Cisco Support Community Expert Series Webcast:

IOS-XR Fundamentals and Architecture

Cisco Support Community – Expert Series Webcast

- Today's featured experts are Cisco Support Engineer Sundeep Valengattil and Sudhir Kumar
- Ask them questions now about **IOS-XR Fundamentals and Architecture**, in the Chat window

Sudeep Valengattil
CSE-HTTS
CCIE 36098

Sudhir Kumar
CSE- HTTS
CCIE 35219

Topic: IOS-XR Fundamentals and Architecture

Event Date: 18-Nov-2014

Panel of Experts

Raj Pathak

CSE-HTTS

CCIE 38760

Thank You For Joining Us Today!

Today's presentation will include audience polling questions

We encourage you to participate!

Thank You For Joining Us Today!

If you would like a copy of the presentation slides, click the PDF file link in the chat box on the right or go to:

https://supportforums.cisco.com/document/12352951/expert-webcast-ios-xr-fundamentals-and-architecture-presentation

Or, https://supportforums.cisco.com/expert-corner/knowledge-sharing

Ask the Expert Event continuing with today's Experts

IOS-XR Fundamentals and Architecture

This is an opportunity to continue asking questions about IOS-XR Fundamentals and Architecture
Service Providers /XR OS and Platforms
Community now through **November 28th, 2014**.

https://supportforums.cisco.com/discussion/12354711/ask-experts-ios-xr-fundamentals-and-architecture

Ask the Expert Events – Current /Upcoming English

Introduction to Cisco Unified Computing System (Cisco UCS) Central

Join Cisco Expert: **Payal Bhaduri**

This is an opportunity to learn and ask questions about basics of Cisco UCS Central tool and its application..

November 10th - 21st, 2014

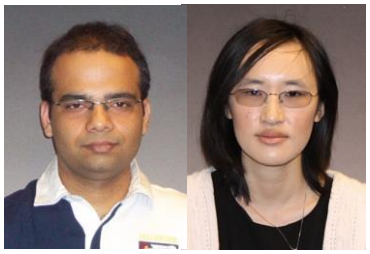

Technical Discussion on UCS-Mini platform

Join Cisco Experts: **Vishal Mehta and Rosalind Lee**

Cisco UCS, originally designed for the data center, is now optimized for branch and remote offices, point-of-sale, and smaller IT environments with Cisco UCS Mini. UCS Mini is for customers who need fewer servers (expandable to 15 servers) but still want the robust management capabilities provided by UCS Manager. s

November 24th – December 5th, 2014

Join the discussion for these Ask The Expert Events:
https://supportforums.cisco.com/community/netpro/expert-corner#view=ask-the-experts

Polling Question 1

What is your level of experience with IOS-XR OS and Platforms

- I have basic working experience on IOS-XR platforms
- I theoretically know IOS-XR, but no practical experience
- I'm playing with it in the lab
- I'm running it in production

Submit your Questions now!

Use the Q & A panel to submit your questions and the panel of experts will respond.

IOS-XR Fundamentals and Architecture

November 2014

Agenda

- High-Level Overview of Cisco IOS XR
- Cisco IOS XR Infrastructure
- Configuration Management
- Cisco IOS XR Monitoring and Operations
- Cisco IOS XR Security
- Introduction to different IOS-XR platforms

Requirements for Carrier Grade NOS

Convergence

- Should have the capability to enable infrastructure and service convergence

Scalability

- Should be able to support the addition of different system components without service disruption

Availability

- HW Redundancy
- Should have the necessary SW capability to enable the system to operate with no or minimal service disruption when such a module fails, and with it is subsequently removed, upgraded, or replaced
- Failure Recovery and Microkernel-Based NOS
- Should be able to contain and recover from most SW failures without service disruption
- Process Restart ability
- Should be able to restart the process with a SW process fails
- Failure Detection
- Should support network features that enable quick failure detection and rerouting of traffic around failed links, modules, or routers
- SW Upgrades and Patching
- Should support SW upgrade and/or patching with no or minimal disruption to service

Security

- Should minimize the impact to date, control, and management plane functions due to such attacks

Service Flexibility

- Should support the addition of new software features, line cards, and/or service modules with no or minimal service disruption

High-Level Overview of Cisco IOS XR

IOS XR distribution models

Localization: processing and storage closer to the resource

- Database specific to the node is located on that node
- Processes are located on a node where they have greater interaction with the resource
- For example, ARP, IM, BFD, and FIB manager are located on the linecard

Load distribution: additional RPs or DRPs are added to the system and processes are distributed across the different RPs or DRP modules

- Routing Protocols incl. BGP, ISIS, PIM, RSVP, LDP, and RIB
- Management Entities incl. SNMP server, SSH, Telnet, and HTTP
- Other processes

IOS XR Architecture: Separation of management, control, and data plane

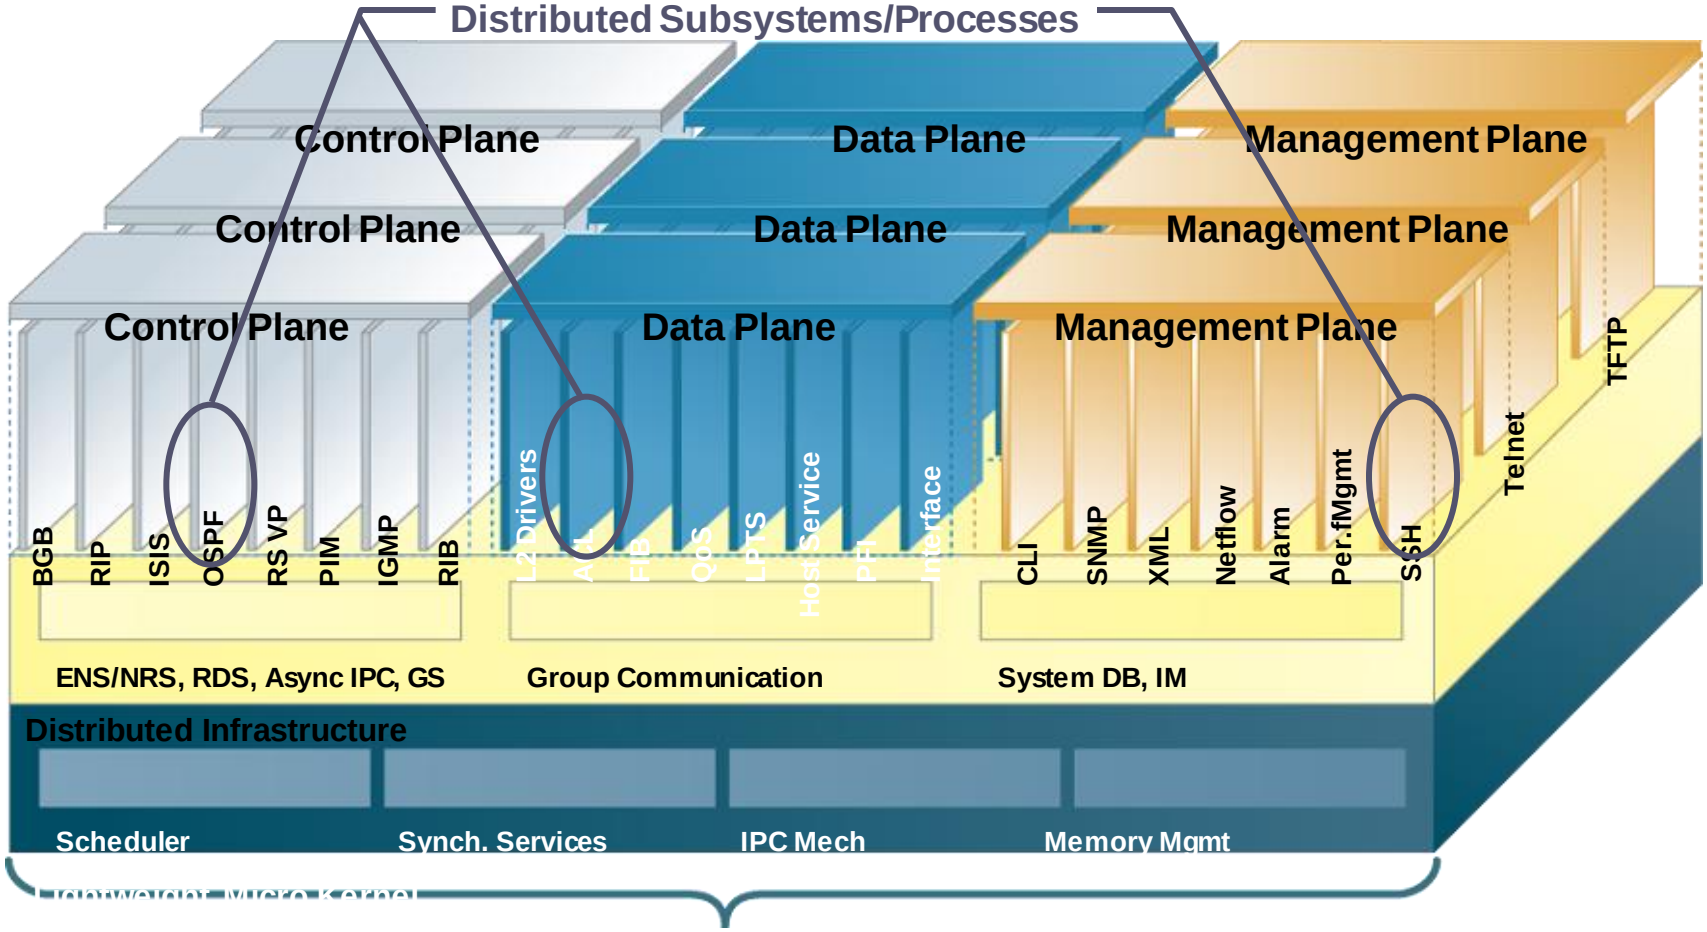

Kernel System Services

Each routing control plane or management plane runs on one or multiple (D)RPs

Data plane processes are located on each node that participates in packet forwarding

