port-asic's on both receive and transmit side
- Exception drops are counted per Asic, not per port.

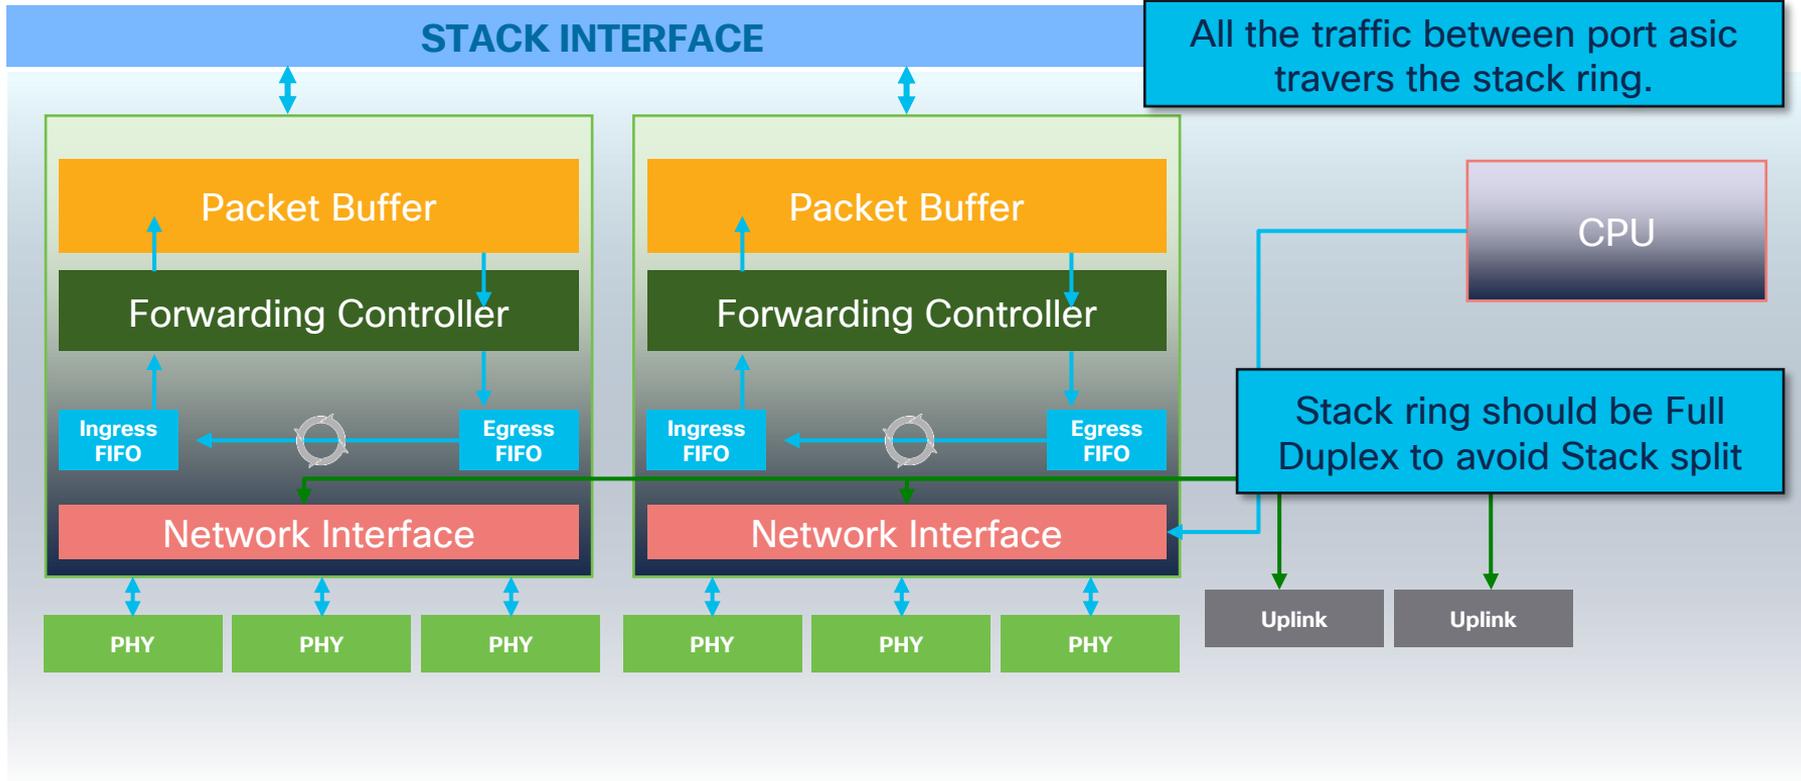
Forwarding Verification

Troubleshooting Forwarding

- Many layers are involved in forwarding
- Troubleshoot from top (IOSd) down toward the hardware layer
- 9500X and 9600 (Supervisor 2) use Silicon One
All other models based on UADP architecture
- Packet processing happens on port asics.
- CPU only involved in programming the asic for forwarding.
- Understand the packet flow through the system



