



Cisco ASR 9000 Architecture and Troubleshooting

CSC-ASR9000


Yong Zhao, Customer Support Engineer

CCIE #23507 RS&SP

yonzhao2@cisco.com

Swiss Army Knife Built for Edge Routing World

- Cisco ASR9000 Market Roles



1. High-End Aggregation & Transport

1. Mobile Backhaul
2. L2/Metro Aggregation
3. CMTS Aggregation
4. Video Distribution & Services

2. DC Gateway Router

1. DC Interconnect
2. DC WAN Edge
3. WEB/OTT

3. Services Router



1. Business Services
2. Residential Broadband
3. Converged Edge/Core
4. Enterprise WAN

Scalable System Architecture and Portfolio

Physical and Virtual

96Tbps System

nV Cluster



9000v,901,903

9001
9001-S

9904

9006

9010

9912

9922

Agenda


- Hardware Architecture
- Operating System & Configuration
- Frame Journey
- Control, Management, & Security
- Troubleshooting

A long-exposure photograph of a city street at night. The foreground is dominated by vibrant, multi-colored light trails from moving vehicles, creating a sense of motion and energy. In the background, modern city buildings are illuminated with various lights, and a pedestrian bridge spans across the street. The overall scene is a dynamic and colorful representation of a bustling urban environment.

Hardware Architecture

ASR 9010 and ASR 9006 Chassis

Shipping since day 1



ASR 9001 Compact Chassis


Shipping since IOS-XR 4.2.1 May 2012

Side-to-Side airflow
2RU

Front-to-back air flow with air flow
baffles, 4RU, require V2 fan

Sub-slot 0 with MPA

Sub-slot 1 with MPA



Redundant
(AC or DC)
Power Supplies
Field Replaceable

Supported MPAs:

- 20x1GE
- 2x10GE
- 4x10GE
- 1x40GE

EoBC
(cluster)

Fixed 4x10G
SFP+ ports

Fan Tray
Field Replaceable

ASR 9001-S Compact Chassis

Shipping since IOS-XR
4.3.1 May 2013

Side-to-Side airflow
2RU

Front-to-back air flow with air flow
baffles, 4RU, require V2 fan

Supported MPAs:


20x1GE
2x10GE
4x10GE
1x40GE

Pay As You Grow

- Low entry cost
- **SW License upgradable to full 9001**

Sub-slot 0 with MPA

Sub-slot 1 with MPA




60G bandwidth are disabled by
software. SW license to enable it

ASR 9922 Large Scale Chassis

Shipping since IOS-
XR 4.2.2 August 2012

Features	Description
Power	4 Power Shelves, 12 16 Power Modules 4.4 2.1 kW DC ; 6.0 3.0 kW AC supplies N+N AC supply redundancy N+N N+1 DC supply redundancy
Fan	4 Fan Trays Front to back airflow
I/O Slots	20 I/O slots
Rack Size	44 RU
RP	1+1 RP redundancy
Fabric	6+1 fabric redundancy.
Bandwidth	Phase 1: 550Gb per Slot Phase 2: 1.6Tb per Slot Future: 2+Tb per Slot
SW	XR 4.2.2 – August 2012




Fully loaded
Engineering testbed



ASR 9912 Large Scale Chassis

Shipping since XR4.3.2 & 5.1.0, Sep 2013


Features	Description
Fan	2 Fan Trays Front to back airflow
I/O Slots	10 I/O slots
Rack Size	30 RU
RP	1+1 RP redundancy
Fabric	6+1 fabric redundancy
Power	3 Power Shelves, 9 12 Power Modules 4.4 2.1 kW DC ; 6.0 3.0 kW AC supplies N+N AC power supply redundancy N+N N+1 DC power supply redundancy
Bandwidth	Phase 1: 550Gb per Slot Phase 2: 1.6Tb per Slot Future: 2+Tb per Slot
SW	XR 4.3.2 & 5.1.0



ASR 9000v “Satellite”

Local or remote 10G to 1G fan out

- Local or remote
- All configuration done on host
- L2/3/4 operation done on host
 - No local switching on satellite
 - Ingress & egress QoS done on host
- 1GE ports can be assigned to specific 10G ports or bundle




Power and Cooling



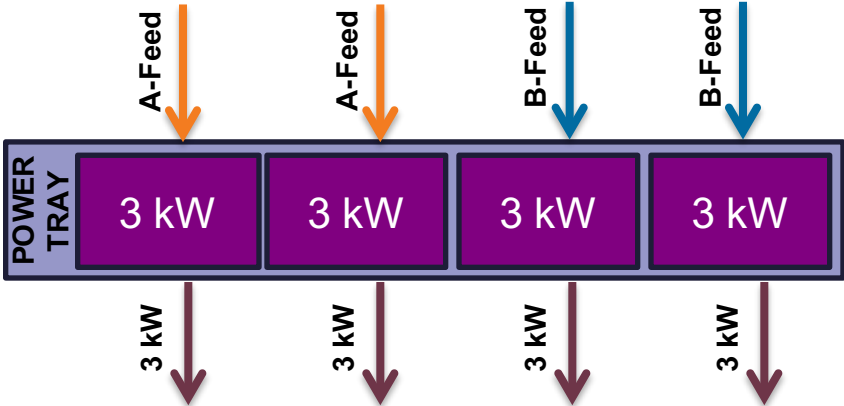
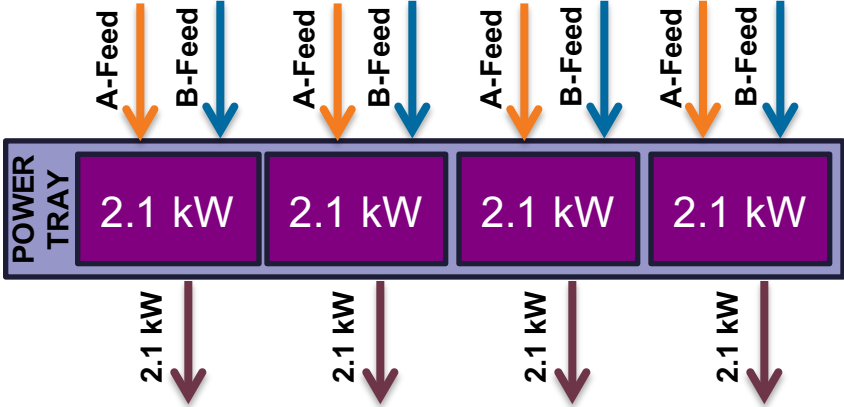
ASR-9010-FAN



ASR-9006-FAN




Power Supply



Version 3 AC Power Supply

New Power Tray and Power Supply


- Three 6kW AC PMs per tray
- Two logically AND'ed single phase 3kW AC power inputs per PM
 - If both inputs are active, output is max 6kW
 - If only one input is active, Power Supply is still working, output will be max 3kW
- Input 16/20A@230V AC & same power cords → easily upgradeable!



Version 3 DC Power Supply

New Power Tray and Power Supply

- Four 4.4kW DC PMs per tray
- Two logically AND'ed DC power feeds per PM
 - If both inputs are active, output is max 4.4kW
 - If only one input is active, Power Supply is still working, output will be max 2.2kW
- DC power requirements & DC lug → “easily” upgradeable if you like copper 😊



New DC Tray

ASR9000 Power Supplies

	AC V2	AC V3	DC V2	DC V3
Max Power	3 kW	6 kW	2.1 kW	4.4 kW
# of Feeds	1	2	2	2
Feed redundancy in PEM	n/a	No	Yes	No
# of PSs per power tray	4	3	4	4
Redundancy scheme ¹⁾	N+N	N+N	N+1	N+N

number of modules required to protect from feed failure (e.g. power grid outage)

ASR 9K Ethernet Line Card Overview

-TR: Transport Optimized
-SE: Service edge Optimized


First-generation LC
(Trident NP)

-L, -B, -E




Second-gen LC
(Typhoon)

-TR, -SE



-TR, -SE

Third-gen LC
(Tomahawk)



ASR 9000 Tomahawk Line Cards

Built with Cutting Edge, Customized, Flexible, Efficient Silicon

Cutting Edge

- Unprecedented Scale – 240Gbps in one ASIC!
- Hardware Offload of CPU Intensive Protocols

Customized

- High Availability Customized Silicon – Hitless FPD Upgrades
- Embedded MACSec for Increased Security

Flexible

- CPAK Delivers 10G, 40G or 100G on any Interface on One LC
- Common Sparring and Onetime LC Qualification

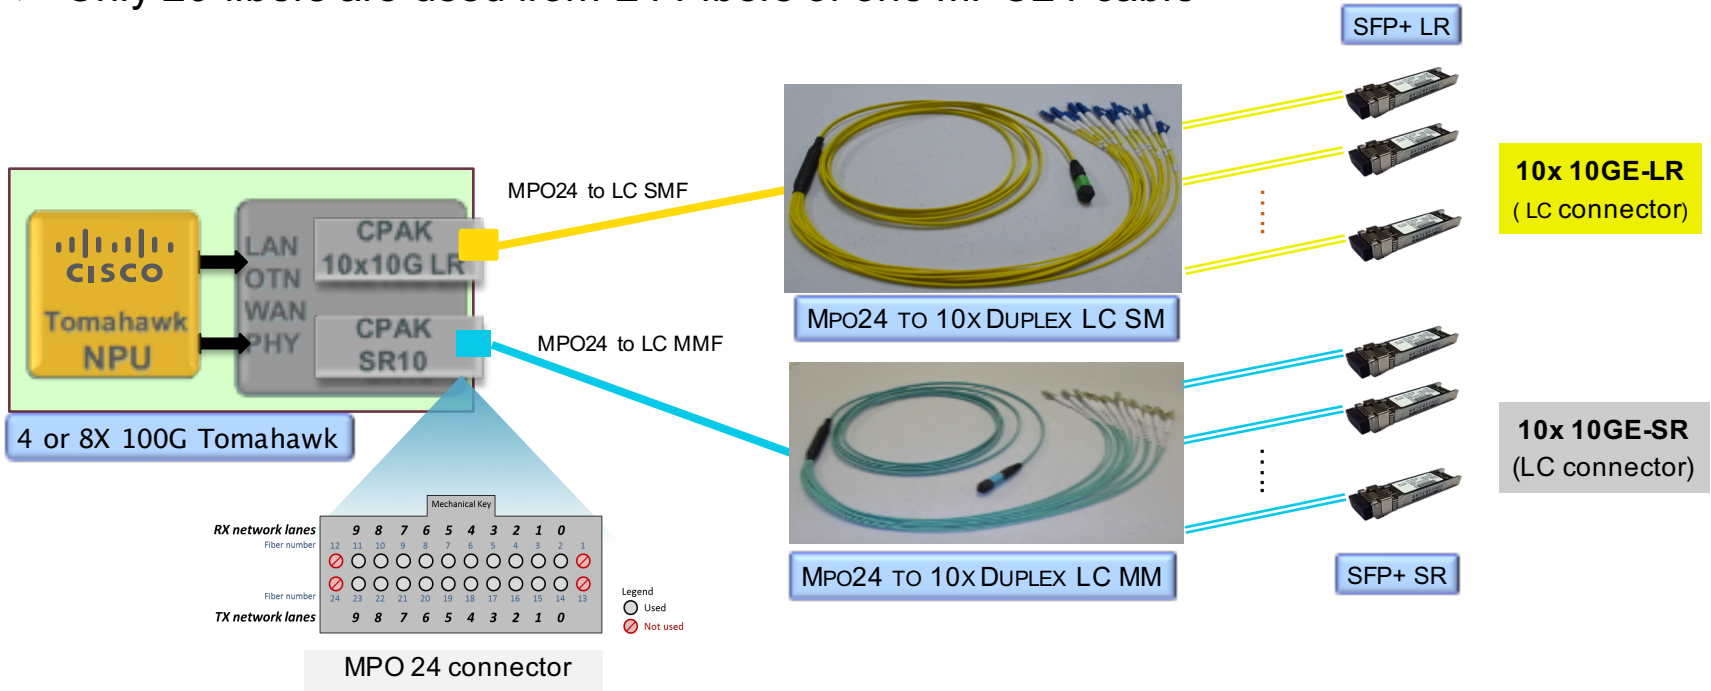
Efficient

- Power Down Unused Linecard Slices to Increase Efficiency
- Lowest Watts per Gbps with CPAK and Optimized Silicon




Breakout Cable Solution – 100G → 10x 10G

- 10 : 1 Breakout Cable with a MPO24 MMF or SMF cable
- Only 20 fibers are used from 24 Fibers of one MPO24 cable

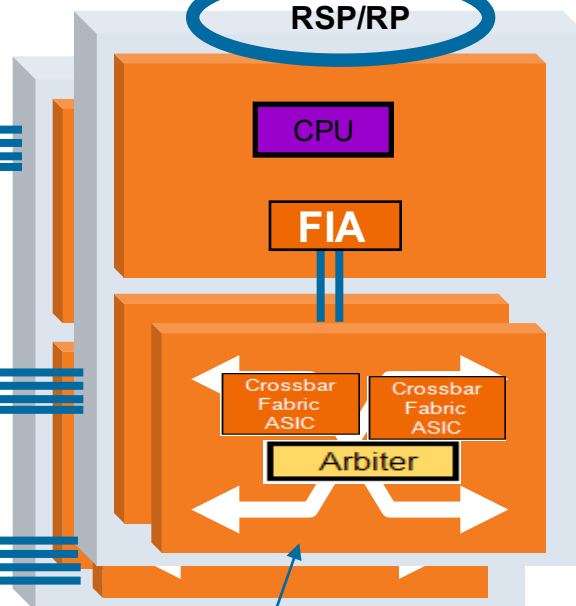


Cisco ASR 9000 Hardware System Components

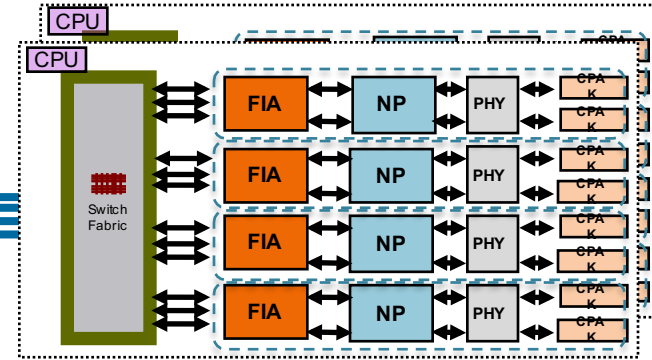
Line Card



RSP/RP



Line Card



Switch Fabric


Integrated on RSP or Separated fabric card

Route Switch Processors and Route Processors

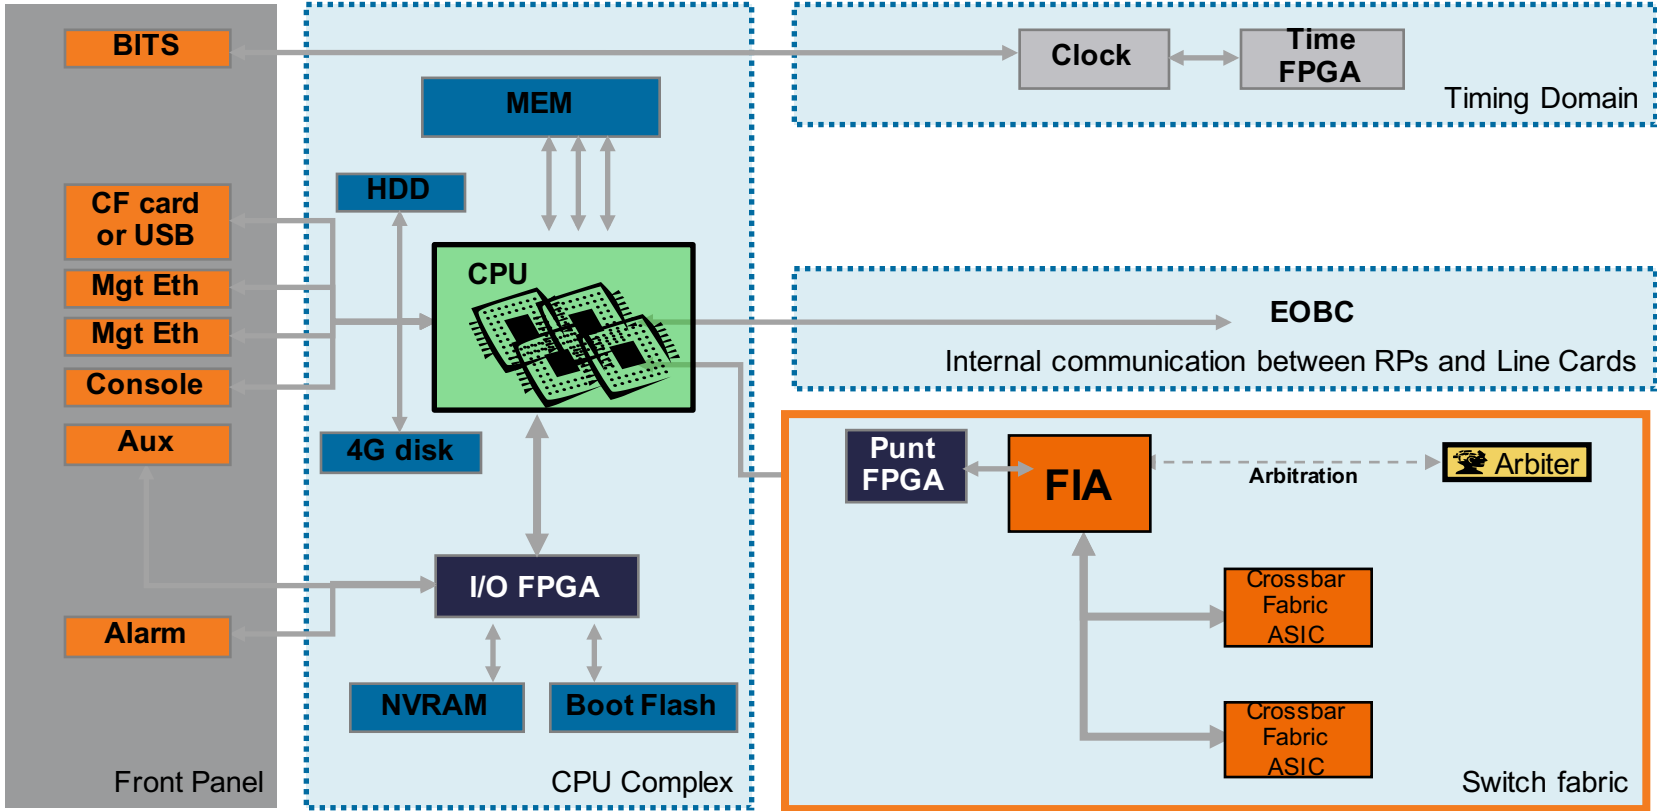
RSP used in ASR9904/9006/9010, RP in ASR9922/9912

	9006/9010 RSP	9904/9006/9010 RSP440	9922/9912 RP1	RSP880	9922/9912 RP2
	1 st Gen RP	2 nd Gen RP and Fabric ASIC		3 rd Gen RP and Fabric ASIC	
Processors	Freescale 8641D 2 x 1.5GHz	Intel x86 4 Core 2.27 GHz	Intel x86 4 Core 2.27 GHz	Intel x86 (Ivy Bridge EP) 6 Core 2GHz	Intel x86 (Ivy Bridge EP) 6 Core 2GHz
RAM	RSP-4G: 4GB RSP-8G: 8GB	RSP440-TR: 6GB RSP440-SE: 12GB	-TR: 6GB -SE: 12GB	-TR: 16GB -SE: 32GB	-TR: 16GB -SE: 32GB
SSD	N/A: 4G / 30G-HDD	2x 16GB Slim SATA	2x 16GB Slim SATA	2x 32GB Slim SATA	2x 32GB Slim SATA
nV EOBC ports	N/A	2 x 1G/10G SFP+	2 x 1G/10G SFP+	4 x 1/10G SFP+	4 x 1/10G SFP+
Punt BW		10GE	10GE	40GE	40GE
Switch fabric bandwidth	92G + 92G (9006/9010)	220G + 220G (9006/9010) 385G + 385G (9904) (fabric integrated on RSP)	660G+110G (separated fabric card)	460G + 460G (9006/9010) 805G + 805G (9904) (fabric integrated on RSP)	1.38Tb + 230G (separated fabric card)