Cisco IOS XR Infrastructure

IOS XR Kernel

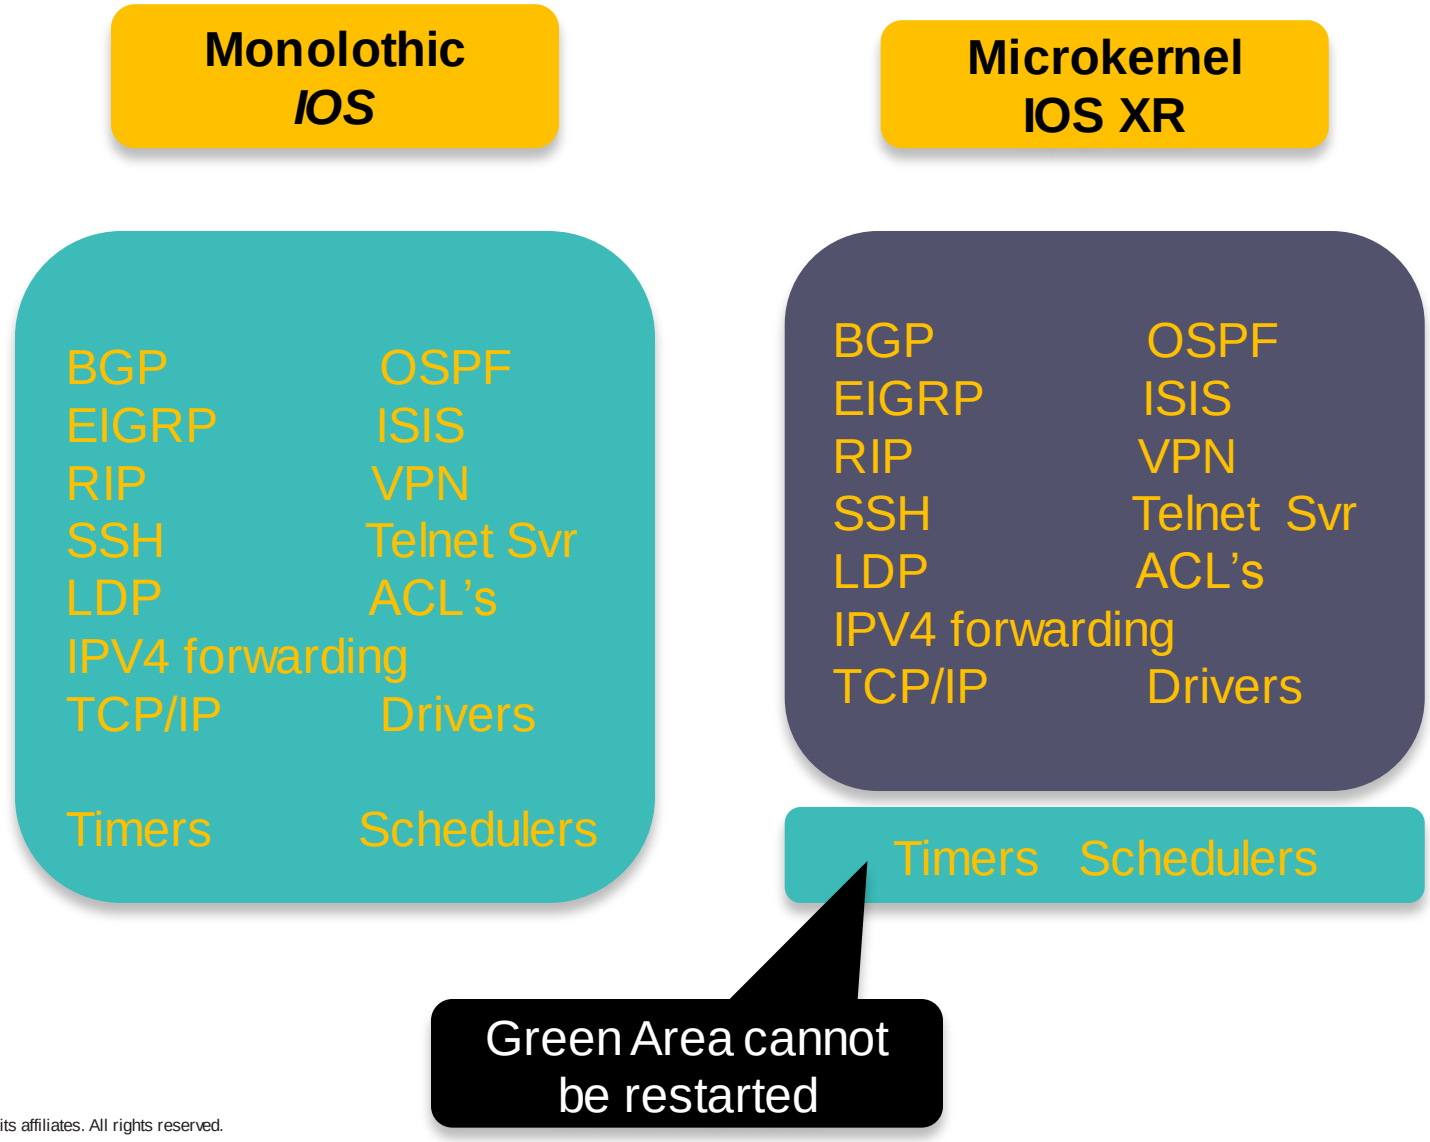

Microkernel-Based OS

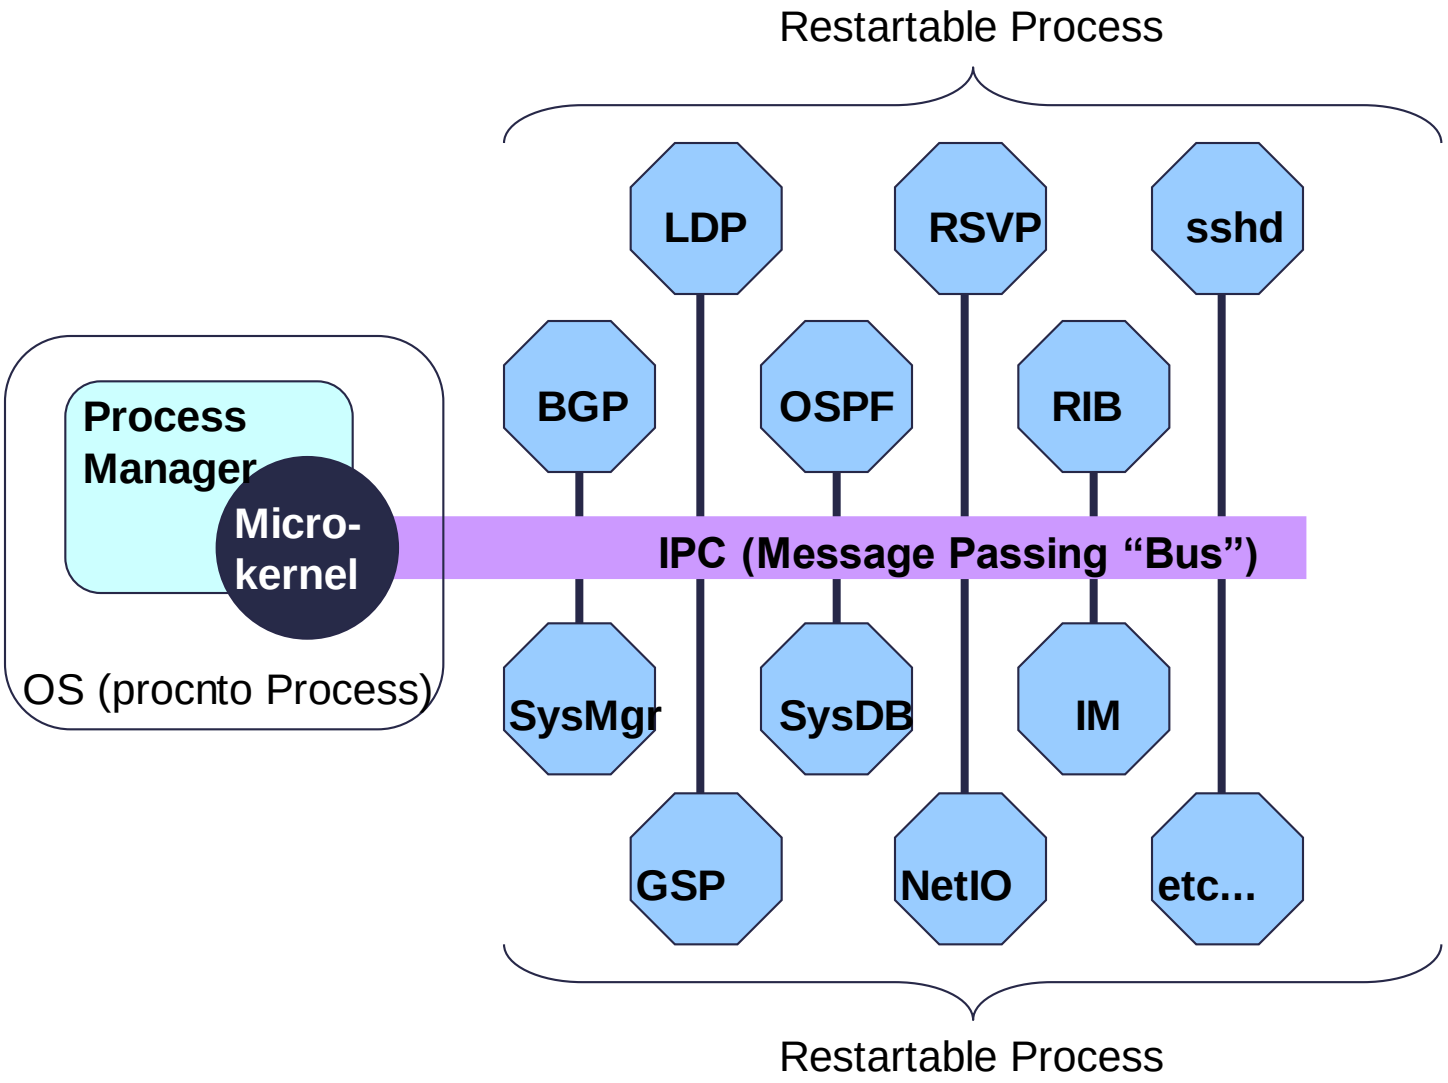

- OS is a group of cooperating processes managed by a small microkernel

Process Management Definitions

Threads

Threads are units of execution, each with an execution context that includes a stack and registers.

A thread is in effect a ‘sub-process’ managed by the parent, responsible for executing a sub portion of the overall process. For example, OSPF has a thread which handles ‘hello’ receipt and transmission.

A thread may only run when the parent process is allocate runtime by the system scheduler. A process with threads is said to be ‘multi-threaded’

Process

A process is group of threads that share virtual address (memory) space

Instance

A process can have more one running instantiation (e.g. sysmgr or isis).

Process Management Definitions

Job ID#

- Each process is allocated a Job ID# or JID when it is first run. The Job ID remains associated with the process even if the process is stop and restarted

Process ID#

- Also known as the PID. Each process is also allocated a ‘PID’ when it is first run. The PID will change if the process is stop and restarted

Thread ID#

- Also known as the TID. If a process is contains threads, each is assigned a TID# associated with the PID/JID

DLL

- Processes may call and access Dynamic Link Libraries which is a method for sharing portions of code. Processes will dynamically load the required DLL at runtime

Threads - example output

```
RP/0/RP0/CPU0:H4#show processes threadname 122
Thu Jul  9 10:37:20.403 JST
JID    TID  ThreadName       pri state    TimeInState        NAME
122    1    io-control       10 Receive        0:00:23:0344 bgp
122    2    chkpt_evm        10 Receive       39:58:21:0119 bgp
122    3    label-thread     10 Receive        0:01:08:0779 bgp
122    4    async            10 Receive       17:29:17:0675 bgp
122    5    io-read          10 Receive        0:00:56:0972 bgp
122    6    io-write         10 Receive        0:00:19:0441 bgp
122    7    router           10 Receive        0:01:08:0799 bgp
122    8    import           10 Receive        0:00:08:0778 bgp
122    9    update-gen       10 Receive        0:00:27:0456 bgp
122    10   crit-event       10 Receive        0:01:08:0779 bgp
122    11   event            10 Receive        0:00:20:0946 bgp
122    12   management       10 Receive        0:01:18:0754 bgp
122    13   rib-update ID 0  10 Receive        0:01:08:0757 bgp
122    14   rib-update ID 1  10 Receive        0:00:08:0756 bgp
```

General Thread Scheduling

- The microkernel provides the following three scheduling algorithms to meet needs for different

- FIFO scheduling
- Round-robin scheduling
- Sporadic scheduling

General Thread Scheduling

- IOS XR microkernel uses a preemptive, priority based, and non-adaptive scheduling algorithm
- Each thread is assigned a priority
- The scheduler is responsible for selecting the next thread to run based on the priority assigned
- The highest priority thread in ready state is selected to run
  - There is a ready state FIFO queue for each priority level

Synchronization Services

- The microkernel provides a message-passing-based synchronous IPC mechanism
  - This message-passing service copies a message directly from the address space of the sender thread to the receiver thread without intermediate buffering
- Other IPC mechanisms use shared memory space is possible to develop
  - Access to the shared memory space must be synchronized to ensure data consistency
  - The microkernel provides mutex, condver, and semaphore synchronization tools

Common Thread States and Transitions

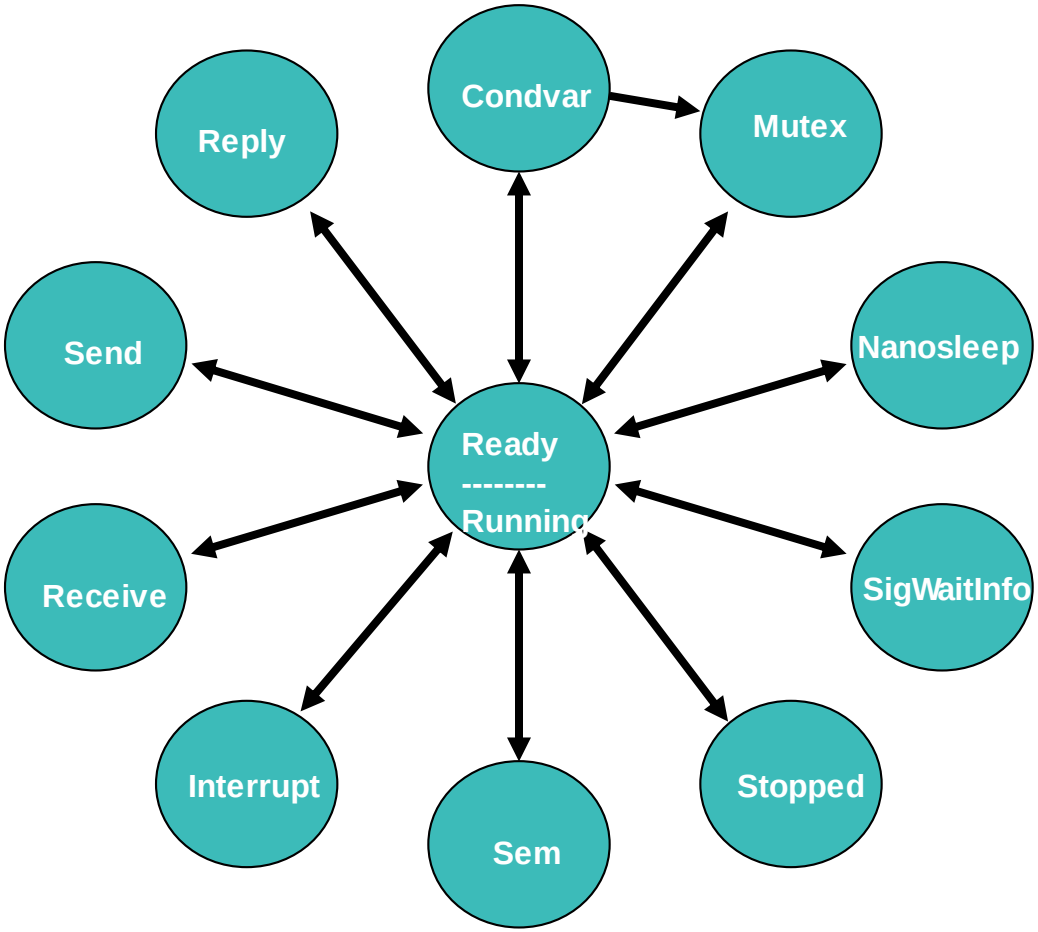

- A thread state can transition from ready to running and vice versa
- A thread in running state may also transition to any f the other states

Thread States

State	Explanation
dead	The kernel is waiting to release the thread's resources
running	Actively running on a CPU
ready	Not running on a CPU but is ready to run
stopped	Suspended (SIGSTOP signal)
send	Waiting for a client to send a message
receive	Waiting for a server to receive a message
reply	Waiting for a server to reply to a message
stack	Waiting for more stack to be allocated
waitpage	Waiting for the process manager to resolve a page fault
sigsuspend	Waiting for a signal
sigwaitinfo	Waiting for a signal
nanosleep	Sleeping for a period of time
mutex	Waiting to acquire a mutex
condvar	Waiting for a conditional variable to be signaled
join	Waiting for the completion of another thread
intr	Waiting for an interrupt
sem	Waiting to acquire a semaphore

System Manager

Functions;

- Start processes during bootup or node reload
- Start processes during route processor (RP) failover
- Start processes in response to user configuration
- Act as a central repository for all process-related information
- Initiate disaster recovery based on the process health
- Invoke dumper to collect a core dump when a process terminates abnormally

System manager runs on each route processor and line card in the system

Two instances of sysmgr process are running on each node

- Primary sysmgr ... responsible for all system manager responsibilities
- Secondary sysmgr ... standby and is ready to assume the primary role if the current primary exits for some reason

Process Attributes

```
RP/0/RP0/CPU0:F34m#more disk0:/hfr-rout-3.6.1/startup/ospf.startup
Wed Jul 15 13:34:55.088 JST
name:ospf
path:/ios/bin
item:/cfg/gl/ipv4-ospf/proc/
copies:10
tuple_dynamic_tag:ON
placement: ON
check_avail: ON
failover_tier: ospf
standby_capable: ON

RP/0/RP0/CPU0:F34m#more disk0:/hfr-fwdg-3.6.1/startup/fib_mgr.startup
Wed Jul 15 13:35:57.810 JST
name:fib_mgr
path:/ios/bin
check_avail: ON
level:99
standby_capable:ON

RP/0/RP0/CPU0:F34m#more disk0:/hfr-base-3.6.1/startup/gsp-rp.startup
Wed Jul 15 13:37:00.467 JST
name:gsp
path:/ios/bin
level: 80
check_avail:on
mandatory: on
```

JID and PID

- System manager assigns a unique job id (JID) to each executable
- The JID is persistent across restarts
- In addition to a JID each process is assigned a unique process ID (PID) when it is started or restarted
- If a process is restarted it is assigned a new PID number, but it retains its original JID number