Catalyst 9200/9300 Stackable Switches



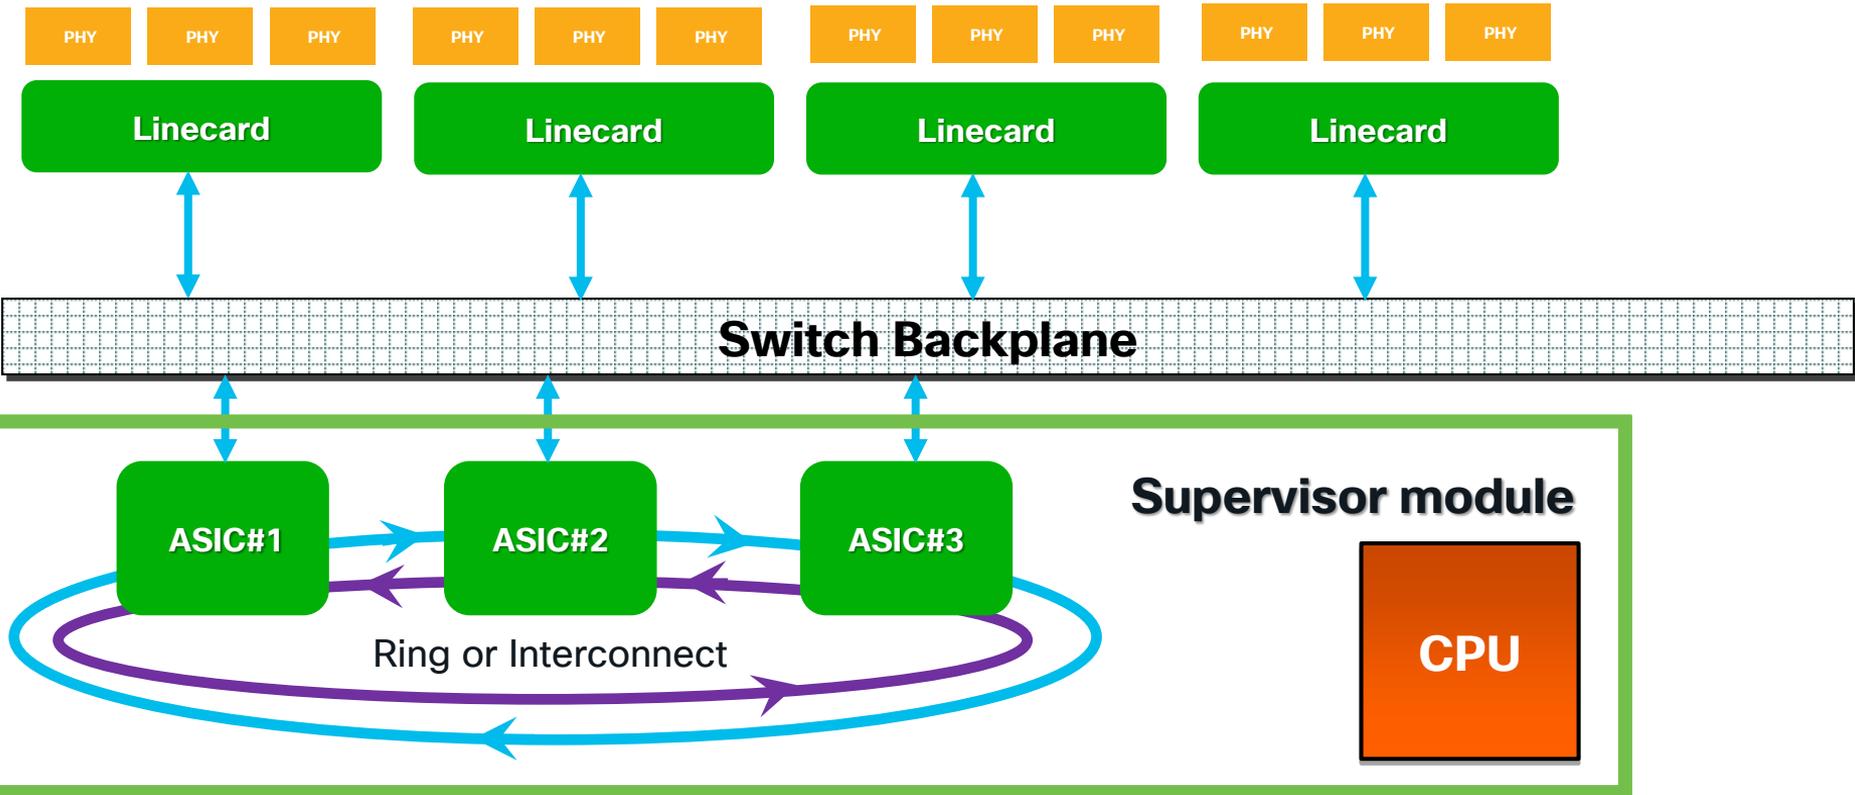
Switch Stacks

```
9300_10#sh switch stack-ports summary
```

Sw#/Port#	Port Status	Neighbor/Port	Cable Length	Link OK	Link Active	Sync OK	#LinkOK	In Loopback
1/1	OK	4/2	100cm	Yes	Yes	Yes	1	No
1/2	OK	2/1	50cm	Yes	Yes	Yes	1	No
2/1	OK	1/2	50cm	Yes	Yes	Yes	2	No
2/2	OK	3/1	50cm	Yes	Yes	Yes	1	No
3/1	OK	2/2	50cm	Yes	Yes	Yes	1	No
3/2	OK	4/1	50cm	Yes	Yes	Yes	2	No
4/1	OK	3/2	50cm	Yes	Yes	Yes	1	No
4/2	OK	1/1	100cm	Yes	Yes	Yes	2	No

