
cisco Live!
ALL IN

The bridge to possible

Troubleshooting BGP

A Primer on Quickly Solving Major Internet Outages

Tarique Shakil
Principal Technical Marketing Engineer, Cloud Networking

BRKENT-3000

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 by the speaker until June 17, 2022.

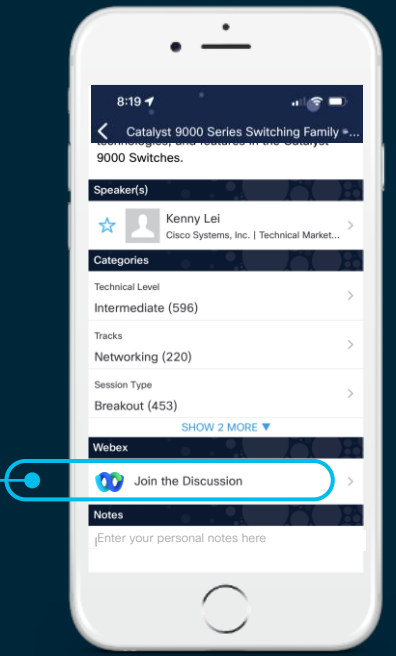

https://ciscolive.ciscoevents.com/ciscolivebot/#BRKENT-3000

Agenda

- Introduction
- BGP Update Generation
- Missing Routes, Unexpected Routes
  - Filtering and Stale Routes
- BGP Table Version & Route Churn
- Troubleshooting with NX-OS
- Conclusion

BGP Update Processing

BGP Update Processing

Route Learning and Propagation Flow

BGP prefixes are injected by explicit configuration

- Network statement – **network <prefix> mask <mask>**
  - Prefix/mask needs to match the RIB exactly
  - Does not enable BGP on an interface like IGP's
- Redistribution – **redistribute ospf <PID>**
  - Injects prefixes from the specified protocol
  - Does not inject 0.0.0.0/0
- Aggregate route – **aggregate-address <prefix> <mask>**
  - Component route must exist in BGP
  - Aggregator attribute is added
- Default route – **default-information originate**

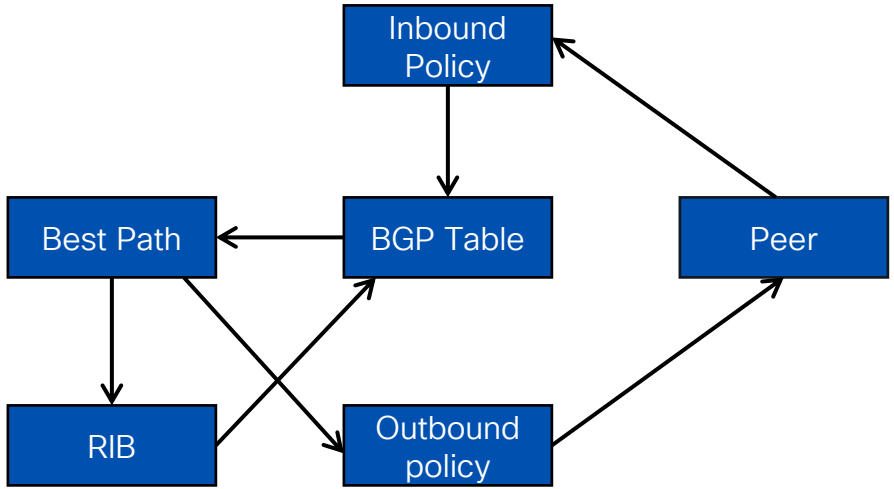

BGP Update Processing – IOS-XE

Update Groups

- Update Group is a collection of peers with identical outbound policy.
- Helps in improving IBGP convergence
  - Update messages are formatted and replicated to all the peers
- A Master is selected in the update group, which is updated first in the group
- Based on the message formatted for the master / Leader, all the peers are then replicated with the same formatted message
  - The message formatting only happens once.

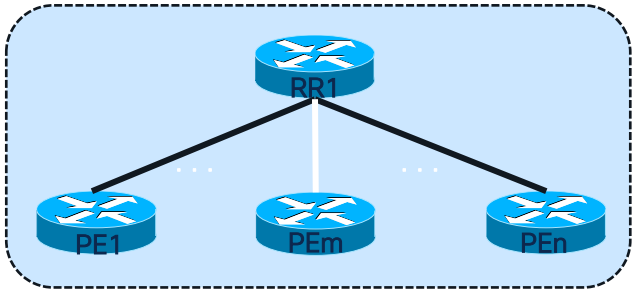

Troubleshooting BGP Convergence

Update Groups

```
R1#show bgp ipv4 unicast update-group
BGP version 4 update-group 2, internal, Address Family: IPv4 Unicast
  BGP Update version : 7/0, messages 0, active RGs: 1
  Route-Reflector Client
  Route map for outgoing advertisements is dummy
  Topology: global, highest version: 7, tail marker: 7
  Format state: Current working (OK, last not in list)
                Refresh blocked (not in list, last not in list)
  Update messages formatted 4, replicated 15, current 0, refresh 0, limit
1000
  Number of NLRIs in the update sent: max 1, min 0
  Minimum time between advertisement runs is 0 seconds
  Has 4 members:
   10.1.12.2        10.1.13.2*        10.1.14.2        10.1.15.2
```