RSP440 – Faceplate and Interfaces



RSP Engine Architecture



ASR 9000 Switch Fabric Overview

Integrated fabric/RP/LC



9001, 2RU, 120G



9001-S, 2RU, 60G

Fabric is integrated on RSP 1+1 redundancy



9904

RSP880: 805G+805G/slot

RSP440: 385G+385G/slot



9006


RSP880: 460G+460G/slot

RSP440: 220G+220G/slot



9010

Separated fabric card 6+1 redundancy



9912

SFC2: 1.38T+230G /slot


SFC1: 660G+110G/slot



9922

Fabric Super-framing Mechanism

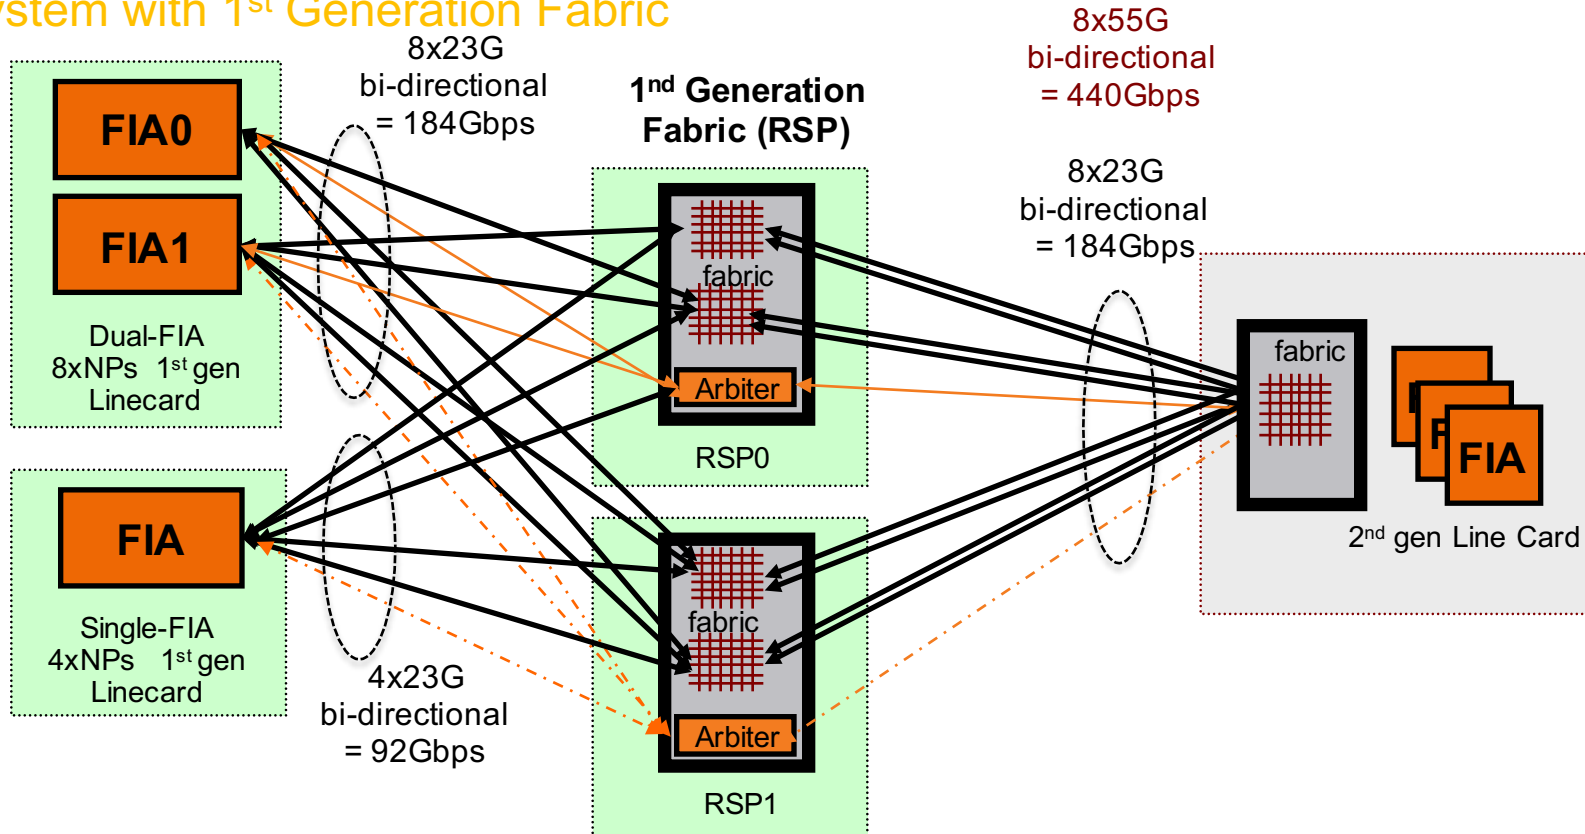
- **Multiple unicast frames** from/to same destinations aggregated into **one super frame**
- Super frame is created if there are frames waiting in the queue, up to 32 frames or when min threshold met, can be aggregated into one super frame
- Super frame only apply to unicast, **not multicast**
- Super-framing significantly **improves total fabric throughput**



- Note that fabric counters are showing super frames not individual packets!!


1st/2nd Generation switch fabric compatibility

System with 1st Generation Fabric



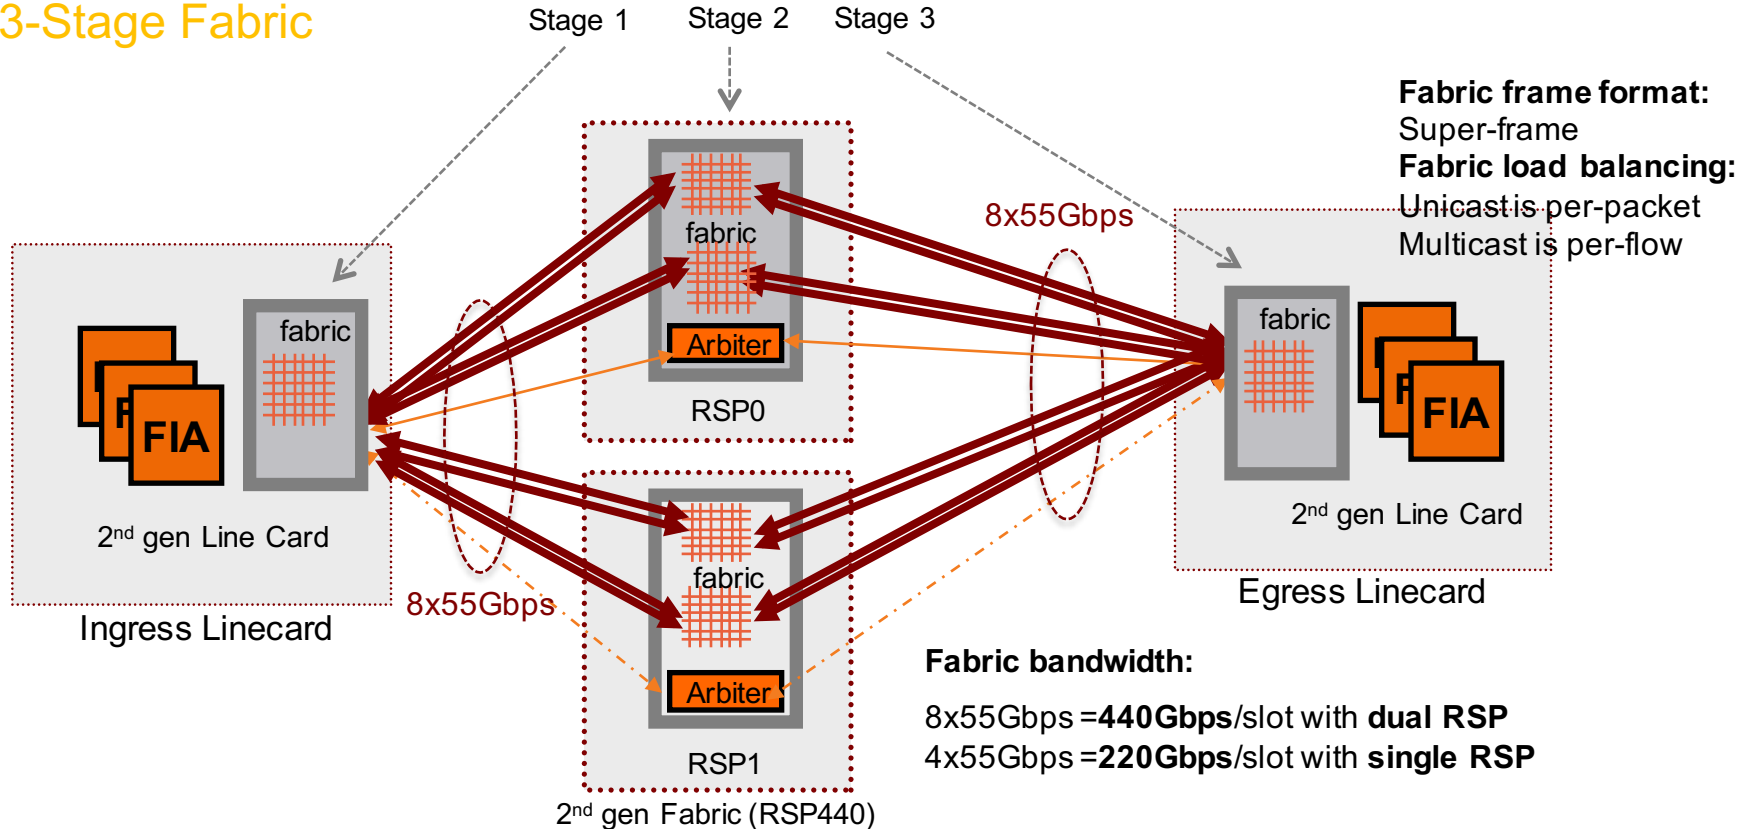
1st/2nd Generation switch fabric compatibility

System With 2nd Generation Fabric



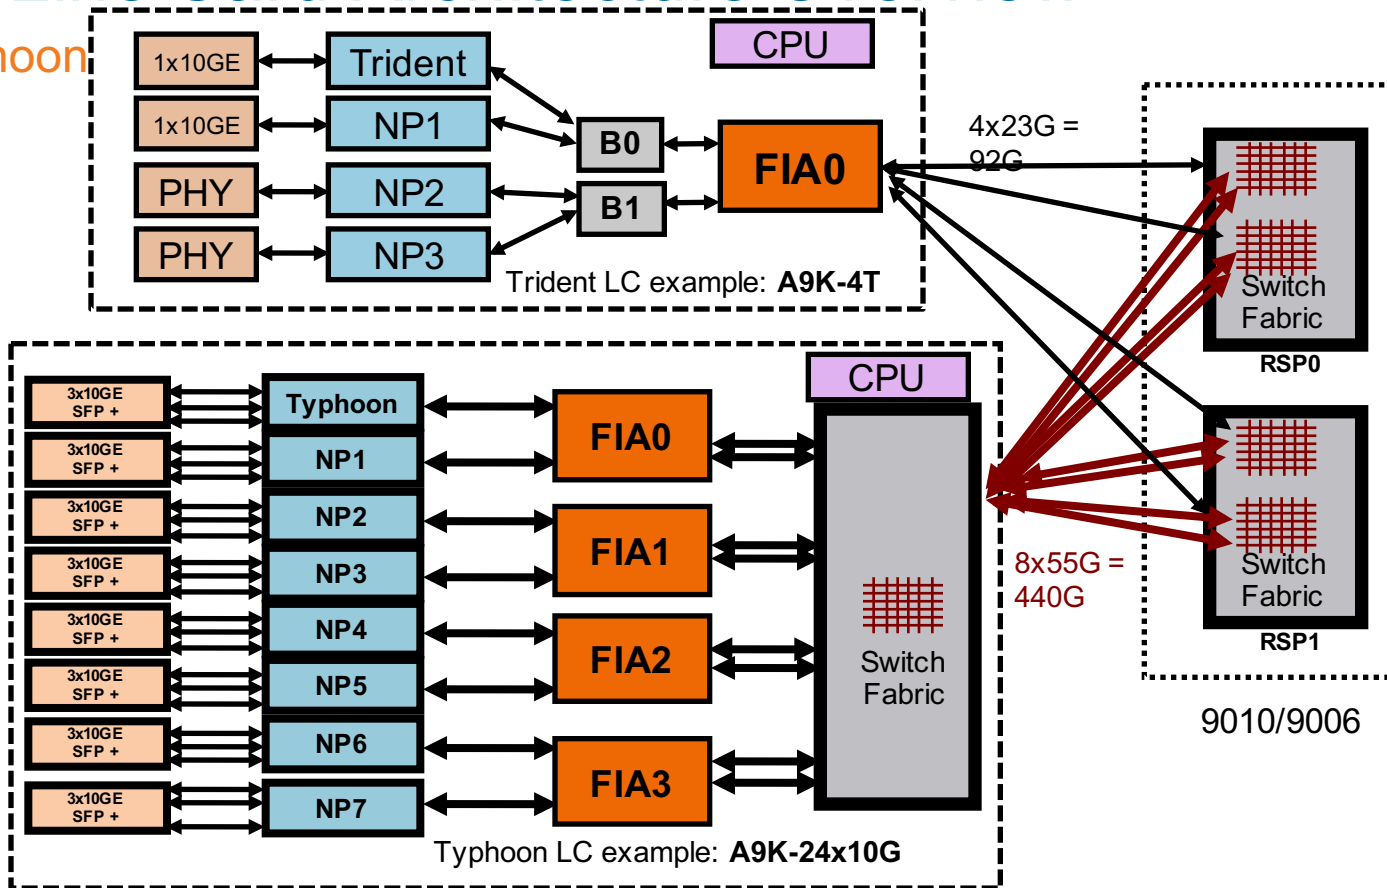
ASR9006/9010 RSP440 Switch Fabric Architecture