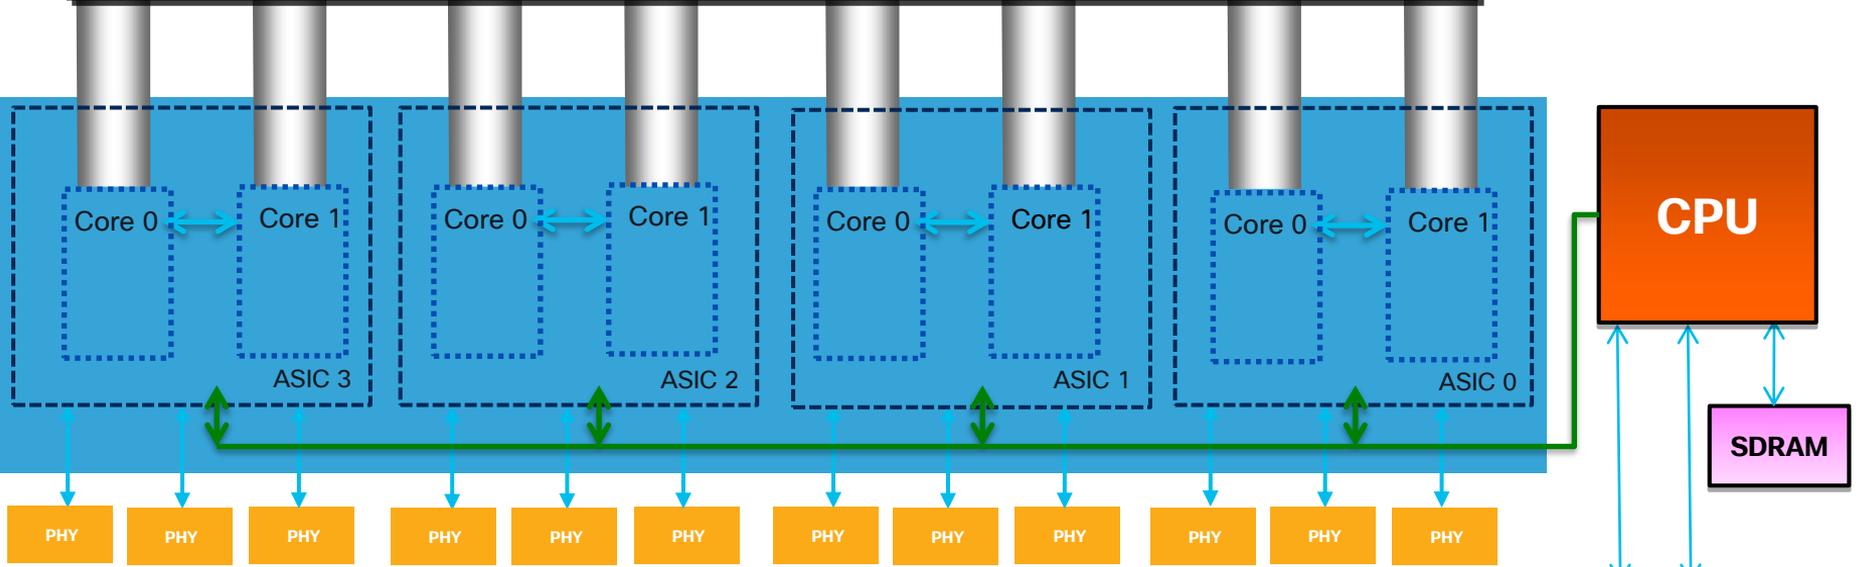
- Ensure all stack ports are connected and Link Ok to ensure stack is running with redundancy and full bandwidth
- High Number of #linkOK or inability to read cable length might indicate bad cable/stack ports
- Show tech-support stack to gather useful commands

Catalyst 9400/9600 Chassis based



Catalyst 9500 Fixed Switches

Internal High Speed Switch Stack Ring



9500 does *not* support stack connectors.
Does support StackWise Virtual

Interface Internal Mappings

Interface to ASIC mapping important to understand data flows

```
9300_1#show platform software fed switch active ifm mappings
```

Interface	IF_ID	Inst	Asic	Core	Port	SubPort	Mac	Cntx	LPN	GPN	Type	Active
GigabitEthernet1/0/1	0x8	1	0	1	0	0	26	6	1	1	NIF	Y
GigabitEthernet1/0/2	0x4c	1	0	1	1	0	6	7	2	2	NIF	Y
GigabitEthernet1/0/3	0x4d	1	0	1	2	0	28	8	3	3	NIF	Y

Internally used interface addressing:

- LPN : Local Port Number
- GPN : Global Port Number
- IF_ID : Interface Identification, used for many fed CLI
- Type : Type of interface, NIF = Network Interface
- Inst : Instance : ASIC + Core
- Port : Asic Ports
- Active : Is Interface Active

Troubleshoot between ports on same ASIC when possible

IFM Mappings logical interface

```
9600_SVL#show platform software fed switch active ifm interfaces vlan
Interface                IF_ID                State
-----
Vlan1                    0x0005ffff          READY
Vlan123                  0x00420012          READY
Vlan900                  0x00420010          READY
9600_SVL#show platform software fed switch active ifm interfaces svi
Interface                IF_ID                State
-----
Vlan123                  0x0000013b          READY
9600_SVL#show platform software fed switch active ifm interfaces tunnel
Interface                IF_ID                State
-----
Tunnel1                  0x00000143          READY
```

