



The bridge to possible

# Introduction to IOS-XR

Rajesh Patki – Technical Leader Customer Delivery

@rpatki

Ramiro Garza Rios – Customer Delivery Architect

@rgarzari

# Cisco Webex App

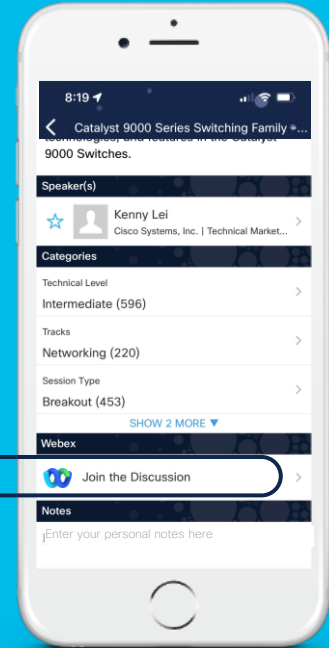
## Questions?

Use Cisco Webex App to chat with the speaker after the session

## How

- 1 Find this session in the Cisco Live 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 until February 24, 2023.

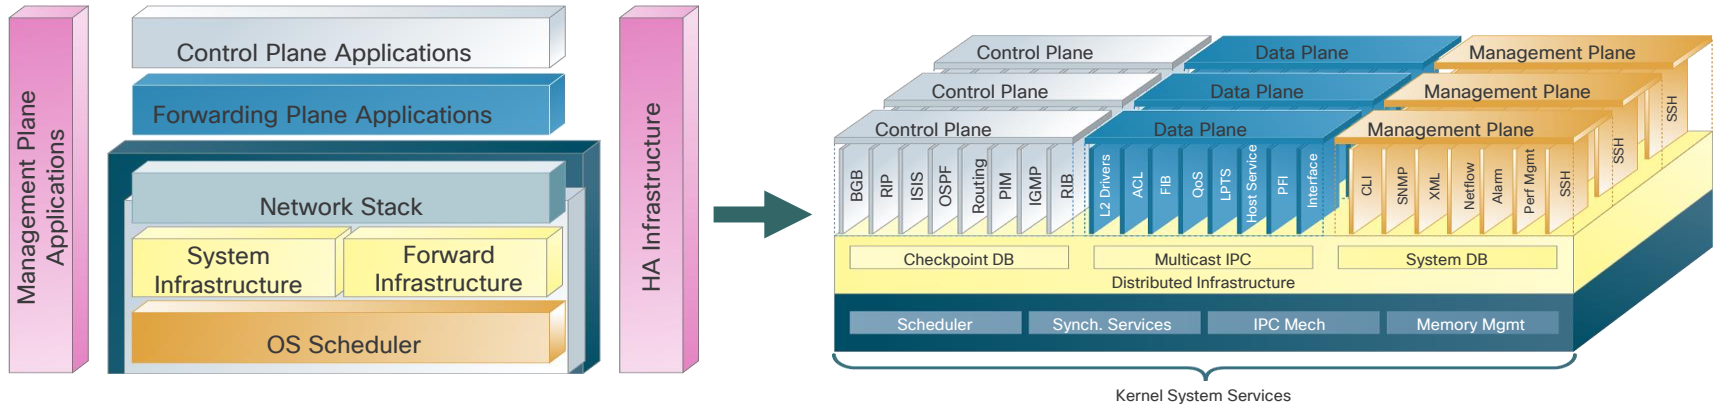




# Agenda

- Introduction
- Architecture Overview
- Software Packages, Installation & Upgrades
- Protocol Configuration Overview
- Route Policy Language (RPL)
- MPLS LDP, MPLS-VPN and MPLS-TE
- Operation and System Monitoring Tools
- Hands-On Lab

# Evolution of Router OS



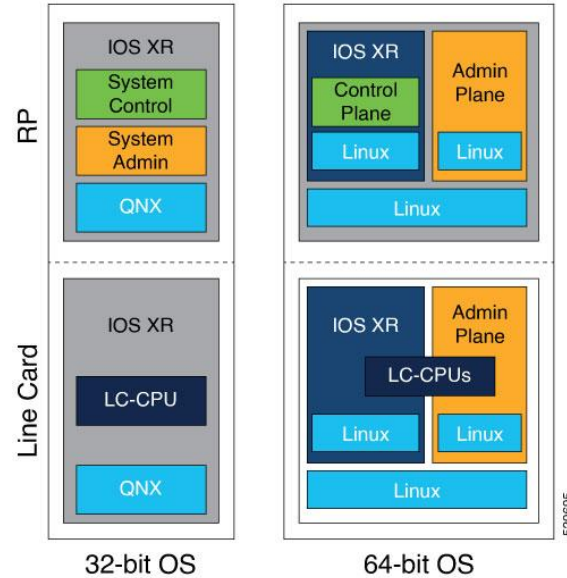
## IOS

- Monolithic Kernel
- Run to Completion Scheduler
- Centralized Infrastructure
- Centralized applications
- Everything has hardware access

## IOS XR

- Micro Kernel ( QNX )
- Preemptive Multitasking
- Distributed Infrastructure
- Distributed applications
- Limited access to hardware

# Evolution of Router OS



## classic-XR ( cXR )

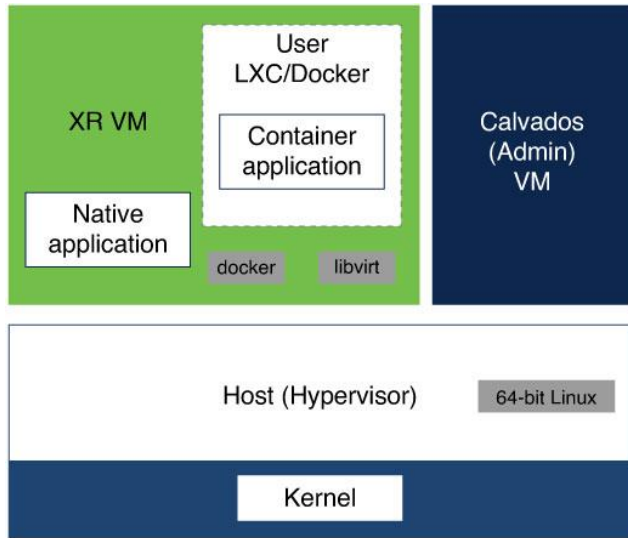
- Micro Kernel ( QNX )
- 32 bit OS / NPU
- No VMs



## evolved XR ( eXR )

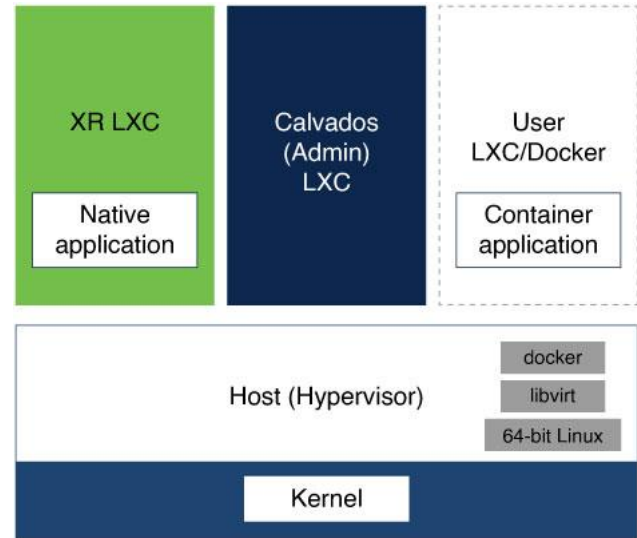
- Linux Kernel
- 64 bit OS /x86 hardware
- 2 VMs on each RP/LC – Control/Admin
- 3<sup>rd</sup> party application hosting and ISSU

# Evolution of Router OS -eXR



## VM based eXR

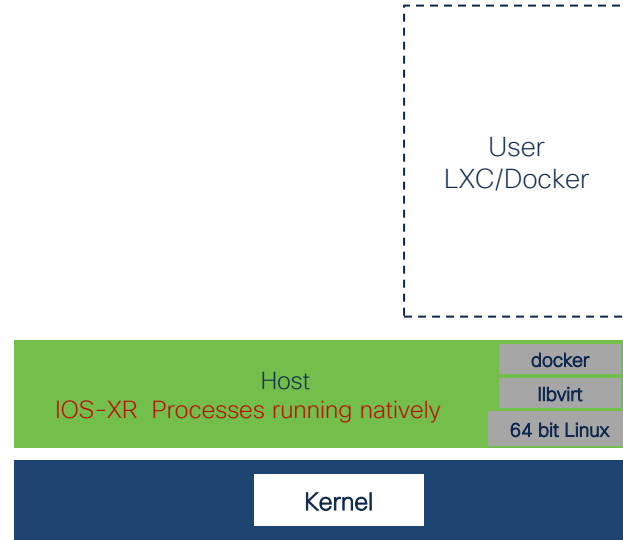
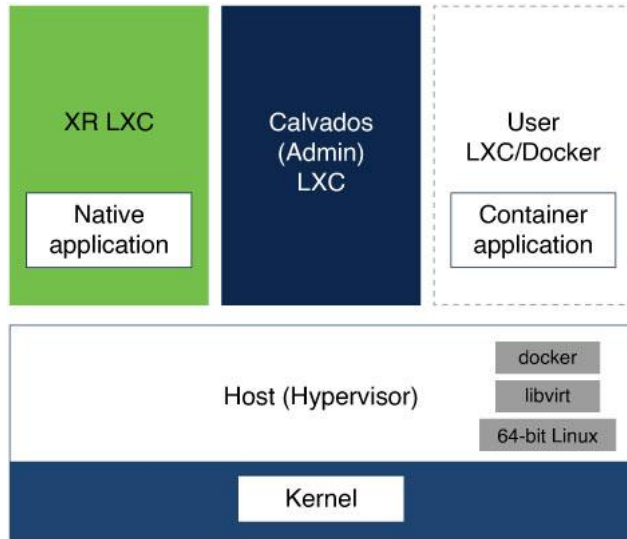
- VM OS can be separate from Host OS
- On ASR9K, NCS 6K
- ISSU support



## Container based eXR

- Container(LXC) OS shares same kernel as Host OS
- On NCS-5500 , NCS-5000
- Lighter OS, no ISSU support.