3-Stage Fabric

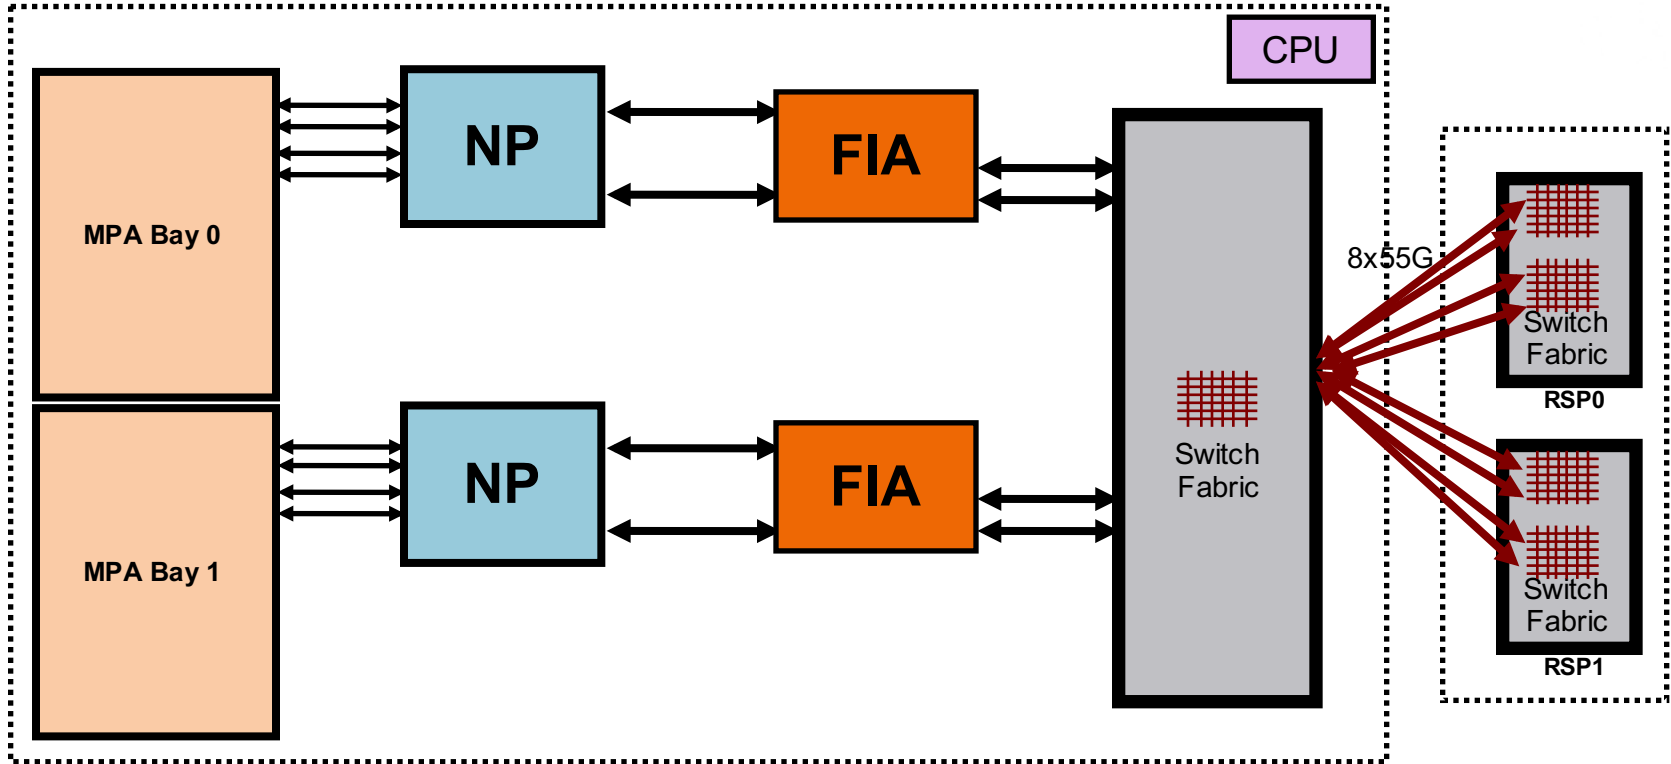


ASR 9000 Line Card Architecture Overview

Trident and Typhoon



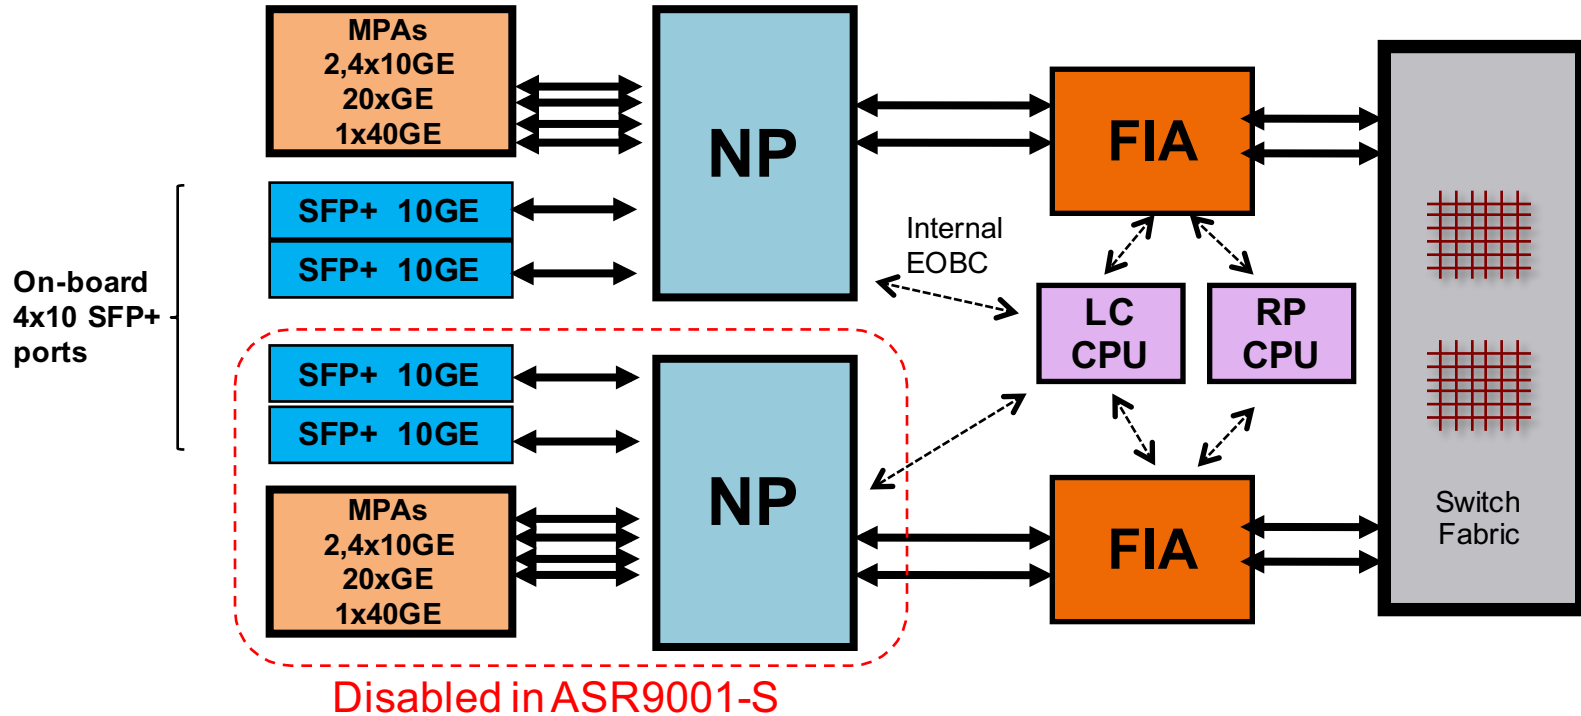
Module Cards – MOD80






ASR 9001 Architecture

Identical HW Components as the Modular Systems



Edge Linecard Silicon Slice Evolution



Network Processor Architecture Details

-TR and -SE have same memory size



-TR and -SE have different memory size

- TCAM: VLAN tag, QoS and ACL classification
- Stats memory: interface statistics, forwarding statistics etc
- Frame memory: buffer, Queues
- Lookup Memory: forwarding tables, FIB, MAC, ADJ
- TR / SE
 - Different TCAM/frame/stats memory size for different per-LC QoS, ACL, logical interface scale
 - Same lookup memory for same system wide scale mixing different variation of LCs doesn't impact system wide scale

-TR: Transport optimized
-SE: Service Edge optimized

A nighttime photograph of a city street. In the background, there are modern buildings with lit windows and a pedestrian bridge with blue lighting. The middle ground shows a road with traffic lights and light trails from vehicles. The foreground is dominated by long, curved light trails in yellow, orange, and red, suggesting a long-exposure shot of light trails from a moving light source.

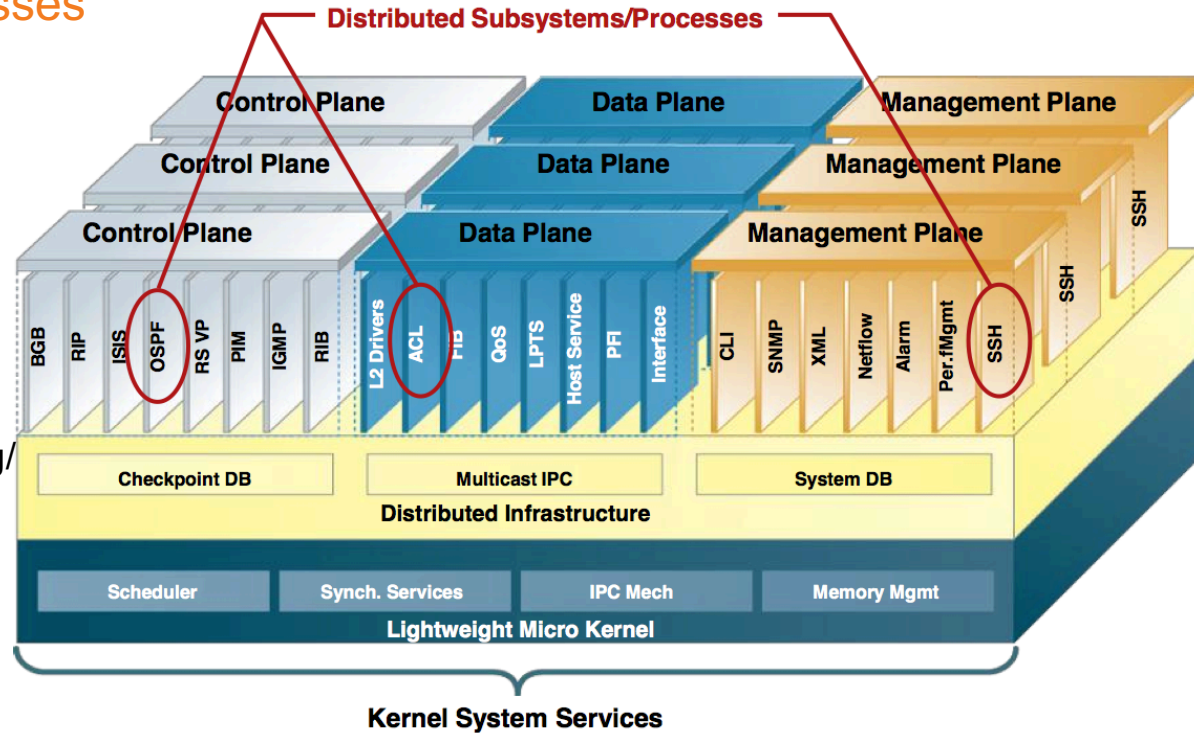
Operating System & Configuration

IOS-XR Architecture

IOS-XR Operating System

Kernel and distributed processes

- Micro-kernel
 - QNX kernel
- Restartable processes
 - A process may start/ terminate based on configuration
 - Scheduler keeps track of process starts/spawning/ priority/path
 - A process can crash/restart/patched
- Distributed processing
 - Processes run on RP and LC CPU's



IOS-XR Operating System

RSP and LC CPU's

- To monitor: a CPU on every card
- Instances of processes running on RSP and LC CPU's

```
RP/0/RSP0/CPU0:rasr9k-1y#show processes cpu location 0/RSP0/CPU0 | exclude ""
0%      0%      0%"
Wed Nov 28 01:36:52.203 UTC
```

CPU utilization for one minute: 26%; five minutes: 25%; fifteen minutes: 22%

PID	1Min	5Min	15Min	Process
94243	3%	3%	3%	spp
254074	23%	22%	19%	netio

```
RP/0/RSP0/CPU0:rasr9k-1y#show processes cpu location 0/0/CPU0 | exclude "" 0%
0%      0%"
Wed Nov 28 01:28:52.281 UTC
```

CPU utilization for one minute: 46%; five minutes: 48%; fifteen minutes: 39%

PID	1Min	5Min	15Min	Process
45085	22%	23%	22%	spp
180316	23%	23%	23%	netio

```
RP/0/RSP0/CPU0:ASR9010-1#show process distribution netio
Mon Jul 27 06:22:43.321 UTC
3 processes found
```

NODE	PID	JID	#THR	TYPE	PROGRAM
0/RSP0/CPU0	532616	333	16	RP	netio
0/RSP1/CPU0	532613	333	16	RP	netio
0/6/CPU0	364639	274	8	LC	netio

Demonstrating Process Restart

Same Job ID, New Process ID

```
RP/0/RSP0/CPU0:rasr9000-2w-a#show processes bgp
wed Jan 15 17:58:34.050 EST
      Job Id: 1048
      PID: 287056
      Executable path: /disk0/iosxr-routing-
4.2.3.CSCuh52959-1.0.0/0x100305/bin/bgp
      Instance #: 1
      Version ID: 00.00.0000
      Respawn: ON
      Respawn count: 1
      Max. spawns per minute: 12
      Last started: Thu Jan  2 09:11:18 2014
      Process state: Run
      Package state: Normal
      started on config: default
      Feature name: ON
      Tag : default
      Process group: v4-routing
      core: MAINMEM
      Max. core: 0
      Placement: Placeable
      startup_path: /pkg/startup/bgp.startup
      Ready: 0.700s
      Available: 85.082s
      Process cpu time: 21.760 user, 2.619 kernel,
24.379 total
JID  TID CPU Stack pri state      TimeInState
HR:MM:SS:MSEC  NAME
1048  1   2  384K  10 Receive      0:00:03:0395
```

```
RP/0/RSP0/CPU0:rasr9000-2w-a#process restart bgp
wed Jan 15 18:03:24.836 EST
RP/0/RSP0/CPU0:Jan 15 18:03:24.874 :
sysmgr_control[65784]: %OS-SYSMGR-4-PROC_RESTART_NAME :
User cisco (con0_RSP0_CPU0) requested a restart of
process bgp at 0/RSP0/CPU0
RP/0/RSP0/CPU0:rasr9000-2w-a#show processes bgp
wed Jan 15 18:03:28.726 EST
      Job Id: 1048
      PID: 3182840
      Executable path: /disk0/iosxr-routing-
4.2.3.CSCuh52959-1.0.0/0x100305/bin/bgp
      Instance #: 1
      Version ID: 00.00.0000
      Respawn: ON
      Respawn count: 2
      Max. spawns per minute: 12
      Last started: Wed Jan 15 18:03:24 2014
      Process state: Run (last exit due to SIGTERM)
      Package state: Normal
      Started on config: default
      Feature name: ON
      Tag : default
      Process group: v4-routing
      core: MAINMEM
      Max. core: 0
      Placement: Placeable
      startup_path: /pkg/startup/bgp.startup
      Ready: 0.225s
```

Process Dumps

Where?

```
RP/0/RSP0/CPU0:rasr9000-2w-a#admin show exception
```

```
Sat Jan 4 00:15:27.885 EST
```

```
Exception path for choice 1 is not configured or removed
Exception path for choice 2 is not configured or removed
Exception path for choice 3 is not configured or removed
Choice fallback one path = dumper_harddisk:/dumper compress = on filename = <process_name>
Choice fallback two path = dumper_disk1a:/dumper compress = on filename = <process_name>
Choice fallback three path = dumper_disk0a:/dumper compress = on filename = <process_name>
Kernel dump not configured
Tftp route for kernel core dump not configured
No config for pakmem tuple
No config for sparse tuple
No config for sprsize tuple
No config for coresize tuple
No config for memory-threshold tuple
No config for core-verification tuple
```

```
RP/0/RSP0/CPU0:rasr9000-2w-a#dir harddisk:/dumper
```

```
Sat Jan 4 00:16:10.138 EST
```

```
Directory of harddisk:/dumper
```

24922	-rw-	216304651	Sat Jan 4 00:16:08 2014	ce_switch.log
24665	-rw-	42408	Tue Nov 5 19:06:35 2013	crashinfo.by.kernel.19070930-173606
24694	-rw-	1586390	Tue Nov 5 19:06:36 2013	kernel_core.by.kernel.19070930-173606.Z
24695	-rw-	1044480	Tue Nov 5 19:06:36 2013	pcds_dump.19070930-173606
24697	-rw-	4813080	Fri Nov 8 17:03:11 2013	first.mpls_lsd_338.node0_RSP0_CPU0.x86.Z
.				

A nighttime photograph of a city street. In the background, there are several tall buildings with lit windows and a pedestrian bridge spanning across the street. The street is illuminated by traffic lights and streetlights. In the foreground, there are long, colorful light trails from moving vehicles, creating a sense of motion and energy. The overall scene is a vibrant urban night scene.

Operating System & Configuration

IOS-XR Install & Upgrade & Config

Reading Installed Packages

Example

```
RP/0/RSP0/CPU0:rasr9000-2w-b#show install active
detail
```

```
Sun Jan  4 23:43:14.325 EST
Secure Domain Router: Owner
```

```
Node 0/RSP0/CPU0 [RP] [SDR: Owner]
```

```
Boot Device: disk0:
```

```
Boot Image: /disk0/asr9k-os-mpi-
```

```
5.1.3/0x100305/mbiasr9k-rsp3.vm
```

```
Active Packages:
```

```
disk0:asr9k-mp1s-px-5.1.3
```

```
disk0:iosxr-mp1s-5.1.3
```

```
disk0:asr9k-mgb1-px-5.1.3
```

```
disk0:asr9k-mgb1-supp-5.1.3
```

```
disk0:iosxr-mgb1-5.1.3
```

```
disk0:asr9k-optic-px-5.1.3
```

```
disk0:asr9k-optics-supp-5.1.3
```

```
disk0:asr9k-k9sec-px-5.1.3
```

```
disk0:iosxr-security-5.1.3
```

```
disk0:asr9k-k9sec-supp-5.1.3
```

```
disk0:asr9k-doc-px-5.1.3
```

```
disk0:asr9k-doc-supp-5.1.3
```

```
disk0:asr9k-fpd-px-5.1.3
```

```
disk0:asr9k-fpd-5.1.3
```

```
disk0:asr9k-mini-px-5.1.3
```

```
disk0:asr9k-scfclient-5.1.3
```

```
disk0:asr9k-os-mpi-5.1.3
```

```
disk0:asr9k-cpp-5.1.3
```

```
disk0:asr9k-ce-5.1.3
```

```
disk0:iosxr-ce-5.1.3
```

```
disk0:asr9k-diags-supp-5.1.3
```

```
disk0:iosxr-diags-5.1.3
```

```
disk0:asr9k-fwding-5.1.3
```

```
disk0:iosxr-fwding-5.1.3
```

```
disk0:iosxr-routing-5.1.3
```

```
disk0:iosxr-infra-5.1.3
```

```
disk0:asr9k-base-5.1.3
```

```
disk0:asr9k-mcast-px-5.1.3
```

```
disk0:asr9k-mcast-supp-5.1.3
```

```
disk0:iosxr-mcast-5.1.3
```

```
Node 0/1/CPU0 [LC] [SDR: Owner]
```

```
Boot Device: mem:
```

```
Boot Image: /disk0/asr9k-os-mpi-5.1.3/1c/mbiasr9k-
```

```
1c.vm
```

```
Active Packages:
```

```
disk0:asr9k-mp1s-px-5.1.3
```

```
disk0:iosxr-mp1s-5.1.3
```

```
disk0:asr9k-optic-px-5.1.3
```

```
disk0:asr9k-optics-supp-5.1.3
```

```
disk0:asr9k-mini-px-5.1.3
```

```
disk0:asr9k-scfclient-5.1.3
```

```
disk0:asr9k-os-mpi-5.1.3
```

```
disk0:asr9k-cpp-5.1.3
```

```
disk0:asr9k-ce-5.1.3
```

```
disk0:iosxr-ce-5.1.3
```

```
disk0:asr9k-diags-supp-5.1.3
```

```
disk0:iosxr-diags-5.1.3
```

Software Maintenance Updates (SMUs)

- Allows for **software package installation**/removal leveraging on Modularity and Process restart
- **Redundant processors are not mandatory** (unlike ISSU) and in many cases is non service impacting and may not require reload.
- Mechanism for
 - delivery of critical bug fixes without the need to wait for next maintenance release
- SMU is named by release, package, and bugid
 - Examples
 - asr9k-px-4.2.3.CSCud37351-1.0.0
 - asr9k-px-4.2.0.CSCtx28089-1.0.0

IOS-XR Version

Defect ID



Cisco Software Manager

Available on CCO in the Downloads Section for ASR9000

The screenshot displays the Cisco Software Manager interface. On the left, the navigation pane shows 'Download Software' and 'Select a Software Type:' with options for 'IOS XR Software' and 'IOS XR Software Maintenance Upgrades (SMU)'. A red box highlights the 'IOS XR Software Maintenance Upgrades (SMU) Bundles' link.

The main content area shows a list of 35 SMUs for 'ASR9K-PX-4.2.3'. The table below summarizes the visible entries:

ST	DDTS	Type	Description	Impact	Functional Areas	SMU ID	SMU Name	
✓	▲	CSCud16470	Optional	OSPFv3 reserved field not zero - RFC 5340	hitless	OSPFV3	AA06838	asr9k-px-4.2.3.CSCud16470
✓	▲	CSCud29892	Optional	bundle replay not processed for a subset of interfaces	hitless	BUNDLE	AA06806	asr9k-px-4.2.3.CSCud29892
✓	▲	CSCud37351	Recommended	423 SMU Pack2 for ASR9k	needs reboot	INFRASTRUCTURE	AA06812	asr9k-px-4.2.3.CSCud37351
✓	▲	CSCud39254	Recommended	NP search memory management failure	needs reboot	INFRASTRUCTURE	AA06844	asr9k-px-4.2.3.CSCud39254
✓	▲	CSCud40419	Optional	Multicast packets destined to 239.x are getting FIB looked-up in 4.2.1	hitless	MCAST	AA06828	asr9k-px-4.2.3.CSCud40419
✓	▲	CSCud41972	Optional	Radius Challenge not working	hitless	AAA	AA06890	asr9k-px-4.2.3.CSCud41972
✓	▲	CSCud54093	Recommended	r423 RO V2 Fan Tray SMU	issu/reload	INFRASTRUCTURE	AA06883	asr9k-px-4.2.3.CSCud54093
✓	▲	CSCud65815	Optional	qos-ea is blocked by prm_server_ty after QoS in-place modification	needs reboot	QOS	AA06911	asr9k-px-4.2.3.CSCud65815
✓	▲	CSCud73764	Optional	ASR9000 SIP-700 Multiple processes blocked on cpp_driver0 process	needs reboot	FORWARDING	AA07177	asr9k-px-4.2.3.CSCud73764
✓	▲	CSCud81249	Optional	ppp_ma crash when	hitless	PPP	AA06909	asr9k-px-4.2.3.CSCud81249
✓	▲	CSCud98419	Recommended	Umbrella ddtts for FP	hitless	INFRASTRUCTURE	AA06927	asr9k-px-4.2.3.CSCud98419

A 'New SMU Alert - ASR9K-PX-4.2.3' dialog box is displayed, stating: '35 new SMU entries have been discovered from Cisco.com on Tuesday, April 16, 2013 1:32:34 PM PDT'.

In the bottom-left foreground, a 'Login Info' dialog box is open, showing a list of devices. The 'Add New Device Group' and 'Add New Device' buttons are circled in red. The device list includes 'sj20lab-as1(Swami)' selected.

Turboboot

A “fresh” start!