- An unique IF_ID is assigned to every logical and physical.
- Every type of interface has its unique IF_ID.
Ex:Layer 2 Vlan has a different IF_ID as the Layer 3 SVI

IF_ID more detail

```
9600_SVL#show platform software fed switch active ifm if-id 0x138
```

```
Interface IF_ID          : 0x00000000000000138
Interface Name          : FortyGigabitEthernet2/1/0/1
Interface Block Pointer : 0x7f33b89fd918
Interface Block State   : READY
Interface State         : Enabled
Interface Status        : ADD, UPD
Interface Type          : ETHER
Port Type               : ROUTE PORT
Port Location          : LOCAL
Slot                   : 15
Unit                   : 0
Slot Unit              : 1
SNMP IF Index          : 148
GPN                    : 577
Port Handle             : 0xb7000104
IPv4 MTU              : 9100
IPv6 MTU               : 0
IPv4 VRF ID          : 0x0
IPv6 VRF ID          : 0x0
Protocol flags         : 0x0007 [ ipv4 ipv6 pim_ipv4 ]
```

Specifying the IF-ID gives
verbose information

Information available varies
depending on Interface Type

One VRF ID per VRF, VRF ID
0 is Global Routing Table

Layer 2 Forwarding. Verifying HW STP state

- Show spanning tree gives IOSd view of Spanning Tree
- Hardware forwarding states can be checked *per switch* on FED layer
- Outputs will show what interface are in forwarding state and if traffic is tagged or untagged
- Flood list indicates what Ports will receive flooded traffic on this switch

```
9300_54#sh platform hardware fed switch active vlan 192 egress
VLAN STP State in hardware
vlan id is:: 192
Interfaces in forwarding state: : Gi1/0/11 (Untagged) , Te1/1/1 (Untagged) , Te1/1/4 (Tagged)
9300_54#sh platform hardware fed switch active vlan 192 ingress
VLAN STP State in hardware
vlan id is:: 192
Interfaces in forwarding state: : Gi1/0/11 (Untagged) , Te1/1/1 (Untagged) , Te1/1/4 (Tagged)
flood list: : Gi1/0/11, Te1/1/1, Te1/1/4
```

Layer 2 Forwarding, IOSd mac address tables

```
9300_54#show mac address vlan 192
      Mac Address Table
-----
Vlan  Mac Address      Type      Ports
-----
  192  0000.cafe.cafe    DYNAMIC   Gi1/0/11
  192  0013.c3c1.0d89    DYNAMIC   Gi1/0/11
  192  0050.5693.5e70    DYNAMIC   Te1/1/4
  192  a0f8.4910.2dd3    STATIC    Vl192
  192  d0ec.35c9.d353    DYNAMIC   Te2/1/1
Total Mac Addresses for this criterion: 5
```

- Show mac address-table command shows overview of all mac addresses dynamic, static and system mac addresses.
- IOSd shows global view.
Each FED maintains its own Mac Address Table

FED MATM Mac Address Table

```

9300_1#sh platform software fed switch 1 matm macTable vlan 100
VLAN  MAC                               Type  Seq#  EC_Bi  Flags  machandle  siHandle  diHandle  *a_time  *e_time  ports
-----
192    0000.cafe.cafe                          0x1   71718    0      0  0x7f2348  0x7f2198  0x7f2478  300      112  GigabitEthernet1/0/11
192    d0ec.35c9.d353                          0x101 71719    0      0  0x7f2328  0x7f2148  0x7f2588  300      14   TenGigabitEthernet2/1/1
192    0013.c3c1.0d89                          0x1   71720    0      0  0x7f23d8  0x7f2378  0x7f4278  300      11   GigabitEthernet1/0/11
192    a0f8.4910.2dd3                          0x8002 0        0      64  0x7f2308  0x7f21c8  0x5234    0        0   Vlan192
192    0050.5693.5e70                          0x1   12772    0      0  0x7f23d8  0x7f2358  0x7f23e8  300      13   TenGigabitEthernet1/1/4

*a_time=aging_time(secs)  *e_time=total_elapsed_time(secs)
Type:
MAT_DYNAMIC_ADDR          0x1  MAT_STATIC_ADDR          0x2  MAT_CPU_ADDR             0x4  MAT_DISCARD_ADDR        0x8
MAT_ALL_VLANS             0x10 MAT_NO_FORWARD            0x20 MAT_IPMULT_ADDR          0x40 MAT_RESYNC                0x80
MAT_DO_NOT_AGE           0x100 MAT_SECURE_ADDR          0x200 MAT_NO_PORT              0x400 MAT_DROP_ADDR            0x800
MAT_DUP_ADDR             0x1000 MAT_NULL_DESTINATION     0x2000 MAT_DOT1X_ADDR          0x4000 MAT_ROUTER_ADDR          0x8000
MAT_WIRELESS_ADDR        0x10000 MAT_SECURE_CFG_ADDR      0x20000 MAT_OPQ_DATA_PRESENT    0x40000 MAT_WIRED_TUNNEL_ADDR    0x80000
MAT_DLR_ADDR             0x100000 MAT_MRP_ADDR              0x200000 MAT_MSRRP_ADDR          0x400000 MAT_LISP_LOCAL_ADDR      0x800000
MAT_LISP_REMOTE_ADDR     0x1000000 MAT_VPLS_ADDR             0x2000000
    
```