# Evolution of Router OS -XR7



## eXR

- VM/ container for Control/Admin
- On ASR9K, NCS 6K, 5K, 5500



## XR7

- No Admin plane, no container
- On 8K , NCS-540, NCS-57B2, NCS1010

# XR7

- Slim OS – modern, lean and trusted
- Removal of Admin plane & system containers
  - Reduced complexity
  - Faster Boot time
  - Memory footprint reduction
- Software integrity
  - Images are signed by Cisco.
  - Tamper resistant self check boot with signature validation
  - 3<sup>rd</sup> Party RPMs ( also cisco signed )
- Automated provisioning using ZTP
- Application hosting
- YANG Modeled API @ every layer of stack .
- Streaming telemetry



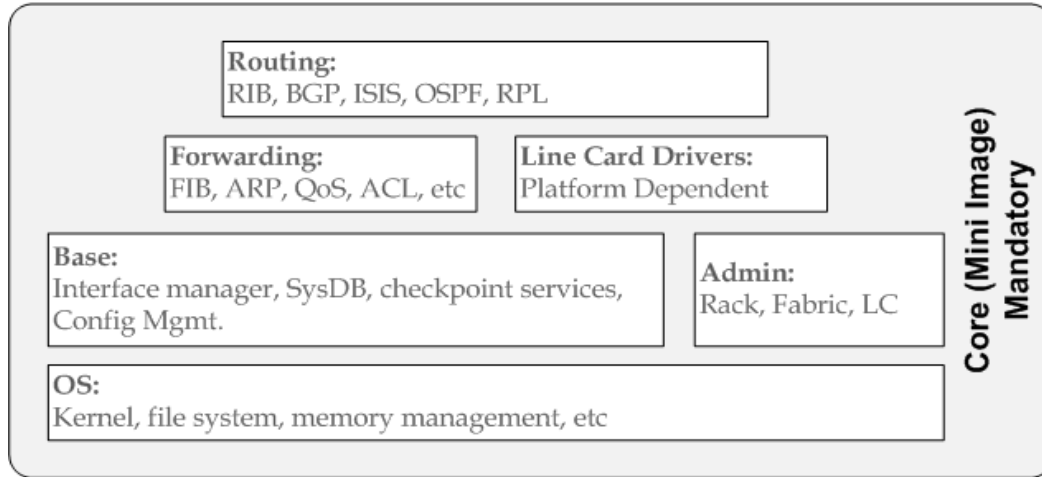
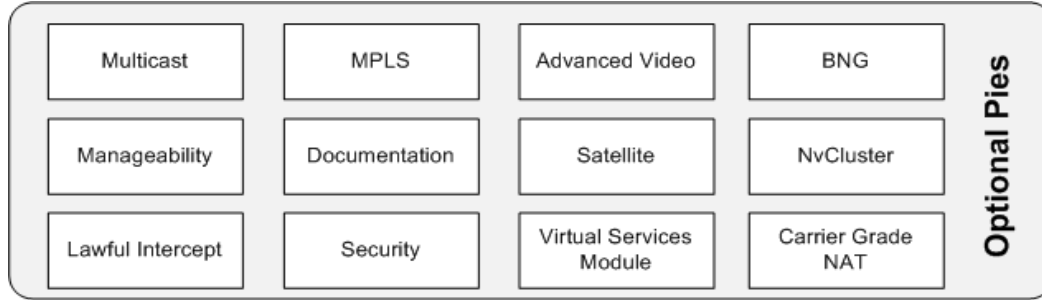
- Available on
  - 8K
  - NCS 57B2
  - NCS 540
  - NCS 1010

Want to learn more?

<https://xrdocs.io/>

# IOS-XR Boot Images and Installation

# IOS-XR Software Packages



## IOS XR Mini:

- Includes the following components:
  - OS
  - Base
  - Admin
  - Line Card Drivers
  - Forwarding Logic
  - Routing Protocols

## IOS XR Versioning

- Major release**  
New functionality (6.2, 7.2 etc.)
- Maintenance release**  
SW fixes (6.2.0, 6.2.1, 6.2.2)
- There are not any special S, T, J, or XT trains; special functionality is added through packages.

# Software Install Terminology

## Software Maintenance Upgrade (SMU)

Provides timely temporary point fixes for urgent issues for a given package version

Fix integrated into the subsequent IOS XR maintenance release.

Implementation changes only. No interface changes (no changes to CLI, APIs, IPC etc.) or new feature content

Ideally does not impact traffic (Hitless, non traffic impacting)

SMU is named by release and bugid - Examples - hfr-px-6.0.5.CSCxy12345.pie



**PIE?**



**Package?**



**Mini?**



**SMU?**

# Identifying packages

```
RP/0/RP0/CPU0:CRS#show install active summary
```

```
Active Packages:
```

```
disk1:hfr-diags-px-6.7.4  
disk1:hfr-doc-px-6.7.4  
disk1:hfr-fpd-px-6.7.4  
disk1:hfr-k9sec-px-6.7.4  
disk1:hfr-mcast-px-6.7.4  
disk1:hfr-mgbl-px-6.7.4  
disk1:hfr-mini-px-6.7.4  
disk1:hfr-mpls-px-6.7.4
```

```
RP/0/RSP0/CPU0:ASR9K#sh install active summary
```

```
Active Packages:
```

```
disk0:asr9k-mgbl-px-6.7.4  
disk0:asr9k-mini-px-6.7.4  
disk0:asr9k-li-px-6.7.4  
disk0:asr9k-k9sec-px-6.7.4  
disk0:asr9k-doc-px-6.7.4  
disk0:asr9k-mcast-px-6.7.4  
disk0:asr9k-optic-px-6.7.4  
disk0:asr9k-fpd-px-6.7.4  
disk0:asr9k-mpls-px-6.7.4
```

```
RP/0/RP0/CPU0:NCS-5501#sh install active
```

```
Node 0/RP0/CPU0 [RP]
```

```
Boot Partition: xr_lv32
```

```
Active Packages: 8
```

```
ncs5500-xr-6.7.4 version=6.7.4 [Boot image]  
ncs5500-mcast-2.2.0.0-r674  
ncs5500-mpls-2.1.0.0-r674  
ncs5500-mgbl-3.0.0.0-r674  
ncs5500-mpls-te-rsvp-2.3.0.0-r674  
ncs5500-ospf-2.0.0.0-r674  
ncs5500-isis-2.2.0.0-r674  
ncs5500-k9sec-3.2.0.0-r674
```

# Identifying packages

```
RP/0/RP0/CPU0:8K#show install active summary
```

```
Active Packages:   XR: 180   All: 1276  
Label:             7.3.15  
Software Hash:    892ad10cc1f79cb419cf15d260bb6ed1
```

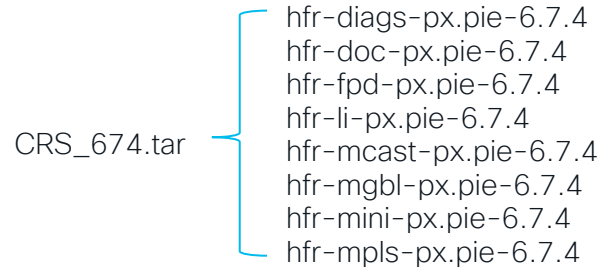
Optional Packages	Version
-----	-----
xr-8000-mcast	7.3.15v1.0.0-1
xr-8000-netflow	7.3.15v1.0.0-1
xr-bgp	7.3.15v1.0.0-1
xr-cdp	7.3.15v1.0.0-1
xr-ipsla	7.3.15v1.0.0-1
xr-is-is	7.3.15v1.0.0-1
xr-k9sec	7.3.15v1.0.0-1
xr-lldp	7.3.15v1.0.0-1
xr-mcast	7.3.15v1.0.0-1
xr-mpls-oam	7.3.15v1.0.0-1
xr-netflow	7.3.15v1.0.0-1
xr-ospf	7.3.15v1.0.0-1
xr-perfmgmt	7.3.15v1.0.0-1
xr-track	7.3.15v1.0.0-1

# PIE Installation Concepts -



- PIE ( Package Installation Envelope )
  - can be added, upgraded, or removed.
- Three phases of PIE installation
  - **Add** – Copies the package file to the local storage and unpacks the files
  - **Activate** – Installs the new code by restarting processes/nodes with new code
  - **Commit** – Locks the activated code to sustain reloads
- Three phases of removing a PIE
  - **DeActivate** – Removes the code from runtime
  - **Commit** – Locks the code change to sustain reloads
  - **Remove** – Removes unused packages from the local storage

# Installation : cXR



Copy tar file to CRS harddisk

- install add tar /harddisk:/674/CRS\_674.tar synchronous
- install activate id 84 prompt-level none synchronous test
- install activate id 84 prompt-level none synchronous
- install verify packages synchronous
- install commit

# Installation : eXR

- Copy tar file to NCS harddisk
- install add source harddisk:/ NCS5500-iosxr-k9-6.6.3.tar
- install prepare id 334
- install activate
- install commit

POST 7.0 we can also upgrade with a single command:

- untar the image and copy the ISO and RPMs to router harddisk
- install source /harddisk:/image/ ncs5500-mini-x-7.3.2.iso noprompt

Other Boot options :

- GISO / IPXE / USB

# Golden ISO ( GISO )

- Custom bootable ISO image for specific customer, which contains
  - ISO image ( is a must )
  - Feature Packages ( optional )
  - SMUs ( optional )
  - Configs ( optional )
- Generated by using the `gisobuild.py` script
  - Script available on GITHUB or on the router itself
  - Created on a server with root access.
  - `./gisobuild.py -i ncs5500-mini-x-6.3.3.iso -r ./633_repo -l LABEL`
- Copy the generated image to router harddisk
- `install update source harddisk:/giso ncs5500-goldenk9-x.iso-6.3.3.LABEL633GISO replace noprompt`

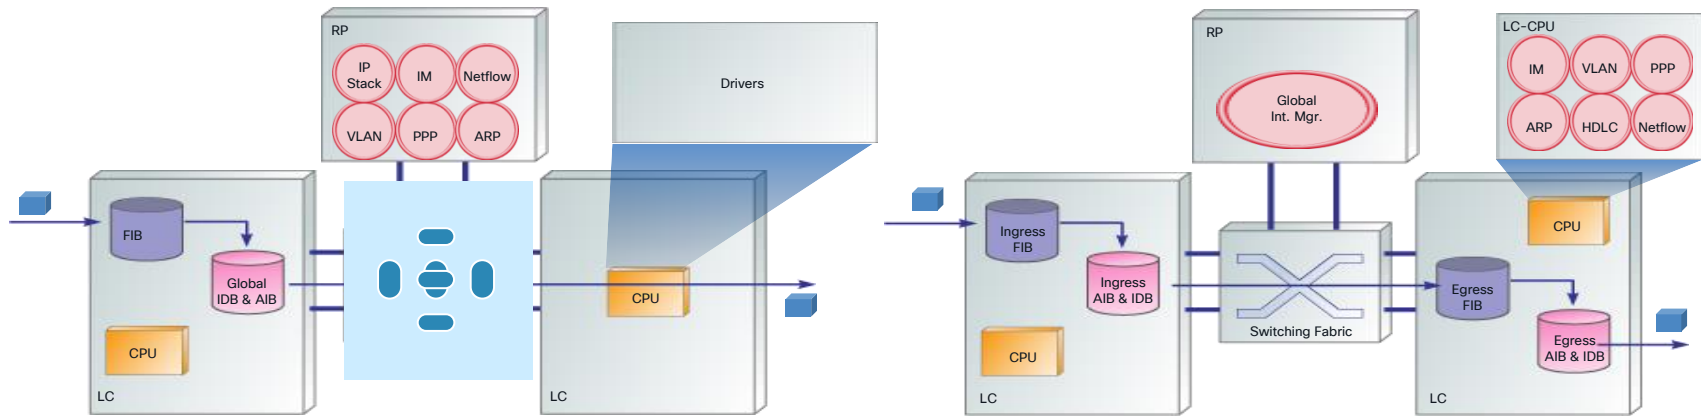
# Installing and Updating Third Party Packages



- Third Party Packages are traditional Linux tools available from the Shell
  - Communication: lighttpd ,openssh, wget, curl, etc.
  - Programing: python, ruby, perl, etc.
  - Utilities: sed, gawk, tar, gzip, vi, etc.
- Additional packages provided by vendors (No Cisco Support)
  - Chef
  - Puppet
- Installed using yum or rpm from the shell.

```
yum-config-manager --add-repo=http://192.168.0.254/XR/6.0.0  
yum install chef -y
```