- Previous disk data is wiped
- A clean “re-image” as a new router
- Potential recovery mechanism
- Not an “upgrade” method

```
Rommon1>unset BOOT
Rommon2>confreg 0x102
Rommon3>sync
Rommon4>reset

Rommon1>IP_ADDRESS=<a.b.c.d>
Rommon2>IP_SUBNET_MASK=<mask>
Rommon3>TFTP_SERVER=<a.b.c.d>
Rommon4>DEFAULT_GATEWAY=<a.b.c.d>
Rommon5>TFTP_RETRY_COUNT=4
Rommon6>TFTP_TIMEOUT=60
Rommon7>TFTP_CHECKSUM=1
Rommon8>priv
Rommon9>diswd
Rommon10>unset BOOT
Rommon11>TURBOBOOT=on, disk0, format
Rommon12>sync
Rommon13>boot tftp://a.b.c.d/path/asr9k-
mini-px.vm-4.2.3
```

PIE Installation Concepts

- PIE install used once system is operational w/ XR
 - .vm files can be used if booting from rommon is required
- Packages can be added or upgraded
- Install from Exec or Admin Mode
- 3 phase install
 - Add – Copy package and unpack
 - Activate – Restart processes/nodes with new code
 - Commit – Lock activated packages through reset



Cards' FPD

Verify/upgrade FPD version

```
RP/0/RSP0/CPU0:rasr9000-2w-b#admin show hw-module fpd location all
```

```
=====
```

Existing Field Programmable Devices							
Location	Card Type	HW Version	Type	Subtype	Inst	Current SW Version	Upg/Dng?
0/RSP0/CPU0	A9K-RSP440-SE	1.0	1c	cbc	0	16.115	No
			1c	fpga1	0	0.09	No
			1c	fpga2	0	1.06	No
			1c	fpga3	0	4.09	No
			1c	rommon	0	0.62	No
0/RSP0/CPU0	ASR-9006-FAN	1.0	1c	cbc	2	5.02	No
0/0/CPU0	A9K-24x10GE-SE	1.0	1c	cbc	0	19.110	No
			1c	fpga2	0	1.02	No
			1c	fpga3	0	1.01	No
			1c	fpga4	0	1.05	No
			1c	rommon	0	1.28	No
0/1/CPU0	A9K-MOD80-SE	1.0	1c	cbc	0	20.116	No
			1c	fpga2	0	1.01	No
			1c	fpga4	0	1.05	No
			1c	rommon	0	1.28	No

```
=====
```

No == good. No change needed

If LC1 needs upgrade

```
RP/0/RSP0/CPU0:rasr9000-2w-a#admin upgrade hw-module fpd all location 0/1/CPU0
```

Cards' FPD

Verify/upgrade FPD version

```
RP/0/RSP0/CPU0:rasr9000-2w-a#admin show running-config
```

```
 fpd auto-upgrade
```

```
RP/0/RSP0/CPU0:rasr9000-2w-a#admin upgrade hw-module fpd all location all
```

```
wed Jan 15 18:57:04.683 EST
```

```
***** UPGRADE WARNING MESSAGE: *****
```

```
* This upgrade operation has a maximum timeout of 160 minutes. *
* If you are executing the cmd for one specific location and *
* card in that location reloads or goes down for some reason *
* you can press CTRL-C to get back the RP's prompt. *
* If you are executing the cmd for _all_ locations and a node *
* reloads or is down please allow other nodes to finish the *
* upgrade process before pressing CTRL-C. *
```

```
% RELOAD REMINDER:
```

- The upgrade operation of the target module will not interrupt its normal operation. However, for the changes to take effect, the target module will need to be manually reloaded after the upgrade operation. This can be accomplished with the use of "hw-module <target> reload" command.
- If automatic reload operation is desired after the upgrade, please use the "reload" option at the end of the upgrade command.
- The output of "show hw-module fpd location" command will not display correct version information after the upgrade if the target module is not reloaded.

```
NOTE: Chassis CLI will not be accessible while upgrade is in progress.
```

```
Continue? [confirm]
```

Auto FPD upgrade
configuration

Manual FPD upgrade

Command Modes

Exec, exec config, admin, admin config

Exec – Normal operations – monitoring interfaces, routing, CEF, and VPNs

```
RP/0/RSP1/CPU0:viking-1#  
show ipv4 interfaces brief      show running-config  
show install active            show cef summary location 0/5/CPU0  
show l2vpn xconnect
```

Config – Normal configuration for router

```
RP/0/RSP1/CPU0:viking-1(config)#  
router bgp 100          l2vpn          policy-map foo  
mpls ldp                ipv4 access-list block-junk
```

Admin – Chassis operations (i.e. fans & power)

```
RP/0/RSP1/CPU0:viking-1(admin)#  
Config-register <0x1922>      show platform
```

Admin Config – admin plane config

```
RP/0/RSP1/CPU0:viking-1(admin-config)#  
username admin-root
```

Rack/Slot/Module/Port

```
RP/0/RSP0/CPU0: asr9001(admin)#show platform
```

Node	Type	State	Config State
0/RSP0/CPU0	ASR9001-RP(Active)	IOS XR RUN	PWR, NSHUT, MON
0/FT0/SP	FAN TRAY	READY	
0/0/CPU0	ASR9001-LC	IOS XR RUN	PWR, NSHUT, MON
0/0/0	A9K-MPA-4X10GE	OK	PWR, NSHUT, MON
0/0/1	A9K-MPA-4X10GE	OK	PWR, NSHUT, MON
0/PM0/SP	A9K-750W-AC	READY	PWR, NSHUT, MON


```
RP/0/RSP0/CPU0: asr9001-nv-Edge(admin)#show platform
```

Node	Type	State	Config State
0/RSP0/CPU0	ASR9001-RP(Active)	IOS XR RUN	PWR, NSHUT, MON
0/FT0/SP	FAN TRAY	READY	
0/0/CPU0	ASR9001-LC	IOS XR RUN	PWR, NSHUT, MON
0/0/0	A9K-MPA-20X1GE	OK	PWR, NSHUT, MON
0/0/1	A9K-MPA-4X10GE	OK	PWR, NSHUT, MON
0/PM0/SP	A9K-750W-DC	READY	PWR, NSHUT, MON
1/RSP0/CPU0	ASR9001-RP(Active)	IOS XR RUN	PWR, NSHUT, MON
1/FT0/SP	FAN TRAY	READY	
1/0/CPU0	ASR9001-LC	IOS XR RUN	PWR, NSHUT, MON
1/0/0	A9K-MPA-20X1GE	OK	PWR, NSHUT, MON
1/0/1	A9K-MPA-2X10GE	OK	PWR, NSHUT, MON
1/PM0/SP	A9K-750W-DC	READY	PWR, NSHUT, MON

A cluster node

XR Configuration Key Concepts

Two Stage Commit, Verification



**Syntax Check
after each line
PASSES**

```
Interface GigabitEthernet0/3/0/0
 ipv4 address 9.9.9.9/24
 taskgroup bgp
 task read bgp
 task write bgp
```

Target Configuration

**Active Configuration
Before Commit**

```
hostname odin
line default
 exec-timeout 1440 0
!
 taskgroup ops
 task read boot
 task write boot
 task execute bgp
!
router static
 address-family ipv4 unicast
 0.0.0.0/0 7.1.9.1
 7.7.7.77/32 7.1.9.1
```

**Active Configuration
After Commit
No Change**

**Semantic Check
during commit
FAILS
BGP cannot be
taskgroup name**



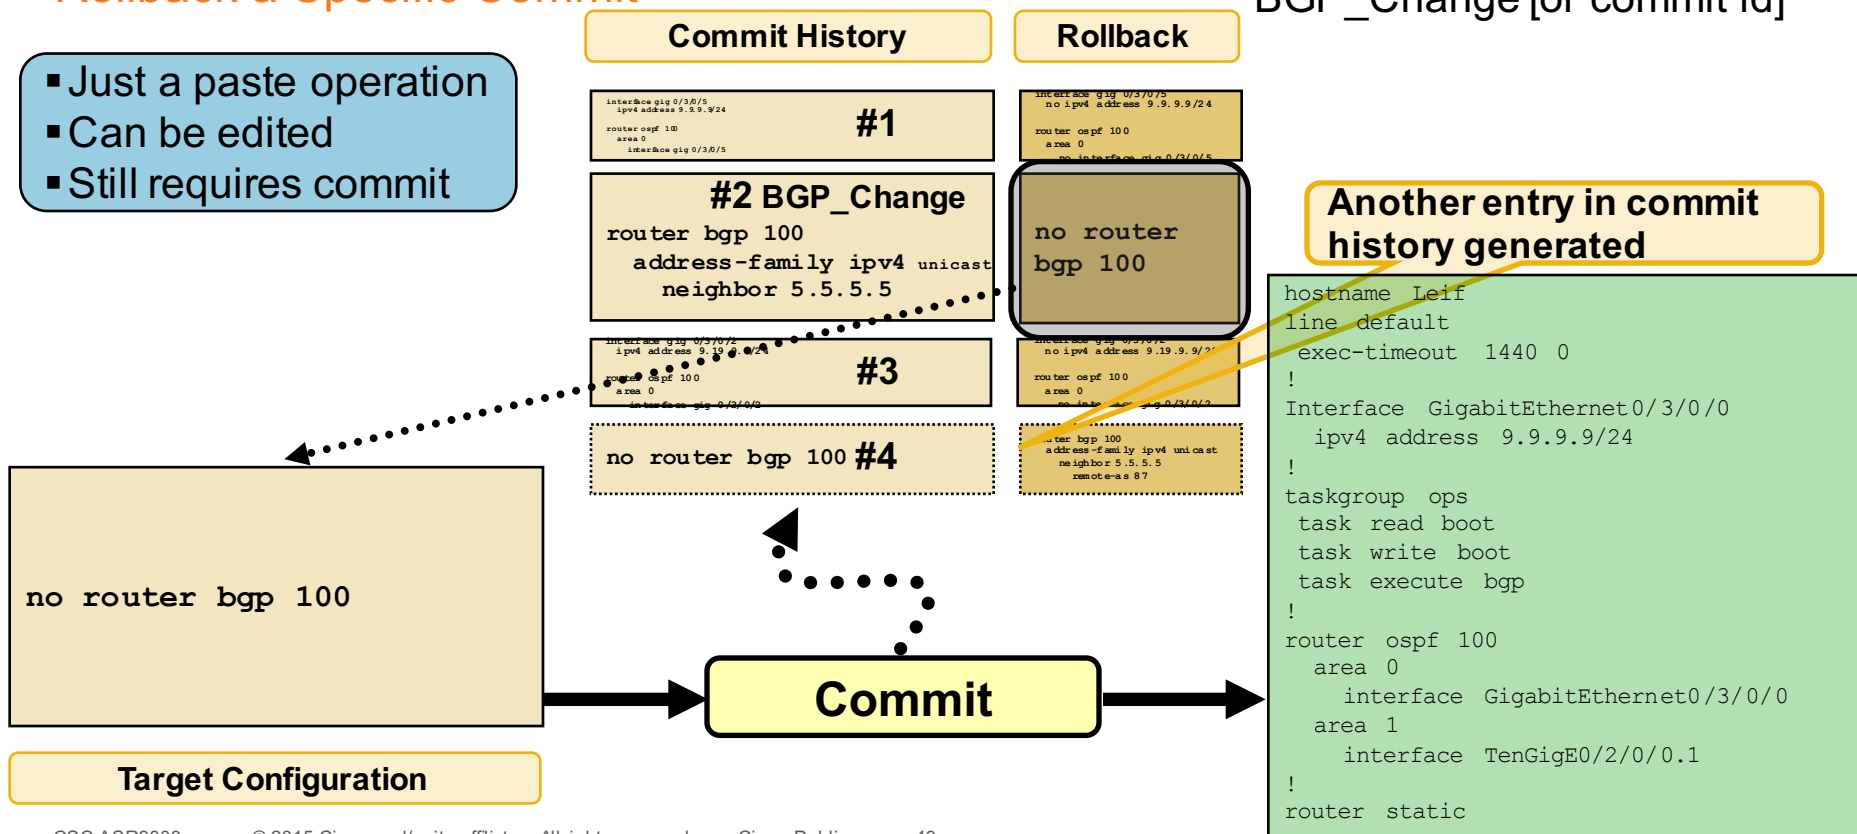
```
hostname odin
line default
 exec-timeout 1440 0
!
 taskgroup ops
 task read boot
 task write boot
 task execute bgp
!
router static
 address-family ipv4 unicast
 0.0.0.0/0 7.1.9.1
 7.7.7.77/32 7.1.9.1
```


XR Configuration Key Concepts

Rollback a Specific Commit

- Just a paste operation
- Can be edited
- Still requires commit

(config)#load rollback changes
BGP_Change [or commit id]



XR Configuration Key Concepts

Concurrent config sessions and exclusive



Enter Proposed Changes

```
interface gig 0/3/0/0
  ipv4 address 9.9.9.9/24

router ospf 100
  area 0
    interface gig 0/3/0/0
  area 1
    interface gig 0/4/0/0
```

First to Commit

Normal Commit
only first user's changes



Enter Proposed Changes

```
interface gig 0/3/0/0
  ipv4 address 9.9.9.7/24

router ospf 100
  area 2
    interface gig 0/3/0/0
  area 4
    interface gig 0/4/0/0
```

Second to Commit

Use **config exclusive** mode to block other users from committing

One or more commits have occurred from other configuration sessions since this session started or since the last commit was made from this session.
You can use the 'show configuration commit changes' command to browse the changes.
Do you wish to proceed with this commit anyway? [no]:

XR Configuration Key Concepts

Monitoring Configuration , commits, changes, sessions

```
RP/0/RSP0/CPU0:rasr9000-2w-b#show running-config
```

```
Mon Jan 5 00:35:19.951 EST
Building configuration...
!! IOS XR Configuration 5.1.3
!! Last configuration change at Tue Dec 2 22:19:25 2014 by cisco
!
service unsupported-transceiver
hostname rasr9000-2w-b
clock timezone EST -5
.
```

```
RP/0/RSP0/CPU0:rasr9000-2w-b#show configuration commit list
```

```
Mon Jan 5 00:35:34.747 EST
SNo. Label/ID User Line Client Time Stamp
~~~~ ~~~~~~ ~~~~ ~~~~~ ~~~~~~ ~~~~~~
1 1000000033 cisco con0_RSP0_CPU0 CLI Tue Dec 2 22:19:25 2014
2 1000000032 cisco con0_RSP0_CPU0 CLI Tue Dec 2 22:15:54 2014
3 1000000031 cisco con0_RSP0_CPU0 CLI Tue Dec 2 22:14:18 2014
.
```

```
RP/0/RSP0/CPU0:rasr9000-2w-b#show configuration commit changes last 5
```

```
Mon Jan 5 00:36:17.813 EST
Building configuration...
!! IOS XR Configuration 5.1.3
interface GigabitEthernet0/1/0/4
 ipv4 address 10.6.7.6 255.255.255.0
!
.
```

```
RP/0/RSP1/CPU0:viking-1# show config sessions
```

Current Configuration Session	Line	User	Date	Lock
00000051-004c4104-00000000	con0_RSP1_	ww	Tue Jul 21 16:58:22 2009	