- Type Field indicates the type of mac address using bitmap.
- Sequence number of an entry changing would indicate relearning
- Owning FED responsible for aging entries. MAT_DO_NOT_AGE bit set in type for non owning FED instances



Layer 3 Forwarding. Routing protocols

```
9300_20#sh ip route 10.48.13.0
Routing entry for 10.48.13.0/24
  Known via "isis", distance 115, metric 30, type level-2
  Redistributing via isis
  Last update from 172.18.255.9 on TenGigabitEthernet1/0/47, 00:40:35 ago
  Routing Descriptor Blocks:
  * 172.18.255.9, from 172.31.254.18, 00:40:35 ago, via TenGigabitEthernet1/0/47
    Route metric is 30, traffic share count is 1

9300_20#show ip arp 172.18.255.9
Protocol Address Age (min) Hardware Addr Type Interface
Internet 172.18.255.9 41 0072.78f8.aed6 ARPA TenGigabitEthernet1/0/47
```

- Verify routing table and next hop information with known correct path
- Confirm the route shown is covered by the correct subnet and learned by the expected routing protocol.

Cisco Express Forwarding

- CEF shows the Forwarding Information Base showing how traffic will be forwarded through the system.
- Sources show where the route is learned from.
RIB(Routing Information Base)
- Use Detail or Internal keywords to get even more detail

```
9300_20#show ip cef 10.48.13.0 internal
10.48.13.0/24, epoch 4, RIB[I], refcnt 6, per-destination sharing
sources: RIB
feature space:
  IPRM: 0x00028000
  Broker: linked, distributed at 4th priority
ifnums:
  TenGigabitEthernet1/0/47(144): 172.18.255.9
  path list 70C9760F41F0, 327 locks, per-destination, flags 0x4D [shble, hvsh, rif, hwn]
  path 70C978450268, share 1/1, type attached nexthop, for IPv4
    nexthop 172.18.255.9 TenGigabitEthernet1/0/47, IP adj out of TenGigabitEthernet1/0/47,
addr 172.18.255.9 70C9761536E8
output chain:
  IP adj out of TenGigabitEthernet1/0/47, addr 172.18.255.9 70C9761536E8
```

Cisco Express Forwarding

```
9300_20#show adjacency 172.18.255.9 detail
Protocol Interface Address
IP TenGigabitEthernet1/0/47 172.18.255.9(174)
0 packets, 0 bytes
sourced in sev-epoch 9
Encap length 14
007278F8AED6007278F813550800
L2 destination address byte offset 0
L2 destination address byte length 6
Link-type after encap: ip
ARP

9300_20#sh ip arp TenGigabitEthernet 1/0/47
Protocol Address Age (min) Hardware Addr Type Interface
Internet 172.18.255.10 - 0072.78f8.1355 ARPA TenGigabitEthernet1/0/47
Internet 172.18.255.9 221 0072.78f8.aed6 ARPA TenGigabitEthernet1/0/47
```

- CEF adjacency displays next hop information.
- Layer 2 rewrite information displayed to verify Destination and Source Mac Address of rewritten packet

Forwarding Manager (FP/RP)

- Forwarding Manager should show the prefix on both Standby/Active RP and on all FP (stack members or supervisor modules)
- aom id shown with FP can be used to display prefix in object manager

```
9300_20#sh platform software ip switch active F0 cef prefix 10.48.13.0/24 detail
Forwarding Table
10.48.13.0/24 -> OBJ_ADJACENCY (0x4a), urpf: 81
Prefix Flags: unknown
aom id: 11562925, HW handle: (nil) (created)
```

```
9300_20#sh platform software object-manager switch active f0 object 11562925
Object identifier: 11562925
Description: PREFIX 10.48.13.0/24 (Table id 0)
Obj type id: 71
Obj type: route-pfx
Status: Done, Epoch: 0, Client data: 0x452c6ae8
```

Forwarding Manager (FP/RP)

- Forwarding Manager should show the prefix on both Standby/Active RP and on all FP (stack members or supervisor modules)
- aom id shown with FP can be used to display prefix in object manager

```
9300_20#show platform software adjacency switch active f0 index 0x4a
Adjacency id: 0x4a (74)
Interface: TenGigabitEthernet1/0/47, IF index: 144, Link Type: MCP_LINK_IP
Encap: 0:72:78:f8:ae:d6:0:72:78:f8:13:55:8:0
Encap Length: 14, Encap Type: MCP_ET_ARPA, MTU: 9100
Flags: no-l3-inject
Nexthop addr: 172.18.255.9
IP FRR MCP_ADJ_IPFRR_NONE 0
aom id: 8386846, HW handle: (nil) (created)
```

Forwarding Manager (FP/RP)

- Looking at downlinks, parents or children of an object reveals related object

```
9300_20#show platform software object-manager switch active f0 object 8386846
Object identifier: 8386846
  Description: adj 0x4a, Flags None
  Obj type id: 40
  Obj type: adj
  Status: Done, Epoch: 0, Client data: 0x4546afd8
9300_20#sh plat soft object-manager switch active f0 object 8386846 downlinks
Object identifier: 987
  Description: PREFIX 0.0.0.0/0 (Table id 0)
  Status: Done
Object identifier: 11562925
  Description: PREFIX 10.48.13.0/24 (Table id 0)
  Status: Done
..
```