# Distributed Forwarding Infrastructure



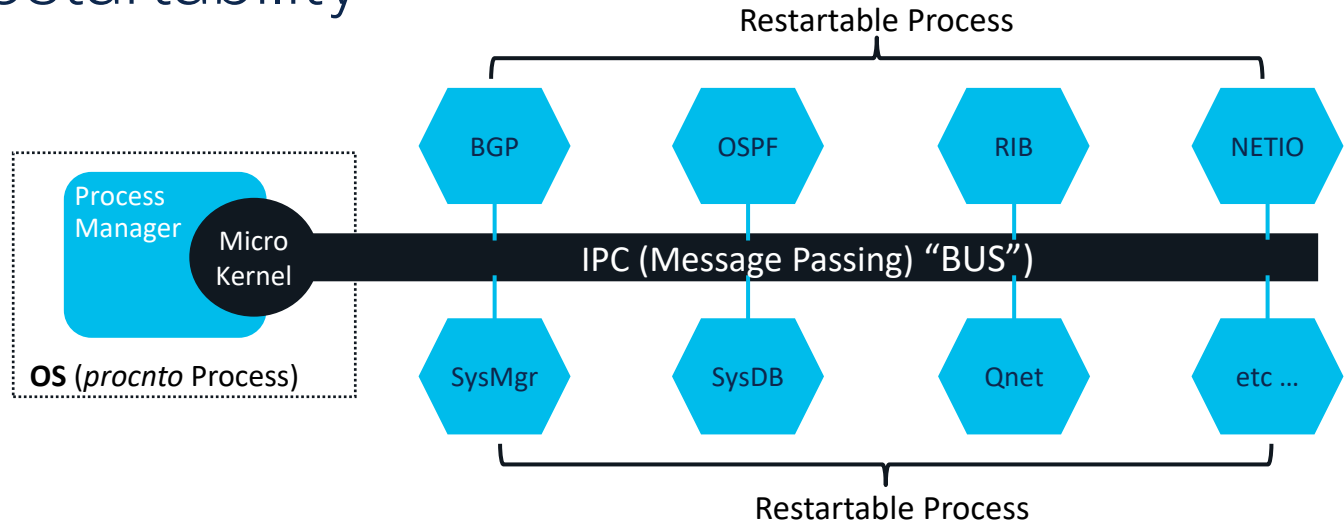
## Single stage forwarding (Centralized)

- Single global Adjacency Information Base (AIB) distributed to all line cards
- Single global Interface Management DB distributed to all line cards
- Only Ingress FIB – forces forwarding features to be run in RP

## Two stage forwarding (Distributed)

- Each line card has independent AIB only for local interfaces
- Each line card has independent Interface DB for local interfaces
- Both Ingress and Egress FIB – allows forwarding features to be independently run in LCs

# Process Restartability

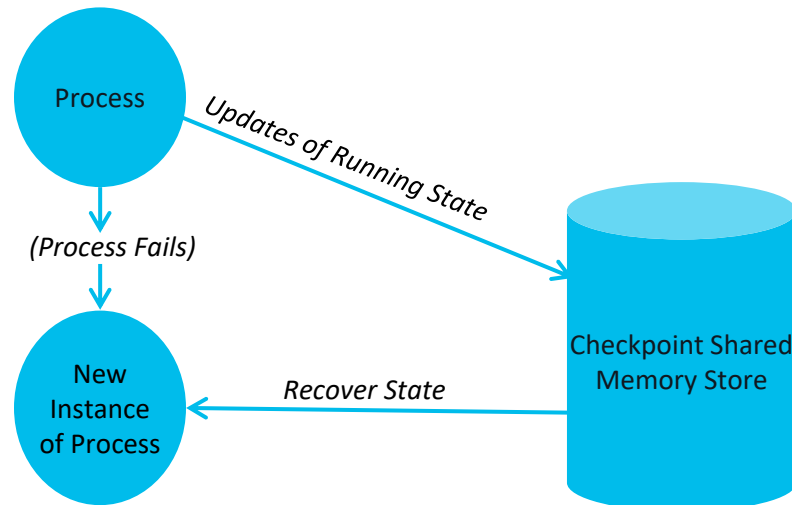


- All processes outside the microkernel (procnto) are individually restartable.
- Certain “Mandatory” processes must always be running. Failure of these can cause RP reload.
- If any of the processes, including SysMgr, SysDB, BGP, or Qnet, is restarted it does not cause the entire system to reload.
- When a process restarts, it recovers its states from persistent storage or peer processes.
  - For example, if RIB process restarts, it will restore the RIB table from its peer processes which are routing protocol processes such as OSPF, BGP, IS-IS, etc. As a result, the RIB table is rebuilt and there’s no traffic disruption if the RIB process is restarted.

# Process Restartability (continued)

## Checkpoint Server

- Checkpoint Server (CS) on the active RP is also mirrored to the CS on the standby RP.
- Useful during failover because processes on the new active RP can recover the states from the local CS.



**Dynamic State Recovery from Checkpoint Server**

# IOS XR Security

- User group and Task groups
  - Task group : defines tasks with permissions.... or ... use predefined task groups.
  - User group : associate one or more task group.
  - Usernames associated with user group.
- AAA authentication/authorization/accounting.
  - AAA server can be configured to map user to task group using optional AAA attributes
- Management Plane Protection.
  - Allows control in-band and out-of-band interfaces.
  - Helps control the interfaces on which network management traffic can access the router.
- Forwarding plane protection:
  - ACLs
  - Unicast RPF
  - LPTS ( Local Packet Transport Service )

# IOS XR Security

- Management Plane Protection.
  - Controls the interface on which Network Management traffic can access the router.
  - In-band interface : n/w mgmt + Internet traffic
  - Out-of-band interface : only n/w mgmt traffic.
  - Management ports on RP: default out-of-band
  - Enhances Security.
  - Reduces device access ACLs.


```
RP/0/RP0/CPU0:8K#show run control-plane
control-plane
management-plane
inband
interface all
allow SSH
!
interface FourHundredGigE0/0/0/0
allow all
!
interface FourHundredGigE0/0/0/2
allow Telnet peer
address ipv4 10.1.0.0/16
!
!
!
!
!
end
```

# Local Packet Transport Protocol (LPTS)

## Interoffice Mail for Data Plane

- Enables delivery of data to distributed processes across the system hardware (RPs, DRPs)
- Used for **'for\_us'** packet prioritization and filtering
  - Sends 'for us' packets only to the nodes that want them
  - Uses HW policers to throttle "for us" traffic
  - Applies to data plane traffic, not IPC
- Integral firewall to protect router resources - Packet forwarding executed in HW - no impact on Line Card CPU
- **Dynamic adjustment** - Flow types reflect both application type (OSPF, BGP, ...) and trust (established, configured or unknown peer) - Additional bandwidth allowed once neighbor up

## Interoffice Mail for Data Plane



## LC/RP CPU guard



# LPTS: Dynamic Control Plane Protection

- DCoPP is an automatic, built in firewall for control plane traffic.
- Every Control and Management packet from the line card is rate limited in hardware to provide flood protect at RP

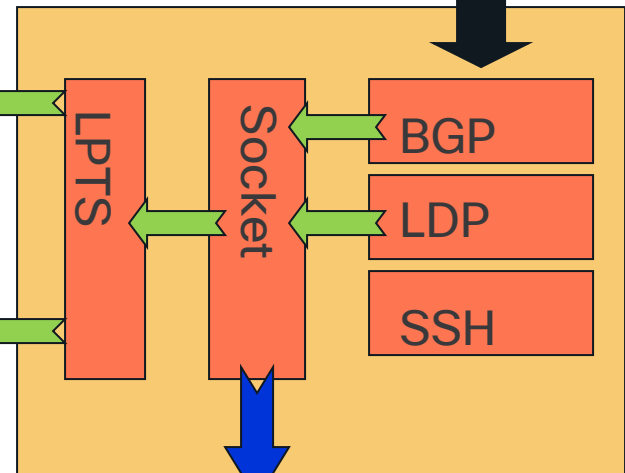
```
router bgp
 neighbor 202.4.48.99
  ttl_security

 mpls ldp
 ...
 !
```

### LC 1 Hardware Forwarding Entries

Local	port	Remote	port	Rate	Priority
Any	ICMP	ANY	ANY	1000	low
any	179	any	any	100	medium
any	179	202.4.48.99	any	1000	medium
202.4.48.1	179	202.4.48.99	2223	10000	medium
200.200.0.2	13232	200.200.0.1	646	100	medium

### LC 2 Hardware Forwarding Entries



# LPTS – sample output

```
RP/0/RP0/CPU0:CRS#show lpts pifib hardware policer loc 0/4/cpu0
```

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

```
Burst = 100ms for all flow types  
-----
```

FlowType	Policer	Type	Cur. Rate	Def. Rate	Accepted	Dropped
BGP-known	113	Static	25000	25000	0	0
BGP-cfg-peer	114	Static	10000	10000	0	0
BGP-default	115	Static	1500	1500	0	0
LDP-TCP-known	122	Static	25000	25000	72532	0
LDP-TCP-cfg-peer	152	Static	10000	10000	0	0
LDP-TCP-default	154	Static	10000	10000	0	0
LDP-UDP	158	Static	2500	2500	421787	0
OSPF-mc-known	107	Static	20000	20000	3636579	0
OSPF-mc-default	111	Static	5000	5000	0	0
OSPF-uc-known	161	Static	5000	5000	0	0
OSPF-uc-default	162	Static	1000	1000	0	0
SNMP	133	Static	2000	2000	0	0
SSH-known	135	Static	1000	1000	0	0
SSH-default	136	Static	1000	1000	0	0

# Config Management using CLI



# IOS-XR CLI: New CLI format

- New CLI reflects the HW position in the system
  - Introduces the Hierarchical location scheme
  - Each linecard has three-level identification: **Shelf/Slot/cpu #**
  - Interfaces have the **Shelf/Slot/Bay/Interface** scheme
- CRS-1 is designed to scale to 72 linecard chassis with a potential of 1296 (1152+144) line cards and RP slots
  - Location identifiers use **R/S/M/I** format
    - R = Rack (applicable in multi-chassis systems)
    - S = Slot (physical slot the module is in)
    - M = module (0 for 'fixed' PLIMs, n for SPAs)
    - I = Interface
- Protocol referenced by address family type - **v4/v6**
- Backward compatible command-set with IOS

```
RP/0/RP0/CPU0:CRS1#show ipv4 interface brief
```

Interface	IP-Address	Status	Protocol	Vrf-Name
Bundle-Ether1000	10.2.66.1	Up	Up	default
HundredGigE0/6/0/2	10.1.14.1	Up	Up	default
HundredGigE0/6/0/3	unassigned	Up	Up	default
HundredGigE0/7/0/0	unassigned	Up	Up	default

# IOS-XR CLI: New CLI format and Configuration Modes

- Config modes include:
  - Admin mode
- Admin mode is newly introduced compared to IOS
- Admin mode allows viewing / configuring shared resources
  - Fabric
  - Logical Router
  - Package installation
  - Online Diagnostics

```
RP/0/RP0/CPU0:CRS# admin
RP/0/RP0/CPU0:CRS(admin)#
```