Troubleshooting BGP Convergence

Update Groups on IOS XR

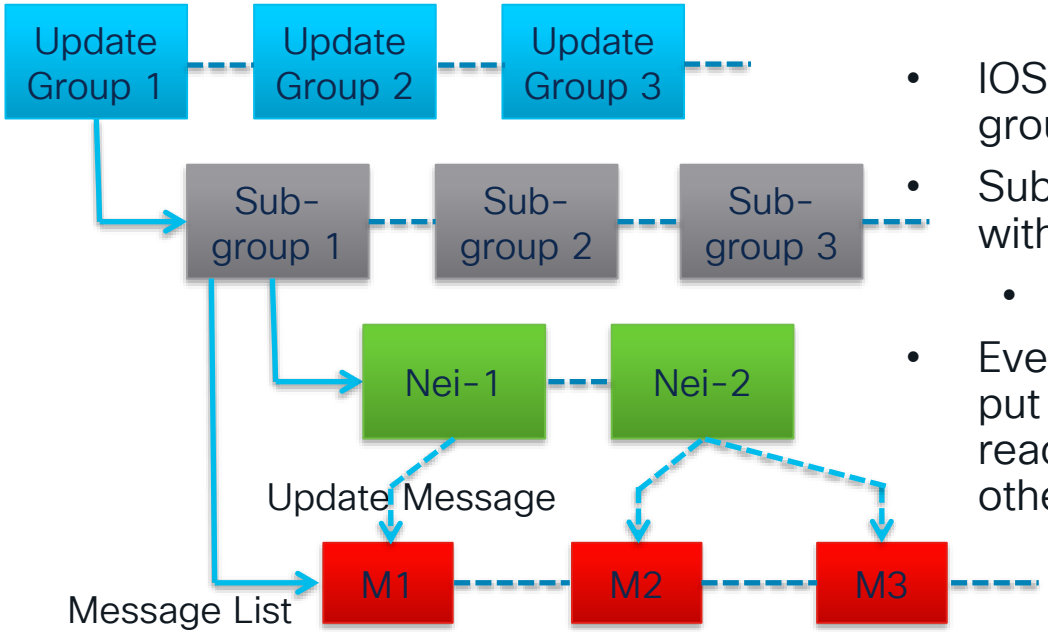

- IOS XR have hierarchical update groups
- Sub-Groups are subset of neighbors within an update Group
  - Neighbors running at same pace
- Even a newly configured neighbor is put in a separate sub-group till it reaches the same table version as other members

Troubleshooting BGP Convergence

Update Groups on IOS XR

```
RP/0/0/CPU0:R10#show bgp update-group
Update group for IPv4 Unicast, index 0.2:
<snip>
Sub-groups merged: 5
Number of refresh subgroups: 0
Messages formatted: 36, replicated: 68
All neighbors are assigned to sub-group(s)
    Neighbors in sub-group: 0.2, Filter-Groups num:3
    Neighbors in filter-group: 0.3(RT num: 3)
        10.1.100.1
    Neighbors in filter-group: 0.1(RT num: 3)
        10.1.100.2
    Neighbors in filter-group: 0.2(RT num: 3)
        10.1.100.8
```

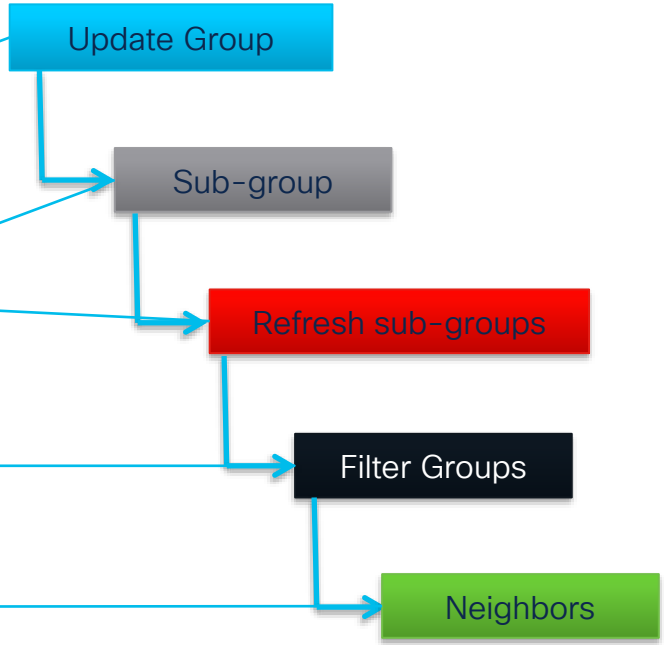

Problems

Stuck BGP Messages

```
R2#show ip bgp sum | begin Neighbor
Neighbor …  MsgRcvd MsgSent    TblVer  InQ OutQ  Up/Down  State/PfxRcd
1.1.1.1  …       53      284    10167    0   97 00:01:20        0
```

The number of packets transmitted is not increasing ☹

The number of packets generated is increasing

At least one BGP keepalive interval apart

```
R2#show ip bgp sum | begin Neighbor
Neighbor …  MsgRcvd MsgSent    TblVer  InQ OutQ  Up/Down  State/PfxRcd
1.1.1.1  …       53      284    10167    0   98 00:02:24        0
```

OutQ is incrementing due to keepalive generation
MsgSent is not incrementing
Something is “stuck” on the OutQ
The keepalives are not leaving R2!!

Troubleshooting Missing Routes / Stale Routes

Missing Routes / Stale Routes

What does it mean?

- Missing Routes
  - The remote peer has not received the route
  - Possible Problem
    - Either speaker didn't advertise the routes or the remote peer didn't receive or process the BGP update
    - Inbound / Outbound Route-maps (Filtering)
- Stale Routes
  - A route present in the BGP table learnt from remote peer but not present on the remote peer BGP table
  - Possible Problem
    - Either remote speaker didn't advertise the withdraw or the local device didn't process the withdraw
    - EOR received

Missing Routes

RPL in IOS XR

- IOS and NX-OS by default install routes in the BGP table for prefixes learnt from eBGP peers
- IOS XR requires a mandatory RPL policy to have them installed in BGP table.
  - The RPL can permit all routes or conditional routes

```
route-policy Inbound-ROUTES
 if destination in A1-Prefix-Set then
    pass
 else
    drop
 endif
end-policy
router bgp 65530
neighbor-group IGW
  remote-as 65535
address-family ipv4 unicast
route-policy Inbound-ROUTES in
```