Error Objects

- Error object that do not get cleared could suggest forwarding issues
- Lookup objects to verify relationship to potential issues

```
9600X_50#show platform software object-manager f0 statistics
Forwarding Manager Asynchronous Object Manager Statistics
Object update: Pending-issue: 2, Pending-acknowledgement: 0
Batch begin:   Pending-issue: 0, Pending-acknowledgement: 0
Batch end:     Pending-issue: 0, Pending-acknowledgement: 0
Command:      Pending-acknowledgement: 0
Total-objects: 1897
Stale-objects: 0
Resolve-objects: 0
Childless-delete-objects: 0
Backplane-objects: 0
Error-objects: 9
Number of bundles: 0
Paused-types: 0
```

Error Objects

- Use keyword `error-object`, `pending-*` to see object id's
- Check objects in error to confirm

```
9600X_50#sh platform software object-manager f0 error-object
Object identifier: 650830
  Description: LB 0x10000417
  Status: Modify/delete-failed
9600X_50#sh platform software object-manager f0 object 650830
Object identifier: 650830
  Description: LB 0x10000417
  Obj type id: 60
  Obj type: loadbalance
  Status: Modify/delete-failed, Epoch: 0, Client data: 0xad0635e8
```

FED Routing tables

```
9300_20#show plat software fed switch active ip route 10.48.13.0/24 detail
vrf   dest                htm                flags  SGT   DGID Last-modified      SecsSinceHit
---   ----                ---                -----
0     10.48.13.0/24         0x78bcb94983a8 0x0    0     0    2025/02/04 20:36:50.190      3
FIB: prefix_hdl:0x2600007a, mpls_ecr_prefix_hdl:0, sgtOverWrite: 0
===== OCE chain =====
ADJ:objid:74 {link_type:IP ifnum:0x90, adj:0x62000050, si: 0x78bcb9199178 IPv4: 172.18.255.9 }
[Output omitted]
Station Index (SI) [0xc3]
RI = 0x1f
DI = 0x53a8
Replication Bitmap: LD
Destination index   = 0x53a8
pmap                = 0x00000000 0x400000000000
pmap_intf : [TenGigabitEthernet1/0/47]
..
```

- FED layer output dependant on platform used (S1 vs Doppler)
- Replication Bitmap (per Asic/Core) shows if traffic is Local, Remote or to other Core.
- Bits in pmap are set if traffic would go out an interface

FED Adjacency table

```
9300_20#sh platform software fed switch 1 ip adj 172.18.255.9 detail
IPV4 Adj entries
dest          if_name          dst_mac          si_hdl  ri_hdl  adj_id  Last-modified
172.18.255.9  TenGig1/0/47    0072.78f8.aed6  0x78178 0x78158 0x4a    2024/12/31 16:31:09
Asic          SI-Index         DI-Index
----          -
0             195              0x53a8
Destination index = 0x53a8
pmap          = 0x00000000 0x400000000000
pmap_intf : [TenGigabitEthernet1/0/47]
Rewrite info:
ASIC#:0 RI:31 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1)
Mapped_rii:L3_UNICAST_IPV4(9)
MAC Addr: MAC Addr: 00:72:78:f8:ae:d6,
L3IF LE Index 38
```

- Rewrites occur on destination Asic.
- Check Adj on Egress switch (when in stack)

FED Tables, Silicon One

```
9600X_50#show plat software fed active ip route 10.48.13.0/24
IPV4ROUTE_ID:id:0x5dd182def298 nobj:(IPNEXTHOP_ID,0x1000001d) 10.48.13.0/24 tblid:0 DA:0
  NPD: device:0 lspa_rec:0 api_type:route(3)
  SDK: is_host:0 l3_dest:0x7b8bedf76590 l3_dest:la_l3_fec_impl_base(oid=2851) vrf(gid/oid):0/0x267
  ADJ:objid:0x1000001d (IPv4: 172.31.250.218) nh_type:NHADJ_NORMAL iif_id:0x4bc ether_type:0x8
#child:105
  srcmac:a453.0e46.22c5 dstmac:4ce1.75eb.0e82
  NPD: fec_oid:2851 was_nor_nh:1 cr_def:0 stale:0 l3port_valid:1
  NPD: device:0 nh_gid/oid:10/0xb1a old_gid/oid:0/0x0 parent_oid:2711
  SDK: cla_nhtype:0
9600X_50#show plat software fed active ip adj 172.31.250.218
Number of IPv4 Adjacency entries: 13
Adjacency ID           Adjacency IPAddress      Dest MacAddr             Nexthop Type            Interface
0x1000001d             172.31.250.218           4ce1.75eb.0e82           NORMAL                   HundredGigE1/0/16
```

- Silicon One outputs differ from Doppler based platform on FED layer but can provide similar outputs
- Note : Not all FED commands work on both Doppler and S1.