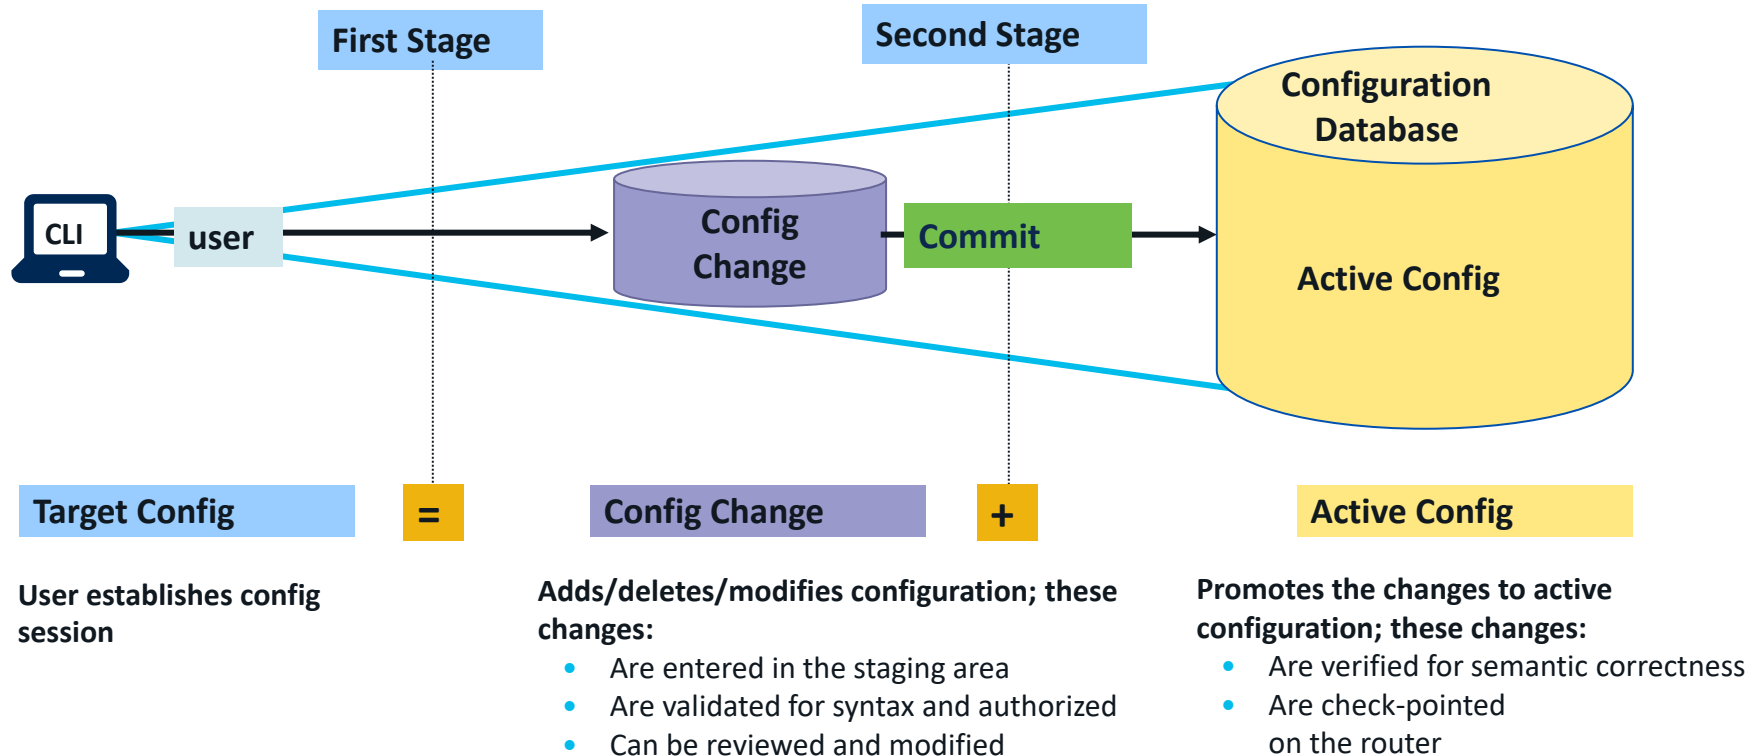
```
RP/0/RP0/CPU0:NCS5501# admin
sysadmin-vm:0_RP0#
```

```
RP/0/RP0/CPU0:8K# admin
```

```
WARNING: Admin mode has been deprecated. Please consult the Command Line Interface Guide.
```

```
RP/0/RP0/CPU0:8K#
```

# IOS-XR CLI: Two Stage Config model



# IOS-XR and IOS Config differences

- IOS-XR configuration is held in binary form which is quicker to parse and process - 'show running-configuration' is just an ASCII representation of the binary data extracted from all nodes in the system
- There is no concept of a startup configuration like in IOS
- If one copies the running config to startup, a backup config with the name "startup" is created
- Router config is based on two stage config model.
- "running" or "active" config can not be modified directly.
- Instead, user config first enters a staging area (first stage)
- Must be explicitly promoted to be part of active config (second stage).

IOS-XR	IOS
Configuration changes do NOT take place after <CR>	Configurations take place immediately after <CR>
Configuration changes must be 'committed' before they take effect	No commit
Allows you to verify your configuration before applying it	No verification required
Two stage configuration model	Not available
Configuration rollback	Not available
Provision to pre-configure	Not available
New config plane – Admin mode	Not available
Feature centric	Interface centric

# IOS-XR CLI: New CLI format and Configuration Modes

- Config modes include:
  - Privileged exec mode
  - Global config mode
- Config sub-mode

```
RP/0/RP0/CPU0:CRS1#config t
RP/0/RP0/CPU0:CRS1 (config)#interface HundredGigE0/6/0/2
RP/0/RP0/CPU0:CRS1 (config-if)#
```

# IOS-XR CLI: Config Commits

```
RP/0/RP0/CPU0:CRS1#show run int gi0/2/0/0
```

```
% No such configuration item(s)
```

```
RP/0/RP0/CPU0:CRS1#conf
```

```
RP/0/RP0/CPU0:CRS1(config)#interface gig0/2/0/0
```

```
RP/0/RP0/CPU0:CRS1(config-if)#ipv4 address 100.12.1.1/24
```

```
RP/0/RP0/CPU0:CRS1(config-if)#commit
```

```
RP/0/RP0/CPU0:Jan 22 00:49:28.119 : config[65691]: %MGBL-CONFIG-6-DB_COMMIT : Configuration committed by user 'root'. Use 'show configuration commit changes 1000000037' to view the changes.
```

```
RP/0/RP0/CPU0:iosxr(config-if)#end
```

```
RP/0/RP0/CPU0:Jan 22 00:49:30.701 : config[65691]: %MGBL-SYS-5-CONFIG_I : Configured from console by root
```

```
RP/0/RP0/CPU0:CRS1#
```

```
RP/0/RP0/CPU0:CRS1#show run int gigabitEthernet 0/2/0/0
```

```
interface GigabitEthernet0/2/0/0
```

```
ipv4 address 100.12.1.1 255.255.255.0
```

# IOS-XR CLI: Config Commit

- Commit keyword writes config into Active Config
- Supplies a commit ID to help in Config Rollback
  - **1000000036** is the commit ID in previous illustration
- List of commits can be viewed
  - History list is maintained upto 100 commits
- Commits can be labeled with user-friendly 'tags'
  - Eliminates the cumbersome IDs
- Configuration commits can be rolled back to a previous commit point.

```
RP/0/RP0/CPU0:CRS1#show configuration commit list
SNo. Label/ID      User      Line      Client      Time Stamp
~~~~ ~~~~~~      ~~~~      ~~~~      ~~~~~~      ~~~~~~
1      1000000037      root      con0_0_CPU      CLI      Sun Jan 22 00:49:28 2023
2      1000000036      john01    vty1:node0_RP0_CPU      CLI      Sun Jan 22 00:11:00 2023
3      1000000035      lina05    vty1:node0_RP0_CPU      CLI      Sat Jan 21 18:01:00 2023
```

# IOS-XR CLI: Config Rollback

```
RP/0/RP0/CPU0:CRS1#conf t
RP/0/RP0/CPU0:CRS1(config)#hostname iox-CL23
RP/0/RP0/CPU0:CRS1(config)#commit
RP/0/RP0/CPU0:Jan 2 01:00:55.302 : config[65691]: %MGBL-CONFIG-6-DB_COMMIT : Configuration
committed by user 'root'. Use 'show configuration commit changes 100000034' to view the changes.
RP/0/RP0/CPU0:iox-CL23(config)#end
RP/0/RP0/CPU0:iox-CL23#
RP/0/RP0/CPU0:iox-CL23#rollback configuration last 1
Loading Rollback Changes.
Loaded Rollback Changes in 1 sec
Committing.
1 items committed in 1 sec (0)items/sec
Updating.RP/0/RP0/CPU0:Jan 2 01:01:07.143 : config_rollback[65691]: %MGBL-CONFIG-6-DB_COMMIT :
Configuration committed by user 'root'. Use 'show configuration commit changes 100000035' to
view the changes.

Updated Commit database in 1 sec
Configuration successfully rolled back 1 commits.
RP/0/RP0/CPU0:CRS1#
```

# IOS-XR CLI: Config error handling

- Two levels of config error handling
- Parser/Syntax error
  - Identified by the parser when the <return> key is entered
- Commit error (aka Semantic error)
  - Syntactically correct but erroneous from config commit standpoint
  - Error details viewed through “show configuration failed” command
  - Common reasons for this error include:
    - Non-atomic config sequence
    - Lack of predecessor config
    - Unsupported config from platform perspective

```
RP/0/RP0/CPU0:CRS1#configure terminal
RP/0/RP0/CPU0:CRS1 (config)#policy-map TEST
RP/0/RP0/CPU0:CRS1 (config-pmap)#class NAME
RP/0/RP0/CPU0:CRS1 (config-pmap-c)#set precedence 0
RP/0/RP0/CPU0:CRS1 (config-pmap-c)#commit
% Failed to commit one or more configuration items during a pseudo-atomic operation. All changes made
have been reverted.Please issue 'show configuration failed [inheritance]' to view the errors
RP/0/RP0/CPU0:ios (config-pmap-c)#show configuration failed
!! SEMANTIC ERRORS: This configuration was rejected by the system due to semantic errors
policy-map TEST
  class NAME
    set precedence 0
!!% Entry not found: class map NAME used in policy map TEST does not exist
```

# Protocol Configuration and Verification

# IOS-XR Routing Protocols

- IOS
  - Protocol configuration exists at:
    - Global Process Level
    - Routing Process
    - Interface Configuration
    - (normally one level deep)
- IOS XR
  - Protocol configuration exists at:
    - Only within the protocol
    - Hierarchical
    - (multiple levels deep is common)

# Static Routes

## IOS

```
ip route 192.1.1.0 255.255.255.0 gi4/0
ip route 223.255.254.0 255.255.255.0 10.13.0.1
ipv6 route 5301::1111/128 fec0::1
```

```
show run | include route
show run | begin ip route
show run | begin ipv6 route
```

## IOS XR

### router static

#### address-family ipv4 unicast

```
43.43.44.0/24 Serial0/5/3/3/0:2
192.1.1.0/24 GigabitEthernet0/4/0/0
223.255.254.254/32 MgmtEth0/1/CPU0/0
```

!

#### address-family ipv6 unicast

```
5301::1111/128 HundredGigE0/5/3/3/0:0
```

!

!

```
show run router static
show run router static ipv4 unicast
show run router static ipv6 unicast
```

# OSPF Configuration

## IOS

```
router ospf 1
 area 0 authentication message-digest
 area 2 authentication message-digest
 network 10.100.1.0 0.0.0.7 area 0
 network 10.200.1.0 0.0.0.15 area 2

interface gi0/0
 ip ospf network point-to-point
 ip ospf message-digest-key 1 md5 CISCO
 ip ospf cost 100
!
interface gi0/1
 ip ospf network point-to-point
 ip ospf message-digest-key 1 md5 CISCO
 ip ospf cost 100
!
interface gi0/2
 ip ospf network point-to-point
 ip ospf message-digest-key 1 md5 CISCO
 ip ospf cost 2000
!
interface gi0/3
 ip ospf network point-to-point
 ip ospf message-digest-key 1 md5 CISCO
 ip ospf cost 9999
```

## IOS XR

```
router ospf 1
 authentication message-digest message-digest-key 1 md5 CISCO
 network point-to-point
 area 0
   cost 100
   interface gi0/0/0/0
   !
   interface gi0/0/0/1
 area 2
   cost 2000
   interface gi0/0/0/2
   !
   interface gi 0/0/0/3
   cost 9999
```

```
show run router ospf
show run router ospf 1 area 2
show run formal router ospf
```

# OSPF Configuration

## IOS XR

```
RP/0/RP0/CPU0:CRS#show run router ospf

router ospf 1
 authentication message-digest message-digest-key 1 md5 CISCO
 network point-to-point
 area 0
   cost 100
   interface gi0/0/0/0
   !
   interface gi0/0/0/1
 area 2
   cost 2000
   interface gi0/0/0/2
   !
   interface gi0/0/0/3
   cost 9999
```

```
RP/0/RP0/CPU0:CRS#show run formal router ospf

router ospf 1 authentication message-digest message-digest-key 1 md5
router ospf 1 network point-to-point
router ospf 1 area 0
router ospf 1 area 0 cost 100
router ospf 1 area 0 interface gi0/0/0/0
router ospf 1 area 0 interface gi0/0/0/1
router ospf 1 area 2
router ospf 1 area 2 cost 2000
router ospf 1 area 2 interface gi0/0/0/2
router ospf 1 area 2 interface gi0/0/0/3
router ospf 1 area 2 interface gi0/0/0/3 cost 9999
```