```
RP/0/0/CPU0:P103#show running-config
router bgp
```

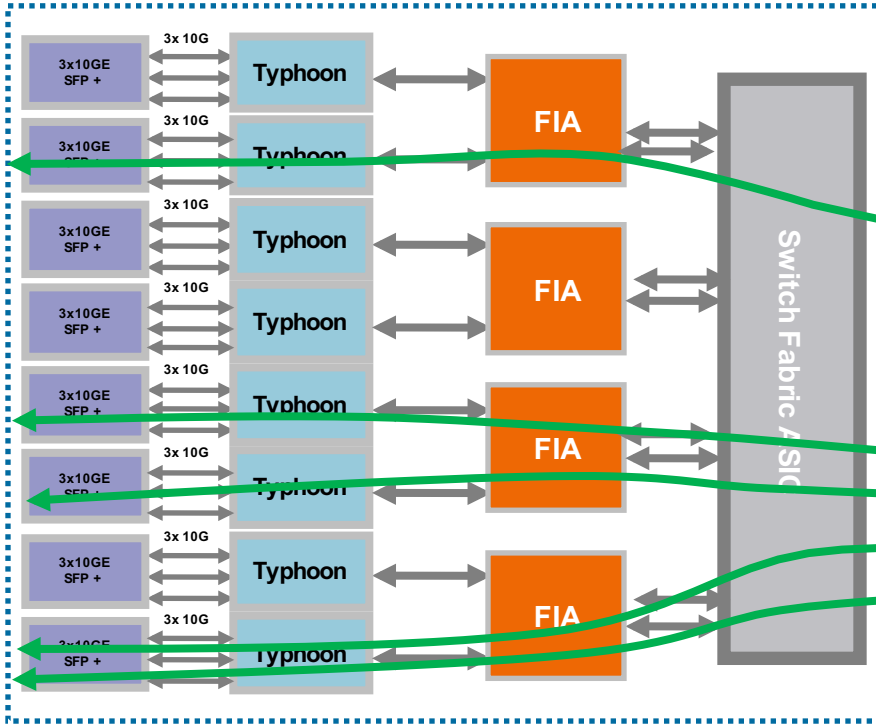
```
Mon Jan 5 15:06:29.774 EST
router bgp 65001
 bgp router-id 10.100.103.1
 bgp cluster-id 10.100.101.1
 bgp graceful-restart
 address-family ipv4 unicast
 !
 address-family vpnv4 unicast
 !
 address-family ipv6 unicast
 !
 address-family l2vpn vpls-vpws
 !
 neighbor-group iBGP
 remote-as 65001
 keychain CISCO
 update-source Loopback0
 address-family ipv4 unicast
 route-reflector-client
 !
 address-family vpnv4 unicast
 route-reflector-client
 !
 address-family ipv6 unicast
 route-reflector-client
```

A nighttime photograph of a city street. In the foreground, there are long, curved light trails from cars, primarily in shades of yellow and orange. In the middle ground, a pedestrian bridge with a blue light strip runs across the street. In the background, there are several tall buildings with lit windows and some flags on poles. The overall scene is illuminated by city lights.

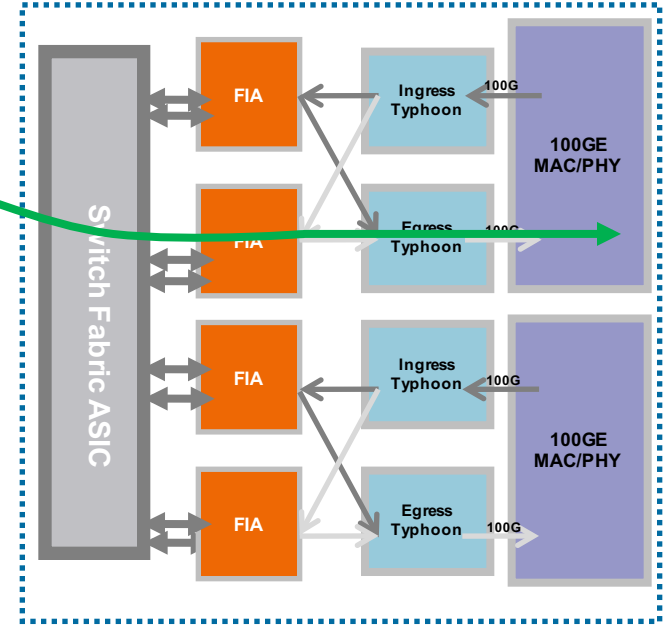


Frame Journey
Unicast Packet Forwarding

Unicast Transit Frame Path

Physical > NP > FIA > Fabric > FIA > NP > Physical



A9K-24X10GE




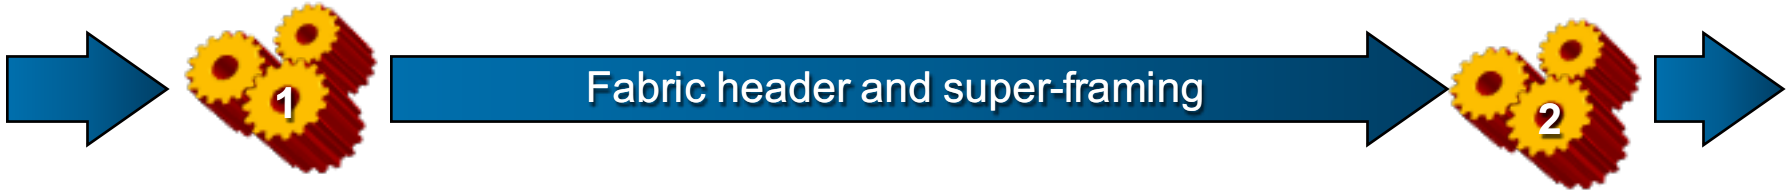
A9K-2X100GE

Unicast Two Stage Forwarding

By ingress NP and egress NP

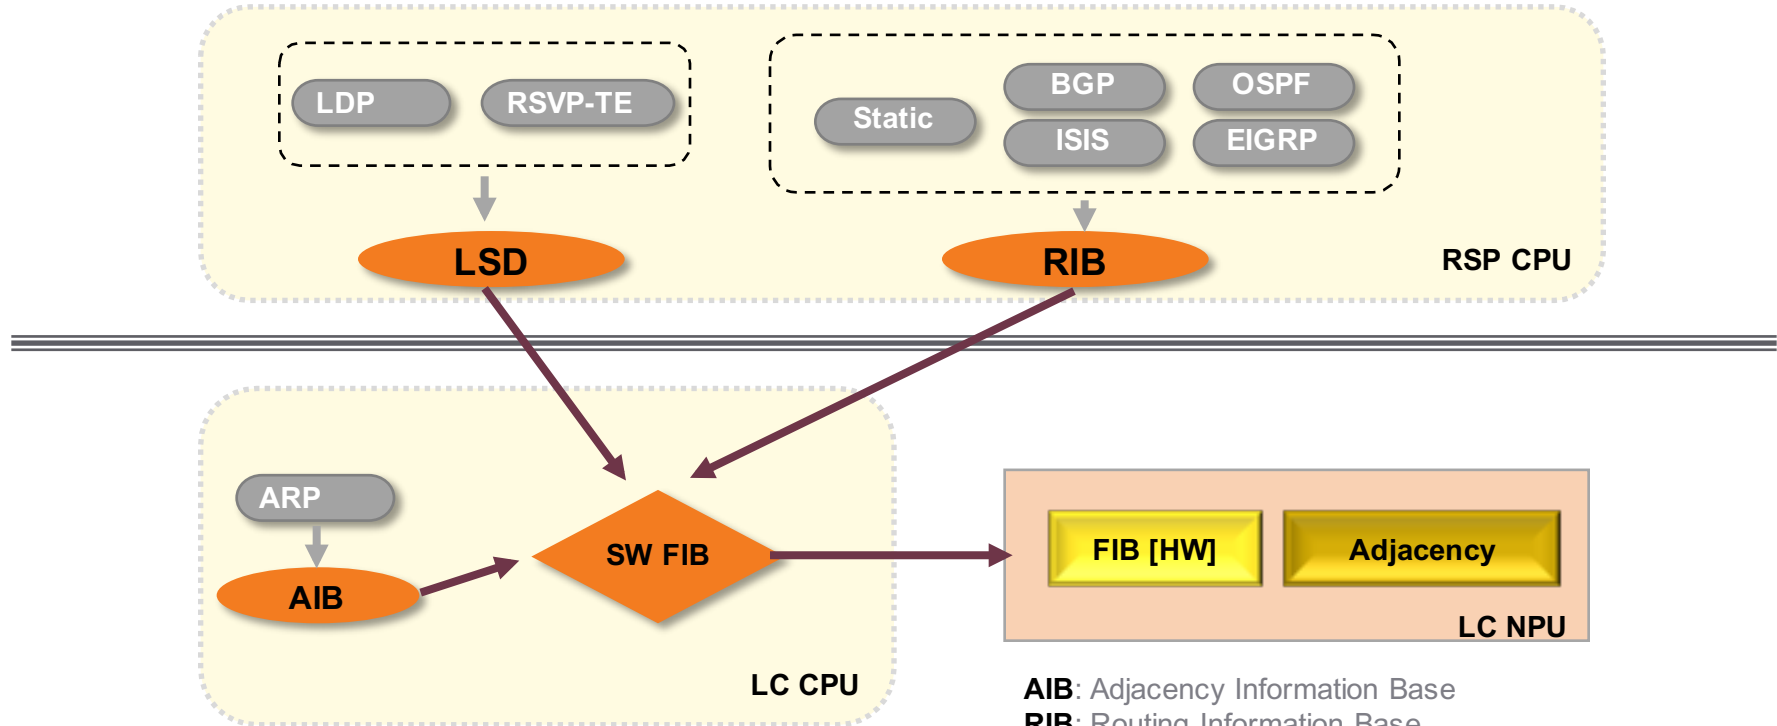
First stage: lookup on ingress NPU →
Egress NPU (or **SFP**: switch fabric port)

Second stage: lookup on egress NPU →
Egress port and rewrite information



The NP FIB

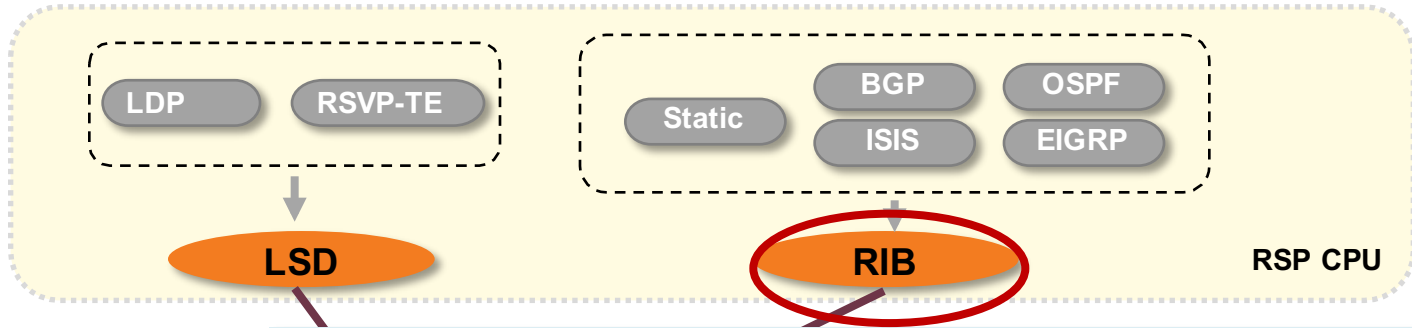
From RP control plane to data plane NP



AIB: Adjacency Information Base
RIB: Routing Information Base
FIB: Forwarding Information Base
LSD: Label Switch Database


The NP FIB

RIB info: example



```
RP/0/RSP0/CPU0:asr#sh route 222.0.0.6/31
```

```
Routing entry for 222.0.0.6/31
  Known via "isis isis1", distance 115, metric 20, type level-1
  Installed Mar  2 17:58:12.251 for 00:00:47
  Routing Descriptor Blocks
    222.0.0.2, from 222.2.2.1, via TenGigE0/1/0/3
    Route metric is 20
  No advertising protos.
```



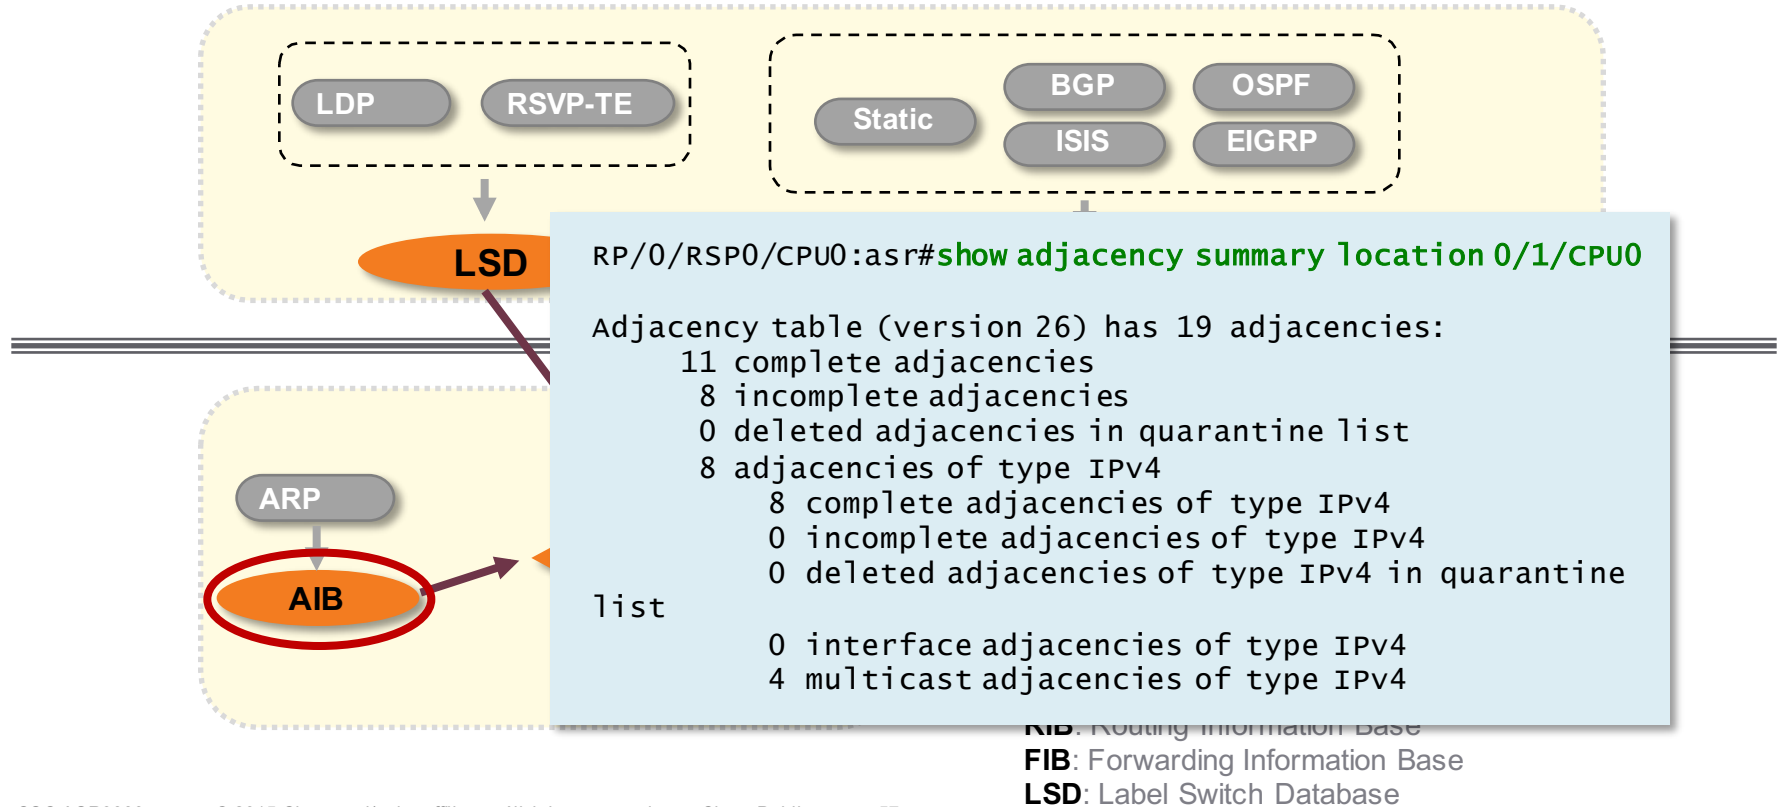
RIB: Routing Information Base

FIB: Forwarding Information Base

LSD: Label Switch Database

The NP FIB

Line card adjacency




The NP FIB

FIB entry in NP: example

```
RP/0/RSP0/CPU0:asr#sh cef 222.0.0.6 hardware ingress to 0/1/CPU0
222.0.0.6/31, version 1, internal 0x40000001 (0xb1d66c6c) [1], 0x0
(0xb1b4f758), 0x0 (0x0)
Updated Mar 2 17:58:11.987
local adjacency 222.0.0.2
Prefix Len 31, traffic index 0, precedence routine (0)
  via 222.0.0.2, TenGigE0/1/0/3, 5 dependencies, weight 0, class 0
  next hop 222.0.0.2
  local adjacency
    EZ:0 Leaf
    =====
Search ctrl-byte0: 0x3 ctrl-byte1: 0x8 ctrl-byte2:0x5
Leaf Action : FORWARD
prefix length : 31
Search Control Flags :
  match      : 1      valid: 1
  done       : 0      ifib_lookup: 0
  ext_lsp_array : 0    match_all_bit: 0
  recursive  : 0      nonrecursive : 1
  default_action: 1

Non Recursive Leaf:
-----
ldi ptr : 10936 (0x2ab8)      igp statsptr:0
rpf ptr : 0x0000
```

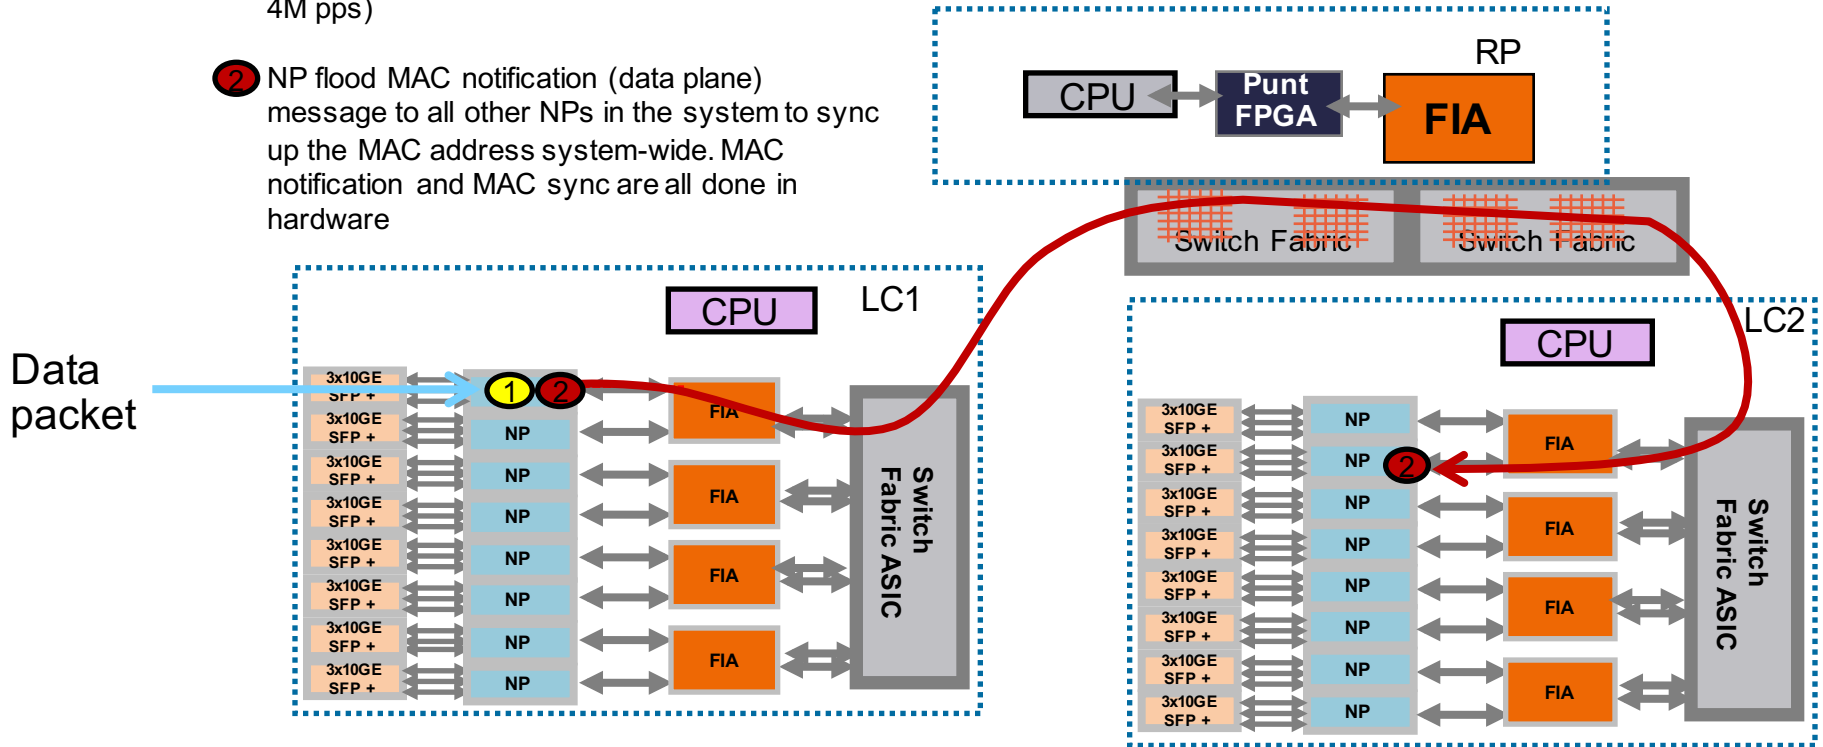


AIB: Adjacency Information Base
RIB: Routing Information Base
FIB: Forwarding Information Base
LSD: Label Switch Database

MAC learning and synchronization

- 1 NP learn MAC address in hardware (around 4M pps)
- 2 NP flood MAC notification (data plane) message to all other NPs in the system to sync up the MAC address system-wide. MAC notification and MAC sync are all done in hardware

Hardware based MAC learning: ~4Mpps/NP



MAC learning and synchronization

show controllers np struct 18 detail all-entries npx location
0/x/cpu0 | i 0000c07acd0 <<< replace your client MAC

```
RP/0/RSP0/CPU0:rasr9k-1y#show 12vpn forwarding bridge-domain BRIDGES:DOMAIN-A mac-address  
hardware ingress location 0/4/CPU0
```

```
Fri Feb 22 18:50:08.433 UTC
```

```
To Resynchronize MAC table from the Network Processors, use the command...
```


```
12vpn resynchronize forwarding mac-address-table location <r/s/i>
```

Mac Address	Type	Learned from/Filtered on	LC learned	Resync Age	Mapped to
0000.c001.0167	dynamic	Gi0/0/1/0.200	0/0/CPU0	0d 0h 0m 18s	N/A
0000.c001.016b	dynamic	Gi0/0/1/0.200	0/0/CPU0	0d 0h 0m 11s	N/A
0000.c001.016c	dynamic	Gi0/0/1/0.200	0/0/CPU0	0d 0h 0m 9s	N/A
0000.c001.016d	dynamic	Gi0/0/1/0.200	0/0/CPU0	0d 0h 0m 18s	N/A
0000.c001.016e	dynamic	Gi0/0/1/0.200	0/0/CPU0	0d 0h 0m 20s	N/A
0000.c001.016f	dynamic	Gi0/0/1/0.200	0/0/CPU0	0d 0h 0m 8s	N/A
0000.c001.0171	dynamic	Gi0/0/1/0.200	0/0/CPU0	0d 0h 0m 17s	N/A
0000.c001.0102	dynamic	Te0/4/0/20.101	0/4/CPU0	0d 0h 0m 16s	N/A
0000.c001.0104	dynamic	Te0/4/0/20.101	0/4/CPU0	0d 0h 0m 20s	N/A
0000.c001.0105	dynamic	Te0/4/0/20.101	0/4/CPU0	0d 0h 0m 8s	N/A
0000.c001.0106	dynamic	Te0/4/0/20.101	0/4/CPU0	0d 0h 0m 9s	N/A
0000.c001.0107	dynamic	Te0/4/0/20.101	0/4/CPU0	0d 0h 0m 18s	N/A
0000.c001.0108	dynamic	Te0/4/0/20.101	0/4/CPU0	0d 0h 0m 15s	N/A
0000.c001.0109	dynamic	Te0/4/0/20.101	0/4/CPU0	0d 0h 0m 3s	N/A
0000.c001.010a	dynamic	Te0/4/0/20.101	0/4/CPU0	0d 0h 0m 4s	N/A