TCAM route utilization

```
9300_20#sh platform software fed switch active ip route summary
Total number of v4 fib entries = 11348
Total number succeeded in hardware = 8186

Mask-Len 0 :- Total-count 1 hw-installed count 1
Mask-Len 4 :- Total-count 1 hw-installed count 1
Mask-Len 8 :- Total-count 2 hw-installed count 2
Mask-Len 24 :- Total-count 10002 hw-installed count 8171
Mask-Len 30 :- Total-count 1331 hw-installed count 0
Mask-Len 32 :- Total-count 11 hw-installed count 11
```

- Hardware installed count should match Total-Count
- Routes not installed into Hardware mostly due to resource issues
- TCAM's use regions to install routes

Verifying TCAM Utilization

```
9300_20#show platform hardware fed switch active fwd-asic resource tcam utilization
```

```
Codes: EM - Exact_Match, I - Input, O - Output, IO - Input & Output, NA - Not Applicable
```

```
CAM Utilization for ASIC [0]
```

Table	Subtype	Dir	Max	Used	%Used	V4	V6	MPLS	Other
Mac Address Table	EM	I	32768	97	0.30%	0	0	0	97
Mac Address Table	TCAM	I	1024	21	2.05%	0	0	0	21
L3 Multicast	EM	I	8192	0	0.00%	0	0	0	0
L3 Multicast	TCAM	I	512	9	1.76%	3	6	0	0
L2 Multicast	EM	I	8192	0	0.00%	0	0	0	0
L2 Multicast	TCAM	I	512	11	2.15%	3	8	0	0
IP Route Table	EM	I	24576	12	0.05%	11	0	1	0
IP Route Table	TCAM	I	8192	8190	99.98%	8177	10	2	1
QOS ACL	TCAM	IO	5120	181	3.54%	52	86	0	43
Security ACL	TCAM	IO	5120	129	2.52%	26	58	0	45
Netflow ACL	TCAM	I	256	42	16.41%	20	20	0	2

- TCAM Full situation can lead to performance issues
- Choose suitable SDM template to optimize TCAM allocation
- TCAM utilization can vary per ASIC
- Check syslog for error messages

Verifying Hardware Forwarding

- Catalyst 9000 based switches allow for packets to be send through the system to verify hardware forwarding.
- Show Platform Hardware switch <id> fed forward command allows building packet from CLI of through pcap to verify hardware forwarding
- Using “exact-route” provides a faster response with less detailed outputs.
- Use packet parameters as it arives on the interface to see where and if it will be forwarded through the system

Verifying Hardware Forwarding

```
9300_20#show platform hardware fed switch 1 forward exact-route interface te 1/0/1 0072.78f8.aed1
0072.78f8.1354 ipv4 10.48.91.151 10.48.13.151 tcp 10001 23 0
Ingress Packet Details:
ether hdr : dest mac: 0072.78f8.1354, src mac: 0072.78f8.aed1
ether hdr : ethertype: 0x0800 (IPv4)
ipv4  hdr : dest ip: 10.48.13.151, src ip: 10.48.91.151
ipv4  hdr : packet len: 40, ttl: 64, protocol: 6 (TCP)
tcp   hdr : dest port: 23, src port: 10001
Ingress Packet Hexdump (first 128 bytes):
007278F813550072 78F8AED608004500 0028000100004006 FD410A305B970A30
0D97271100170000 00000000000005000 2000EB2E00000000 0000000000000000
Egress Interface: TenGigabitEthernet1/0/47
Egress Packet Details:
ether hdr : dest mac: 0072.78f8.aed6, src mac: 0072.78f8.1355
ether hdr : ethertype: 0x0800 (IPv4)
ipv4  hdr : dest ip: 10.48.13.151, src ip: 10.48.91.151
ipv4  hdr : packet len: 40, ttl: 63, protocol: 6 (TCP)
tcp   hdr : dest port: 23, src port: 10001
Egress Packet Hexdump (first 128 bytes):
007278F8AED60072 78F8135508004500 0028000100003F06 FE410A305B970A30
0D97271100170000 00000000000005000 2000EB2E00000000 0000000000000000
```

Ingress packet

Egress packet

Verifying Hardware Forwarding

```
9300_20#show platform hardware fed switch 1 forward interface te 1/0/47 0072.78f8.aed1
0072.78f8.1354 ipv4 10.48.91.151 10.48.13.151 tcp 10001 23 0
Show forward is running in the background. After completion, syslog will be generated.
*Feb 10 14:52:48.864: %SHFWD-6-PACKET_TRACE_DONE: Switch 1 F0/0: fed: Packet Trace
Complete: Execute (show platform hardware fed switch <> forward last summary|detail)
```

Using packet captures with show platform forward reduces complexity

```
9300_20#show platform hardware fed switch active forward interface te 1/0/47 pcap
flash:inpacket.pcap number 2 data
```

Results:

- show platform hardware fed switch <id> forward last summary
- show platform hardware fed switch <id> forward last detail
- shfwd.log files in flash:

Before needing to troubleshoot

- Get familiar with the CLI's
- Download this pdf for reference

Webex App

Questions?

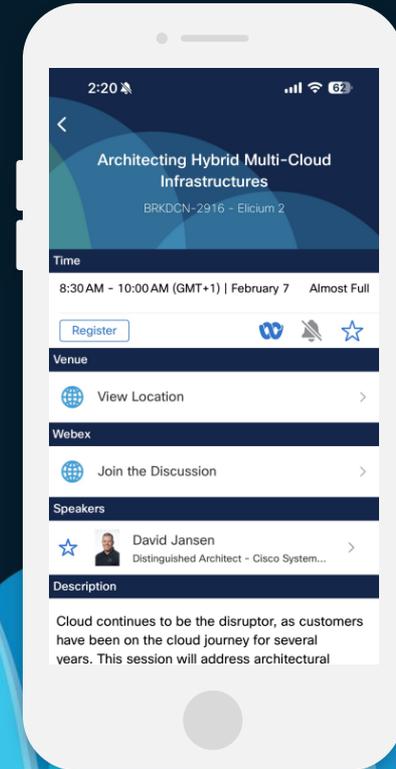
Use the Webex app to chat with the speaker after the session

How

- 1 Find this session in the Cisco Events mobile app
- 2 Click “Join the Discussion”
- 3 Install the Webex app or go directly to the Webex space
- 4 Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until February 28, 2025.