```
show run router ospf
show run router ospf 1 area 2
show run formal router ospf
```

# ISIS CLI Differences

## IOS

### router isis ios

```
net 47.1111.1111.0001.0000.0c00.0006.00
domain-password mydomainpasswd authenticate snp validate
area-password myareapasswd authenticate snp validate
metric-style wide
log-adjacency-changes
nsf ietf
maximum-paths 6
!
```

### interface TenGigabitEthernet0/0/0

```
ip address 201.1.1.2 255.255.255.0
ip router isis ios
isis metric 30
```

### interface TenGigabitEthernet0/0/1

```
ip address 201.1.2.2 255.255.255.0
ip router isis ios
isis metric 100
```

## IOS-XR

### router isis iosxr

```
net 47.1111.1111.0001.0000.0c00.0007.00
nsf ietf
log adjacency changes
lsp-password mydomainpasswd level 1 send-only snp send-only
lsp-password myareapasswd level 2 send-only snp send-only
```

### address-family ipv4 unicast

```
metric-style wide
maximum-paths 6
```

!

### interface TenGigE0/2/0/0

```
address-family ipv4 unicast
metric 30
```

### interface TenGigE0/2/0/1

```
hello-interval 60
address-family ipv4 unicast
metric 100
```

# BGP

## IOS

```
router bgp 65000
  bgp log-neighbor-changes
  no bgp default ipv4-unicast
  neighbor 172.16.2.5 remote-as 651111
  neighbor 192.168.1.2 remote-as 65000
  neighbor 192.168.1.2 update-source Loopback 0
!
address-family ipv4
  network 192.168.1.1 mask 255.255.255.255
  neighbor 172.16.2.5 activate
  neighbor 192.168.1.2 activate
  maximum-paths 8
```

## IOS XR

```
router bgp 65000
  address-family ipv4 unicast
  network 192.168.1.1/32
  maximum-paths 8
!
  neighbor 192.168.1.2
  remote-as 65000
  update-soruce Loopback 0
  address-family ipv4 unicast
!
  neighbor 172.16.2.25
  remote-as 651111
  address-family ipv4 unicast
  route-policy PASS-ALL in
  route-policy PASS-ALL out
```

**Will not accept routes from  
EBGP peers by default**


# Route Policy Language



How many of you  
get confused by  
route-maps?

# RPL brings clarity to Route-Maps

```
router bgp 100
  bgp log-neighbor-changes
  neighbor 100.64.1.1 remote-as 65000
  neighbor 100.64.1.1 password bgn2dcx
  neighbor 100.64.1.1 remove-private-as
  neighbor 100.64.1.1 soft-reconfiguration inbound
  neighbor 100.64.1.1 prefix-list INTERNET-IN in
  neighbor 100.64.1.1 prefix-list INTERNET-OUT out
  neighbor 100.64.1.1 route-map INTERNET-IN in
  neighbor 100.64.1.1 route-map INTERNET-OUT out
  neighbor 100.64.1.1 filter-list 3 in
  neighbor 100.64.1.1 filter-list 7 out
  neighbor 100.64.1.1 distribute-list 1 in
  neighbor 100.64.1.1 distribute-list 2 out
```



What is the order of processing?

# What is RPL ?

- Route Policy Language
- Used to filter routing information
  - Remove routes
  - Change attributes



# Why use RPL ?

- Clarity
  - Easy to understand if-then-else logic
- Scalability – reuse the logic
  - Parameterization
  - Nesting of route-policies
- “inline” keyword shows us expanded view of all included policies.
- Can view the effect of the route-policy before applying it.
  - Test/Verify the proposed route-policy behavior before its actually applied.
- RPLs can be used in BGP, OSPF, ISIS, EIGRP, RIP
  - Not restricted just to BGP




# Actions in a RPL



There is an implicit drop at the end of RPL processing.

A route must be given a **'ticket'** to ensure that it has been inspected by the RPL

- **Pass** – prefix allowed if not later dropped
  - **pass** grants a ticket to defeat default drop
  - Execution **continues** after pass
- **Set** – value changed, prefix allowed if not later dropped
  - Any **set** at any level grants a ticket
  - Execution **continues** after **set**
  - Values can be set more than once
- **Drop (default)** – prefix is discarded 
  - Explicit drop **stops** policy execution
  - Implicit drop (if policy runs to end without getting a ticket)
- **Done** – accepts prefix and **stops** processing
  - **done** grants a ticket to defeat default drop



# Basic RPL Examples

- Basic **Pass** Policy

## Example Configuration

```
route-policy PASS-ALL
  pass
end-policy
```

- Basic **Drop** Policy

## Example Configuration

```
route-policy DROP-ALL
  drop
end-policy
```

Somewhat redundant due to implicit drop

# RPL Examples

- Basic conditional statement

Logic	Example Configuration
<pre>if match-condition-one then   action-one endif</pre>	<pre>if med eq 150 then   pass endif</pre>

Conditional Match

Action

- Branching options

Logic	Example Configuration
<pre>if match-condition-one then   action-one else   action-two endif</pre>	<pre>if destination in (10.0.0.0/8 ge 8) then   pass else   drop endif</pre>

Comparison operator

Notice we are matching networks directly in the RPL.  
Supports Prefix Matching or WildCard

# RPL Examples

(continued)

- Multiple Branching options

Logic	Example Configuration
<pre>if match-condition-one then   action-one elseif match-condition-two then   action-two else   action-three endif</pre>	<pre>if destination in (10.0.0.0/8 ge 8) then   set tag 1 elseif destination in (172.16.0.0/12 ge 12) then   set tag 2 else   drop endif</pre>

# Nested Conditions

- *If* statements within other *if/elseif/else* statements
  - Method used for multiple conditions
- Nesting can be any depth

Logic	Example Configuration
<pre>if match-condition-one then   if match-condition-two then     action-one   endif endif</pre>	<pre>if as-path passes-through '100' then   if destination in PREFIX-SET-RFC1918 then     pass   endif endif</pre>

# Original Value

- A conditional match does not occur on intermediary values during the route policy processing.

```
Conditional Matches on Original Value
route-policy ORIGINAL-VALUES
  if med eq 100 then
    set med 200
  endif
  if med eq 200 then
    drop
  endif
end-policy
```

- In the example, only the original routes with a MED of 200 are dropped and the routes with values set to 200 are not dropped.

# Simplifying BGP AS-Path Conditions

AS Path Selection Criteria	Route-Map AS-Path ACL Logic ( <i>ip as-path access-list 1</i> )	RPL Logic
Local Routes	permit ^\$	if as-path is-local
Only Routes From Neighbor AS 200	permit ^200_	if as-path neighbor-is '200'
Only Routes Originating From AS 200	permit _200\$	if as-path originates-from '200'
Passes Through AS200	permit _200_	if as-path passes-through '200'
Routes From 3 ASes or less away	permit ^[0-9]+ [0-9]+ [0-9]+?	if as-path length le 3

# RPL Examples

## Bad RPL Logic

```
route-policy METRIC-MODIFICATION
  if destination in (10.0.0.0/8 ge 8) then
    set med 100
  endif
  set med 200
end-policy
```

Overwrites  
Setting

## Great RPL Logic – Option #1

```
route-policy METRIC-MODIFICATION
  if destination in (10.0.0.0/8 ge 8) then
    set med 100
  else
    set med 200
  endif
end-policy
```

## Good RPL Logic – Option #2

```
route-policy METRIC-MODIFICATION
  if destination in (10.0.0.0/8 ge 8) then
    set med 100
  done
  endif
  set med 200
end-policy
```

Stops all  
processing on  
matched  
prefixes

# RPL Policy Sets

- Prefix-lists, ACLs, AS-Path ACLs can be confusing because of permit/deny actions
- IOS XR uses policy sets to store the same information:  
Prefix Set, AS-Path Set, Community Set, Extended Community Set
- There is not a deny in a Policy Set  
More later on how to fix that 😊
- Processing occurs until the first match is made

# Named and Inline Set (Same behavior)

## Inline Set Example Configuration

```
route-policy RFC1918-INLINE-SET
  if destination in (10.0.0.0/8 ge 8, 172.16.0.0/12 ge 12, 192.168.0.0/16 ge 16) then
    pass
  else
    drop
  endif
end-policy
```

## Named Set Example Configuration

```
route-policy RFC1918-PREFIX-SET
  if destination in PREFIX-SET-RFC1918 then
    pass
  else
    drop
  endif
end-policy
!
prefix-set PREFIX-SET-RFC1918
  10.0.0.0/8 ge 8,
  172.16.0.0/12 ge 12,
  192.168.0.0/16 ge 16
end-set
```

# Viewing Set Based RPLS

Keyword required to see sets in the RPL

## Inline Example Configuration

```
RP/0/RP0/CPU0:XR1# show rpl route-policy RFC1918-PREFIX-SET inline

route-policy RFC1918-PREFIX-SET
  if destination in (10.0.0.0/8 ge 8, 172.16.0.0/12 ge 12, 192.168.0.0/16 ge 16) then
    pass
  else
    drop
  endif
end-policy
```

Avoids having to scroll back and forth in the configuration

# RPL Existence

- **Question:** Non IOS XR systems have a different behavior to a route if there is reference to a non-existent ACL, prefix-list, or route-map. How does IOS XR process these occurrences?
- **Answer:** IOS XR RPL does not use access-lists or prefix-lists. It uses policy-sets. If a policy-set or RPL is referenced and does not exist, the configuration **will fail on the commit!**

# Nesting of RPLs

## Example Configuration

```
route-policy CHILD-ONE
  set weight 100
end-policy

route-policy CHILD-TWO
  set community (2:1234) additive
end-policy

route-policy PARENT
  apply CHILD-ONE
  apply CHILD-TWO
  pass
end-policy
```

Can go multiple levels deep

# Parameter Passing

```
Single Parameter
route-policy PARAM($MED)
  set med $MED
end-policy

router bgp 300
address-family ipv4 unicast
!
neighbor 192.1.1.2
  remote-as 400
  address-family ipv4 unicast
    route-policy PARAM (50) in
    route-policy PASS-ALL out
```

List of policy parameters

Accessing the passed parameter

Calling policy and passing parameter

# Using Multiple Parameters in BGP

## Multiple Parameter

```
prefix-set PREFIX-SET-RFC1918
 10.0.0.0/8 ge 8,
 172.16.0.0/12 ge 12,
 192.168.0.0/16 ge 16
end-set

route-policy SP-PEER($AS, $PREFIX)
 if destination in $PREFIX and as-path originates-from '$AS' then
   pass
 endif
end-policy

router bgp 300
 address-family ipv4 unicast
 !
 neighbor 192.1.1.2
  remote-as 400
  address-family ipv4 unicast
   route-policy SP-PEER (50, PREFIX-SET-RFC1918) in
   route-policy PASS-ALL out
```

Can reference a  
policy-set

# Traditional Boolean Logic Operators

## Negation

```
if not destination in PREFIX-SET-RFC1918 then
    pass
endif
```

## Conjunction

```
if destination in PREFIX-SET-RFC1918 and as-path passes-through '100' then
    pass
endif
```

## Disjunction

```
if destination in PREFIX-SET-RFC1918 or as-path passes-through '100' then
    pass
endif
```

# Boolean Logic Operations

- Compound Boolean expressions are supported
- Boolean expressions are evaluated in order of operator precedence, from left to right
- The highest precedence operator is NOT, followed by AND, and then OR

med eq 10 and not destination in (10.1.3.0/24) or community is (56:78)

# Boolean Logic Order of Operations

- The highest precedence operator is **NOT**, followed by **AND**, and then **OR**
  - Parentheses may be used to change the order of operations and improve readability
- `med eq 10 and (not destination in (10.1.3.0/24) or community is (56:78))`
- Conditional match that requires a route to not pass through AS 100 or AS 200, and must be within the 192.168.0.0/16 network range

## Use of parentheses

if not (as-path passes-through '100' or as-path passes-through '200') and destination in (192.168.0.0/16 ge 16)

# Prefix List Functionality in an RPL

- Prefix List