System Manager and Process Lifecycle

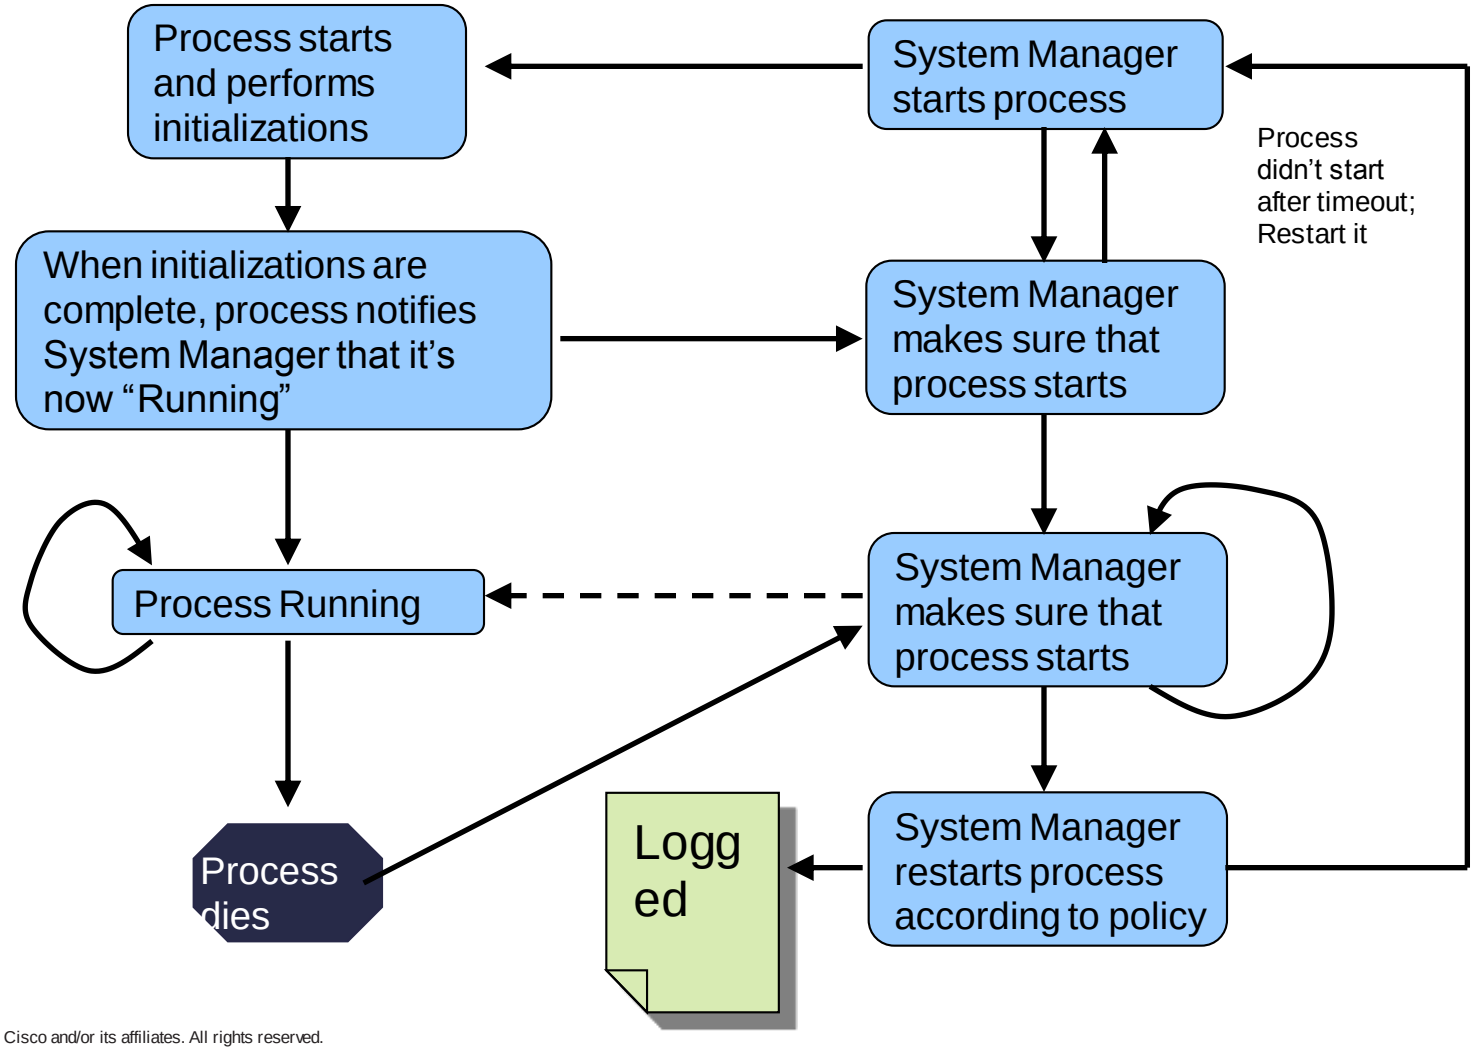

Process States

You should really only see Run and Killed, other states are transitional and process should not be stuck here unless we're having problems!

- Run – Normal state.
- Exited – Normal exit (not run), cleaning up.
- Hold – Process was restarted but respawn was inhibited.
- Wait – Wait for item to be set (sysdb notification).
- Restart - Waiting for respawn timer to expire.
- Initializing – Waiting for EOI timer after intialization. This only happens when a conditional process is spawned … not for restarts, failovers, etc.
- Killed – was running and then forcibly stopped.
- Queued – queueing until node is "active".
- Tuple-set – Infrastructure ready (sysdb tuple found). Process can start. An intermediary state after boot.
- Unknown, Error, None – Software problem catch-all states.

show process

```
RP/0/RP0/CPU0:F34m#show processes ospf
Wed Jul 15 15:55:42.033 JST
                  Job Id: 302
                     PID: 487663
         Executable path: /disk0/hfr-rout-3.6.1/bin/ospf
              Instance #: 1
              Version ID: 00.00.0000
                 Respawn: ON
           Respawn count: 1
  Max. spawns per minute: 12
            Last started: Tue Jul  7 10:44:00 2009
           Process state: Run
           Package state: Normal
       Started on config: cfg/gl/ipv4-ospf/proc/v4_TEST-Policy/ord_f/default/ord_a/routerid
                    core: TEXT SHAREDMEM MAINMEM
                Max. core: 0
                Placement: ON
             startup_path: /pkg/startup/ospf.startup
                   Ready: 6.805s
               Available: 2.440s
        Process cpu time: 10.115 user, 0.433 kernel, 10.548 total
JID   TID CPU Stack pri state        TimeInState    HR:MM:SS:MSEC   NAME
302     1   1   68K  10 Receive        0:00:00:0801    0:00:09:0509 ospf
302     2   1   68K  10 Receive        0:01:17:0640    0:00:00:0410 ospf
302     3   0   68K  10 Receive        0:49:31:0405    0:00:00:0021 ospf
302     4   1   68K  10 Receive        0:00:00:0100    0:00:00:0101 ospf
302     5   1   68K  10 Receive        0:00:17:0639    0:00:00:0039 ospf
302     6   0   68K  10 Receive        0:00:17:0639    0:00:00:0029 ospf
302     7   0   68K  10 Condvar        0:49:28:0009    0:00:00:0003 ospf
302     8   1   68K  10 Receive      197:11:35:0329    0:00:00:0000 ospf
302     9   1   68K  10 Receive      197:11:05:0222    0:00:00:0001 ospf
RP/0/RP0/CPU0:F34m#
```

Inter Process Communication (IPC)

Synchronous vs. Asynchronous

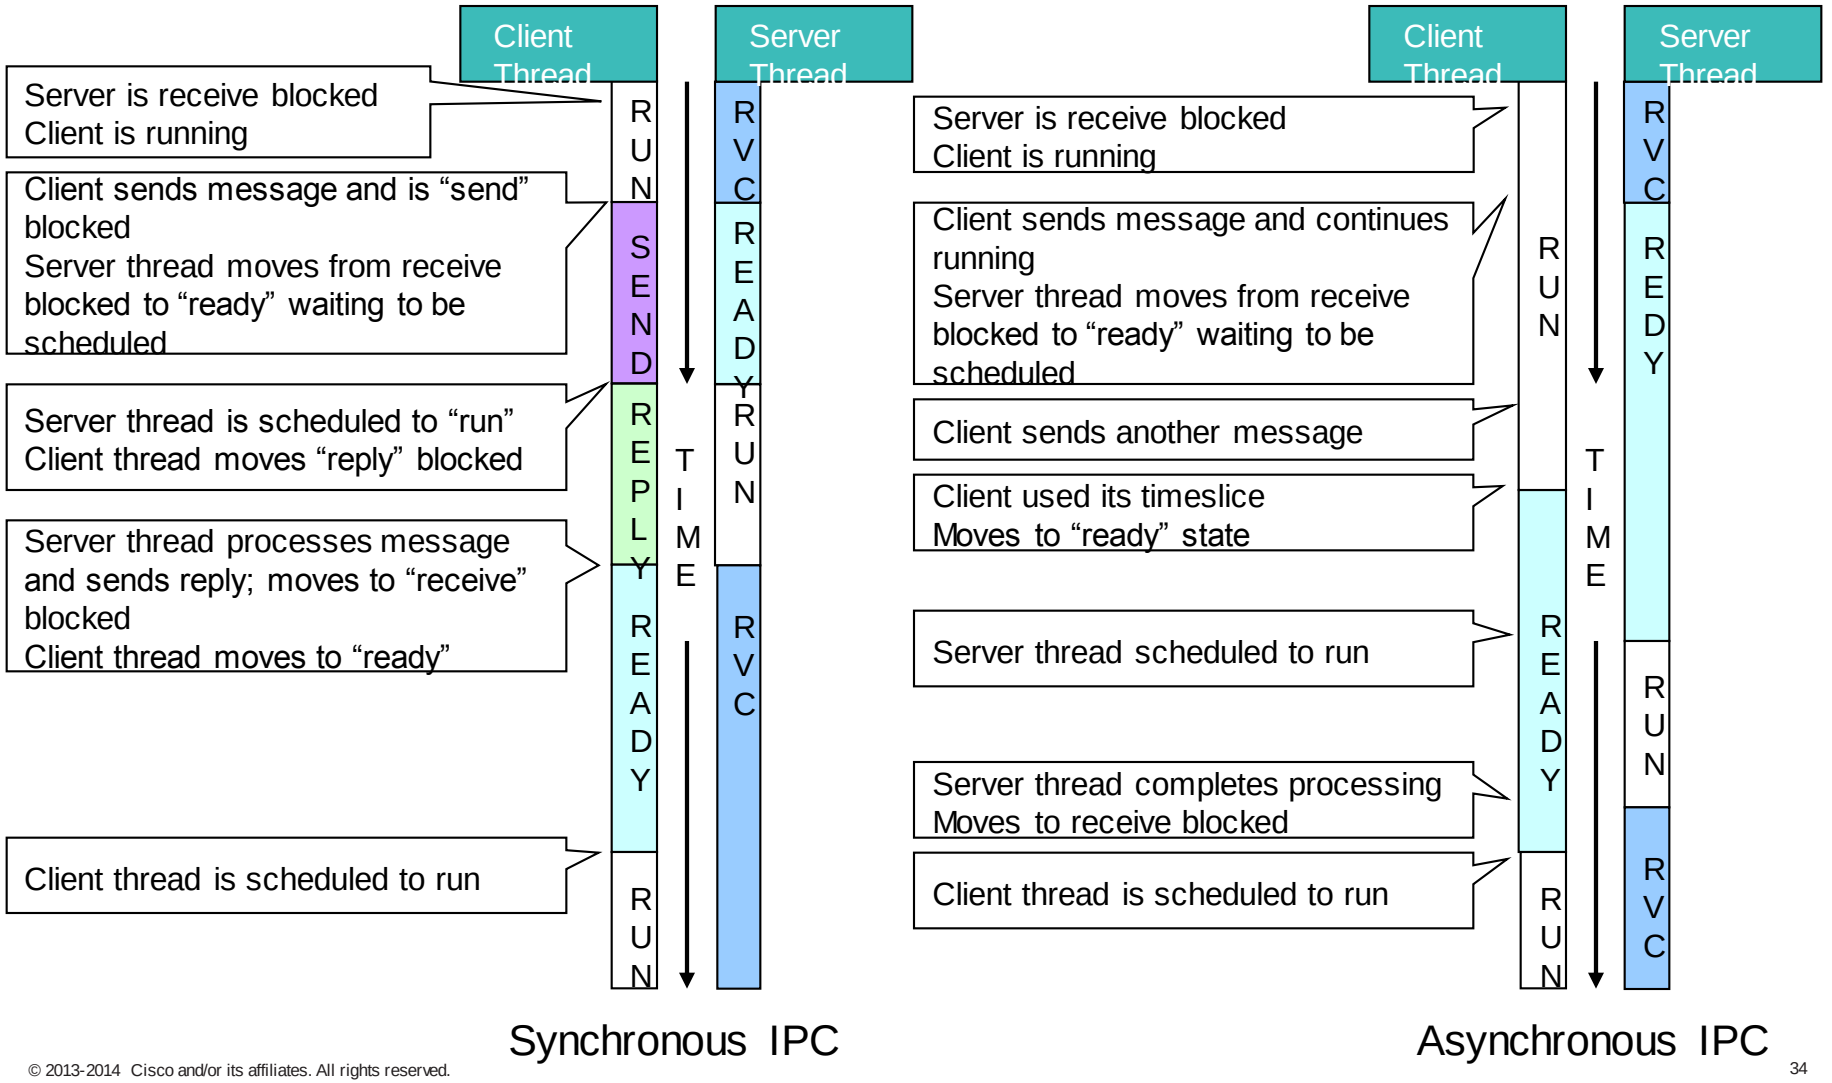

Intra-node IPC vs. Inter-node IPC

- Some IPC mechanisms can be used only in a single CPU or SMP environment, which is referred to here as a node
- A node is not necessarily equivalent to a linecard or RP/DRP
  - Some linecard and RP/DRPs have multiple CPUs or SMPs ... therefore, they have multiple nodes.
  - A CRS-1 linecard (MSC) has two nodes, CPU0 and SP, corresponding to the main CPU and service processor (SP), respectively
- An intra-node IPC mechanism can be used only within a single node, whereas an inter-node IPC mechanism can be used between different nodes
  - Generally the inter-node IPC mechanism supports both inter-node and intra-node communication