Missing Routes

BGP not in read-write mode

- May not see the routes in BGP table, incase BGP remains in read-only mode
  - To have the BGP routes installed, BGP should be in read-write mode
- On XR, use the below commands to verify BGP in read-write mode
  - **Show bgp**
  - **Show bgp process performance-statistics detail**
    - At the very bottom of this output, you will see the below lines, if the device entered the read-write mode

```
First neighbor established:  Jan 23 20:15:45
Entered DO_BESTPATH mode:    Jan 23 20:15:49
Entered DO_IMPORT mode:      Jan 23 20:15:49
Entered DO_RIBUPD mode:      Jan 23 20:15:49
Entered Normal mode:         Jan 23 20:15:49
Latest UPDATE sent:          Jan 23 20:18:39
```

Unexpected Routes

Route-Map Problem

```
route-map OSPF2BGP permit 10
 match ip prefix-list FILTERv4
!
router bgp 100
 address-family ipv4 unicast
   redistribute ospf 1 route-map OSPF2BGP
```

- What is the outcome of the above redistribution ?

Unexpected Routes

Route-Map Problem

```
route-map OSPF2BGP permit 10
 match ip prefix-list FILTERv4
route-map OSPF2BGP permit 20
 match ipv6 prefix-list FILTERv6
!
router bgp 100
 address-family ipv4 unicast
   redistribute ospf 1 route-map OSPF2BGP
 address-family ipv6 unicast
   redistribute ospfv3 1 route-map OSPF2BGP
```

- What is the outcome of the above redistribution ?

Unexpected Routes

Route-Map Problem

```
route-map OSPF2BGP permit 10
 match ip prefix-list FILTERv4
route-map OSPF2BGP permit 20
 match ipv6 prefix-list FILTERv6
!
router bgp 100
 address-family ipv4 unicast
   redistribute ospf 1 route-map OSPF2BGP
 address-family ipv6 unicast
   redistribute ospfv3 1 route-map OSPF2BGP
```

- What is the outcome of the above redistribution ?

Unexpected Routes

Route-Map Behavior

- A route map processes routes or IP packets in a linear fashion, that is, starting from the lowest sequence number.
- If referred policies (for example, prefix lists) within a match statement of a route-map entry return either a no-match or a deny-match, Device fails the match statement and processes the next route-map entry.
- Without any match statement in a route-map entry, the permission (permit or deny) of the route-map entry decides the result for all the routes or packets.

Troubleshooting Filtering

Topology

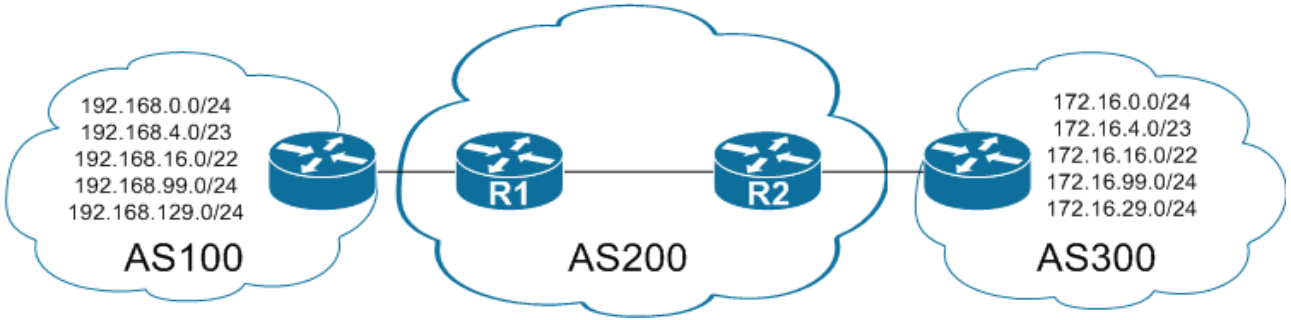


```
R2#show bgp ipv4 unicast
Network          Next Hop       Metric LocPrf Weight Path
*> 172.16.0.0/24    192.168.200.3       0             0 300 80 90 21003 2100 i
*> 172.16.4.0/23    192.168.200.3       0             0 300 1080 1090 1100 1110 i
*> 172.16.16.0/22   192.168.200.3       0             0 300 11234 21234 31234 i
*> 172.16.99.0/24   192.168.200.3       0             0 300 40 i
*> 172.16.129.0/24  192.168.200.3       0             0 300 10010 300 30010 30050 i
*>i192.168.0.0      10.12.1.1           0    100      0 100 80 90 21003 2100 i
*>i192.168.4.0/23   10.12.1.1           0    100      0 100 1080 1090 1100 1110 i
*>i192.168.16.0/22  10.12.1.1           0    100      0 100 11234 21234 31234 i
*>i192.168.99.0     10.12.1.1           0    100      0 100 40 i
*>i192.168.129.0    10.12.1.1           0    100      0 100 10010 300 30010 30050 i
```

Troubleshooting Filtering

Regex Query Modifiers

Modifier	Description
_ (Underscore)	Matches a space
^ (Caret)	Indicates the start of the string
$ (Dollar Sign)	Indicates the end of the string
[] (Brackets)	Matches a single character or nesting within a range
- (Hyphen)	Indicates a range of numbers in brackets
[^] (Caret in Brackets)	Excludes the characters listed in brackets
() (Parentheses)	Used for nesting of search patterns
| (Pipe)	Provides 'or' functionality to the query
. (Period)	Matches a single character, including a space
* (Asterisk)	Matches zero or more characters or patterns
+ (Plus Sign)	One or more instances of the character or pattern
? (Question Mark)	Matches one or no instances of the character or pattern.