```
ipv4 prefix-list PREFIX-LIST deny 192.168.1.1/32  
ipv4 prefix-list PREFIX-LIST permit 192.168.0.0/16 ge 16
```

- RPL Logic

```
if ( (not destination in (192.168.1.1) ) and destination in (192.168.0.0/16 ge 16) ) then  
    pass  
endif
```

# RPL Show Commands

Only display prefixes matching policy – filter show command

```
RP/0/0/1:CRS1# show bgp route-policy SAMPLE
BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 729
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network Next Hop Metric LocPrf Weight Path
* 10.13.0.0/16 192.168.40.24 0 1878 704 701 200 ?
* 10.16.0.0/16 192.168.40.24 0 1878 704 701 i
```

# IOS XR: MPLS LDP, TE and VPN Configuration and Verification



# LDP

```
RP/0/RP0/CPU0:CRS1 (config)# mpls ldp
RP/0/RP0/CPU0:CRS1 (config-ldp)# ?
  backoff          Configure session backoff parameters
  discovery        Configure discovery parameters
  explicit-null    Configure explicit-null advertisement
  graceful-restart Configure LDP graceful restart feature
  holdtime        Configure session holdtime
  interface        Enable LDP on an interface and enter interface submode
  label           Configure label advertisement control
  log             Configure log information
  neighbor        Configure neighbor parameters
  router-id       Configure router Id
  signalling      Configure LDP signalling parameters
```

Must be  
explicitly  
configured

- LDP is enabled on an interface via the interface submode underneath the LDP submode

```
RP/0/RP0/CPU0:CRS1 (config-ldp)# interface HundredGigE0/6/0/2
RP/0/RP0/CPU0:CRS1 (config-ldp-if)# ?
  discovery        Configure interface LDP discovery parameter
```

# show mpls forwarding/show mpls interfaces

```
RP/0/RP0/CPU0:CRS1# show mpls forwarding
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched	T
17	Pop Label	192.168.1.1/32	PO0/6/0/2	192.168.6.1	0	
18	Pop Label	192.168.1.3/32	PO0/6/0/0	192.168.7.3	0	
19	Unlabelled	192.168.1.4/32	PO0/6/0/1	192.168.8.5	0	
20	Unlabelled	192.168.1.5/32	PO0/6/0/1	192.168.8.5	0	
21	29	192.168.1.6/32	PO0/6/0/2	192.168.6.1	0	
22	30	192.168.1.7/32	PO0/6/0/0	192.168.7.3	0	
	Unlabelled	192.168.1.7/32	PO0/6/0/1	192.168.8.5	0	

```
RP/0/RP0/CPU0:CRS1# show mpls interfaces
```

Interface	LDP	Tunnel	Enabled	
HundredGigE0/6/0/0		Yes	Yes	Yes
HundredGigE0/6/0/1		No	Yes	Yes
HundredGigE0/6/0/2		Yes	Yes	Yes

# IOS XR L3VPN

# VRF Configuration

## Configuring the VRF in the Global Mode

```
vrf vpn1
  description foo
  address-family ipv4 unicast
    import route-target
    100:1
  export route-target
    200:1
```

## Assigning interface to VRF

```
interface g0/1/0/2
  vrf vpn1
  ipv4 address 1.1.1.2/24
```

**Prior to assigning the VRF to the interface, the IP address must be removed, or a semantic error will be seen on commit**

```
!!% 'RSI' detected the 'fatal' condition 'The interface's numbered and unnumbered IPv4/IPv6 addresses or IPv6 enable must be removed prior to changing or deleting the VRF'
```

# VRF-Lite & PE-CE Configuration

## router static

```
vrf vpn1
  address-family ipv4 unicast
    10.1.1.1/32 g0/1/0/2
```

## router eigrp 100

```
vrf vpn1
  address-family ipv4
    router-id 100.100.100.100
    redistribute bgp 100 route-policy policy1
  interface g0/1/0/2
    site-of-origin 100:1
```

## router ospf 100

```
vrf vpn1
  router-id 100.100.100.100
  domain-id type 0005 value 000102030405
  domain-tag 101
  redistribute bgp 100 route-policy policy1
  area 0
  interface g0/1/0/2
```

## router bgp 100

```
bgp router-id 100.100.100.100
  address-family vpnv4 unicast
    vrf vpn1
      rd [auto | 100:1]
      label-allocation-mode [per-ce | per-vrf]
  address-family ipv4 unicast
    neighbor 1.1.1.1
      remote-as 65523
  address-family ipv4 unicast
    route-policy vpn1-in in
    route-policy vpn1-out out
```

**VPNv4 Address-Family  
Required for BGP in VRF**

**RD configured under the  
BGP VRF instance**

# BGP VPNv4 Configuration

## Configuring PE to RR VPNv4 iBGP Neighbors

```
router bgp 100
  bgp router-id 100.100.100.100
  address-family vpnv4 unicast
!
neighbor 192.168.1.1
  remote-as 100
  update-source loopback 0
  address-family vpnv4 unicast
```

# IOS XR Ethernet VPN (EVPN)

# Attachment Circuit (AC) Configuration

## Configuring the AC

```
interface GigabitEthernet0/0/0/4.30 l2transport  
  encapsulation dot1q 30
```

## Assigning AC to EVPN Virtual Private Virtual Service (VPWS)

```
l2vpn xconnect group evpn_vpws  
  p2p evpn_vpws_5-6  
  interface GigabitEthernet0/0/0/4.30  
    neighbor evpn evi 1 target 6 source 5
```

**l2transport keyword  
denotes AC/Layer 2  
subinterface**

# BGP VPNv4 Configuration

## Configuring PE to RR L2VPN EVPN iBGP Neighbors

```
router bgp 100

address-family l2vpn evpn
  neighbor 192.168.100.100
  remote-as 100
address-family l2vpn evpn
```

# Full EVPN VPWS Configuration

## XR1

```
interface GigabitEthernet0/0/0/4.30 l2transport
encapsulation dot1q 30
int GigabitEthernet0/0/0/4
```

```
l2vpn xconnect group evpn_vpws
p2p evpn_vpws_5-6
interface GigabitEthernet0/0/0/4.30
neighbor evpn evi 1 target 6 source 5
```

```
! L2VPN BGP session to RR
router bgp 100
address-family l2vpn evpn
neighbor 192.168.100.100
remote-as 100
address-family l2vpn evpn
```

## XR2

```
interface GigabitEthernet0/0/0/4.30 l2transport
encapsulation dot1q 30
int GigabitEthernet0/0/0/4
```

```
l2vpn xconnect group evpn_vpws
p2p evpn_vpws_5-6
interface GigabitEthernet0/0/0/4.30
neighbor evpn evi 1 target 5 source 6
```

```
! L2VPN BGP session to RR
router bgp 100
address-family l2vpn evpn
neighbor 192.168.100.100
remote-as 100
address-family l2vpn evpn
```

# EVPN Show Commands example

Displays l2vpn connection details

```
RP/0/RP0/CPU0:XR1# show l2vpn xconnect detail

Group evpn_vpws, XC evpn_vpws_5-6, state is up; Interworking none
AC: GigabitEthernet0/0/0/4.30, state is up
  Type VLAN; Num Ranges: 1
  Rewrite Tags: []
  VLAN ranges: [30, 30]
  MTU 1504; XC ID 0x2; interworking none
  Statistics:
    packets: received 21, sent 23
    bytes: received 2208, sent 2336
    drops: illegal VLAN 0, illegal length 0
EVPN: neighbor 192.168.2.2, PW ID: evi 1, ac-id 6, state is up ( established )
  XC ID 0xa0000003
  Encapsulation MPLS
  Encap type Ethernet, control word disabled
  Sequencing not set
  Ignore MTU mismatch: Enabled
  Transmit MTU zero: Enabled
  LSP : Up
```

# EVPN Show Commands example (continued)

Displays I2vpn connection details

```

EVPN          Local          Remote
-----
Label         24004          24007
MTU           1518          unknown
Control word disabled      disabled
AC ID         5              6
EVPN type     Ethernet      Ethernet

-----

Create time: 01/05/2022 02:25:24 (14:23:47 ago)
Last time status changed: 01/05/2022 16:34:22 (00:14:49 ago)
Statistics:
  packets: received 23, sent 21
  bytes: received 2336, sent 2208
RP/0/RP0/CPU0:XR1#
```

# Traffic Engineering Configuration and Verification



# Traffic Engineering Configuration

- **TE has its own configuration mode**
  - Specifies which physical interfaces TE should be enabled on.
  - Various TE timers / Affinity-map / DiffServ are all configured in this mode.
- **RSVP also has its configuration mode**
  - Specifies interface and bandwidth associated with that interface.
- **Traffic Engineering needs to be enabled in the IGP**
  - This enables TE extensions in that area ( OSPF sends type 10 LSA)
- **Build the Tunnel-TE interface.**
  - Configure the destination IP address.
  - Path options/ initial signaled b/w / auto route announce are configured here.

# MPLS-TE Submode

```
RP/0/RP0/CPU0:CRS1(config)# mpls traffic-eng

RP/0/RP0/CPU0:CRS1(config-mpls-te)# ?
bfd                Configure BFD parameters
fast-reroute       Fast-reroute config parameters
interface         Enable MPLS-TE on an interface
link-management    MPLS Link Manager subcommands
maximum            Maximum number of configurable tunnels
path-selection     Path Selection Configuration
reoptimize         Reoptimize timers frequency
signalling         Signalling options
topology           Topology Database Configuration
RP/0/RP0/CPU0:CRS1(config-mpls-te)# interface HundredGigE0/6/0/0
```

```
RP/0/RP0/CPU0:CRS1# show running-config mpls traffic-eng
mpls traffic-eng
 interface HundredGigE0/6/0/0
 !
 interface HundredGigE0/6/0/1
 !
 interface HundredGigE0/6/0/2
```

# RSVP Configuration

```
RP/0/RP0/CPU0:CRS1# show running-config rsvp
rsvp
  interface HundredGigE0/6/0/0
    bandwidth 155520
  !
  interface HundredGigE0/6/0/1
    bandwidth 155520
  !
  interface HundredGigE0/6/0/2
    bandwidth 155520
  !
!
```

# Prepare the IGP (OSPF / ISIS)

## OSPF

```
router ospf 1
 area 0
  mpls traffic-eng
 !
 mpls traffic-eng router-id Loopback0
 !
```

## ISIS

```
router isis 1
 address-family ipv4 unicast
  mpls traffic-eng level-1
  mpls traffic-eng router-id Loopback0
  metric-style wide [ transition ] [ level { 1 | 2 } ]
 !
 !
```

# Tunnel Head-End Configuration

1. Create a 'tunnel-te <num>' interface
2. Under this, add your head-end interface configuration

```
interface tunnel-te6
  ipv4 unnumbered Loopback0
  autoroute announce
  destination 192.168.1.6
  path-option 5 explicit name foo
  path-option 10 dynamic
!
explicit-path name foo
  index 1 next-address ipv4 unicast 192.168.1.3
  index 2 next-address ipv4 unicast 192.168.1.8
  index 3 next-address ipv4 unicast 192.168.1.6
```

# Tunnel Head-End Configuration

```
RP/0/RP0/CPU0:CR3-CRS#show mpls traffic-eng tunnels tabular
```

Tunnel Name	LSP ID	Destination Address	Source Address	Tun State	FRR State	LSP Role	Path Prot
tunnel-te1	15	10.122.0.1	10.122.0.3	up	Inact	Head	Inact
tunnel-te2	5	10.122.0.2	10.122.0.3	up	Inact	Head	Inact
tunnel-te7	4	10.122.0.7	10.122.0.3	up	Inact	Head	Inact
tunnel-te8	4	10.122.0.8	10.122.0.3	up	Inact	Head	Inact
CR1-CRS_t3	2	10.122.0.3	10.122.0.1	up	Inact	Tail	
CR2-CRS_t3	514	10.122.0.3	10.122.0.2	up	Inact	Tail	
CR7-CRS_t3	576	10.122.0.3	10.122.0.7	up	Inact	Tail	
CR8-CRS_t3	265	10.122.0.3	10.122.0.8	up	Inact	Tail	
CR9-CRS_t4	122	10.122.0.4	10.122.0.9	up	Inact	Mid	
CR5-CRS_t9	6	10.122.0.9	10.122.0.5	up	Inact	Mid	

# IOS-XR Operation and Monitoring Tools

# Process Management

- Process
  - An executable portion of code run within its own memory space
- Threads
  - A process may contain one or more threads or a ‘sub-process’ e.g. OSPF process has a thread which handles ‘hellos’
  - A thread may only run when the parent process is allocated runtime by the system scheduler.
- Job ID (JID) and Process ID (PID)
  - Each process is allocated a Job ID# or JID when it is first run. **Remains associated** with the process even if the process is stopped & restarted - The PID **changes** if the process is stopped and restarted
- Thread ID# (TID)
  - If a process contains threads, each is assigned a TID# associated with the PID/JID.