CISCO *Live!*



Fill Out Your Session Surveys



Participants who fill out a minimum of 4 session surveys and the overall event survey will get a unique Cisco Live t-shirt.

(from 11:30 on Thursday, while supplies last)



All surveys can be taken in the Cisco Events mobile app or by logging in to the Session Catalog and clicking the 'Participant Dashboard'



Content Catalog

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at ciscolive.com/on-demand. Sessions from this event will be available from March 3.



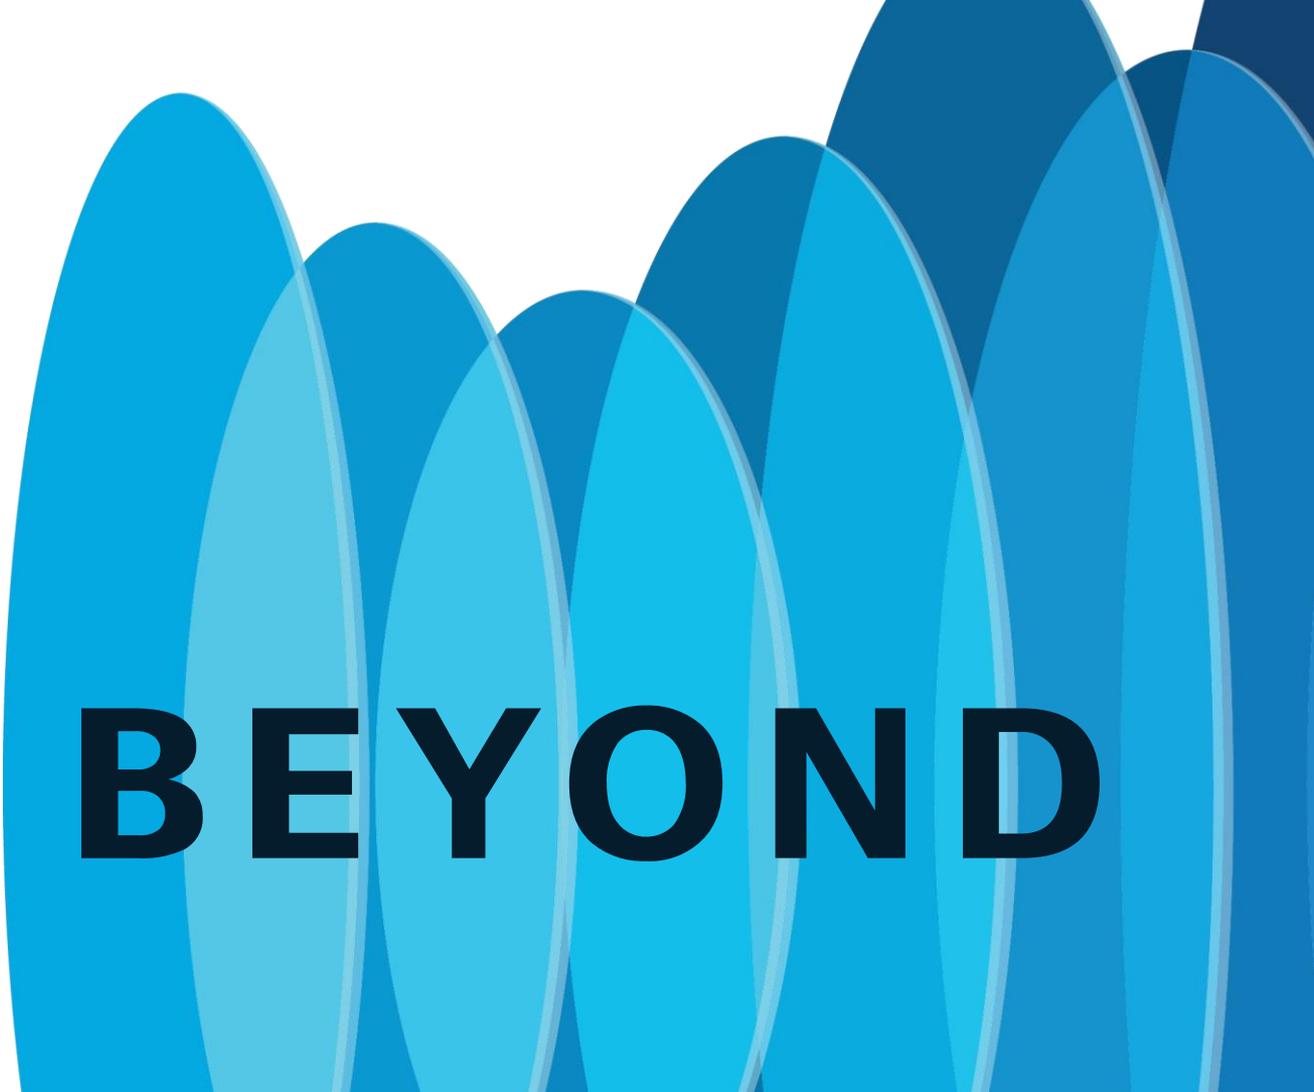
Thank you

CISCO *Live!*



CISCO *Live!*

GO BEYOND

A series of overlapping, vertically-oriented ovals in various shades of blue, ranging from light to dark, positioned on the right side of the image.