IPC
Connection-oriented vs. Rendezvous

- With connection-oriented scheme the receiver (typically a server) creates and advertises a channel, and the sender (client) connects to the channel before any message is sent by the sender
  - Both the sender and receiver are aware of the identity of each other
  - If either the sender or receiver restarts, the connection must be reestablished
- With the rendezvous scheme, the data, not the receiver or its channel, is the primary focus of the IPC
  - The communicating parties find each other through the name of the date structure
  - The rendezvous scheme is suitable for a producer-consumer model of communication
- In a producer-consumer model, a producer process generates data while one or more consumers process the data
  - The consumers are not interested in the identity of the producer; they are interested in the identity of the data produced

IPC
Point-to-point vs. Point-to-multipoint

- Point-to-point communication is used between two threads
- It is necessary to send the same message to several threads within a single node or to multiple nodes in some cases
  - However, the use of point-to-point communication for this scenario is inefficient
  - Moreover, the sender has to know the identity of all the receivers
- The use of point-to-multipoint communication solves both problems

Light Weight Messaging

- LWM is a point-to-point, synchronous, and connection-oriented communication built on top of QNX synchronous IPC
  - The microkernel is involved in copying the message from the sender thread memory space to the receiver's memory space
- Because the client is blocked after sending each message, there are at least two context switching instances per message
- LWM can be used both for intra-node and inter-node communication transparently
  - The sender does not need to know whether the receiver is on a local or remote node

Qnet

- Qnet is a QNX Neutrino protocol for communication between processes residing on different nodes

- Qnet enables IPC to work across nodes
  - LWM uses Qnet transparently to enable inter-node communication

- The use of a symbolic link (symlink) enables location transparency to the Qnet protocol
  - When a process needs to communicate with another process, it uses the symlink associated with the service and does not need to know where the service is located
  - A server process register with Qnet symlink manager (QSM) and publishes its service using symlink

Group Service Protocol

- GSP is a point-to-multipoint, connectionless, and asynchronous communication mechanism

- GSP communication identifies members by a logical group ID (GID) that can be mapped to a multicast address
  - The multicast address could be an IP multicast address, multicast MAC address, or fabric group ID (FGID)

- With GSP the sender does not need to know the identity or location of receivers
  - GSR provides the sender with a mechanism to send to all members of a GID, a subset of members, or to a single member

Bulk Content Downloader (BCDL)

- BCDL is used to download large data tables using GSP
- It is primarily used to download routing tables and MPLS label bindings to linecards and FGID tables to fabric cards
- Uses the producer/consumer model
  - RIB and LSD on the RP are producers
  - FIB manager processes on the linecards and RP are consumers of the data generated by the producer processes

IOS XR System Database

- SysDB provides a common mechanism for applications to store, modify, and/or access system information
- SysDB provides fully distributed in-memory data storage organized as a hierarchical namespace
  - This is accomplished by partitioning the SysDB namespace into three distinct planes: admin, shared, and local
  - Each SysDB server process is associated with exactly one of the shared-plane, the local-plane, or the admin-plane and can access items only in its own plane
- SysDB is used to store both configuration (cfg) and operational status (oper) items

Polling Question 2

Do you find configuration on XR platforms to be a complicated task?

a. Yes, I do

b. No, I do not

c. I'm not certain

Configuration Management

Distributed Configuration Management

- IOS XR has different processes, running on every modular service card (MSC) or node, which are responsible for the configuration of that particular node
- These processes are responsible for applying configuration and managing operations for their node
- RP is responsible only for the summary state information of all interfaces in the system
- The MSC-specific configuration is separated from the control plane-specific configuration

Configuration Planes

Admin plane

- Global configuration that is applicable to the entire physical system (SDR carving and Admin user access info)
- Accessible from the owner SDR, which has admin privileges

Shared plane

- Global configuration that is applicable to an entire SDR (routing protocol, QoS, ACL, multicast, MPLS, and etc.)
- Accessible from throughout the SDR to which it applies

Local plane

- Configuration that is applicable to an individual node within an SDR
- Accessible only from the SDR in which the node resides
- If a node is brought down or removed from the SDR, the local plane configuration relevant to that node is moved to an preconfiguration area from which it will later be restored if the node is added or brought back up again

Configuration Manager on RP and LC

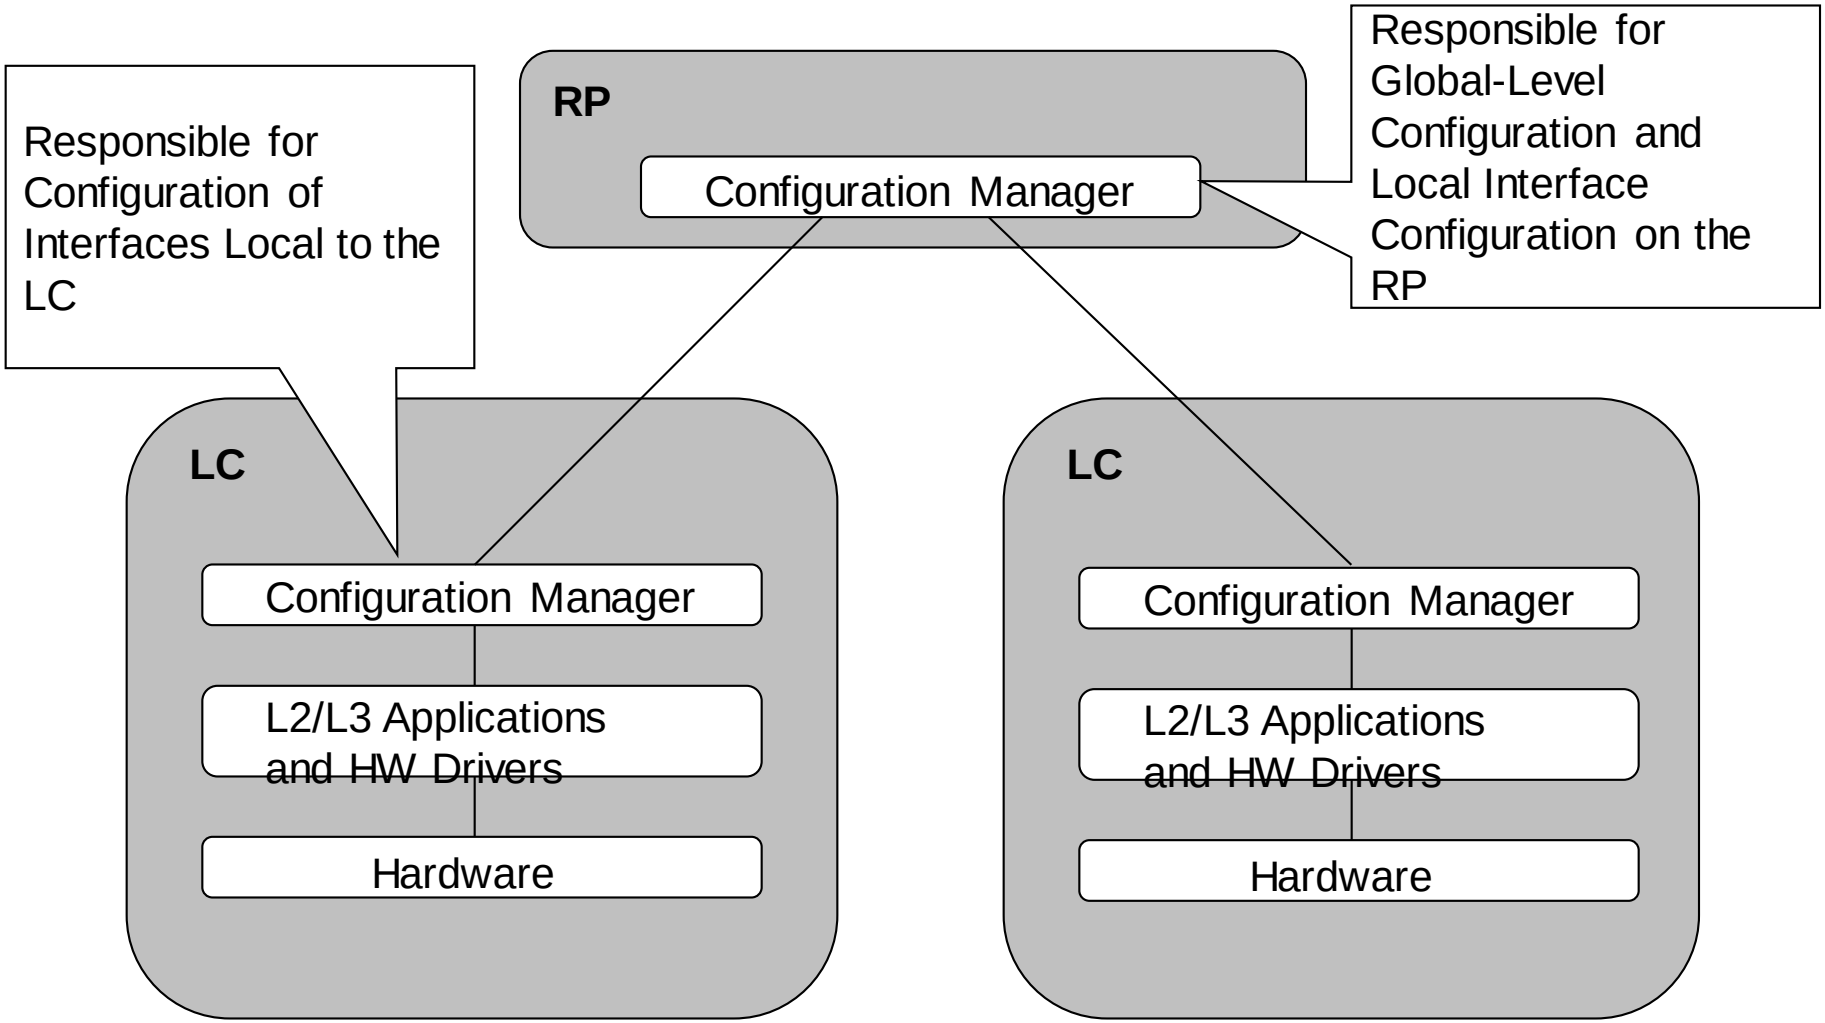

Configuration File System

- CFS is a set of files and directories used to store the primary persistent configuration in binary format
- CFS is managed by Cfgmgr and is replicated to other management nodes (RP and DRP) within the SDR by RDSFS
- It also has the persistent ASCII backup config, configuration history, commits/checkpoints, startup failed configuration files, versioning info, commitdb metadata files, and so on
- The configuration files are arranged in a series of directories under a single CFS root directory on a management node ({media}:/config/)
- CFS is constantly synchronized between Active and Standby RPs to preserve the router's configuration state in the event of an RP failover
- CFS is also created and synched in all Management nodes except the MSCs/nodes (because no disk exits there)

CFS on disk0

```
RP/0/RP0/CPU0:F34m#dir disk0:/config
Mon Jul 27 11:29:54.335 JST

Directory of disk0:/config

9850        drwx  16384       Thu Jul 23 14:07:50 2009  running
9851        drwx  16384       Thu Nov 16 17:55:51 2006  .startup
9852        drwx  16384       Thu Nov 16 17:56:01 2006  failed
9853        drwx  16384       Thu Nov 16 17:55:51 2006  pkgoir
1843        drwx  16384       Thu Jul 23 14:05:46 2009  admin
9855        drwx  16384       Thu Nov 16 17:55:52 2006  econf
9856        drwx  16384       Thu Nov 16 17:55:54 2006  user
9857        drwx  16384       Thu Nov 16 17:55:52 2006  archive
9861        drwx  16384       Thu Nov 16 17:55:53 2006  history
9873        drwx  16384       Thu Nov 16 17:55:56 2006  removed_cfg
3           drwx  16384       Thu Jul 23 14:05:39 2009  ad
1787        drwx  16384       Thu Jul 23 14:07:48 2009  lr
13302       drwx  16384       Thu Oct 23 10:24:36 2008  snmp

1004994560 bytes total (574701568 bytes free)
```