- Basic command
  - **show process**

## Process restart-ability

- **process restart <process name | number> <option>**
- Monitor commands:
  - **monitor processes**
  - **monitor threads (or top)**
- Troubleshooting commands:
  - **show exception, exception**
  - **dumpcore**
  - **show context, clear context**
  - **process core**
  - **follow process**

# Process Restartability

```
RP/0/RP1/CPU0:CRS1#process shutdown snmpd
```

```
RP/0/RP1/CPU0:CRS1#show processes snmpd
```

```
Job Id: 288
```

```
PID: 143532
```

```
Executable path: /disk0/hfr-base-6.6.2/bin/snmpd
```

```
Instance #: 1
```

```
Version ID: 00.00.0000
```

```
Respawn: ON
```

```
Respawn count: 1
```

```
Max. spawns per minute: 12
```

```
Last started: Mon Jan 18 10:32:22 2021
```

```
Process state: Killed (last exit status:
```

```
15)
```

```
Package state: Normal
```

```
Registered item(s): cfg/gl/snmp/  
cfg/gl/udpsnmp/  
cfg/gl/mibs/
```

```
core: TEXT SHARED MEM MAIN MEM
```

```
Max. core: 0
```

```
startup_path:
```

```
/pkg/startup/snmpd.startup
```

```
Ready: 11.636s
```

Process state reported as 'killed'

JID# remains constant,  
PID# changed on restart

Respawn counter incremented with  
process restart

```
RP/0/RP1/CPU0:CRS1#process restart snmpd
```

```
RP/0/RP1/CPU0:CRS1#show processes snmpd
```

```
Job Id: 288
```

```
PID: 8528114
```

```
Executable path: /disk0/hfr-base-6.6.2/bin/snmpd
```

```
Instance #: 1
```

```
Version ID: 00.00.0000
```

```
Respawn: ON
```

```
Respawn count: 2
```

```
Max. spawns per minute: 12
```

```
Last started: Fri Jan 22 11:46:38 2021
```

```
Process state: Run (last exit status :
```

```
15)
```

```
Package state: Normal
```

```
Started on config:
```

```
cfg/gl/snmp/admin/community/ww
```

```
core: TEXT SHARED MEM MAIN MEM
```

```
Max. core: 0
```

```
startup_path: /pkg/startup/snmpd.startup
```

```
Ready: 6.657s
```

```
Process cpu time: 0.721 user, 0.145 kernel,  
0.866 total
```

# 'Monitor process' command

- Command provide Unix 'top' like information
- Displays details on number of running processes, CPU and memory utilization
- Automatically updates every 10 seconds
- Can specify the location of the node that you wish to monitor, for example 0/RP0/CPU0 or 0/2/CPU0
- To change the parameters displayed by monitor processes, enter one of the interactive commands eg. ? to get help, n for the number of entries, t – sorted on cpu time, q to quit

```
233 processes; 788 threads; 4663 channels, 5906 fds
CPU states: 94.8% idle, 4.1% user, 1.0% kernel
Memory: 4096M total, 3599M avail, page size 4K
```

t - Sort on CPU time

JID	TIDS	Chans	FDs	Tmrs	MEM	HH:MM:SS	CPU	NAME
1	26	236	183	1	0	67:18:56	1.06%	procnto-600-smp-cisco...
256	5	39	21	4	292K	0:02:44	0.79%	packet
69	10	454	9	3	2M	0:33:07	0.62%	qnet
331	8	254	21	13	2M	0:15:20	0.52%	wdsysmon
55	11	23	15	6	36M	0:31:18	0.50%	eth_server
241	12	96	83	13	1M	0:04:54	0.37%	netio
171	15	97	44	9	2M	0:03:33	0.12%	gsp

m - Sort on memory usage

JID	TIDS	Chans	FDs	Tmrs	MEM	HH:MM:SS	CPU	NAME
55	11	23	15	6	36M	0:00:00	0.00%	eth_server
155	1	7	18	4	12M	0:00:00	0.00%	fgid_server
100	2	11	16	3	11M	0:00:00	0.00%	fgid_aggregator
257	8	16	36	3	8M	0:00:00	0.00%	parser_server
65554	7	16	3	3	7M	0:00:00	0.00%	devb-ata
53	5	237	633	0	4M	0:00:00	0.00%	dllmgr
121	11	48	67	19	3M	0:00:00	0.00%	bgp

# 'Show memory compare' command

- Process how to use the command:
  1. `show memory compare start`  
Takes the initial snapshot of heap usage
  2. `show memory compare end`  
Takes the second snapshot of heap usage
  3. `show memory compare report`  
Displays the heap memory comparison report

```
RP/0/RP1/CPU0:CRS1#show memory compare start
Successfully stored memory snapshot /harddisk:/malloc_dump/memcmp_start.out
RP/0/RP1/CPU0:CRS1#show memory compare end
Successfully stored memory snapshot /harddisk:/malloc_dump/memcmp_end.out
RP/0/RP1/CPU0:CRS1#show memory compare report
```

JID	name	mem before	mem after	difference	mallocs	restart
57	i2c_server	11756	11916	160	1	
121	bgp	2522256	2522208	-48	-1	
234	lpts_pa	408536	407632	-904	-14	
224	isis	3089108	3087900	-1208	0	
314	tcp	247196	245740	-1456	-9	
241	netio	808136	806464	-1672	-46	

# System Monitoring



# Multipipe support

- Support multiple pipes on the command line so that the output can be processed by enhanced utility set to filter and format the output of any show command.
- Upto 8 pipes are supported. This limit will be superseded by the limit of characters that can be typed on the single line (1024) if the individual commands specified with pipes are long enough.

```
RP/0/RP0/CPU0:CRS1# show platform | in FP-X | in 100GbE | utility wc -l
```

```
82
```

```
RP/0/RP0/CPU0:CRS1# show log start Jan 3 07:00:00 | in LDP | in " UP|DOWN" | ex "10.2[0-5]"
```

```
RP/6/RP0/CPU0:Jan 3 17:10:18.428 : mpls_ldp[1038]: %ROUTING-LDP-5-NBR_CHANGE : Neighbor 10.80.1.1:0, DOWN (Interface state down)
```

```
RP/6/RP0/CPU0:Jan 3 17:10:58.362 : mpls_ldp[1038]: %ROUTING-LDP-5-NBR_CHANGE : Neighbor 10.80.1.1:0, UP
```

```
RP/0/RP0/CPU0:CRS1#
```

```
RP/0/RP0/CPU0:CRS1# show proc cpu | ex " 0% ."
```

```
CPU utilization for one minute: 2%; five minutes: 2%; fifteen minutes: 2%
```

```
PID 1Min 5Min 15Min Process
```

```
131105 1% 1% 1% ce_switch
```

```
131106 1% 1% 1% eth_server
```

```
RP/0/RP0/CPU0:CRS1#
```

# Persistent Syslog

- Router can be configured to have persistent syslogs on the harddisk:
- Logs not written to harddisk in real-time, they are first staged on disk0: then flushed on harddisk after about every 4 hours.
- Useful when syslog server is inaccessible for any reason.

## logging archive

```
device harddisk      ← Archived on harddisk.  
file-size 500       ← 500 MB (single log file)  
archive-size 500    ← 500 MB (total size of syslogs)  
archive-length 4     ← 4 weeks (logs older than this are removed)
```

```
RP/0/RP0/CPU0:CRS1#dir harddisk:/var/log/2020/11/03  
Directory of harddisk:/var/log/2020/11/03  
10008690    -rwx  1023890      Tue Nov  3 15:48:26 2020  syslog.03.1  
15426650112 bytes total (14786479104 bytes free)  
RP/0/RP0/CPU0:CRS1#
```

# 'Describe' command

- Details of a command and associated process/files

```
RP/0/RP1/CPU0:CRS#describe show controllers pse summary
```

```
The command is defined in metro_driver.parser
```

```
Node 0/RP1/CPU0 has file metro_driver.parser for boot package /disk0/hfr-os-mbi-3.2.90/mbihfr-rp.vm from hfr-lc
```

```
Package:
```

```
hfr-lc
```

```
hfr-lc V3.2.90[3I] linecard package for ppc
```

```
Vendor : Cisco Systems
```

```
Desc : linecard package for ppc
```

```
Build : Built on Tue May 24 23:46:10 CEST 2005
```

```
Source : By edde-bld1 in /vws/afz/production/3.2.90.3I/hfr/workspace for c2.95.3-p8
```

```
Component:
```

```
hfr-metro-driver V0.0.0[main/204] Driver for Metro ASIC
```

```
File: metro_driver.parser
```

```
User needs ALL of the following taskids:
```

```
interface (READ)
```

```
drivers (READ)
```

```
It will take the following actions:
```

```
Spawn the process:
```

```
metro_cli -t 0x1
```

Package command is located in

Version of component code

Permissions required for execution of command

KSH command that CLI is calling – only visible to 'cisco-support' users from 3.2.0

# Trace functionality

- 'trace' functionality is like a form of 'always-on' debug

```
RP/0/RP1/CPU0:CRS1#show ospf trace
OSPF Trace Summary (2, RP/1/RP0/CPU0:CRS1, 0M)

  Trace Name      Size      Count  Description
  -----
1.  adj           65536     6291   adjacency
2.  adj_cycle     65536    893383 dbd/flood events/pkts
3.  config        2048       486   config events
4.  errors        8192    868816 errors
5.  events        4096       255   mda/rtrid/bfd/vrf
6.  ha            8192       485   startup/HA/NSF
7.  hello         2048    3982447 hello events/pkts
8.  idb          8192        973   interface
9.  pkt          2048    1927767 I/O packets
10. rib          65536    52190  rib batching
11. spf          65536    93138  spf/topology
12. spf_cycle    65536   352143 spf/topology detail
13. te          4096     3893  mpls-te
14. test        1024    20052  testing info
15. mq          65536         5  message queue info
RP/0/RP0/CPU0:CRS1#show ospf trace hello
Traces for OSPF 2 (Fri Jan 22 08:55:38)
Traces returned/requested/available: 2048/2048/2048
Trace buffer: hello

1   Jan 22 08:49:45.305* ospf_send_hello: area 0.0.0.80 intf MADJ: BE1008 from 0.0.0.0
2   Jan 22 08:49:45.546 ospf_rcv_hello: intf BE1009 area 0.0.0.74 from 10.1.0.9 10.1.9.2
3   Jan 22 08:49:45.546 ospf_check_hello_events: intf MADJ: BE1009 area 0.0.0.74 from 0.0.0.0
4   Jan 22 08:49:45.573* ospf_send_hello: area 0.0.0.74 intf MADJ: BE1008 from 0.0.0.0
5   Jan 22 08:49:45.845* ospf_rcv_hello: intf BE1009 area 0.0.0.80 from 10.1.0.9 10.1.9.2
6   Jan 22 08:49:45.845* ospf_check_hello_events: intf MADJ: BE1009 area 0.0.0.80 from 0.0.0.0
7   Jan 22 08:49:45.917* ospf_send_hello: area 0.0.0.80 intf Te0/5/0/7 from 10.1.80.1
8   Jan 22 08:49:46.232 ospf_rcv_hello: intf BE1008 area 0.0.0.74 from 10.1.0.8 10.1.8.2
```