```
.
```

L3 Unicast Packet Journey

Mapping the port to NP and FIA



! Example: Path from GigabitEthernet0/0/1/0 192.3.1.2 TO TenGigE0/4/0/20.6 192.6.1.2

```
RP/0/RSP0/CPU0:rasr9k-1y#show controllers NP ports all location 0/0/CPU0
```

```
Fri Feb 22 15:57:32.307 UTC
```

```
Node: 0/0/CPU0:
```

NP	Bridge	Fia	Ports
0	--	0	TenGigE0/0/0/0, TenGigE0/0/0/1, TenGigE0/0/0/2, TenGigE0/0/0/3
1	--	1	GigabitEthernet0/0/1/0 - GigabitEthernet0/0/1/19

Map the port to NP and FIA

```
RP/0/RSP0/CPU0:rasr9k-1y#show controllers NP ports all location 0/4/CPU0
```


```
Fri Feb 22 15:55:22.370 UTC
```

```
Node: 0/4/CPU0:
```

NP	Bridge	Fia	Ports
0	--	0	TenGigE0/4/0/0, TenGigE0/4/0/1, TenGigE0/4/0/2
1	--	0	TenGigE0/4/0/3, TenGigE0/4/0/4, TenGigE0/4/0/5
2	--	1	TenGigE0/4/0/6, TenGigE0/4/0/7, TenGigE0/4/0/8
3	--	1	TenGigE0/4/0/9, TenGigE0/4/0/10, TenGigE0/4/0/11
4	--	2	TenGigE0/4/0/12, TenGigE0/4/0/13, TenGigE0/4/0/14
5	--	2	TenGigE0/4/0/15, TenGigE0/4/0/16, TenGigE0/4/0/17
6	--	3	TenGigE0/4/0/18, TenGigE0/4/0/19, TenGigE0/4/0/20
7	--	3	TenGigE0/4/0/21, TenGigE0/4/0/22, TenGigE0/4/0/23

L3 Unicast Packet Journey

The egress interface identifier



```
RP/0/RSP0/CPU0:rasr9k-1y#show controllers pm interface tenGigE 0/4/0/20.6
```

```
Fri Feb 22 16:45:22.404 UTC
```

```
Ifname(1): TenGigE0_4_0_20.6, ifh: 0xc001340 :
```

```
iftype 0x19
```

```
egress_uidb_index 0x1d
```

```
ingress_uidb_index 0x1d
```

```
port_num 0x14
```

```
subslot_num 0x0
```

```
phy_port_num 0x14
```

```
channel_id 0x6
```

```
channel_map 0x0
```

```
lag_id 0x0
```

```
virtual_port_id 0x0
```

```
switch_fabric_port 0x136
```

```
in_tm_qid_fid0 0x0
```

```
in_tm_qid_fid1 0x0
```

```
in_qos_drop_base 0x0
```

```
out_tm_qid_fid0 0x0 0x0 0x0 0x0 0x0 0x0 0x0
```

```
0x0
```

```
out_tm_qid_fid1 0x0 0x0 0x0 0x0 0x0 0x0 0x0
```

```
0x0
```

```
out_qos_drop_base 0x0
```

```
bandwidth 10000000 kbps
```

```
.
```

Get internal identifiers

Interface handle: unique to logical sub-interface

Fabric port, shared between all port sub-interfaces. The fabric destination.

L3 Unicast Packet Journey

Ingress NP FIB

```
RP/0/RSP0/CPU0:rasr9k-1y#show cef ipv4 192.6.1.2
hardware ingress location 0/0/CPU0
Fri Feb 22 17:40:35.887 UTC
192.6.1.0/24, version 364, attached, connected,
internal 0xc0000c1 (ptr 0x8856b534) [1], 0x0
(0x873dde50), 0x0 (0x0)
Updated Feb 22 16:09:42.862
remote adjacency to TenGigE0/4/0/20.6
Prefix Len 24, traffic index 0, precedence
routine (0), priority 0
via TenGigE0/4/0/20.6, 2 dependencies, weight
0, class 0 [flags 0x8]
path-idx 0 [0x8a60a7bc 0x0]
remote adjacency
LEAF - HAL pd context :
sub-type : IPV4, ecd_marked:0,
has_collapsed_ldi:0, collapse_bwalk_required:0,
ecdv2_marked:0
Leaf H/W Result:

Physical Result: 0x11dd0600 (LE)

Raw Data0: 0x91ad1000 00000001 360c0013
40000000
Raw Data1: 0x00000000 00000136 00180000
00000000
```

```
RX H/W Result on NP:1 [Adj ptr:0x3a (BE)]:

Raw Data0: 0x91000000 00000136 0c001340 00000000
adj_resolve_control_byte0
match: 1
valid: 1
iptunl_adj: 0
remote_rack: 0

adj_resolve_control_byte1
adj_down: 0
mgscp_en: 0
rx_lag_hash_en: 0
rx_lag_adj: 0

adj_resolve_control_byte2
rx_lag_adj: 0
rx_adj_null0: 0
rp_destined: 0
rx_punt: 0
rx_drop: 0
sfp/vqi : 0x136
if_handle : 0xc001340
```

L3 Unicast Packet Journey

Egress NP FIB

```
RP/0/RSP0/CPU0:rasr9k-1y#show cef ipv4 192.6.1.2
hardware egress location 0/4/CPU0
Fri Feb 22 17:55:28.494 UTC
192.6.1.2/32, version 0, internal 0x4080001 (ptr
0x8efc2704) [1], 0x0 (0x8e0f2210), 0x0 (0x0)
Updated Feb 22 16:13:35.351
local adjacency 192.6.1.2
Prefix Len 32, traffic index 0, Adjacency-
prefix, precedence routine (0), priority 0
via 192.6.1.2, TenGigE0/4/0/20.6, 3
dependencies, weight 0, class 0 [flags 0x0]
path-idx 0 [0x91a2cef8 0x0]
next hop 192.6.1.2
local adjacency
LEAF - HAL pd context :
sub-type : IPV4, ecd_marked:0,
has_collapsed_ldi:0, collapse_bwalk_required:0,
ecd_v2_marked:0
Leaf H/W Result:

Physical Result: 0x11e80300 (LE)

Raw Data0: 0x91ad1000 8a030001 360c0013
40400000
.
```

```
TX H/W Result for NP:6 (index: 0x38a (BE)):
Raw Data0: 0x91080000 1d000000 dc050000 400b5f00
Raw Data1: 0x0000c006 01020000 00000000 00000000
adj_resolve_control_byte0
reserved: 0
egr_uidb_internal: 1 match: 1
valid: 1 iptun1_adj: 0
adj_resolve_control_byte1
tx_adj_null0: 0
tx_punt: 0 tx_drop: 0
default_action: 1
spare: 0
adj_resolve_control_byte2
spare: 0
spare_cb: 0
flags
gre_adj : 0
uidb_index : 0x1d00 (LE)
reserve_pad_word: 0
13_mtu : 1500
reserve_pad_1 : 0
adj_stats_index : 0x400b5f00
dest_mac : 0x0000.c006.0102
ether reserved : 00000000000000000000
.
```

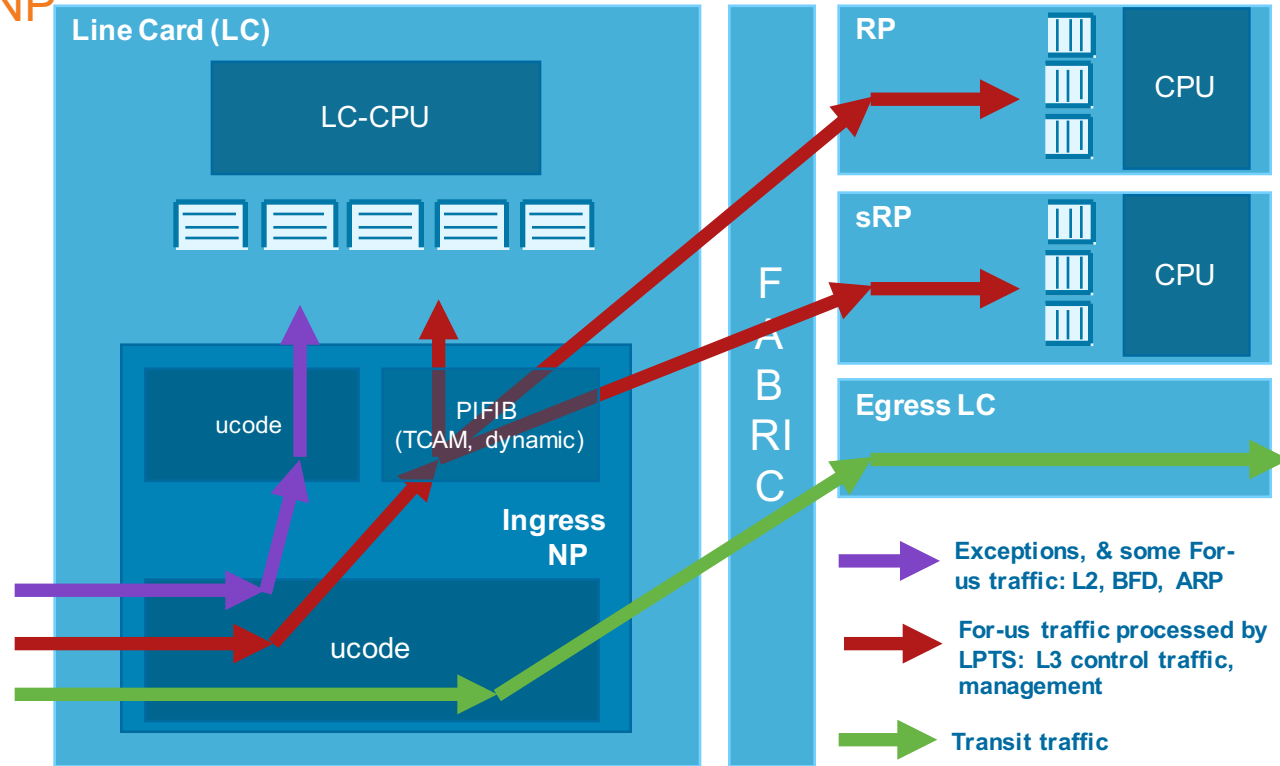

A nighttime photograph of a city street. In the background, there are modern buildings with lit windows and a pedestrian bridge with blue lighting. The middle ground shows a road with traffic lights and light trails from vehicles. The foreground is dominated by long, curved light trails in yellow, orange, and red, suggesting a long-exposure shot of a busy intersection.

Control, Management, & Security
For Us, Exceptions, & Resource Protection

Traffic: Transit, For us, and Exceptions

Differentiate on ingress NP

- Transit
 - Look up, re-write, forward
- For us
 - Destined to RP, or link local scope
 - Punt to RP or ingress LC CPU
- Exception
 - MTU failure, TTL failure, etc. Should have been transit
 - Punt to LC CPU

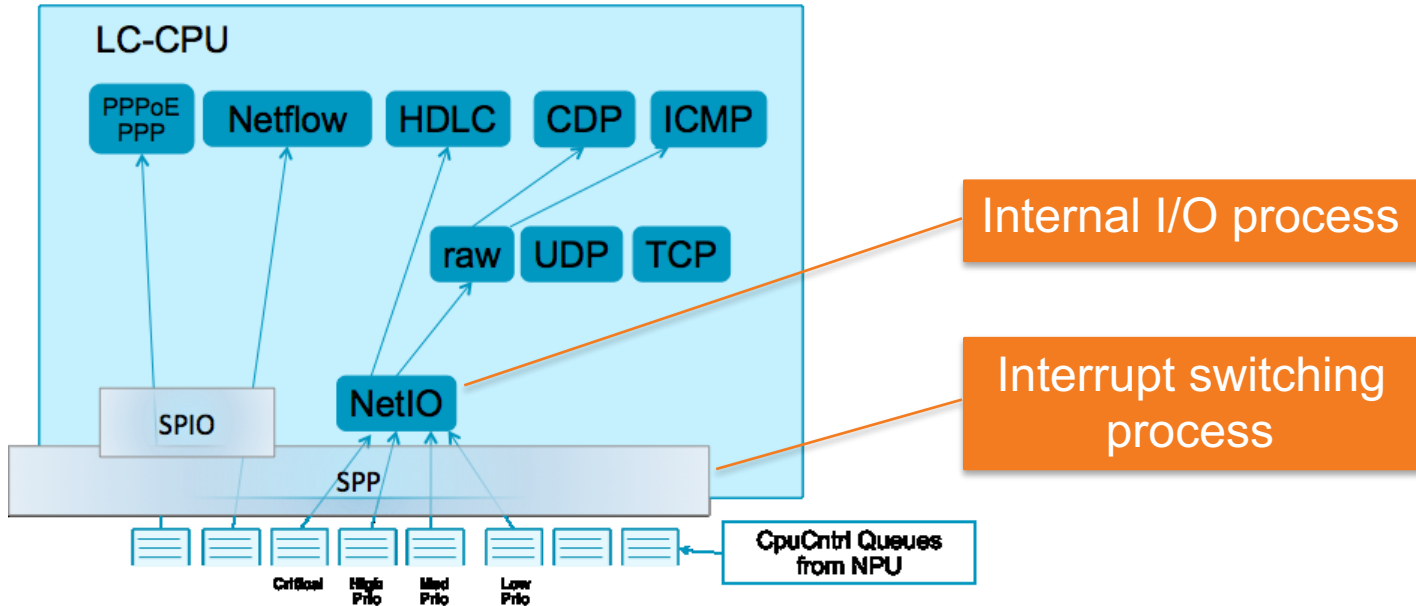


For Us Frame Path

From ingress NP to RP CPU or LC CPU

Inside LC CPU

NetIO ~ IP INPUT, process switching
SPP ~ rx interrupt handler



Control Plane Protection

LPTS flow type policers

#clear controller np counters all location ...

```
RP/0/RSP0/CPU0:rasr9k-1y#show lpts pifib hardware police location 0/0/CPU0
Sun Dec  2 06:07:36.931 UTC
-----
Node 0/0/CPU0
Burst = 100ms for all flows
-----
```

FlowType	Police ID	Type	Cur. Rate	Def. Rate	Accepted	Dropped
unconfigured-default	100	Static	2500	2500	0	0
Fragment	101	Local	0	2500	0	0
OSPF-mc-known	102	Static	2000	2000	0	0
OSPF-mc-default	103	Local	0	1500	53	26
OSPF-uc-known	104	Static	2000	2000	0	0
OSPF-uc-default	105	Local	0	1500	0	0
ISIS-known	143	Static	0	0	0890	0
ISIS-default	144	Local	0	0	0	0
BGP-known	106	Static	2500	2500	4070	0
BGP-cfg-peer	107	Static	2000	2000	17	0
BGP-default	108	Local	400000	1500	138918630	3848639925
PIM-mcast-default	109	Local	0	0	0	0
PIM-mcast-known	176	Static	2500	2500	0	0
PIM-ucast	110	Static	1500	1500	0	0
IGMP	111	Static	3000	3000	0	0

Established session packets

Configured peer packets

BGP packets from unknown


Management

Management Plane Protocols

- Simple Network Management Protocol (SNMP)
- Telnet / SSH
- SFTP / FTP / TFTP
- Secure Copy Protocol (SCP)
- TACACS+ / RADIUS
- NetFlow (also used by the Data Plane as that is where the traffic comes from)
- Network Time Protocol (NTP)
- Syslog

Management Traffic

Management interfaces: Out-of-band, in-band, and “global”



- No communication permitted between inband and out-of-band
- Management VRF is not necessary, but looks cleaner

Management Traffic

Out-of-Band: Virtual address, interfaces, and protocols

```
RP/0/RSP0/CPU0:rasr9k-1y#show running-config
control-plane management-plane out-of-band
Wed Dec  5 00:45:07.132 UTC
control-plane
management-plane
out-of-band
vrf MGMT
interface MgmtEth0/RSP0/CPU0/0
  allow SSH peer
  address ipv4 172.16.1.0/24
!
  allow SNMP peer
  address ipv4 172.16.1.98
!
!
interface MgmtEth0/RSP1/CPU0/0
  allow SSH peer
  address ipv4 172.16.1.0/24
!
  allow SNMP peer
  address ipv4 172.16.1.98
!
!
```

```
RP/0/RSP0/CPU0:rasr9k-1y#show mgmt-plane
Wed Dec  5 00:46:26.162 UTC

Management Plane Protection

inband interfaces
-----
interface - TenGigE0_0_0_2/
  ssh configured -
  peer v4 allowed - 192.168.1.0/24

outband interfaces
-----
interface - MgmtEth0_RSP0_CPU0_0/
  ssh configured -
  peer v4 allowed - 172.16.1.0/24
  snmp configured -
  peer v4 allowed - 172.16.1.98
interface - MgmtEth0_RSP1_CPU0_0/
  ssh configured -
  peer v4 allowed - 172.16.1.0/24
  snmp configured -
  peer v4 allowed - 172.16.1.98
```

Management Traffic

In-band: If OOB is not available

```
RP/0/RSP0/CPU0: rasr9k-1y#show running-config
control-plane management-plane inband
Tue Dec 11 23:05:11.597 UTC
control-plane
management-plane
inband
interface TenGigE0/0/0/2
allow SSH peer
address ipv4 192.168.1.0/24
!
!
!
!
```

```
RP/0/RSP0/CPU0: rasr9k-1y#show mgmt-plane
Wed Dec 11 00:46:26.162 UTC

Management Plane Protection

inband interfaces
-----
interface - TenGigE0_0_0_2/
ssh configured -
peer v4 allowed - 192.168.1.0/24

outband interfaces
-----
interface - MgmtEth0_RSP0_CPU0_0/
ssh configured -
peer v4 allowed - 172.16.1.0/24
snmp configured -
peer v4 allowed - 172.16.1.98
interface - MgmtEth0_RSP1_CPU0_0/
ssh configured -
peer v4 allowed - 172.16.1.0/24
snmp configured -
peer v4 allowed - 172.16.1.98
```


A nighttime photograph of a city street. In the background, there are several tall buildings with lit windows and balconies. A pedestrian bridge with a glass railing spans across the street. In the middle ground, a traffic light pole with green lights is visible. The foreground is dominated by long, curved light trails from moving vehicles, creating a sense of motion and energy. The overall scene is illuminated by city lights, creating a vibrant and dynamic atmosphere.

Troubleshooting System Diagnostics

Background Diagnostics

RSP default diagnostics

```
RP/0/RSP0/CPU0:rasr9000-2w-a#admin show diagnostic content location 0/RSP0/CPU0
wed Dec 11 19:44:32.957 EST
```