CFS directory

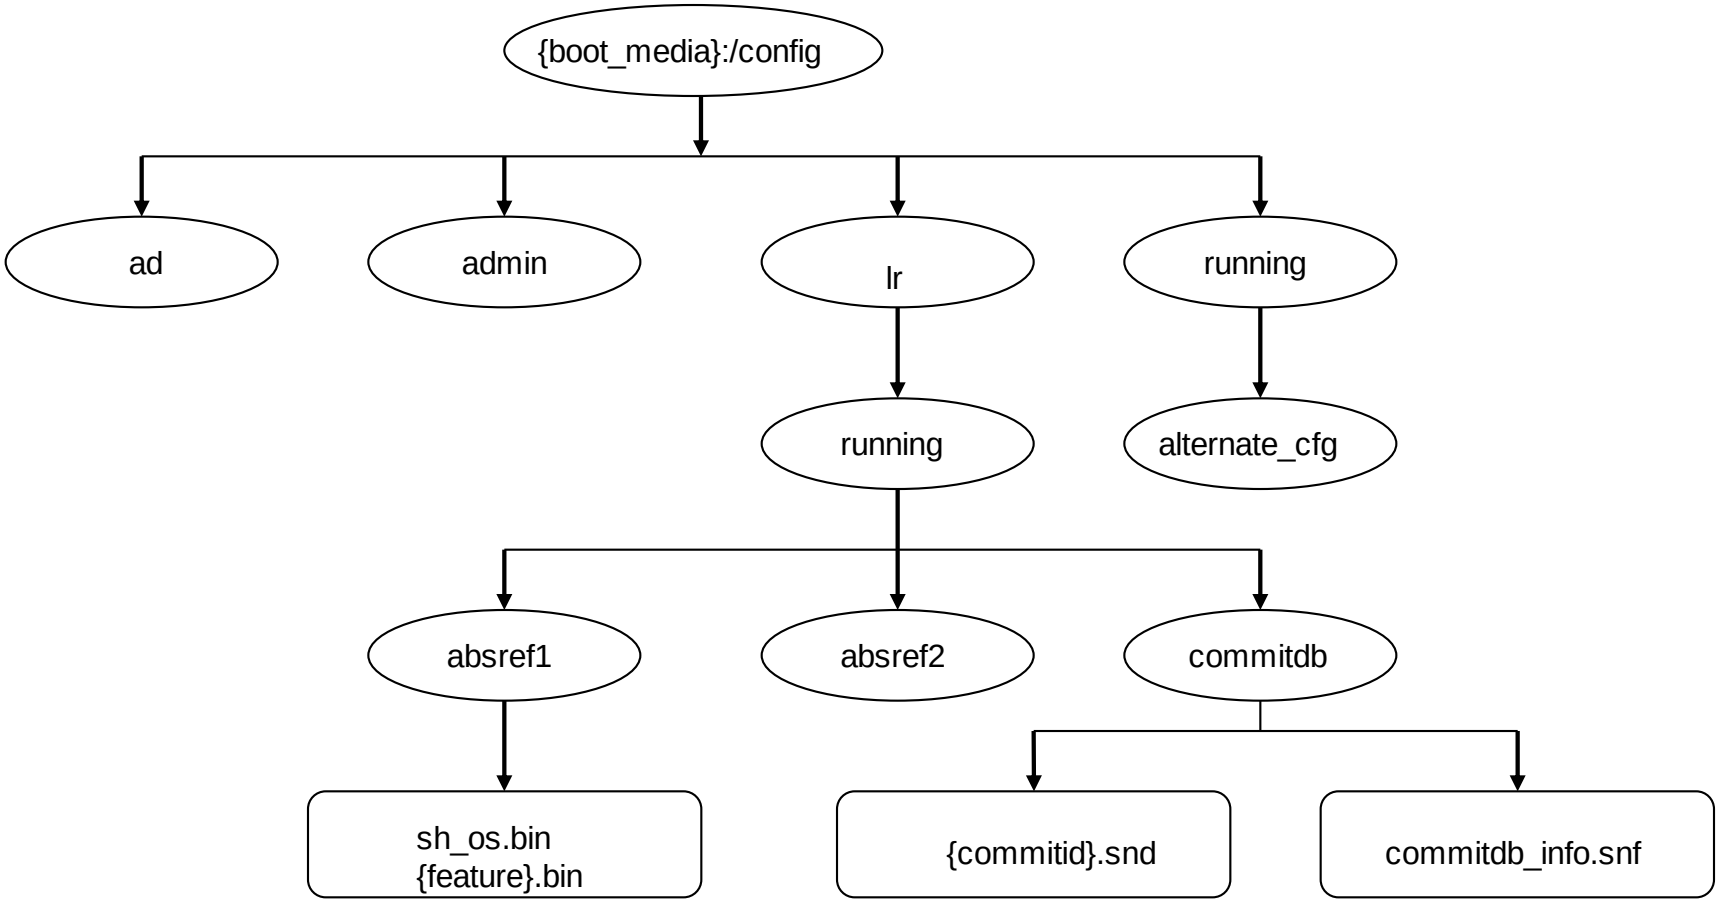

Two-Stage Configuration

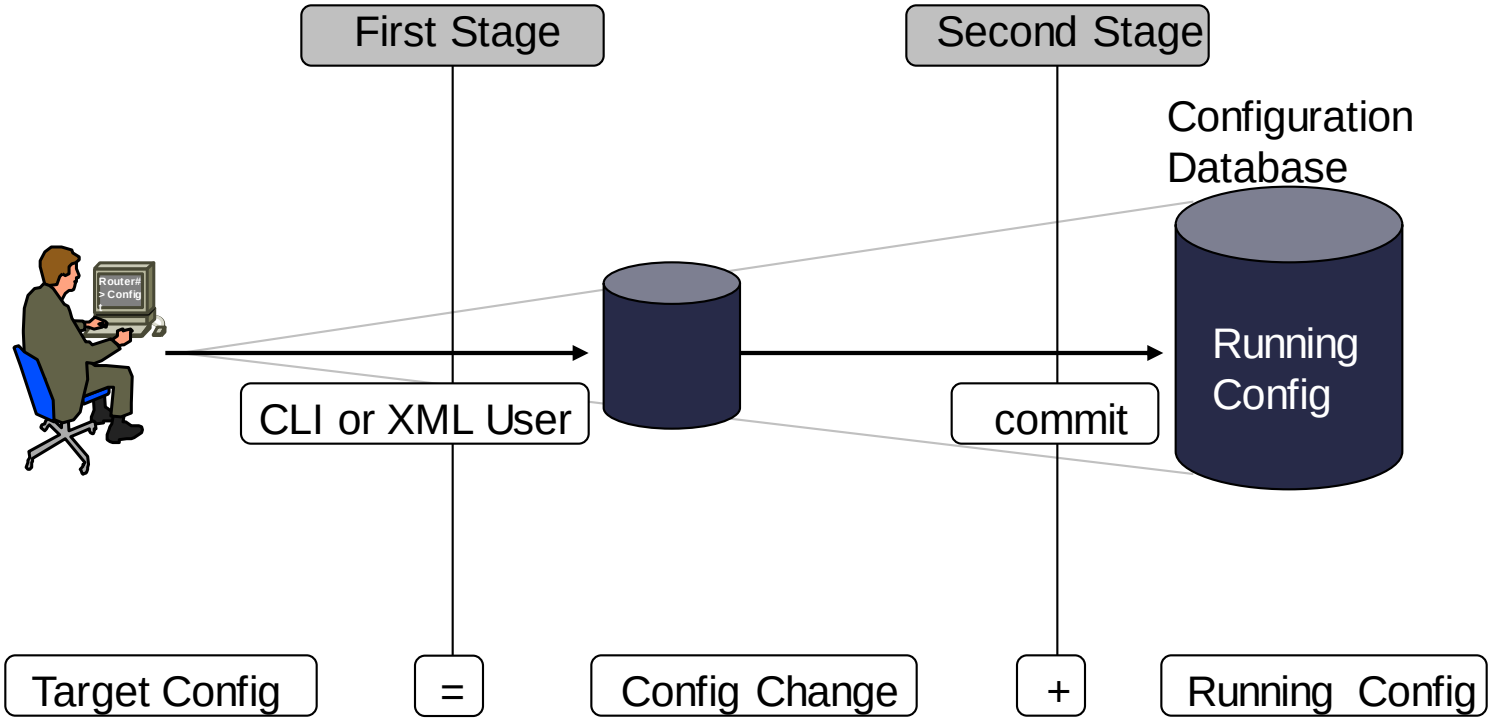

Building Target Configuration

```
RP/0/RP0/CPU0:F34m(config)#show running-config interface GigabitEthernet0/0/1/0
Tue Jul 28 11:48:40.917 JST
interface GigabitEthernet0/0/1/0
description TEST
!

RP/0/RP0/CPU0:F34m(config)#interface GigabitEthernet0/0/1/0
RP/0/RP0/CPU0:F34m(config-if)#ipv4 address 192.168.20.5 255.255.255.0
RP/0/RP0/CPU0:F34m(config-if)#show config
Tue Jul 28 11:49:53.032 JST
Building configuration...
interface GigabitEthernet0/0/1/0
 ipv4 address 192.168.20.5 255.255.255.0
!
end

RP/0/RP0/CPU0:F34m(config-if)#show config merge interface GigabitEthernet0/0/1/0
Tue Jul 28 11:50:14.091 JST
interface GigabitEthernet0/0/1/0
description TEST
 ipv4 address 192.168.20.5 255.255.255.0
```

Configuration Log and changes

```
RP/0/RP0/CPU0:H4#conf t
Tue Jul 28 13:41:08.184 JST
RP/0/RP0/CPU0:H4(config)#hostname R4
RP/0/RP0/CPU0:H4(config)#commit
RP/0/RP0/CPU0:Jul 28 13:41:22.992 : ike[202]: %SECURITY-IKE-4-WARNING : You may want to
    configure a domain-name
RP/0/RP0/CPU0:Jul 28 13:41:24.184 : config[65791]: %MGBL-CONFIG-6-DB_COMMIT :
    Configuration committed by user 'admin'. Use 'show configuration commit changes
    1000001059' to view the changes.
RP/0/RP0/CPU0:R4(config)#end
RP/0/RP0/CPU0:Jul 28 13:41:30.140 : config[65791]: %MGBL-SYS-5-CONFIG_I : Configured from
    console by admin on vty0 (172.16.0.1)
RP/0/RP0/CPU0:R4#show configuration commit changes 1000001059
Tue Jul 28 13:41:46.805 JST
Building configuration...
hostname R4
end
```

Commit options

```
RP/0/RP0/CPU0:H4(config)#commit ?
  best-effort  Commit the configuration changes via best-effort operation
  comment      Assign a comment to this commit
  confirmed    Rollback this commit unless there is a confirming commit
  force        Override the memory checks
  label        Assign a label to this commit
  replace      Replace the contents of running configuration
  <cr>         Commit the configuration changes via atomic operation
```

- The target configuration can be committed using either pseudo-atomic or best-effort options
- In case of any verification errors with pseudo-atomic option, the commit operation is rolled back and the error are returned to the configuration agent
- In the case of the commit with the best-effort option, only that part of the target configuration that has gone through successful verification gets committed, and not necessarily everything in the target buffer

Commit options
Commit Confirmed

```
RP/0/RP0/CPU0:Q3#conf t
Wed Jul 29 00:43:23.725 UTC
RP/0/RP0/CPU0:Q3(config)#hostname R3
RP/0/RP0/CPU0:Q3(config)#commit confirmed ?
  <30-65535>  Seconds until rollback unless there is a confirming commit
  minutes     Specify the rollback timer in the minutes
  <cr>        Commit the configuration changes via atomic operation
RP/0/RP0/CPU0:Q3(config)#commit confirmed 30
RP/0/RP0/CPU0:Jul 29 00:43:54.879 : config[65801]: %MGBL-CONFIG-6-DB_COMMIT :
    Configuration committed by user 'admin'. Use 'show configuration commit changes
    1000000206' to view the changes.
RP/0/RP0/CPU0:R3(config)#
RP/0/RP0/CPU0:R3(config)#RP/0/RP0/CPU0:Jul 29 00:44:25.241 :
    cfgmgr_trial_confirm[65803]: %MGBL-CONFIG-6-DB_COMMIT : Configuration committed by
    user 'admin'. Use 'show configuration commit changes 1000000207' to view the changes.

RP/0/RP0/CPU0:Q3(config)#
```

- The ‘commit confirmed’ command helps to commit the target configuration on a trial basis and not as a permanent configuration
- A user can provide the amount of time to wait before the configuration rollback to the previous configuration

Configuration Navigation

```
RP/0/RP0/CPU0:Q3#show running-config router ospf 1
Wed Jul 29 01:35:35.006 UTC
router ospf 1
 router-id 172.20.3.3
 nsf ietf
 area 0
  mpls traffic-eng
  interface Loopback0
  !
  interface GigabitEthernet0/1/3/0
  !
  interface GigabitEthernet0/1/3/1
  !
!
 mpls traffic-eng router-id Loopback0
!

RP/0/RP0/CPU0:Q3#show running-config router ospf 1 area 0 interface GigabitEthernet0/1/3/0
Wed Jul 29 01:36:08.163 UTC
router ospf 1
 area 0
  interface GigabitEthernet0/1/3/0
  !
 !
!
```

- The running configuration can be viewed as a series of functional units
- You can view the BGP or OSPF configuration individually or you can parse using include, exclude, or begin options

Configuration Management behavior for OIR

- Case 1: Insert or Remove a Node
  - When the same SPA is OIRed, interface configuration will be automatically restored when the SPA or node is up
- Case 2: Replace the Node with a Different Node (different media type)
  - All the interface configuration are deleted
  - However control plane configuration such as the OSPF routing protocol will not be changed by this action (user must manually change)
- Case 3: Replace the Node with a Higher Density Node (same media type)
  - Configurations will be reapplied only to port that had been configured before
  - The user has to configure the additional ports and activate the control plane for the same as required
- Case 4: Replace the Node with a Lower Density Node (same media type)
  - Interface-level configuration will be reapplied to those ports that are now in existence and had been previously configured
  - The missing ports will move to preconfigured state
  - Control plane configuration will remain unaffected and therefore active on the new ports

Configuration Management during Package Activation and Deactivation

- Package activation/deactivation results in a configuration name space version change, and IOS XR removes the affected configuration from the router's running configuration prior to the package being deactivated
- IOS XR saves the removed configuration to a filename with date and timestamp
- '`show configuration removed`' shows removed configuration

Interface Preconfiguration

- IOS XR introduces a new concept called interface preconfiguration where you can start configuring the hardware and the interfaces even if they are not present in the system
- All the CLI commands specific to the interface can be provisioned much before the node or SPA is installed
- After inserting the MSC all the preconfiguration gets applied automatically, then it becomes part of running configuration
- Preconfiguration applies to MSC and interface commands only

```
RP/0/RP0/CPU0:Q3#conf t
Thu Jul 30 02:42:00.948 UTC
RP/0/RP0/CPU0:Q3(config)#interface preconfigure POS
    0/2/0/1
RP/0/RP0/CPU0:Q3(config-if-pre)#encapsulation ppp
RP/0/RP0/CPU0:Q3(config-if-pre)#ipv4 address 10.1.1.1
    255.255.255.0
RP/0/RP0/CPU0:Q3(config-if-pre)#no shut
RP/0/RP0/CPU0:Q3(config-if-pre)#pos crc 32
RP/0/RP0/CPU0:Q3(config-if-pre)#exit
RP/0/RP0/CPU0:Q3(config)#controller preconfigure SONET
    0/2/0/1
RP/0/RP0/CPU0:Q3(config-sonet)#clock source internal
RP/0/RP0/CPU0:Q3(config-sonet)#framing sonet
RP/0/RP0/CPU0:Q3(config-sonet)#no shut
RP/0/RP0/CPU0:Q3(config-sonet)#end
Uncommitted changes found, commit them before
    exiting(yes/no/cancel)? [cancel]:yes
RP/0/RP0/CPU0:Jul 30 02:45:17.407 :
    config[65801]: %MGBL-CONFIG-6-DB_COMMIT :
    Configuration committed by user 'admin'. Use 'show
    configuration commit changes 1000000214' to view
    the changes.
```

```
RP/0/RP0/CPU0:Jul 30 02:45:17.465 :
    config[65801]: %MGBL-SYS-5-CONFIG_I : Configured
    from console by admin on vty0 (172.16.0.1)
RP/0/RP0/CPU0:Q3#sho running-config
Thu Jul 30 02:45:37.619 UTC
!
interface preconfigure POS0/2/0/1
 ipv4 address 10.1.1.1 255.255.255.0
 encapsulation ppp
 pos
  crc 32
 !
!
controller preconfigure SONET0/2/0/1
 framing sonet
 clock source internal
```

Polling Question 3

How often to you contact TAC for IOS-XR troubleshooting?

a. I work with TAC on a regular basis
b. I sometimes contact TAC
c. I rarely contact TAC
d. I have never contacted TAC