# 'Show tech' command

```
RP/0/RP0/CPU0:CRS1# show tech-support snmp file harddisk:sh_tech_snmp
++ Show tech start time: 2021-Jan-22.090643.UTC ++
Fri Jan 22 09:06:44 UTC 2021 Waiting for gathering to complete
.....
Fri Jan 22 09:10:24 UTC 2021 Compressing show tech output
Show tech output available at 0/RP0/CPU0 : harddisk:/rpatki_demo_sh_tech.tgz
++ Show tech end time: 2011-Jan-22.091025.UTC ++

RP/0/RP0/CPU0:CRS1# dir harddisk: | in sh_tech_snmp
Fri Jan 22 09:10:46.951 UTC
58948      -rw-   709261      Fri Jan 22 09:10:25 2014  rpatki_sh_tech_snmp.tgz
RP/0/RP0/CPU0:CRS1#

RP/0/RP1/CPU0:CRS1#show tech-support cef ipv4 location ?
0/2/CPU0    Fully qualified location specification
0/3/CPU0    Fully qualified location specification
0/5/CPU0    Fully qualified location specification
0/RP0/CPU0  Fully qualified location specification
0/RP1/CPU0  Fully qualified location specification
WORD        Fully qualified location specification
```

# 'Monitor interface' command

- overall reporting (requires MGBL package)

```
CRS1-CRS                               Monitor Time: 00:00:37                SysUptime: 245:59:24

Protocol:General
Interface                               In (bps)           Out (bps)           InBytes/Delta      OutBytes/Delta
BE1003                                  0/ --%             0/ --%              964.1M/0           893.5M/0
BE1007                                  12000/ 0%          6000/ 0%            1.3G/1003          787.2M/963
BE1008                                  5000/ 0%           8000/ 0%            1.5G/1382          1.4G/1471
BE1009                                  9000/ 0%           11000/ 0%           957.4M/372         1.4G/854
BE1010                                  0/ --%             0/ --%               0/0                0/0
BE1067                                  4000/ 0%           5000/ 0%            484.4M/324         622.4M/248
Mg1/RP0/CPU0/0                          1000/ 0%           0/ 0%               173.8M/954         79086/0
Mg1/RP1/CPU0/0                          1000/ 0%           0/ 0%               173.9M/954         94122/0
Mg0/RP1/CPU0/0                          1000/ 0%           0/ 0%               173.5M/954         230160/0
Mg0/RP0/CPU0/0                          2000/ 0%           10000/ 0%           186.3M/1614        18.0M/6720
Te1/6/0/0                                0/ 0%              0/ 0%               0/0                0/0
Te1/6/0/1                                7000/ 0%           2000/ 0%            748.4M/479         276.0M/478
Te1/6/0/2                                0/ 0%              0/ 0%               0/0                0/0

--- snip

Gi1/15/3/5                              0/ 0%              0/ 0%               0/0                0/0
Gi1/15/3/6                              0/ 0%              0/ 0%               0/0                0/0

Quit='g',      Clear='c',      Freeze='f', Thaw='t',
Next set='n',  Prev set='p',  Bytes='y',   Packets='k'
(General='g',  IPv4 Uni='4u', IPv4 Multi='4m', IPv6 Uni='6u', IPv6 Multi='6m')
```

# 'Monitor interface'

- per interface reporting  
(requires MGBL package)

```
RP/0/RP1/CPU0:CRS1# monitor interface Bundle-ether 1008
CRS1                Monitor Time: 00:00:18          SysUptime: 246:02:20

Bundle-Ether1008 is up, line protocol is up
Encapsulation ARPA

Traffic Stats:(2 second rates)
Input  Packets:                6489005                Delta
Input  pps:                    8
Input  Bytes:                  1507217455            1274
Input  Kbps (rate):            5                    ( 0%)
Output Packets:                7079943                15
Output pps:                    9
Output Bytes:                  1490126647           2024
Output Kbps (rate):            8                    ( 0%)

Errors Stats:
Input  Total:                  0                    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'
```

# Structured command outputs

- Customers may want a structured output of XR commands.
- XML API exist on XR for a long time.
- [https://www.cisco.com/c/en/us/td/docs/ios\\_xr\\_sw/iosxr\\_r4-1/xml/programming/guide/xl41apidoc/xl41over.html](https://www.cisco.com/c/en/us/td/docs/ios_xr_sw/iosxr_r4-1/xml/programming/guide/xl41apidoc/xl41over.html)
- Command output can be seen in XML format with the “show operational” command as shown.
- Latest images & platforms have REST API and model driven API support.
- More details on <https://developer.cisco.com/site/ydk/>  
<https://developer.cisco.com/video/net-prog-basics/>

```
RP/0/0/CPU0:XR1#sh operational Interfaces InterfaceBriefTable
InterfaceBrief/InterfaceName=Loopback0 xml
<?xml version="1.0"?>
<Response MajorVersion="1" MinorVersion="0">
  <Get>
    <Operational>
      <Interfaces MajorVersion="1" MinorVersion="6">
        <InterfaceBriefTable>
          <InterfaceBrief>
            <Naming>
              <InterfaceName>
                Loopback0
              </InterfaceName>
            </Naming>
            <Interface>
              Loopback0
            </Interface>
            <ParentInterface>
              None
            </ParentInterface>
            <Type>
              IFT_LOOPBACK
            </Type>
            <State>
              IM_STATE_UP
            </State>
            <ActualState>
              IM_STATE_UP
            ...
```

# Lab Access

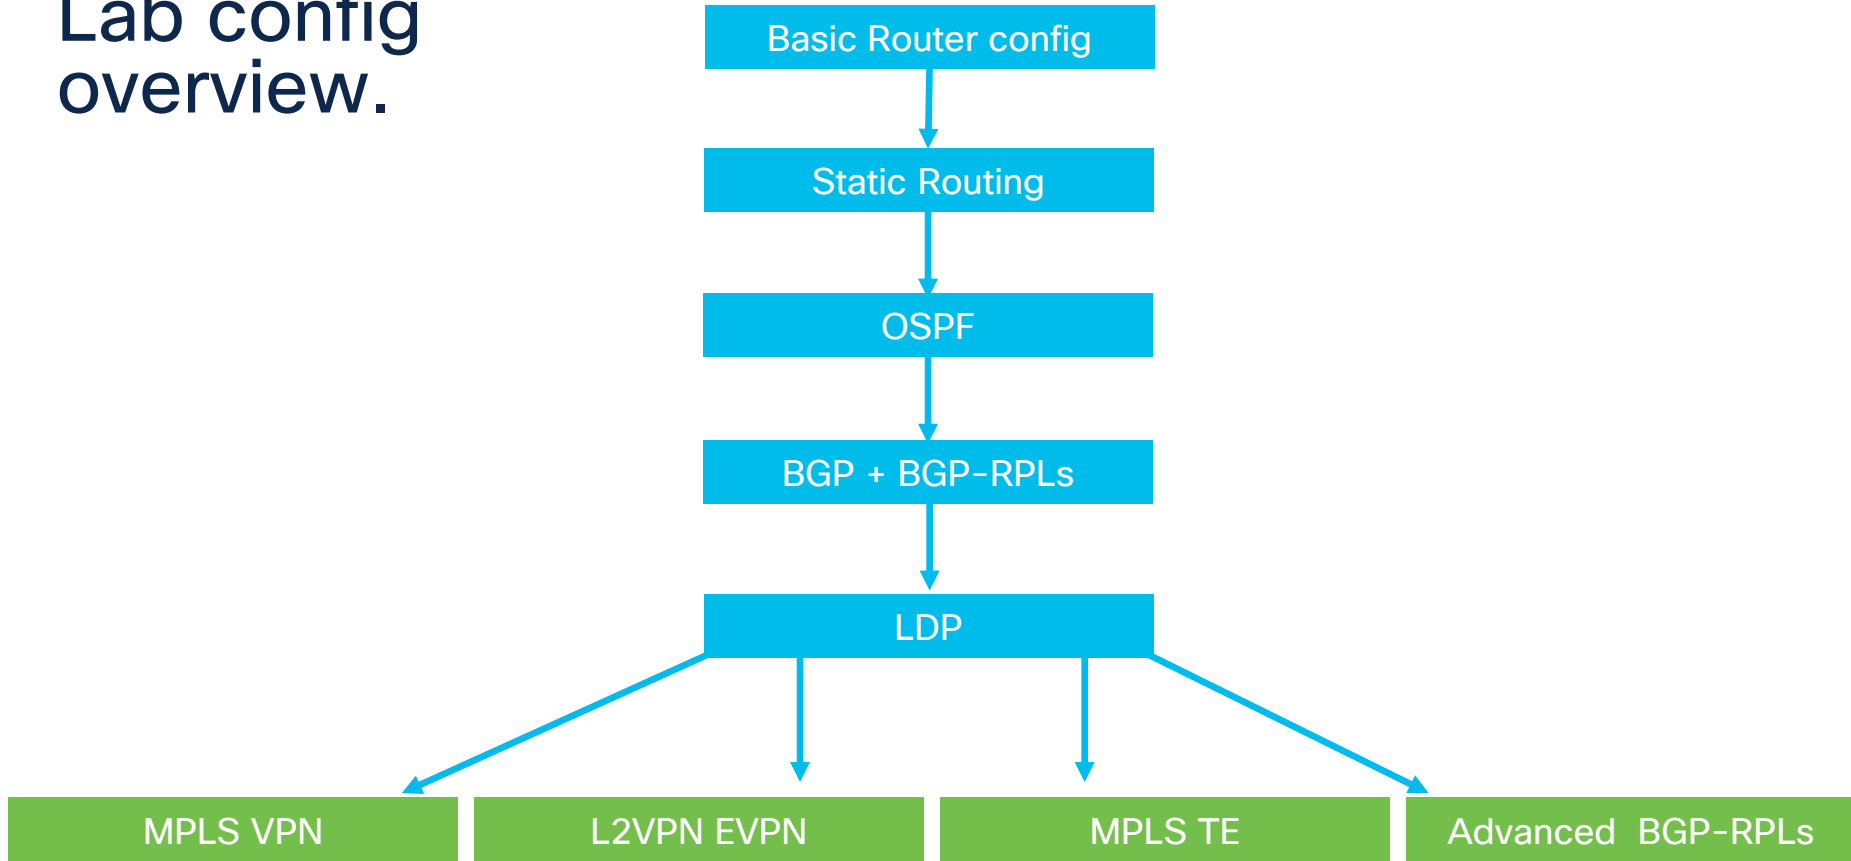


# Lab Orientation

- Every student works by themselves. Don't have to wait on others to proceed!
- Lab activities involve only changes on the IOS XR router configurations.
- Please complete the Lab Guide's lab setup and lab access sections.



# Lab config overview.



# Technical Session Surveys

- Attendees who fill out a minimum of four session surveys and the overall event survey will get Cisco Live branded socks!
- Attendees will also earn 100 points in the Cisco Live Game for every survey completed.
- These points help you get on the leaderboard and increase your chances of winning daily and grand prizes.



# Complete your Session Survey

- Please complete your session survey after each session. Your feedback is important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (open from Thursday) to receive your Cisco Live t-shirt.
- All surveys can be taken in the Cisco Events Mobile App or by logging in to the Session Catalog and clicking the "Attendee Dashboard" at <https://www.ciscolive.com/emea/learn/sessions/session-catalog.html>



# Continue Your Education



Visit the Cisco Showcase for related demos.



Book your one-on-one Meet the Engineer meeting.



Attend any of the related sessions at the DevNet, Capture the Flag, and Walk-in Labs zones.



Visit the On-Demand Library for more sessions at [ciscolive.com/on-demand](https://ciscolive.com/on-demand).



The bridge to possible

Thank you

CISCO *Live!*

CISCO *Live!*

ALL IN