```
RP 0/RSP0/CPU0:
```

```
Diagnostics test suite attributes:
```

- M/C/* - Minimal bootup level test / Complete bootup level test / NA
- B/O/* - Basic ondemand test / not Ondemand test / NA
- P/V/* - Per port test / Per device test / NA
- D/N/* - Disruptive test / Non-disruptive test / NA
- S/* - Only applicable to standby unit / NA
- X/* - Not a health monitoring test / NA
- F/* - Fixed monitoring interval test / NA
- E/* - Always enabled monitoring test / NA
- A/I - Monitoring is active / Monitoring is inactive

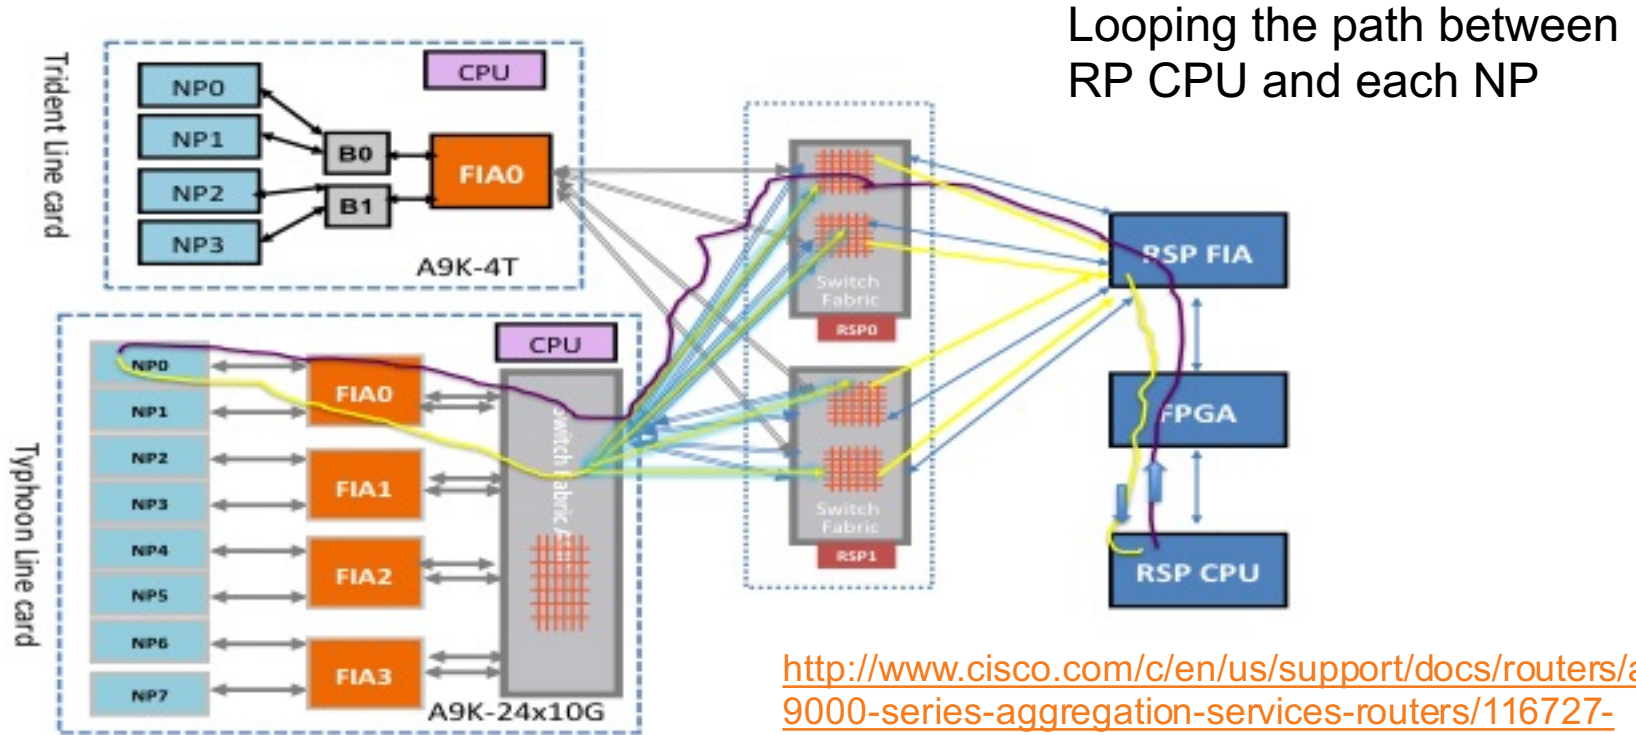
Every minute or every 5 seconds

ID	Test Name	Attributes	Test Interval (day hh:mm:ss.ms)	Thre- shold
1)	CPUctrlScratchRegister	***N***A	000 00:01:00.000	1
2)	ClkCtrlScratchRegister	***N***A	000 00:01:00.000	1
3)	ZenJfScratchRegister	***N***A	000 00:01:00.000	1
4)	FabSwitchIdRegister	*B*N***A	000 00:01:00.000	1
5)	SrspStandbyEobcHeartbeat	*B*NS***A	000 00:00:05.000	3
6)	SrspActiveEobcHeartbeat	*B*NS***A	000 00:00:05.000	3
7)	FabricLoopback	MB*N***A	000 00:01:00.000	3
8)	PuntFabricDataPath	*B*N***A	000 00:01:00.000	3

Error threshold (consecutive)

Background Diagnostics

Test example: PuntFabricDataPath



<http://www.cisco.com/c/en/us/support/docs/routers/asr-9000-series-aggregation-services-routers/116727-troubleshoot-punt-00.html>

Background Diagnostics

LC default diagnostics

```
RP/0/RSP0/CPU0:rasr9000-2w-a#admin show diagnostic content location 0/0/CPU0
wed Dec 11 20:32:08.842 EST
```

```
A9K-24x10GE-SE 0/0/CPU0:
```

```
Diagnostics test suite attributes:
```

- M/C/* - Minimal bootup level test / Complete bootup level test / NA
- B/O/* - Basic ondemand test / not Ondemand test / NA
- P/V/* - Per port test / Per device test / NA
- D/N/* - Disruptive test / Non-disruptive test / NA
- S/* - Only applicable to standby unit / NA
- X/* - Not a health monitoring test / NA
- F/* - Fixed monitoring interval test / NA
- E/* - Always enabled monitoring test / NA
- A/I - Monitoring is active / Monitoring is inactive

ID	Test Name	Attributes	Test Interval (day hh:mm:ss.ms)	Thre- shold
1)	CPUCtrlScratchRegister	*B*N***A	000 00:01:00.000	1
2)	PHYCtrlScratchRegister	*B*N***A	000 00:01:00.000	1
3)	PortCtrlScratchRegister	*B*N***A	000 00:01:00.000	1
4)	FIAScratchRegister	*B*N***A	000 00:01:00.000	1
5)	LcEobcHeartbeat	*B*N***A	000 00:00:05.000	3
6)	NPULoopback	*B*N***A	000 00:01:00.000	3

Background Diagnostics

Errors and clears

- Set: threshold exceeded
 - Path to LC 2 NP 0 failed in this example
- Clear: test previously failing, now passed
 - Indication of “transient” fault. Keep watching
- “show pfm location all” shows platform errors reported

```
RP/0/RSP0/CPU0:Feb 5 05:05:44.051 :  
pfm_node_rp[354]:%PLATFORM-DIAGS-3-PUNT_FABRIC_DATA_PATH_FAILED :  
Set|online_diag_rsp[237686]|System Punt/Fabric/data Path  
Test(0x2000004)|failure threshold is 3, (slot, NP)failed:  
(0/2/CPU0, 0)
```

```
RP/0/RSP0/CPU0:Feb 5 05:05:46.051 :  
pfm_node_rp[354]:%PLATFORM-DIAGS-3-PUNT_FABRIC_DATA_PATH_FAILED :  
Clear|online_diag_rsp[237686]|System Punt/Fabric/data Path  
Test(0x2000004)|failure threshold is 3, (slot, NP)failed:  
(0/2/CPU0, 0)
```

LC & NP Resources

L3 forwarding resources

```
RP/0/RSP0/CPU0:rasr9000-2w-a#show cef resource hardware ingress detail location 0/0/CPU0
```

```
Wed Dec 11 22:27:57.319 EST
```

```
CEF resource availability summary state: GREEN
```

```
CEF will work normally
```

```
  ipv4 shared memory resource:
```

```
    CurrMode GREEN, CurrAvail 1609412608 bytes, MaxAvail 1683308544 bytes
```

```
  ipv6 shared memory resource:
```

```
    CurrMode GREEN, CurrAvail 1609412608 bytes, MaxAvail 1683308544 bytes
```

```
  mpls shared memory resource:
```

```
    CurrMode GREEN, CurrAvail 1609412608 bytes, MaxAvail 1683308544 bytes
```

```
  common shared memory resource:
```

```
    CurrMode GREEN, CurrAvail 1609412608 bytes, MaxAvail 1683308544 bytes
```

```
  DATA_TYPE_TABLE_SET hardware resource: GREEN
```

```
  DATA_TYPE_TABLE hardware resource: GREEN
```

```
  DATA_TYPE_IDB hardware resource: GREEN
```

```
  DATA_TYPE_IDB_EXT hardware resource: GREEN
```

```
  DATA_TYPE_LEAF hardware resource: GREEN
```

```
  DATA_TYPE_LOADINFO hardware resource: GREEN
```

```
  DATA_TYPE_PATH_LIST hardware resource: GREEN
```

```
  DATA_TYPE_NHINFO hardware resource: GREEN
```

```
  DATA_TYPE_LABEL_INFO hardware resource: GREEN
```

```
  DATA_TYPE_FRR_NHINFO hardware resource: GREEN
```

```
  DATA_TYPE_ECD hardware resource: GREEN
```

```
.
```

LC & NP Resources

L2 service resources

```
RP/0/RSP0/CPU0:rasr9000-2w-a#show l2vpn forwarding resource hardware ingress detail location  
0/0/CPU0
```

```
Wed Dec 11 22:26:07.070 EST
```

```
L2VPN forwarding resource availability summary state: GREEN
```

```
shared memory resource:
```

```
    CurrMode GREEN, CurrAvail 1609412608 bytes, MaxAvail 1650212864 bytes
```

```
AC hardware resource: GREEN
```

```
MPLS hardware resource: GREEN
```

```
PBB PORT hardware resource: GREEN
```

```
NHOP hardware resource: GREEN
```

```
L2TP hardware resource: GREEN
```

```
L2TP_SESSION hardware resource: GREEN
```

```
VFI hardware resource: GREEN
```

```
BRIDGE hardware resource: GREEN
```

```
BRIDGE SHG hardware resource: GREEN
```

```
BRIDGE PORT hardware resource: GREEN
```

```
BRIDGE MAC hardware resource: GREEN
```

```
MSTI MAIN PORT hardware resource: GREEN
```

```
BRIDGE MAIN PORT hardware resource: GREEN
```

```
MCAST TABLE hardware resource: GREEN
```

```
MCAST LEAF hardware resource: GREEN
```

```
MCAST XID hardware resource: GREEN
```

```
PBB BMAC SA hardware resource: GREEN
```

```
.
```

TCP Resources

TCP connections states

```
RP/0/RSP0/CPU0: rasr9000-2w-a#show tcp brief
```

```
Fri Dec 13 22:29:33.189 EST
```

PCB	VRF-ID	Recv-Q	Send-Q	Local Address	Foreign Address	State
0x10174ad8	0x60000000	0	0	:::179	:::0	LISTEN
0x101745f8	0x60000001	0	0	:::179	:::0	LISTEN
0x10174798	0x60000002	0	0	:::179	:::0	LISTEN
0x10174938	0x60000003	0	0	:::179	:::0	LISTEN
0x101cce2c	0x6000000b	0	0	:::179	:::0	LISTEN
0x1016fefc	0x00000000	0	0	:::179	:::0	LISTEN
0x101ac7ac	0x60000000	0	0	10.101.111.1:179	10.100.101.1:20100	ESTAB
0x1017bbf4	0x60000000	0	0	10.101.111.1:646	10.101.188.1:30687	ESTAB
0x10182b38	0x60000000	0	0	10.101.111.1:179	10.100.103.1:59214	ESTAB
0x1002e004	0x6000000d	0	0	10.100.111.1:17514	10.100.111.100:13680	ESTAB
0x10161e18	0x60000000	0	0	0.0.0.0:23	0.0.0.0:0	LISTEN
0x101cc968	0x00000000	0	0	0.0.0.0:23	0.0.0.0:0	LISTEN
0x101594b4	0x60000000	0	0	0.0.0.0:646	0.0.0.0:0	LISTEN
0x1016f7f0	0x60000000	0	0	0.0.0.0:179	0.0.0.0:0	LISTEN
0x10165d74	0x60000001	0	0	0.0.0.0:179	0.0.0.0:0	LISTEN
0x1016f4b0	0x60000002	0	0	0.0.0.0:179	0.0.0.0:0	LISTEN
0x1016f650	0x60000003	0	0	0.0.0.0:179	0.0.0.0:0	LISTEN
0x101ade54	0x6000000b	0	0	0.0.0.0:179	0.0.0.0:0	LISTEN
0x1016276c	0x00000000	0	0	0.0.0.0:179	0.0.0.0:0	LISTEN
0x1015e304	0x00000000	0	0	0.0.0.0:0	0.0.0.0:0	CLOSED

TCP Resources

TCP connections parameters

```
RP/0/RSP0/CPU0:rasr9000-2w-a#show tcp detail pcb 0x10182b38
Wed Dec 11 22:47:18.708 EST
```

```
=====
=
Connection state is ESTAB, I/O status: 0, socket status: 0
Established at Thu Dec 5 04:13:28 2013
```

```
PCB 0x10182b38, SO 0x10182968, TCPCB 0x101cc2b4, vrfid
0x60000000,
Pak Prio: Medium, TOS: 192, TTL: 255, Hash index: 717
Local host: 10.101.111.1, Local port: 179 (Local App PID:
287053)
Foreign host: 10.100.103.1, Foreign port: 59214
```

```
Current send queue size in bytes: 0 (max 24576)
Current receive queue size in bytes: 0 (max 32768) mis-
ordered: 0 bytes
Current receive queue size in packets: 0 (max 0)
```

Timer	Starts	Wakeups	Next(msec)
Retrans	9789	0	0
Sendwnd	0	0	0
Timewait	0	0	0
AckHold	9817	9605	0
KeepAlive	1	0	0
PmtuAger	0	0	0
GiveUp	0	0	0
Throttle	0	0	0

```
iss: 2038437204 snduna: 2038628753 sndnxt: 2038628753
sndmax: 2038628753 sndwnd: 31856 sndcwnd: 3648
irs: 249565735 rcvnxt: 2495854116 rcvwnd: 31894
rcvad: 2495886010
```

```
SRTT: 217 ms, RTTO: 300 ms, RTV: 11 ms, KRTT: 0 ms
minRTT: 1 ms, maxRTT: 289 ms
```

```
ACK hold time: 200 ms, Keepalive time: 0 sec, SYN waittime: 30
sec
Giveup time: 0 ms, Retransmission retries: 0, Retransmit
forever: FALSE
Connect retries remaining: 0, connect retry interval: 0 secs
```

```
State flags: none
Feature flags: MD5, win Scale, Nagle
Request flags: win Scale
```

```
Datagrams (in bytes): MSS 1216, peer MSS 1216, min MSS 1240,
max MSS 1240
```

```
Window scales: rcv 0, snd 0, request rcv 0, request snd 0
Timestamp option: recent 0, recent age 0, last ACK sent 0
Sack blocks {start, end}: none
Sack holes {start, end, dups, rxmit}: none
```

```
Socket options: SO_REUSEADDR, SO_REUSEPORT, SO_NBIO
Socket states: SS_ISCONNECTED, SS_PRIV
Socket receive buffer states: SB_DEL_WAKEUP
Socket send buffer states: SB_DEL_WAKEUP
Socket receive buffer: Low/High watermark 1/32768
Socket send buffer : Low/High watermark 2048/24576, Notify
threshold 0
```

```
PDU information:
#PDU's in buffer: 0
FIB Lookup Cache: IFH: 0x134e0 PD ctx: size: 8 data: 0x0
0xb1494a74
Num Labels: 0 Label Stack:
```

A long-exposure photograph of a city street at night. The image shows light trails from cars moving across the road, with buildings and streetlights in the background. The scene is illuminated by city lights, creating a vibrant and dynamic atmosphere. The text "Troubleshooting Forwarding Path" is overlaid on the center of the image.

Troubleshooting Forwarding Path

Monitor Interface

See interface stats in almost real time

```
RP/0/RSP0/CPU0:rasr9000-2w-b#monitor interface tenGigE 0/1/1/1

rasr9000-2w-b          Monitor Time: 00:00:22          SysUptime: 501:59:18

TenGigE0/1/1/1 is up, line protocol is up
Encapsulation ARPA

Traffic Stats:(2 second rates)                                Delta
Input  Packets:                2495245669613                14890408
Input  pps:                    7441113
Input  Bytes:                  164703177204108                982758522
Input  Kbps (rate):            3928857                    ( 39%)
Output Packets:                3017277633655                13261227
Output pps:                    6626897
Output Bytes:                  205177835436607                901762428
Output Kbps (rate):            3605031                    ( 36%)

Errors Stats:
Input  Total:                  1                    0
Input  CRC:                    0                    0
Input  Frame:                  0                    0
Input  Overrun:                0                    0
Output Total:                  0                    0
Output Underrun:               0                    0

Quit='q', Freeze='f', Thaw='t', Clear='c', Interface='i',
Next='n', Prev='p'

Brief='b', Detail='d', Protocol(IPv4/IPv6)='r'
```

The Physical

Checking on port physical: SFP/XFP, levels

```
RP/0/RSP0/CPU0:rasr9000-2w-b#show controllers TenGigE 0/0/0/0 phy
```

```
Mon Dec 9 13:53:37.848 EST
```

```
SFP EEPROM port: 0
```

```
  Xcvr Type: SFP
```

```
  Xcvr Code: SFP-10G-SR
```

```
  Encoding: 64B66B
```

```
  Bit Rate: 10300 Mbps
```

```
  Link Reach 50u fiber: 80 meter
```

```
  Link Reach 62.5u fiber: 20 meter
```

```
  Vendor Name: CISCO-FINISAR
```

```
  Vendor OUI: 00.90.65
```

```
  Vendor Part Number: FTLX8571D3BCL-C2 (rev.: A )
```

```
  Laser wavelength: 850 nm (fraction: 0.00 nm)
```

```
  Optional SFP Signal: Rate Sel, LOS
```

```
  Vendor Serial Number: FNS164018G7
```

```
  Date Code (yy/mm/dd): 12/10/06 lot code:
```

```
Thresholds:
```

```
  Temperature: Alarm High
```

```
                +75.000 C
```

```
  Voltage:      3.630 Volt
```

```
  Bias:         11.800 mAmps
```

```
  Transmit Power: 1.479 mW (1.70 dBm)
```

```
  Receive Power: 1.585 mW (2.00 dBm)
```

```
Temperature: 26.684
```

```
Voltage: 3.301 Volt
```

```
TX Bias: 7.612 mAmps
```

```
TX Power: 0.613 mW (-2.13 dBm)
```

```
RX Power: 0.567 mW (-2.46 dBm)
```

```
Oper. Status/Control:
```

```
Warning High
```

```
  +70.000 C
```

```
  3.465 Volt
```

```
  10.800 mAmps
```

```
  0.741 mW (-1.30 dBm)
```

```
  0.794 mW (-1.00 dBm)
```

```
Warning Low
```

```
  +0.000 C
```

```
  3.135 Volt
```

```
  5.000 mAmps
```

```
  0.186 mW (-7.30 dBm)
```

```
  0.102 mW (-9.90 dBm)
```

```
Alarm Low
```

```
  -5.000 C
```

```
  2.970 Volt
```

```
  4.000 mAmps
```

```
  0.074 mW (-11.30 dBm)
```

```
  0.041 mW (-13.90 dBm)
```

The Physical

Reading the controller counters: In, out, invalid, unicast, mcast, frame sizes