Cisco IOS XR Monitoring and Operations

Using SNMP

- IOS XR has a built-in SNMP agent, whose job is to retrieve and update MIB that describe the state of the device or protocols and features enabled in the device
- IOS XR supports the following versions of SNMP:
  - SNMPv1
  - SNMPv2c
  - SNMPv3
- See below for CRS-1 MIB support list
  - ftp://ftp-sj.cisco.com/pub/mibs/supportlists/crs1/crs1-supportlist.html

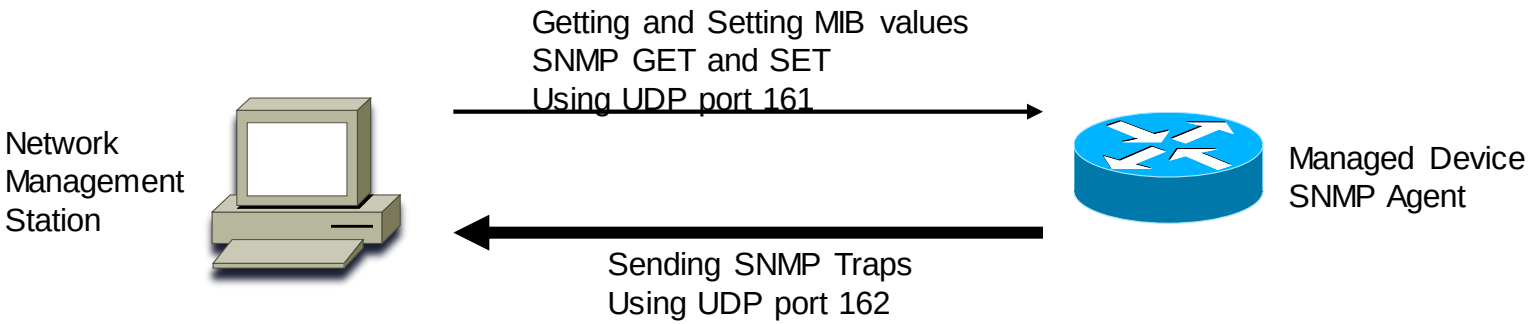

SNMP in the context of a VRF

- A network scenario might require SNMP management stations to be present in L3 VPN contexts

```
vrf NET_MGMT
 address-family ipv4 unicast
 import route-target
   100:111
 export route-target
   100:111
!
snmp-server commumity public RO
snmp-server community private RW
snmp--server traps snmp
snmp-server trap-source Loopback101
snmp-server vrf NET_MGMT
 host 10.10.20.100 traps version 2c public
!
interface Loopback101
 vrf NET_MGMT
 ipv4 address 10.10.21.1 255.255.255.255
!
interface MgmtEth0/0/CPU0/0
 vrf NET_MGMT
 ipv4 address 10.10.20.21 255.255.255.0
```

Syslog

- IOS XR router sends syslog messages to a syslog daemon
- The syslog daemon controls the dispatch of syslog messages to the logging buffer, terminal lines, or an external syslog server

```
LC/0/7/CPU0:Oct 31 00:15:25.294 : ifmgr[146]: %PKT_INFRA-LINEPROTO-
5-UPDOWN : Line protocol on Interface GigabitEthernet0/7/0/4,
changed state to Down
```

Message	Field
`LC/0/7/CPU0:Oct 31 00:15:25.294 :`	NodeID: Time Stamp:
`ifmgr[146]:`	process-name [pid]:
`%PKT_INFRA-LINEPROTO-5-UPDOWN :`	%message-group-severity-message-code :
`Line protocol on Interface GigabitEthernet0/7/0/4, changed state to Down`	message-text

Logging Destination

- Logging messages to the console are on by default and this behavior can be changed with the ‘logging console disable’
- A user that accesses a router through a vty line can enable logging messages by ‘terminal monitor’ or disable it by ‘terminal monitor disable’
- Logging messages can be directed to a syslog server using the configuration command ‘logging {svr_address | svr_name}’
- Logging messages can be sent to an internal circular buffer using the configuration command ‘logging buffered [size] [level]’
- IOS XR also allows the logging messages to be sent to an SNMP NMS station using the configuration ‘snmp-server enable traps syslog’

Local archiving of Logging Messages

- The syslog archive can be created on a local storage device and can be particularly useful for having an onboard backup in case communication is lost with an external syslog server
- The logging archive is typically created under /var/log followed by a directory name based on a date

```
logging archive
 device harddisk
 file-size 1
 frequency weekly
 archive-size 10
 archive-length 4
```

In the above example:
The archive length of 4 weeks specifies that the log will be kept for a month before getting removed
The frequency specifies whether the log will be collected on a daily or a weekly basis
The archive size specifies the total size that limits the archive and the file size is the maximum size for an individual file
When the file size is exceeded, a new file is created

Embedded Event Manager

- EEM detects events related to fault occurrence and fault recovery as well as maintains process and reliability statistics
  - A crucial resource such as memory or CPU oversteps normal operation thresholds
  - OIR of hardware such as a fabric or MCS
  - A process crash
  - Application defined events
- When an event occurs in the system, EEM gets notified by event detectors, and then EEM may take a predetermined corrective actions specified by a user-defined policy

EEM Event Detector and Events processing

- System Manager Event Detector
- Timer Services Event Detector
- Syslog Event Detector
- None Event Detector ... publishes an event when the Event Manager run command executes an EEM policy
- Watchdog System Monitor Event Detector
- Distributed Event Detectors ... publishing of EEM events for processes that are local to a node

Registering and Using Onboard EEM Policies

- IOS XR ships with some onboard TCL scripts that can be registered and used
- These TCP scripts can be browsed by going to the directory /pkg/lib/tcl/fm_scripts/

```
RP/0/RP0/CPU0:F5#show event manager policy available system
No.  Type    Time Created              Name
1    system  Thu Aug  6 08:54:56 2009  ospf_sysmgr_abort.tcl
2    system  Thu Aug  6 08:54:56 2009  ospf_sysmgr_user.tcl
3    system  Thu Aug  6 09:05:15 2009  periodic_diag_cmds.tcl
4    system  Thu Aug  6 09:05:15 2009  periodic_proc_avail.tcl
5    system  Thu Aug  6 09:05:15 2009  periodic_sh_log.tcl
6    system  Thu Aug  6 09:14:24 2009  sl_sysdb_timeout.tcl
7    system  Thu Aug  6 09:05:15 2009  tm_cli_cmd.tcl
8   system  Thu Aug  6 09:05:15 2009  tm_crash_hist.tcl

RP/0/RP0/CPU0:F5(config)#aaa authorization event manager default local
RP/0/RP0/CPU0:F5(config)#event manager environment _cron_entry_diag 55 8,20 * * *
RP/0/RP0/CPU0:F5(config)#event manager policy periodic_diag_cmds.tcl username cisco

RP/0/RP0/CPU0:F5#show event manager policy registered
Mon Aug 17 16:34:26.523 JST
No.  Type    Event Type          Time Registered           Name
1   system  timer cron          Fri Nov 14 16:32:54 2008  periodic_diag_cmds.tcl
 name {diag_cmds} cron entry {55 8,20 * * *}
 nice 0 priority normal maxrun_sec 900 maxrun_nsec 0
persist_time: 3600 seconds, username: cisco
```

User-Defined EEM Policy

- In addition the onboard TCL scripts that come with the IOS XR system, users may write their own TCL-based policies

WDSYSMON

- WDSYSMON stands for Watchdog System Monitor which runs on RPs and linecards
- WDSYSMON monitors the processes and memory on each node for memory and CPU usage, as well as boot device because it is considered a critical resource for the continuous operation of IOS XR system
- The following example shows the memory states of a given node and the default values at which WDSYSMON generates its alarm messages
- After determining the memory state of the node to be SEVERE, WDSYSMON finds top memory consumer process and restarts the process to avoid further memory depletion

```
RP/0/RP0/CPU0:F5#show watchdog memory-state location 0/RP0/CPU0
Mon Aug 17 17:55:22.260 JST
Memory information:
    Physical Memory: 4096     MB
    Free Memory:     2374.390 MB
    Memory State:         Normal
RP/0/RP0/CPU0:F5#show watchdog threshold memory defaults location 0/RP0/CPU0
Mon Aug 17 17:55:58.460 JST
 Default memory thresholds:
 Minor:     409      MB
 Severe:    327      MB
 Critical:  204.799 MB
Memory information:
    Physical Memory: 4096     MB
    Free Memory:     2374.390 MB
    Memory State:         Normal
```

Monitoring Memory

Memory comparison snapshot

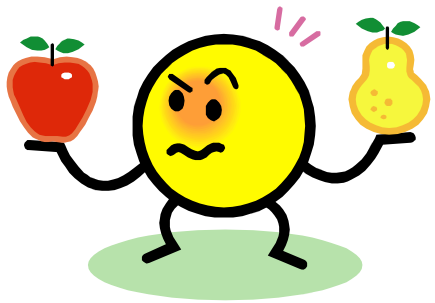

```
RP/0/RP0/CPU0:F5#show memory compare start
Mon Aug 17 18:19:40.959 JST
Successfully stored memory snapshot /harddisk:/malloc_dump/memcmp_start.out
RP/0/RP0/CPU0:F5#show memory compare end
Mon Aug 17 18:25:57.406 JST
Successfully stored memory snapshot /harddisk:/malloc_dump/memcmp_end.out
RP/0/RP0/CPU0:F5#show memory compare report
Mon Aug 17 18:26:07.863 JST
JID   name                mem before   mem after    difference mallocs
restart/exit/new
---   ----                ----------   ---------    ---------- ------- ------
----------
328   parser_server       10897312     10898776     1464       1
376   statsd_manager      1135772      1137080      1308       0
432   wdsysmon            1638684      1639812      1128       4
384   sysdb_svr_admin     1723684      1724096      412        9
365   shelfmgr            723356       723584       228        3
65746 exec                182808       182888       80         3
385   sysdb_svr_local     2609684      2607448      -2236      -53
```

Show system verify

- IOS XR has a built-in utility that performs a high-level system verification:
  - show system verify [start | report | detail]
- It performs various checks with regard to memory, blocked or aborted processes, interface counters, and infrastructure health
- It launches a TCL script that takes a snapshot of the system state at a certain point in time and then takes another snapshot at a later point in time and compares the two reports

System Monitoring

- Monitoring
  - show tech
  - show system verify
  - monitor interface
- Online Help
  - describe
  - man

Show tech variations

```
RP/0/RP0/CPU0:CRS#show tech ?
  bcdl        Output show commands of interest for bcdl debugging
  cef         Output show commands of interest for CEF debugging
  file        Specify a valid file name (e.g. disk0:tmp.log)
  gsp         Output show commands of interest for gsp debugging
  lrd         Output show commands of interest for lrd debugging
  ospfv3      Output show commands for OSPFv3 debugging
  password    Include password in output
  pfi         Output show commands of interest for PFI debugging
  placement   Gather lots of information about process placement
  platform    show tech-support platform output
  rib         Show IP RIB related information
  routing     Routing show tech-support output
  terminal    Send output to terminal
```

Monitor One Interface
monitor interface command

```
RP/0/RP1/CPU0:CRS1#monitor interface mgmtEth 0/rp1/CPU0/0

CRS1                Monitor Time: 00:00:52          SysUptime: 16:33:16
MgmtEth0/RP1/CPU0/0 is up, line protocol is up Encapsulation ARPA


Traffic Stats:(5 minute rates)                                  Delta
  Input  Packets:                    313326                       17
  Input  pps:                             5
  Input  Bytes:                    34467898                     1033
  Input  Kbps:                            3
  Output Packets:                    37633                        4
  Output pps:                             0
  Output Bytes:                    2034463                      244
  Output Kbps:                            0

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',
Detail='d', Brief='b', Next='n', Prev='p'
```

Monitor All Interfaces
monitor interface all command

```
RP/0/RP1/CPU0:CRS1#monitor interface all

CRS1                Monitor Time: 00:00:55          SysUptime: 16:35:50

Interface  Encap                Input      Output        Input          Output
                                  pps         pps         Kbps            Kbps
MgmtEth0/RP0/CPU0/0 ARPA            4           0            2               0
MgmtEth0/RP1/CPU0/0 ARPA            4           0            2               0
POS0/4/0/0          HDLC         1230           0       542210               0
POS0/4/0/3          HDLC            0           0            0               0
POS0/4/0/6          PPP             0           0            0               0
POS0/4/0/7          HDLC            0        1230            0          542210


Quit='q', Freeze='f', Thaw='t', Clear='c', Next set='n', Prev set='p'
```

Online Help
Requires –doc package

```
RP/0/RP1/CPU0:CRS#man keyword mpls

Following Commands Matched for 'mpls'

    FEATURE              COMMAND
      is-is                mpls traffic-eng
      is-is                mpls traffic-eng router-id (is-is)
      is-is                show isis mpls traffic-eng adjacency-log
      is-is                show isis mpls traffic-eng advertisements
      is-is                show isis mpls traffic-eng tunnel
      ospf                 mpls traffic-eng area (ospf)
      ospf                 mpls traffic-eng router-id (ospf)
      ospf                 show ospf mpls traffic-eng
      quality-of-service   match mpls experimental topmost
      quality-of-service   set mpls experimental topmost

RP/0/RP1/CPU0:CRS#man command show cef ipv4

COMMAND

    show cef ipv4

DESCRIPTION

To display the IPv4 Cisco Express Forwarding (CEF) table, use the show cef ipv4
  command in EXEC mode.

show cef ipv4 [prefix [mask] | type instance] [detail] [location node-id]
```

Details about commands
describe command

```
RP/0/5/CPU0:GSR1#describe show install active
Package:
    c12k-base
        c12k-base V3.5.0[00]  Base Package for 124xx
        Vendor : Cisco Systems
        Desc   : Base Package for 124xx
        Build  : Built on Tue Jun 12 20:36:49 UTC 2007
        Source : By zamboni in
/auto/srcarchive2/production/3.5.0/c12k/workspace for c2.95.3-p8
        Card(s): RP, DRP, DRPSC, OC3-POS-4, OC12-POS, GE-3, OC12-POS-4, OC48-
POS, E3-OC48-POS, E3-OC12-POS-4, E3-OC3-POS-16, E3-OC3-POS-8, E3-OC3-POS-4,
E3-OC48-CH, E3-OC12-CH-4, E3-OC3-CH-16, E3-GE-4, E3-OC3-ATM-4, E3-OC12-ATM-4,
E5-CEC, E5-CEC-v2, SE-SEC
        Restart information:
          Default:
            parallel impacted processes restart

Component:
    installmgr V[main/472]  On the box installation program



User needs ALL of the following taskids:

        pkg-mgmt (READ)
```