Troubleshooting Filtering

Regex

```
R2#show bgp ipv4 unicast regexp _300_
! Output omitted for brevity
     Network          Next Hop      Metric LocPrf Weight Path
*> 172.16.0.0/24    192.168.200.3         0             0 300 80 90 21003 455 i
*> 172.16.4.0/23    192.168.200.3         0             0 300 878 1190 1100 1010 i
*> 172.16.16.0/22   192.168.200.3         0             0 300 779 21234 45 i
*> 172.16.99.0/24   192.168.200.3         0             0 300 145 40 i
*> 172.16.129.0/24  192.168.200.3         0             0 300 10010 300 1010 40 50 i
*>i192.168.129.0    10.12.1.1             0    100      0 100 10010 300 1010 40 50 i
```

```
R2#show bgp ipv4 unicast regexp ^300_
! Output omitted for brevity
     Network          Next Hop      Metric LocPrf Weight Path
*> 172.16.0.0/24    192.168.200.3         0             0 300 80 90 21003 455 i
*> 172.16.4.0/23    192.168.200.3         0             0 300 878 1190 1100 1010 i
*> 172.16.16.0/22   192.168.200.3         0             0 300 779 21234 45 i
*> 172.16.99.0/24   192.168.200.3         0             0 300 145 40 i
*> 172.16.129.0/24  192.168.200.3         0             0 300 10010 300 1010 40 50 i
```

Troubleshooting Filtering

Regex

```
R2#show bgp ipv4 unicast regexp [4-8]0_
! Output omitted for brevity
     Network          Next Hop    Metric LocPrf Weight Path
*> 172.16.0.0/24    192.168.200.3      0             0 300 80 90 21003 455 i
*> 172.16.99.0/24   192.168.200.3      0             0 300 145 40 i
*> 172.16.129.0/24  192.168.200.3      0             0 300 10010 300 1010 40 50 i
*>i192.168.0.0      10.12.1.1          0    100      0 100 80 90 21003 455 i
*>i192.168.99.0     10.12.1.1          0    100      0 100 145 40 i
*>i192.168.129.0    10.12.1.1          0    100      0 100 10010 300 1010 40 50 i
```

```
R2#show bgp ipv4 unicast regexp ^[13]00_[^3-8]
! Output omitted for brevity
     Network          Next Hop    Metric LocPrf Weight Path
*> 172.16.99.0/24   192.168.200.3      0             0 300 145 40 i
*> 172.16.129.0/24  192.168.200.3      0             0 300 10010 300 1010 40 50 i
*>i192.168.99.0     10.12.1.1          0    100      0 100 145 40 i
*>i192.168.129.0    10.12.1.1          0    100      0 100 10010 300 1010 40 50 i
```

Troubleshooting Filtering

Prefix-List Blocking Prefixes

```
XE-RTR#debug bgp ipv4 unicast updates in
BGP updates debugging is on (inbound) for address family: IPv4 Unicast

XE-RTR#clear bgp ipv4 unicast 10.1.45.4 soft in
! Output omitted for brevity
* 18:59:42.515: BGP(0): process 10.1.12.0/24, next hop 10.1.45.4, metric 0 from 10.1.45.4
* 18:59:42.515: BGP(0): Prefix 10.1.12.0/24 rejected by inbound filter-list.
* 18:59:42.515: BGP(0): update denied
```

```
NXOS5# debug bgp updates
NXOS5# clear bgp ipv4 unicast 10.1.45.4 soft in
! Output omitted for brevity
19:02:54 bgp: 300 [8449] UPD: [IPv4 Unicast] 10.1.45.4 Inbound as-path-list 1, action permit
 19:02:54 bgp: 300 [8449] UPD: [IPv4 Unicast] 10.1.45.4 Inbound as-path-list 1, action deny
 19:02:54 bgp: 300 [8449]  UPD: [IPv4 Unicast] Dropping prefix 10.1.12.0/24 from peer 10.1.45.4,
due to attribute policy rejected
```

Troubleshooting Filtering

IOS XR BGP RPL Debugging

```
route-policy R4-IN
  if destination in (10.0.0.0/8 le 32) then
    pass
  endif
  if destination in (172.16.0.0/12 le 32) then
    set med 20
  endif
end-policy
```

```
RP/0/0/CPU0:XR#debug bgp policy-execution events
RP/0/0/CPU0:XR#clear bgp ipv4 unicast 10.1.45.4 soft
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]: --Running policy 'R4-IN':---
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]:   Attach pt='neighbor-in-dflt'
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]:   Attach pt inst='default-IPv4-Uni-10.1.45.4'
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]: Input route attributes:
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]:   as-path: 200 100 600
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]:   as-path-length: 3
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]:   as-path-unique-length: 3
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]:   community: No Community Information
. . .
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]:   path-type: ebgp
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]:   aigp-metric: 0
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]:   validation-state: not-found
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]: Policy execution trace:
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]:   Condition: destination in (10.0.0.0/8 ...)
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]:   Condition evaluated to FALSE
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]:   Condition: destination in (172.16.0.0/12 ...)
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]:   Condition evaluated to FALSE
RP/0/0/CPU0: 06:19:10.000 : bgp[1053]:   End policy: result=DROP
```

Stale Routes

Symptoms and Possible Causes

- Symptoms
  - Stale Entry to BGP Peer
  - Traffic Black-Hole
  - Outage
- Possible Causes
  - BGP Slow Peer
  - Sender didn't process the updates
  - Receiver didn't process the update