```
RP/0/RSP0/CPU0:rasr9k-1y#show controllers TenGigE0/4/0/20
stats
```

```
Sun Feb 24 14:44:18.899 UTC
```

```
Statistics for interface TenGigE0/4/0/20 (cached values):
```

```
Ingress:
```

```
Input total bytes      = 3081227904920
Input good bytes       = 3081227904920
```

```
Input total packets   = 23220024479
Input 802.1Q frames   = 0
Input pause frames    = 0
Input pkts 64 bytes   = 7143534733
Input pkts 65-127 bytes = 2888766549
Input pkts 128-255 bytes = 13124923916
Input pkts 256-511 bytes = 62799261
Input pkts 512-1023 bytes = 0
Input pkts 1024-1518 bytes = 0
Input pkts 1519-Max bytes = 0
```

```
Input good pkts       = 23220024479
Input unicast pkts    = 23220023458
Input multicast pkts  = 62
Input broadcast pkts  = 959
```

```
Input drop overrun    = 0
```

```
Egress:
```

```
Output total bytes    = 1345771624
Output good bytes     = 1345771624
```

```
Output total packets  = 21895707
Output 802.1Q frames  = 0
Output pause frames   = 0
Output pkts 64 bytes  = 21665536
Output pkts 65-127 bytes = 21179
Output pkts 128-255 bytes = 168767
Output pkts 256-511 bytes = 40225
Output pkts 512-1023 bytes = 0
Output pkts 1024-1518 bytes = 0
Output pkts 1519-Max bytes = 0
```

```
Output good pkts      = 21895707
Output unicast pkts   = 21870499
Output multicast pkts = 25195
Output broadcast pkts = 13
Output drop underrun  = 0
Output drop abort     = 0
Output drop other     = 0

Output error other    = 0
```

The Internal Path

Troubleshooting NP Forwarding

1. Identify interface in questions with problem
2. Identify the mapping from interface to NPU
3. Examine NP counters
4. Look for rate counter that match lost traffic rate
 - If none of the counter match the expect traffic, check drops at interface controller
5. Lookup the counter description
6. If required capture the packets hitting the counter (typhoon only)
7. If packets are forwarded to the fabric, run fabric troubleshooting steps
8. Identify egress NP and repeat 3 to 6.

Inside Typhoon NP

Reading pipeline counters

```
show controllers np descriptions location  
0/0/cpu0 | i PARSE_DROP_IN_UIDB_TCAM_MISS
```

```
RP/0/RSP0/CPU0:rasr9000-2w-a#show controllers NP counters np0 location 0/1/CPU0
```

```
wed Nov 27 21:09:07.635 EST
```

```
Node: 0/1/CPU0:
```

```
Show global stats counters for NP0, revision v2
```

```
Read 64 non-zero NP counters:
```

Offset	Counter	FrameValue	Rate (pps)
16	MDF_TX_LC_CPU	6722114	10
17	MDF_TX_WIRE	1826039	3
21	MDF_TX_FABRIC	1635541	2
29	PARSE_FAB_RECEIVE_CNT	1837406	3
33	PARSE_INTR_RECEIVE_CNT	5083364	7
37	PARSE_INJ_RECEIVE_CNT	1228130	2
499	RSV_ING_L2_SMAC_MISS	60	0
502	RSV_ING_L2_LEARN	60	0
541	RSV_REFRESH_FROM_NOTIFY_CNT	62	0
584	RSV_L2BC_BVI	2	0
604	RESOLVE_REMOTE_RACK_PREP_CNT	5539915	8
708	LRN_PERIODIC_AGING_DELETE_ENTRY	60	0
774	ARP	119	0
848	PUNT_ADJ	2	0
852	PUNT_ACL_DENY	161	0
900	PUNT_STATISTICS	5083356	7
902	PUNT_DIAGS_RSP_ACT	11419	0
904	PUNT_DIAGS_RSP_STBY	11427	0

List of NP counters:

<https://supportforums.cisco.com/docs/DOC-26566>

Typhoon Counters and Rates

NP drops, rate and direction

```
RP/0/RSP0/CPU0:rasr9000-2w-b#show controllers NP counters np0 location 0/0/CPU0
Tue Dec 10 14:18:39.195 EST
Node: 0/0/CPU0:
-----
Show global stats counters for NP0, revision v2
Read 59 non-zero NP counters:
Offset Counter FrameValue Rate (pps)
-----
 16 MDF_TX_LC_CPU 11004363 9
 17 MDF_TX_WIRE 8712222364719 29761820
 21 MDF_TX_FABRIC 11063035007386 27714366
 29 PARSE_FAB_RECEIVE_CNT 8712222113330 29761820
 33 PARSE_INTR_RECEIVE_CNT 9401470 9
 37 PARSE_INJ_RECEIVE_CNT 832185 1
 41 PARSE_FNET_RECEIVE_CNT 11070653296959 27714366
 45 PARSE_TM_LOOP_RECEIVE_CNT 8437075 5
359 PARSE_MAC_NOTIFY_RCVD 183 0
 367 PARSE_FAST_DISCARD_LOW_PRIORITY_DROP_0 106211394050 883832
 368 PARSE_FAST_DISCARD_LOW_PRIORITY_DROP_1 106210662138 883856
 369 PARSE_FAST_DISCARD_LOW_PRIORITY_DROP_2 106211061617 883943
 370 PARSE_FAST_DISCARD_LOW_PRIORITY_DROP_3 106211474043 883922
 373 DBG_RSV_EP_L_RSV_ING_L3_IFIB 3707021673 0
 830 PUNT_NO_MATCH 4746 0
 831 PUNT_NO_MATCH_EXCD 464963896 0
 849 PUNT_ADJ_EXCD 273406 0
 852 PUNT_ACL_DENY 1479378 0
 853 PUNT_ACL_DENY_EXCD 1163570900 0
```

To egress

To fabric

From fabric

From interface

NP catching up

FIA Counters

FIA counts, drops and direction

<https://supportforums.cisco.com/document/12135016/asr9000xr-understanding-and-troubleshooting-fabric-issues-a9k>

```
RP/0/RSP0/CPU0:rasr9000-2w-b#show controllers fabric
fia instance 0 stats location 0/0/CPU0
Tue Dec 10 14:49:58.704 EST

***** FIA-0 *****
Category: count-0
  From Unicast xbar[0]          733461306331
  From Unicast xbar[1]          733460650405
  From Unicast xbar[2]          0
  From Unicast xbar[3]          0
  From MultiCast xbar[0]        233068
  From MultiCast xbar[1]        0
  From MultiCast xbar[2]        0
  From MultiCast xbar[3]        0
  To Unicast xbar[0]            0
  To Unicast xbar[1]            0
  To Unicast xbar[2]            0
  To Unicast xbar[3]            0
  To MultiCast xbar[0]          457138023968
  To MultiCast xbar[1]          11117127781061
  To MultiCast xbar[2]          489302108080
  To MultiCast xbar[3]          0
  To Line Interface[0]          0
  To Line Interface[1]          0
  From Line Interface[0]        0
  From Line Interface[1]        0
  Ingress drop:                  97191712670
  Egress drop:                    0
  Total drop:                    97191712670
```

```
RP/0/RSP0/CPU0:rasr9000-2w-b#show controllers fabric fia
instance 0 drops ingress location 0/0/CPU0
Tue Dec 10 15:33:37.655 EST

***** FIA-0 *****
Category: in_drop-0
  From Spau Drop-0              0
  acpt tbl-0                    0
  ctl len-0                     0
  short pkt-0                   0
  max pkt len-0                 0
  min pkt len-0                 0
  From Spau Drop-1              0
  acpt tbl-1                    0
  ctl len-1                     0
  short pkt-1                   0
  max pkt len-1                 0
  min pkt len-1                 0
  Tail drp                       125787328841
  Vqi drp                       0
  Header parsing drp            0
  pw to ni drp                  0
  ni from pw drp                0
  sp0 crc err                   0
  sp0 bad align                 0
  sp0 bad code                  0
  sp0 align fail                3
  sp0 prot err                  0
  sp1 crc err                   0
  sp1 bad align                 0
```

Enhancement TS commands

- The “show drops location” was available before 5.3.0 but it only listed the NP and fabric drops so HW drops. With the “**all**” keyword being introduced, we now have the ability to check both HW and SW drops.
- the command is extensible live on the router by editing a grammar file saved under “disk0a:/usr/packet_drops.list”, default path “/pkg/etc/packet_drops.list”

```
RP/0/RSP0/CPU0:ASR9010-1#show drops all location 0/6/cpu0  
Sun Jul 26 23:14:14.374 UTC
```

```
=====  
Checking for drops on 0/6/CPU0  
=====
```

```
show arp traffic:  
[arp:ARP] IP Packet drop count for node 0/6/CPU0: 4
```

```
show cef drops:  
[cef:0/6/CPU0] No route drops      packets : 69  
[cef:0/6/CPU0] Discard drops      packets : 5
```

```
show controllers fabric fia drops egress:  
[fabric:FIA-0] Uc eq pkt-len-crc/lookup-drp: 7  
[fabric:FIA-1] Uc eq pkt-len-crc/lookup-drp: 3
```

```
show controllers fabric fia drops ingress:  
[fabric:FIA-0] sp0 align fail: 5  
[fabric:FIA-0] sp1 crc err: 21
```

A long-exposure photograph of a city street at night. The image shows light trails from cars moving through the frame, creating a sense of motion. In the background, there are modern buildings with lit windows and a pedestrian bridge. The overall scene is illuminated by city lights, creating a vibrant and dynamic atmosphere.

Troubleshooting Packet Capture

Packet Capture: Problem Packets

Example: incrementing drops

```
RP/0/RSP0/CPU0:rasr9000-2w-b#show controllers NP counters np0 location 0/0/CPU0 | include DROP
Sat Jan 18 18:46:52.618 EST
```

370	RSV_DROP_XID_NO_MATCH	209680463	0
404	RSV_ING_VPWS_ERR_DROP	3719838164404	11160601
411	RSV_L2_SHG_DROP	27390624	0
1171	MDF_PUNT_POLICE_DROP	7924962278163	23809032
1178	MODIFY_PUNT_REASON_MISS_DROP	1	0
1246	VIRTUAL_IF_GENERIC_INPUT_DROP	1	0

```
RP/0/RSP0/CPU0:rasr9000-2w-b#show controllers NP counters np0 location 0/0/CPU0 | include DROP
Sat Jan 18 18:46:56.297 EST
```

370	RSV_DROP_XID_NO_MATCH	209680463	0
404	RSV_ING_VPWS_ERR_DROP	3719879236984	11161027
411	RSV_L2_SHG_DROP	27390624	0
1171	MDF_PUNT_POLICE_DROP	7925049898728	23809936
1178	MODIFY_PUNT_REASON_MISS_DROP	1	0
1246	VIRTUAL_IF_GENERIC_INPUT_DROP	1	0

incrementing

Rate [PPS] or
increments from
last command run

Packet Capture: Problem Packets

Example: incrementing drops

```
RP/0/RSP0/CPU0:rasr9000-2w-b#monitor np counter RSV_ING_VPWS_ERR_DROP np0 count 3 location 0/0/CPU0
```

```
Sat Jan 18 19:02:36.386 EST
```

Warning: Every packet captured will be dropped! If you use the 'count' option to capture multiple protocol packets, this could disrupt protocol sessions (eg, OSPF session flap). So if capturing protocol packets, capture only 1 at a time.

Warning: A mandatory NP reset will be done after monitor to clean up. This will cause ~50ms traffic outage. Links will stay down.

```
Proceed y/n [y] >
```

```
Monitor RSV_ING_VPWS_ERR_DROP on NP0 ... (Ctrl-C to quit)
```

```
Sat Jan 18 19:02:44 2014 -- NP0 packet
```

From TenGigE0/0/0/0: 157 byte packet, bytes[0-3] invalid!

```
0000: 00 00 02 01 61 90 00 00 c0 02 01 02 81 00 00 0a .....a...@.....
0010: 08 00 45 00 00 8b 00 00 00 00 40 3d f8 30 c0 01 ..E.....@=x0@.
0020: 01 01 c0 01 01 02 00 00 00 00 00 00 00 00 00 ..@.....
0030: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
0040: 10 f3 11 05 00 00 00 00 00 00 00 00 00 00 00 .....
0050: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
0060: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
0070: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
0080: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
0090: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
```

Ignore
[internal]

UP to 300 B
No CRC



Packet Capture: Problem Packets

Decoding

```
00 00 02 01 61 90 00 00 c0 02 01 02 81 00 00 0a
08 00 45 00 00 8b 00 00 00 00 40 3d f8 30 c0 01
01 01 c0 01 01 02 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
10 f3 11 05 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

https://scripts.cisco.com/ui/u/se/xr_monitor_np_counter

Use a decoder

```
Ethernet II, Src: WesternD_02:01:02 (00:00:c0:02:01:02), Dst: Xerox_01:61:90 (00:00:02:01:61:90)
  Destination: Xerox_01:61:90 (00:00:02:01:61:90)
  Source: WesternD_02:01:02 (00:00:c0:02:01:02)
  Type: 802.1Q Virtual LAN (0x8100)
802.1Q Virtual LAN, PRI: 0, CFI: 0, ID: 10
  000. .... = Priority: 0
  ...0 .... = CFI: 0
  ... 0000 0000 1010 = ID: 10
  Type: IP (0x0800)
Internet Protocol, Src: 192.1.1.1 (192.1.1.1), Dst: 192.1.1.2 (192.1.1.2)
  Version: 4
  Header length: 20 bytes
  Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
  Total Length: 139
```

Packet Capture: Transit Packets

Example: IPv4 L3VPN ingress

```
RP/0/RSP0/CPU0:rasr9000-2w-b#show running-config ipv4 access-list CAPTURE
```

```
Sat Jan 18 20:13:35.941 EST
```

```
ipv4 access-list CAPTURE
```

```
10 permit ipv4 192.4.1.0/24 10.10.6.0/24 capture
```

```
20 permit ipv4 any any
```

```
!
```

Count in NP

```
RP/0/RSP0/CPU0:rasr9000-2w-b#show running-config interface TenGigE 0/0/0/2
```

```
Sat Jan 18 20:13:50.654 EST
```

```
interface TenGigE0/0/0/2
```

```
vrf TRAFFIC
```

```
ipv4 address 192.4.1.1 255.255.255.0
```

```
ipv4 access-group CAPTURE ingress
```

```
!
```

Let all else go!

Apply to transit

```
RP/0/RSP0/CPU0:rasr9000-2w-b#show controllers NP counters np0 location 0/0/CPU0 | include
```

```
ACL_CAPTURE_NO_SPAN
```

```
Sat Jan 18 20:14:26.109 EST
```

```
477 ACL_CAPTURE_NO_SPAN
```

```
6802507
```

```
38003
```

```
RP/0/RSP0/CPU0:rasr9000-2w-b#show controllers NP counters np0 location 0/0/CPU0 | include
```

```
ACL_CAPTURE_NO_SPAN
```

```
Sat Jan 18 20:14:28.819 EST
```

```
477 ACL_CAPTURE_NO_SPAN
```

```
6905417
```

```
38003
```

```
RP/0/RSP0/CPU0:rasr9000-2w-b#show controllers NP counters np0 location 0/0/CPU0 | include
```

```
ACL_CAPTURE_NO_SPAN
```

```
Sat Jan 18 20:14:34.597 EST
```

```
477 ACL_CAPTURE_NO_SPAN
```

```
7124969
```

```
37991
```

NP ACL "capture"
counter
incrementing

Packet Capture: Transit Packets

Example: IPv4 L3VPN ingress

```
RP/0/RSP0/CPU0:rasr9000-2w-b#monitor np counter ACL_CAPTURE_NO_SPAN np0 count 3 location 0/0/CPU0
Sat Jan 18 20:31:53.311 EST
```

Warning: Every packet captured will be dropped! If you use the 'count' option to capture multiple protocol packets, this could disrupt protocol sessions (eg, OSPF session flap). So if capturing protocol packets, capture only 1 at a time.

Warning: A mandatory NP reset will be done after monitor to clean up. This will cause ~50ms traffic outage. Links will stay up.

Proceed y/n [y] >

Monitor ACL_CAPTURE_NO_SPAN on NP0 ... (Ctrl-C to quit)

```
Sat Jan 18 20:32:34 2014 -- NP0 packet
```

From TenGigE0/0/0/2: 250 byte packet, bytes[0-3] invalid

```
0000: 00 11 0b 00 61 92 00 00 c0 04 01 02 08 00 45 60  . . . . a . . . . @ . . . . E `
0010: 00 ec 00 00 00 00 40 3d a8 08 c0 04 01 02 0a 0a  . ] . . . . @ = ( . @ . . . .
0020: 06 5d 00 00 00 00 00 00 00 00 00 00 00 00 00 00  . ] . . . . . . . . . . . . . .
0030: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  . . . . . . . . . . . . . .
0040: 10 f3 11 05 00 00 00 00 00 00 00 00 00 00 00 00  . s . . . . . . . . . . . . . .
0050: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  . . . . . . . . . . . . . .
0060: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  . . . . . . . . . . . . . .
0070: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  . . . . . . . . . . . . . .
0080: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  . . . . . . . . . . . . . .
0090: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  . . . . . . . . . . . . . .
.
```

Those 3 packets are dropped!



SPP packet captures

```
RP/0/RSP0/CPU0#run attach 0/0/CPU0
```

```
attach: Starting session 1 to node 0/0/CPU0
```

```
# spp_ui  
spp-ui>
```

```
spp-ui> trace filter node client/punt  
Node "client/punt" set for trace filtering. Index: 11
```

```
spp-ui> trace filter set 52 4 0xD4000001  
Modified filter for offset 52 successfully  
spp-ui> trace filter set 56 4 0xD4000002  
Modified filter for offset 56 successfully
```

```
RP/0/RSP0/CPU0:A9K-BNG#packet-trace spp platform  
protocol arp start-capture count 5 location 0/0/cpu0
```

```
Wed Mar 12 16:28:30.176 EDT
```

```
Sending command: trace filter
```

```
Trace filter set for protocol: ARP
```

```
Sending command: trace start 5
```

```
Started capture for 5 packets
```

```
Wrote ASCII trace to /tmp/spp_
```

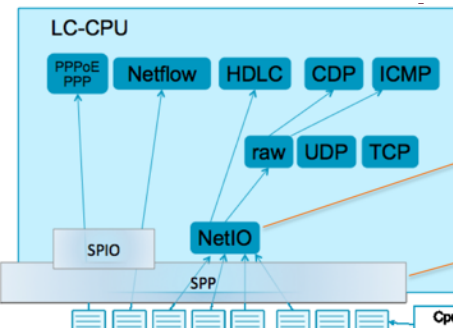
```
Sending command: trace print
```

```
Packet serial 3
```

```
client/inject:
```

```
length 110 phys_int_index -1 next_ctx 0xdeadbeef time  
16:28:30.512
```

```
00: 00 65 7a 00 00 00 00 70 72 00 00 02 00 5e 00 00  
10: 80 00 00 00 00 00 0f 8c 40 c1 0c c8 50 00 00 00  
20: 00 00 0d 34 3f ff f2 90 20 04 fe 03 01 04 00 05  
30: 00 00 00 00 5e 00 00 00 00 00 00 00 00 04 00 02  
40: 40 00 10 34 ff ff ff ff ff 66 66 44 44 22 22
```



Decoder https://scripts.cisco.com/ui/use/xr_spp_ui_to_pcap



Thank you.



CISCO