Verifying the System

1. admin
2. show platform
3. show version
4. show running-config
5. show logging
6. show environment
7. show context
8. exit
9. show context
10. show memory summary detail location all
11. show memory heap summary {job-id | all}
12. top processes
13. show running-config
14. show system verify start
15. show system verify report
16. show {ipv4 | ipv6} interface brief
17. show install active

Cisco IOS XR Security

Topics

- Secure Operating System
- Securing Access to the Router
- Securing the forwarding plane

Secure Operating System

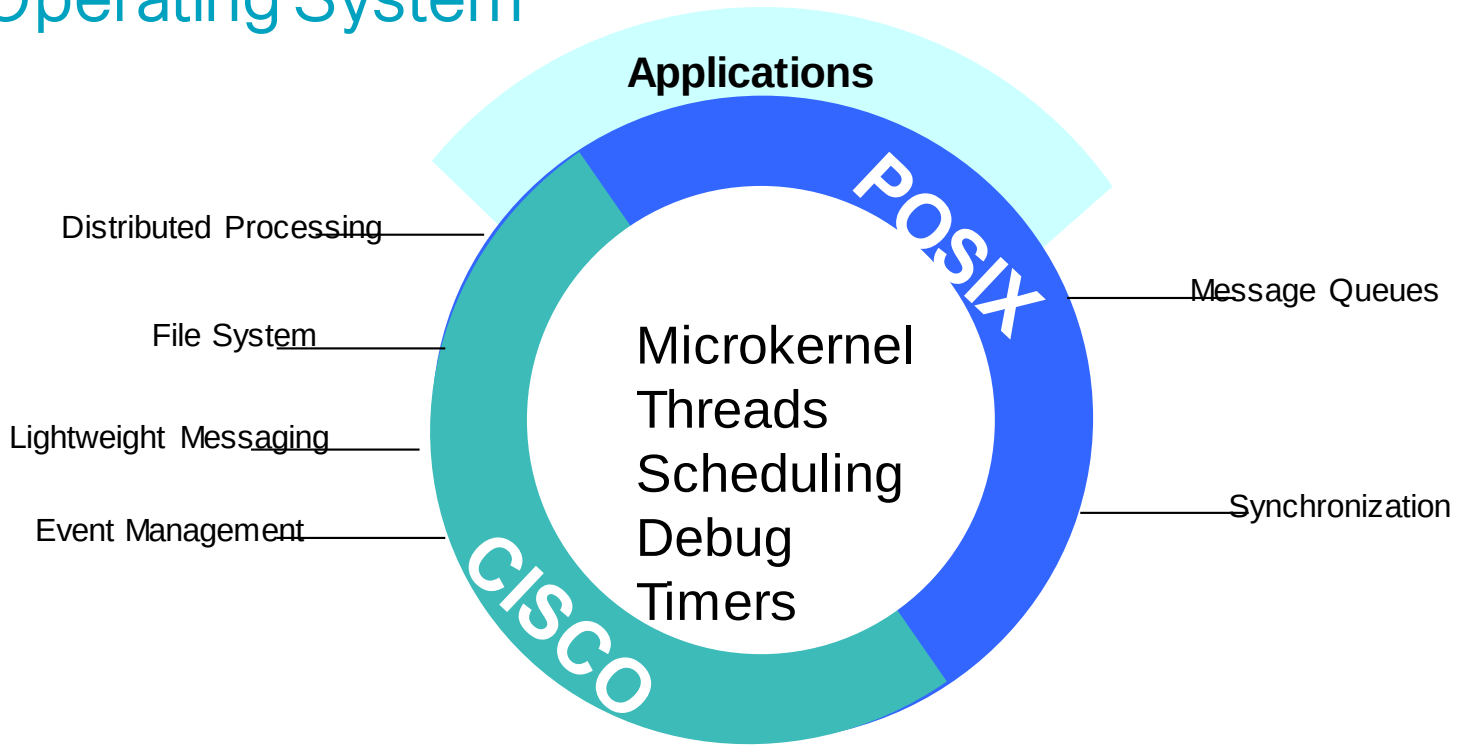

- IOS XR is a microkernel-based operating system offering memory protection and fault tolerance.
- All basic OS and router functionality is implemented as process. All the distributed services run on top of the microkernel
- The microkernel is protected from faults occurring in the protocol or device driver software due to the layered model

Securing Access to the Router –Admin/SDR Plane

- Admin Plane:
  - maintains responsibility for the owner SDR, and certain administrative responsibilities for all other non-owner SDRs
  - is accessible only to a type of user known as the root-system user
- SDR Plane:
  - The root-system user has ability to create SDRs and root SDR user (root-lr)
  - The root-lr user is the equivalent of root-system user from an SDR perspective and has jurisdiction only for the particular SDR on which it is defined

Securing Access to the Router –User Groups and Task Groups

- A user that logs in to an IOS XR router may have one or more preconfigured user groups assigned to it
- Some user groups are pre-created by default and others may be defined via configuration

User Groups and Task Groups	Purpose
cisco-support	Used by Cisco Support Team. Provides access to troubleshooting commands
netadmin	Provides the ability to control and monitor all system- and network-related parameters
operator	Provides very basic user priviledges
root-lr	Provides the ability to control and monitor the specific SDR
root-system	Provides the ability to control and monitor the entire system
sysadmin	Provides the ability to control and monitor all system parameters but cannot configure network protocols
serviceadmin	Provides the ability to administer session border controllers

Securing Access to the Router – Task Groups and Task IDs

- Task groups contain a collection of task IDs that define actions such as READ, WRITE, EXECUTE, or DEBUG (R/W/E/D)

```
RP/0/RP0/CPU0:F5#show running-config taskgroup igp-admin
Mon Aug 24 15:57:51.997 JST
taskgroup igp-admin
 task read ospf
 task read mpls-te
 task write ospf
 task execute ospf
 task debug ospf
 task debug bundle
 description OSPF Administrator
!

RP/0/RP0/CPU0:F5#show running-config usergroup igp-admin
Mon Aug 24 15:58:07.340 JST
usergroup igp-admin
 taskgroup igp-admin
!

RP/0/RP0/CPU0:F5#show running-config username igpadmin
Mon Aug 24 15:58:24.639 JST
username igpadmin
 group igp-admin
 secret 5 $1$kN1M$9I71jeLthzYrKlBE6wWVW.
!

RP/0/RP0/CPU0:F5#show aaa taskgroup igp-admin
Mon Aug 24 15:58:50.516 JST
Task group 'igp-admin'

Task IDs included directly by this group:
Task:               bundle  :                              DEBUG
Task:              mpls-te  : READ
Task:                 ospf  : READ    WRITE    EXECUTE    DEBUG

Task group 'igp-admin' has the following combined set
  of task IDs (including all inherited groups):
Task:               bundle  :                              DEBUG
Task:              mpls-te  : READ
Task:                 ospf  : READ    WRITE    EXECUTE    DEBUG
```

Securing Access to the Router – Task Groups and Task IDs

- Task groups contain a collection of task IDs that define actions such as READ, WRITE, EXECUTE, or DEBUG (R/W/E/D)

```
RP/0/RP0/CPU0:F5#show running-config taskgroup igp-admin
Mon Aug 24 15:57:51.997 JST
taskgroup igp-admin
 task read ospf
 task read mpls-te
 task write ospf
 task execute ospf
 task debug ospf
 task debug bundle
 description OSPF Administrator
!

RP/0/RP0/CPU0:F5#show running-config usergroup igp-admin
Mon Aug 24 15:58:07.340 JST
usergroup igp-admin
 taskgroup igp-admin
!

RP/0/RP0/CPU0:F5#show running-config username igpadmin
Mon Aug 24 15:58:24.639 JST
username igpadmin
 group igp-admin
 secret 5 $1$kN1M$9I71jeLthzYrKlBE6wWVW.
!

RP/0/RP0/CPU0:F5#show aaa taskgroup igp-admin
Mon Aug 24 15:58:50.516 JST
Task group 'igp-admin'

Task IDs included directly by this group:
Task:               bundle  :                              DEBUG
Task:              mpls-te  : READ
Task:                 ospf  : READ    WRITE    EXECUTE    DEBUG

Task group 'igp-admin' has the following combined set
  of task IDs (including all inherited groups):
Task:               bundle  :                              DEBUG
Task:              mpls-te  : READ
Task:                 ospf  : READ    WRITE    EXECUTE    DEBUG
```

Securing Access to the Router
User –User Group – Task Group – Task ID hierarchy

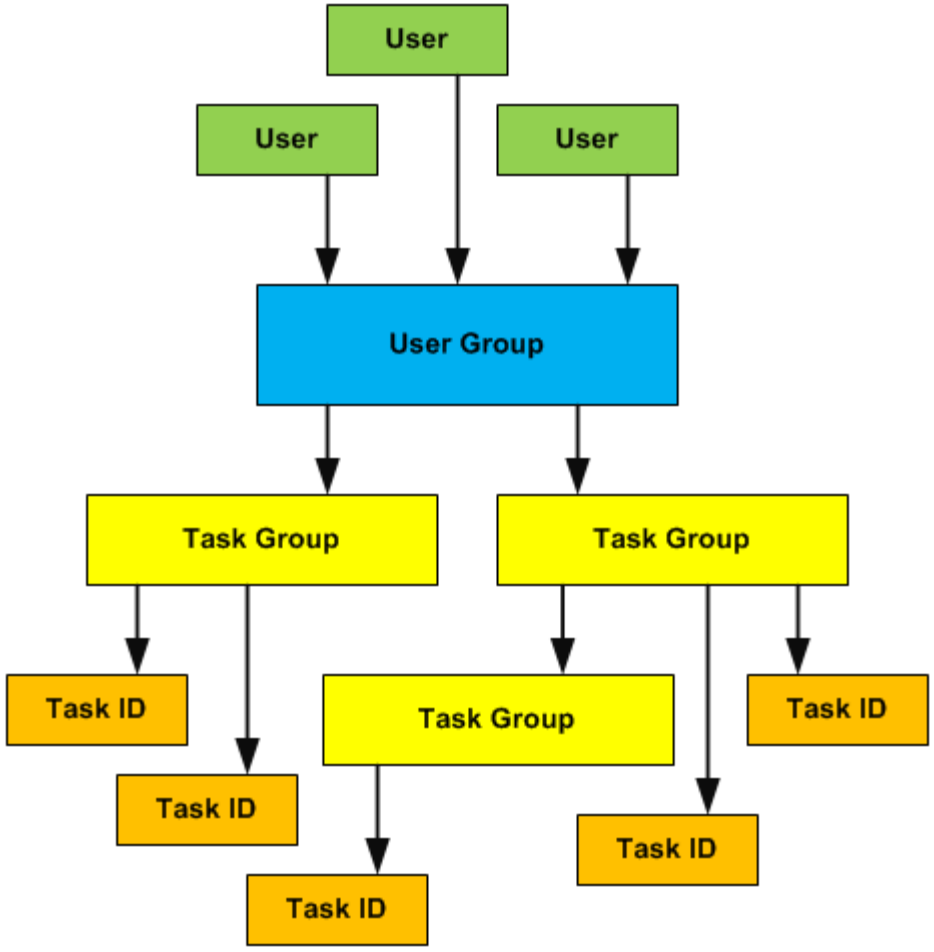

Securing Access to the Router – External AAA

- IOS XR supports external AAA such as TACACS+ and RADIUS
- AAA server and client are identified by IP address and a secret shared key is configured between them
- The notion of a user group on IOS XR local AAA is unrelated to a user group on an ACS server
- IOS XR task groups are identified as optional attributes on the ACS server
- IOS XR AAA supports mapping between privilege levels that can be defined for a given user in the external AAA server file
  - The local user group on the router needs to be configured with a user group with a name that matches the privilege level
  - Privilege level 1 to 13 may be mapped as so
  - Privilege level 15 maps to the root-system and privilege level 14 maps to root-lr

Securing Access to the Router – Configuring SSH

- IOS XR supports two versions of SSH:
  - SSHv1 uses RSA key
  - SSHv2 uses DSA
- Enabling SSH on IOS XR requires hfr-k9sec PIE to be installed
- In addition to installing the hfr-k9sec PIE, RSA or DSA key must be generated before SSH runs in server mode

```
RP/0/RP1/CPU0:CRS1-1(admin)#show install active | include k9sec
disk0:hfr-k9sec-3.6.0
!
! The following command generates DSA key pairs
!
RP/0/RP1/CPU0:CRS1-1#crypto key generate dsa
The name for the keys will be: the_default
Choose the size of your DSA key modulus. Modulus size can be 512, 768, or 1024 bits.
    Choosing a key modulus
How many bits in the modulus [1024]: 1024
Generating DSA keys ...
Done w/ crypto generate keypair
[OK]
!
RP/0/RP1/CPU0:CRS1-1(config)#ssh server v2
RP/0/RP1/CPU0:CRS1-1(config)#commit
```

Securing Access to the Router – Management Plane Protection (MPP)

- MPP enhances the manageability and security aspects of IOS XR
- MPP helps alleviate the need to configure more access lists in controlling router access
- Management ports on (D)RPs are not configurable under MPP because they are out of band by default
- MPP controls incoming traffic for protocols such as TFTP, telnet, SNMP, SSH, and HTTP
- MPP allows control for both of in-band and out-of-band interfaces
- MPP can specify a peer IPv4 or IPv6 address or subnet from which traffic is allowed

Securing Access to the Router
MPP in-band and out-of-band interface

- An in-band management interface is an interface that receives and processes management packets as well as forwards Internet traffic
  - This interface may be referred to as a shared management interface

- An out-of-band interface allows only management protocol traffic to be forwarded or processed
  - This type of interface does not process or receive any customer or Internet traffic

Securing the forwarding plane – uRPF & ACL

uRPF

- Strict uRPF
  - If the reverse path back to the source address of incoming packet is not learned via the interface on which strict uRPF is enabled, the packet is dropped
- Loose uRPF
  - If the route for the source interface is not in the routing table, the packet is dropped
  - It is useful when a case of asymmetric routing might be present on the network

ACL

- Ability to filter on TTL, packet length, fragments
- Interface level statistics in hardware
- Interface ACL processing happens before LPTS processing
- Logging gives ability for forensics and is rate limited on number of packets sent to the CPU to avoid over running CPU resources