Topology Reference

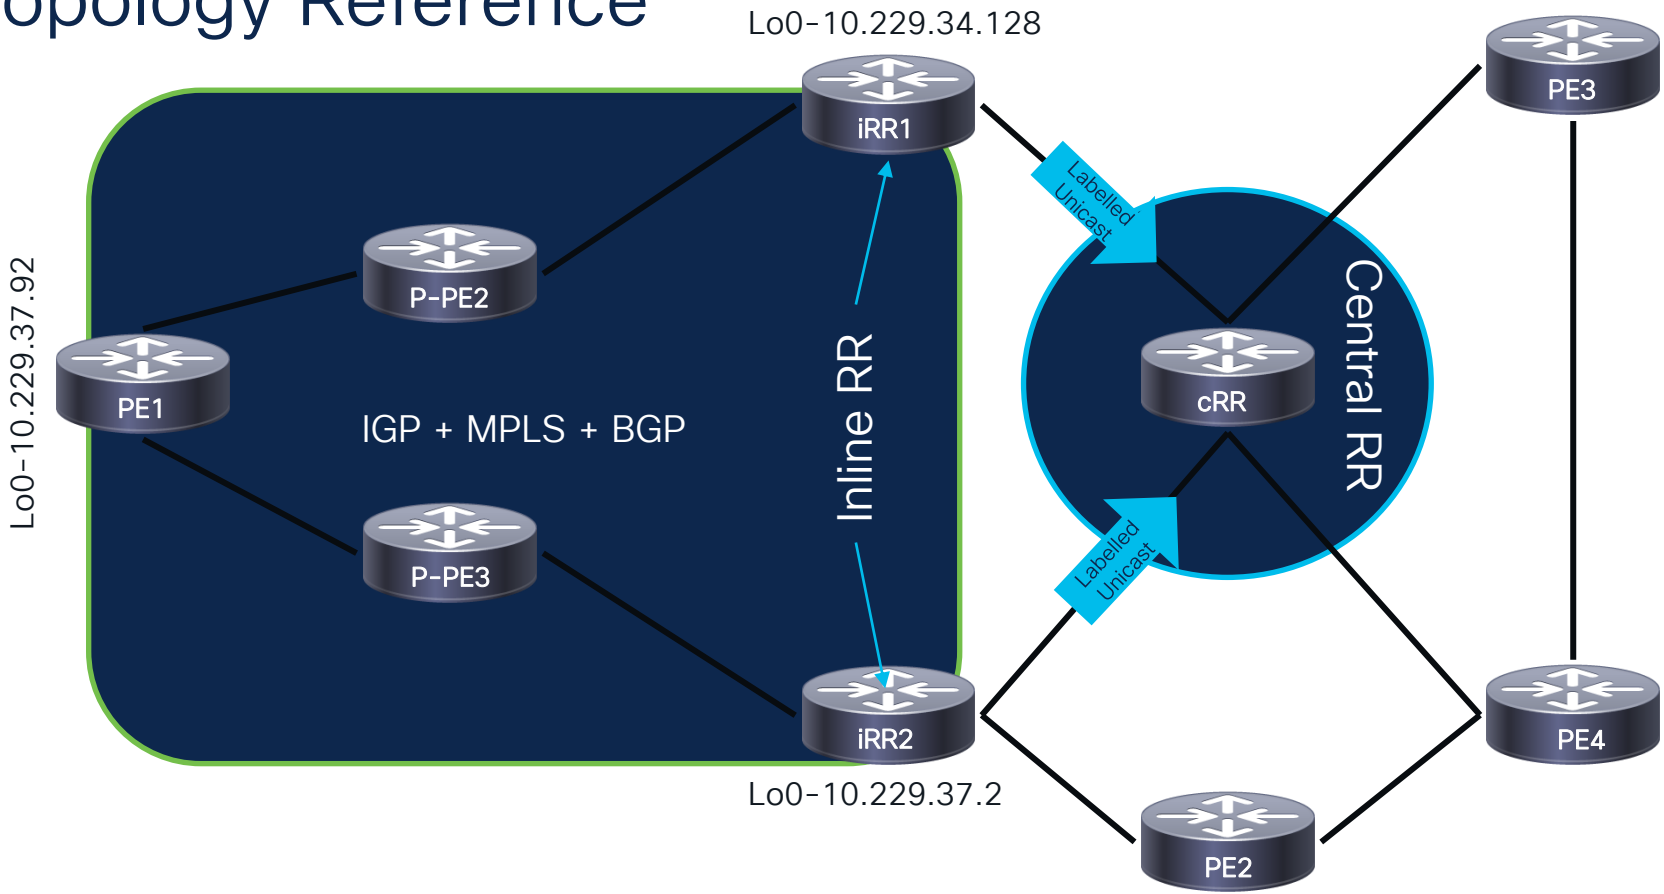

Stale Routes

Example – Route on BGP Speaker

```
RP/0/RSP0/CPU0:RR2#show bgp ipv4 labeled-unicast 10.229.37.92
BGP routing table entry for 10.229.37.92/32
  Local Label: 25528
Last Modified: Jan 13 10:20:52.424 for 11:45:15
Paths: (1 available, best #1)
   Path #1: Received by speaker 0
 Advertised to update-groups (with more than one peer):
  0.1 0.2 0.3 0.7
 Local
  10.229.34.128 (metric 5) from 192.168.53.9 (10.229.37.92)
    Received Label 26596
    Origin IGP, metric 0, localpref 100, valid, internal, best, group-best
    Received Path ID 1, Local Path ID 0, version 301642
    Community: 65080:109
    Originator: 10.229.37.92, Cluster list: 0.0.254.56, 10.229.34.128
```

Stale Routes

Example – Stale Entry on Receiving Router

```
Central-RR#show bgp ipv4 unicast 10.229.37.92
BGP routing table entry for 10.229.37.92/32, version 290518
BGP Bestpath: deterministic-med
Paths: (3 available, best #2, table default)
  Refresh Epoch 1
  Local, (Received from a RR-client)
    10.229.34.128 (metric 116) from 10.229.34.128 (10.229.34.128)
      Origin IGP, metric 0, localpref 100, valid, internal, best2
      Community: 65080:109
      Originator: 10.229.37.92, Cluster list: 10.229.34.128
      mpls labels in/out nolabel/26596
      rx pathid: 0x1A, tx pathid: 0x1
  Local, (Received from a RR-client)
    10.229.37.2 (metric 113) from 10.229.37.2 (10.229.37.2)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Community: 65080:109
      Originator: 10.229.37.92, Cluster list: 10.229.37.2
      mpls labels in/out nolabel/27183
      rx pathid: 0x7, tx pathid: 0x0
```

Stale Routes

How to Troubleshoot?

- On IOS, its difficult to get to the root cause after the problem has occurred.
  - Enable conditional debugs and wait for the issue to happen again
  - Reproduce the problem in lab environment (hard but not impossible)
- On IOS XR, use **show bgp trace** and **bgp debugs** to understand if the advertisement has been sent/received or not
  - Debug
- On NX-OS, use **show bgp event-history events | errors** to figure out if the prefix has been received / advertised or not

Stale Routes or Missing Routes / Advertisements

Conditional Debugs

```
IOS-1#show access-list 99
Standard IP access list 99
     permit 10.1.1.0 0.0.0.255

IOS-1#debug ip bgp 2.2.2.2 update 99
```

```
IOS-XR
route-policy DEBUG_BGP
 if destination in BGP_PREFIX then
  pass
 else
  drop
 endif
end-policy
 prefix-set BGP_PREFIX
  100.1.1.0/24
 end-set
debug bgp update ipv4 unicast [in | out] route-policy DEBUG_BGP
```

BGP Route Churn and Troubleshooting with BGP Table Version

cisco Live!

Route Churn

Symptom - High CPU?

```
Router#show process cpu
  CPU utilization for five seconds: 100%/0%; one minute: 99%; five minutes: 81%
....
139     6795740   1020252       6660 88.34% 91.63% 74.01%    0 BGP Router
```

- Define “High”
  - Know what normal CPU utilization is for the router in question
  - Is the CPU spiking due to “BGP Scanner” or is it constant?
- Look at the scenario
  - Is BGP going through “Initial Convergence”?
- If not then route churn is the usual culprit
  - Illegal recursive lookup or some other factor causes bestpath changes for the entire table

Route Churn

High CPU due to BGP Router

- How to identify route churn?
  - Do "**`sh ip bgp summary`**", note the table version
  - Wait 60 seconds
  - Do "**`sh ip bgp summary`**", compare the table version from 60 seconds ago
- You have 150k routes and see the table version increase by 300
  - This is probably normal route churn
  - Know how many bestpath changes you normally see per minute
- You have 150k routes and see the table version fluctuating by 20K – 50k
  - This is bad and is the cause of your high CPU

Route Churn

```
Router#Show ip bgp all sum | in tab
BGP table version is 936574954, main routing table version 936574954
BGP table version is 429591477, main routing table version 429591477
Router#


Router#Show ip bgp all sum | in tab
BGP table version is 936576768, main routing table version 936575068
BGP table version is 429591526, main routing table version 429591526
Router#

Router#show ip route | in 00:00:0
B       187.164.0.0 [200/0] via 218.185.80.140, 00:00:00
B       187.52.0.0 [200/0] via 218.185.80.140, 00:00:00
B       187.24.0.0 [200/0] via 218.185.80.140, 00:00:00
B       187.68.0.0 [200/0] via 218.185.80.140, 00:00:00
B       186.136.0.0 [200/0] via 218.185.80.140, 00:00:00
. . . .
```

Over 1800 prefixes flapped

< 4 seconds later

Route Churn

Table Version Changes?

- What causes massive table version changes?
- Flapping peers
  - Hold-timer expiring?
  - Corrupt UPDATE?
- Route churn
  - Don’t try to troubleshoot the entire BGP table at once
  - Identify one prefix that is churning and troubleshoot that one prefix
  - Will likely fix the problem with the rest of the BGP table churn

Route Churn

Flapping Routes in BGP

- Figuring out flapping routes from routing table is easy (even in vrf)
  - **`Show ip route vrf * | in 00:00:0|VRF`**
- How about identifying flapping routes on the VPNv4 Route Reflector?
  - **`Show bgp vpnv4 unicast all summary | in table`**
  - Use the table version as the marker in the below command to see the routes which flapped after the last command that was executed
  - **`Show bgp vpnv4 unicast all version`** [*`version-num`* | **`recent`** *`version-num`*]
  - Use the next-hop of the prefixes from the above command, to see why the prefixes are flapping

Route Churn

Flapping Routes in BGP

```
R1#show bgp ipv4 unicast version recent 6
BGP table version is 12, local router ID is 192.168.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
     Network          Next Hop            Metric LocPrf Weight Path
r>i 192.168.2.2/32   192.168.2.2              0    100      0 i
r>i 192.168.3.3/32   192.168.3.3              0    100      0 i
*mi 192.168.200.200/32
                     192.168.3.3              0    100      0 200 i
*>i                  192.168.2.2              0    100      0 200 i
```

Route Churn

Flapping Routes in BGP on IOS XR

- IOS XR has more interesting command for table version updates
  - **Show bgp** *afi safi* **version** *<start-version> <end-version>*

```
RP/0/0/CPU0:XR1#show bgp ipv4 unicast version 5 7
VRF: default
------------
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop            Metric LocPrf    Version Path
*>i192.168.2.2/32     192.168.2.2              0    100          6
i*>i192.168.3.3/32     192.168.3.3             0    100          7
i*>i192.168.200.200/32 192.168.2.2             0    100          5 200 i
i                      192.168.3.3             0    100          5 200 i

Processed 3 prefixes, 4 paths
```

Route Churn

Which AFI?