Securing the forwarding plane – Local Packet Transport Service (LPTS)

- LPTS provides software architecture to deliver locally destined traffic to the correct node on the router and provides security against overwhelming the router resources with excessive traffic
- LPTS achieves security by policing flows of locally destined traffic to a value that can be easily sustained by the CPU capabilities of the platform
- LPTS provides sort of a built-in firewall for an IOS XR router by taking preemptive measures for traffic flows destined to the router

Introduction to different IOS-XR platforms

CRS Family – Investing in the Future
10X Capacity Gains over 10 Years

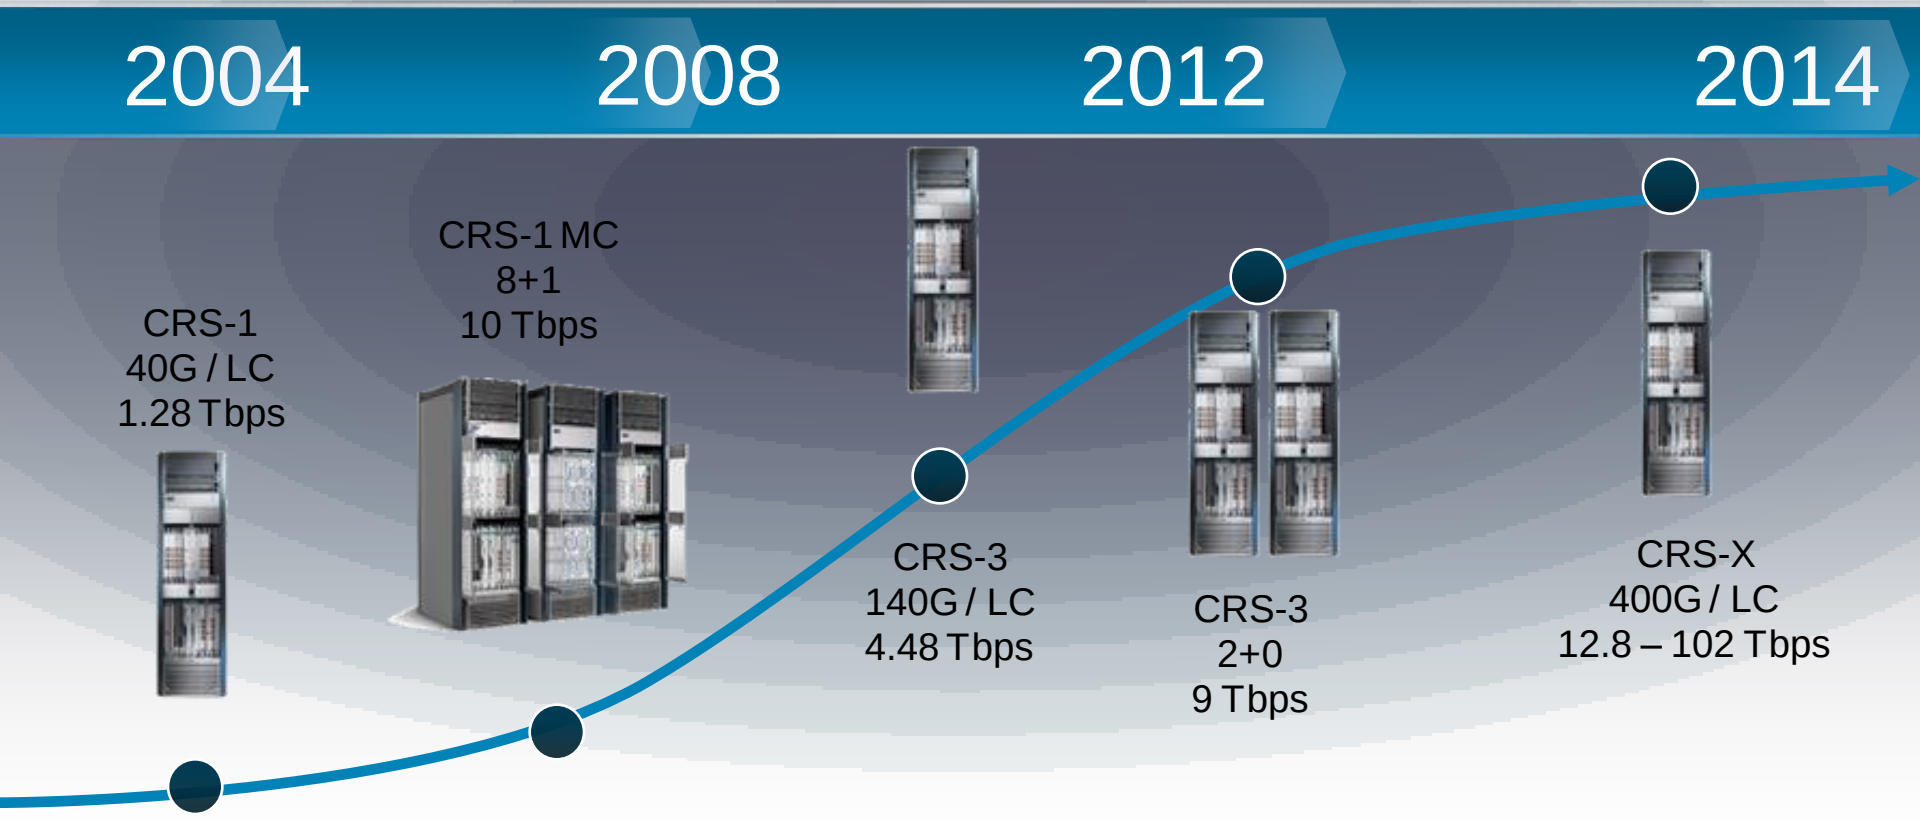

ASR 9000 - Reinventing Edge Routing

Cisco NCS Family – Programmable, Converged, Elastic, Scalable

Cisco NCS 6000

Cisco NCS 4000

Cisco NCS 2000

Submit Your Questions Now!

Use the Q & A panel to submit your questions and our expert will respond

Please be sure to fill out the survey at the end to let us know the quality of the presentation and how this webcast filled your business needs

Join the Cisco Support Community

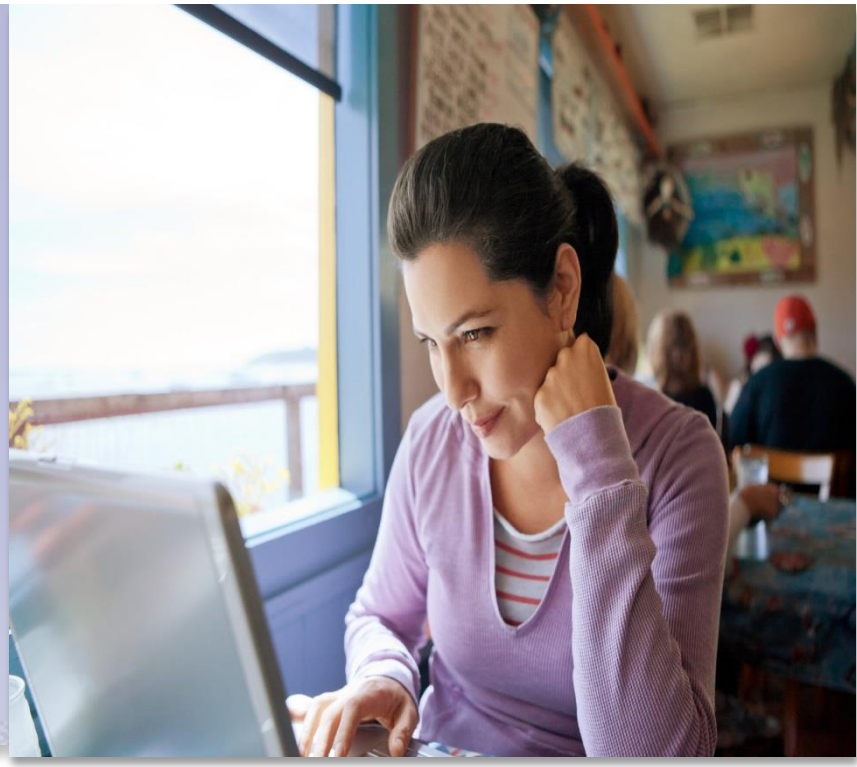

Participate in Live Interactive Technical Events and much more

http://bit.ly/1jll93B

Event Top Contributors

Class 2014

External Contributors

Leonardo Oliveira
Portuguese - VoIP

Bruno Rangel
Portuguese - Collaboration Voice & Video

Guilherme Henrique Villarinho
Portuguese - VoIP

Cisco Contributors

Akhil Behl
Collaboration, Voice, and Video

Hector Carranza Contreras
Routing and Swtiching -Spanish

Néstor González
Wireless -Spanish

Aashish Jolly
Collaboration, Voice, and Video

Hitesh Kumar
Service Providers

Vinay Kumar
Service Providers

Rafael Lima
Routing and Swtiching - Portuguese

Yasuhiro Nakajima
Network Infrastructure - Japanese

Oleksander Nesterov
Data Centers - Russian

Sergey Oliferov
Collaboration, Voice, and Video - Russian

Ricardo Prado
Routing and Swtiching - Spanish

Vignesh Rajendran Praveen
Network Infrastructure

Oleg Tipisov
Security - Russian

Vivek Ruhil
Service Providers

Mike Shchekotilov
Collaboration, Voice, and Video - Russian

Amit Singh
Collaboration, Voice, and Video

Vinayak Sudame
Data Centers

Ryota Takao
Network Infrastructure - Japanese

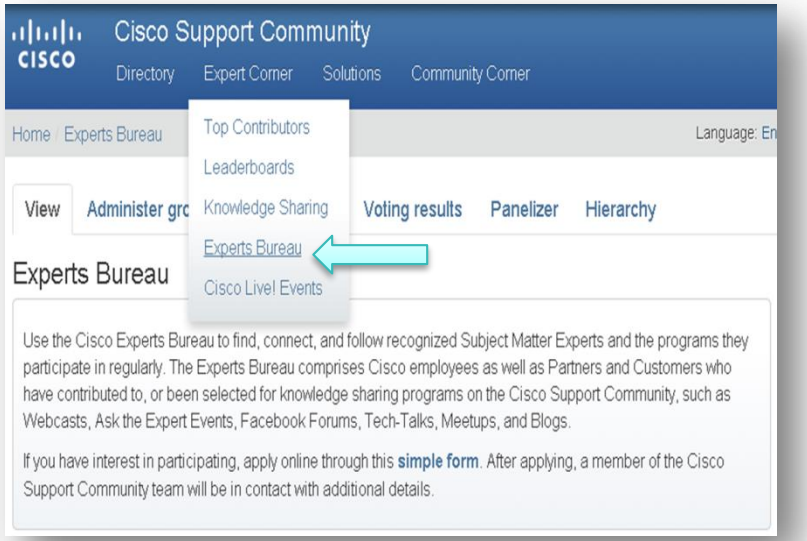

https://supportforums.cisco.com/expert-corner/top-contributors

We invite you to actively collaborate in the Cisco Support Community & Social Media

http://www.facebook.com/CiscoSupportCommunity

http://twitter.com/#!/cisco_support

http://www.youtube.com/user/ciscosupportchannel

https://plus.google.com/110418616513822966153?prsrc=3#110418616513822966153/posts

http://www.linkedin.com/groups/CSC-Cisco-Support-Community-3210019

Newsletter Subscription
https://tools.cisco.com/gdrp/coiga/showsurvey.do?surveyCode=589&keyCode=146298_2&PHYSICAL%20FULFILLMENT%20Y/N=NO&SUBSCRIPTION%20CENTER=YES

We have communities in other languages!

If you speak Spanish, Portuguese, Japanese, Russian or Chinese we invite you to participate and collaborate in your language

Spanish → https://supportforums.cisco.com/community/spanish

Portuguese → https://supportforums.cisco.com/community/portuguese

Japanese → https://supportforums.cisco.com/community/csc-japan

Russian → https://supportforums.cisco.com/community/russian

New Chinese Community!

Chinese → http://www.csc-china.com.cn/

Rate Support Community Content

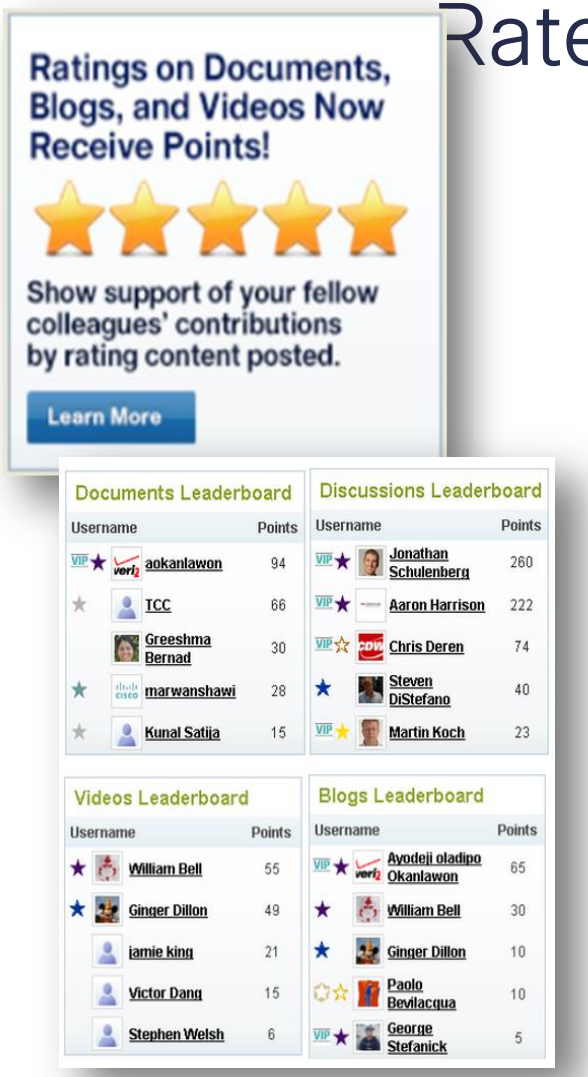

Now your ratings on documents, videos, and blogs count give points to the authors!!!

So, when you contribute and receive ratings you now get the points in your profile.

Help us to recognize the quality content in the community and make your searches easier. Rate content in the community.

https://supportforums.cisco.com/blog/154746

More IT Training Videos & Tech Seminars
on the Cisco Learning Network

View Upcoming Sessions Schedule
cisco.com/go/techseminars

Thank you for Your Time!

Please take a moment to complete the evaluation

Thank you.