- If there are too many updates coming onto the router, one way to identify it would be
  - **`Show ip traffic | section TCP`**
- Symptom – TCP traffic increasing rapidly, but table version for IPv4 and VPNv4 AFI is only increasing by 200 or 300 or a smaller value
- Check for different AFI's enabled on the router and checking for the table version changes in those AFI's
  - Especially IPv6 or VPNv6 as those can have more impact with fewer prefixes flapping

Embedded Event Manager (EEM)

- Serves as a powerful tool for high CPU troubleshooting
- Triggered based on event and thresholds
- Multiple actions can be set based on events

```
event manager applet HIGHCPU
event snmp oid "1.3.6.1.4.1.9.9.109.1.1.1.1.3.1" get-type exact entry-op gt entry-val "90"
exit-op lt exit-val "70" poll-interval 5 maxrun 200
action 1.0 syslog msg "START of TAC-EEM: High CPU"
action 1.1 cli command ”show clock"
action 1.3 cli command ”show ip bgp all summary | append disk0:proc_CPU"
action 2.0 cli command "sh clock | append disk0:proc_CPU"
action 2.1 cli command "show process cpu sorted | append disk0:proc_CPU“
action 2.2 cli command "show proc cpu history | append disk0:proc_CPU"
action 2.3 cli command " show ip bgp all summary | append disk0:proc_CPU"
action 3.1 cli command "show log | append disk0:proc_CPU"
action 4.0 syslog msg "END of TAC-EEM: High CPU"
```

Troubleshooting with NX-OS

Troubleshooting with NX-OS

BGP Event-History

- NX-OS event-history capability is alternate to running debugs
- Event-History Buffer Sizes:
  - Large
  - Medium
  - Small
- Event-History maintained for:
  - Events
  - Errors
  - Detail
  - Msgs
  - CLI

Troubleshooting with NX-OS

Processing an Incoming Update – show bgp event-history detail

- Manually enable Detail Event-History using the command "**event-history detail size [large | medium | small]**"

```
05:28:12.515623: (default) UPD: Received UPDATE message from 10.1.23.2
05:28:12.515616: (default) BRIB: [IPv4 Unicast] (192.168.1.1/32 (10.1.23.2)): returning from
bgp_brib_add, new_path: 0, change: 0, undelete: 0, history: 0, force: 0, (pflags=0x28), reeval=0
05:28:12.515608: (default) BRIB: [IPv4 Unicast] 192.168.1.1/32 from 10.1.23.2 was already in BRIB
with same attributes
05:28:12.515600: (default) BRIB: [IPv4 Unicast] (192.168.1.1/32 (10.1.23.2)): bgp_brib_add:
handling nexthop
05:28:12.515593: (default) BRIB: [IPv4 Unicast] Path to 192.168.1.1/32 via 192.168.2.2 already
exists, dflags=0x8001a
05:28:12.515580: (default) BRIB: [IPv4 Unicast] Installing prefix 192.168.1.1/32 (10.1.23.2) via
10.1.23.2  into BRIB with extcomm
05:28:12.515557: (default) UPD: [IPv4 Unicast] Received prefix 192.168.1.1/32 from peer
10.1.23.2, origin 0, next hop 10.1.23.2, localpref 0, med
005:28:12.515524: (default) UPD: 10.1.23.2 Received attr code 2, length 10, AS-Path: <200 100 >
05:28:12.515503: (default) UPD: Attr code 3, length 4, Next-hop: 10.1.23.2
05:28:12.515454: (default) UPD: Attr code 1, length 1, Origin: IGP
05:28:12.515446: (default) UPD: 10.1.23.2 parsed UPDATE message from peer, len 52 , withdraw len
0, attr len 24, nlri len 5
```

Troubleshooting with NX-OS

Update Generation – show bgp event-history detail

```
05:28:11.478903: (default) UPD: [IPv4 Unicast] 10.1.23.2 Created UPD msg (len 52) with prefix
192.168.1.1/32 ( Installed in HW) path-id 1 for peer
05:28:11.478886: (default) UPD: 10.1.23.2 Sending attr code 3, length 4, Next-hop: 10.1.23.3
05:28:11.478880: (default) UPD: 10.1.23.2 Sending attr code 2, length 10, AS-Path: <300 100 >
05:28:11.478870: (default) UPD: 10.1.23.2 Sending attr code 1, length 1, Origin: IGP
05:28:11.478856: (default) UPD: [IPv4 Unicast] consider sending 192.168.1.1/32 to peer 10.1.23.2,
path-id 1, best-ext is off
. . .
05:28:11.478717: (default) EVT: [IPv4 Unicast] soft refresh out completed for 1 peers
05:28:11.478690: (default) EVT: [IPv4 Unicast] Adding peer 10.1.23.2 for update gen
05:28:11.478686: (default) BRIB: [IPv4 Unicast] Group setting SRM for dest 192.168.3.3/32
05:28:11.478682: (default) BRIB: [IPv4 Unicast] Group setting SRM for dest 192.168.2.2/32
05:28:11.478678: (default) BRIB: [IPv4 Unicast] Group setting SRM for dest 192.168.1.1/32
05:28:11.478666: (default) EVT: [IPv4 Unicast] 1 peer(s) being soft refreshed out
05:28:11.478661: (default) EVT: [IPv4 Unicast] 10.1.23.2 [peer index 2]
05:28:11.478638: (default) EVT: [IPv4 Unicast] Doing soft out BGP table walk for peers
05:28:10.478332: (default) EVT: [IPv4 Unicast] Scheduling peer 10.1.23.2 for soft refresh out
05:28:10.478321: (default) EVT: Received ROUTEREFRESH message from 10.1.23.2
```

Troubleshooting with NX-OS

Conditional Debugging and URIB

- Conditional Debugging

```
debug logfile bgp
debug bgp events updates rib brib import
debug-filter bgp vrf vpn1
debug-filter bgp address-family ipv4 unicast
debug-filter bgp neighbor 10.1.202.2
debug-filter bgp prefix 192.168.2.2/32
```

- Troubleshooting URIB

```
Show routing internal event-history ufdm
Show routing internal event-history ufdm-summary
Show routing internal event-history recursive
```

Troubleshooting with NX-OS

Route Policy Manager

- Route-map functionality is provided by a new process in DC-OS called Route Policy Manager (RPM)
- RPM handles route-maps, AS path access lists, community lists and prefix lists
- The route-maps are configured the same way as they are configured in Cisco IOS, but are managed by RPM
  - If there are any issues seen with route-maps not functioning

Troubleshooting with NX-OS

Route Policy Manager

```
NX-1# show system internal sysmgr service name rpm
Service "rpm" ("rpm", 203):
        UUID = 0x131, PID = 5265, SAP = 348
        State: SRV_STATE_HANDSHAKED (entered at time Mon Jan
30 03:07:59 2017).
        Restart count: 1
        Time of last restart: Mon Aug 22 03:07:57 2016.
        The service never crashed since the last reboot.
        Tag = N/A
        Plugin ID: 1
```

Troubleshooting with NX-

Route Policy Manager

```
template peer-policy PP-Test1
   send-community
   route-map RM-Test1 out
!
neighbor 192.168.2.2 remote-as 65000
     inherit peer-session ps-ebgp-peer-to-
mpls-core
       address-family ipv4 unicast
       inherit peer-policy PP-Test1 5
       send-community
       prefix-list pl-nab-core-devl-routes in
       no prefix-list pl-cloud-routes out
       route-map RM-Test2 out
       soft-reconfiguration inbound
. . . .
```

```
NX-1#  sh route-map RM-Test1
route-map RM-Test1, permit, sequence 10
  Match clauses:
    ip address prefix-lists: sy3-routes
  Continue: sequence 20
  Set clauses:
    community 65135:999
route-map RM-Test1, permit, sequence 999
  Match clauses:
  Set clauses:
!
NX-1# sh route-map RM-Test2
route-map RM-Test1, permit, sequence 10
  Match clauses:
    ip address prefix-lists: pl-cloud-routes
  Set clauses:
route-map RM-Test1, permit, sequence 20
  Match clauses:
    as-path (as-path filter): as-me1-o365-ext-
routes
  Set clauses:
```

Troubleshooting with NX-OS

```
NX-2# show system internal rpm route-map
Policy name: RM-Test1                Type: route-map
Version: 6                           State: Ready
Ref. count: 1                        PBR refcount: 0
Stmt count: 5                        Last stmt seq: 999
Set nhop cmd count: 0                Set vrf cmd count: 0
Set intf cmd count: 0                Flags: 0x00000003
PPF nodeid: 0x00000000               Config refcount: 0
PBR Stats: No
Clients:
    bgp-65136 (Route filtering/redistribution)    ACN version: 0
```

Troubleshooting with NX-OS

```
# show system internal rpm event-history rsw

Routing software interaction logs of RPM
1) Event:E_DEBUG, length:88, at 96760 usecs after Sun Apr 23 22:19:12 2017
    [120] [3959]: Bind ack sent - client bgp-65136 uuid 0x0000011b for policy RM-Test2  <<<<<  Outbound
route-map bound to BGP client
2) Event:E_DEBUG, length:83, at 96717 usecs after Sun Apr 23 22:19:12 2017
    [120] [3959]: Bind request - client bgp-65136 uuid 0x0000011b policy RM-Test2
3) Event:E_DEBUG, length:88, at 782159 usecs after Sun Apr 23 21:51:06 2017
    [120] [3959]: Bind ack sent - client bgp-65136 uuid 0x0000011b for policy RM-Test2
<snip>
    [120] [3959]: UnBind request succesfull - client bgp-65136 policy RM-Test1 <<<<<Unbind for route-map
referenced in peer-policy
6) Event:E_DEBUG, length:99, at 781950 usecs after Sun Apr 23 21:51:06 2017
    [120] [3959]: UnBind request - client bgp-65136 uuid 0x0000011b policy RM-Test1
7) Event:E_DEBUG, length:102, at 344591 usecs after Sun Apr 23 21:47:39 2017
    [120] [3959]: Bind ack sent - client bgp-65136 uuid 0x0000011b for policy RM-Test1 <<<<< Route-map
referenced in peer-policy
8) Event:E_DEBUG, length:97, at 344557 usecs after Sun Apr 23 21:47:39 2017
    [120] [3959]: Bind request - client bgp-65136 uuid 0x0000011b policy RM-Test1
```

Troubleshooting with NX-OS

Route Policy Manager

- Use RPM Event-history when troubleshooting any misbehavior of route policy / redistribution / missing routes / routes not learnt
- In case of issues, collect "**`show tech rpm`**"
- Use the below commands to troubleshoot RPM issues
  - `Show system internal rpm event-history events (For RPM Events)`
  - `Show system internal rpm event-history errors  (For errors with RPM)`
  - `Show system internal rpm event-history rsw  (RPM Interaction with RPM software)`
  - `Show system internal rpm event-history msgs (RPM Message logs)`
  - `Show system internal rpm event-history trace (RPM Traces)`

Cisco Learning and Certifications

From technology training and team development to Cisco certifications and learning plans, let us help you empower your business and career. www.cisco.com/go/certs

Pay for Learning with Cisco Learning Credits

(CLCs) are prepaid training vouchers redeemed directly with Cisco.

Learn

Cisco U.
IT learning hub that guides teams and learners toward their goals

Cisco Digital Learning
Subscription-based product, technology, and certification training

Cisco Modeling Labs
Network simulation platform for design, testing, and troubleshooting

Cisco Learning Network
Resource community portal for certifications and learning

Train

Cisco Training Bootcamps
Intensive team & individual automation and technology training programs

Cisco Learning Partner Program
Authorized training partners supporting Cisco technology and career certifications

Cisco Instructor-led and Virtual Instructor-led training
Accelerated curriculum of product, technology, and certification courses

Certify

Cisco Certifications and Specialist Certifications
Award-winning certification program empowers students and IT Professionals to advance their technical careers

Cisco Guided Study Groups
180-day certification prep program with learning and support

Cisco Continuing Education Program
Recertification training options for Cisco certified individuals

Here at the event? Visit us at **The Learning and Certifications lounge at the World of Solutions**

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

cisco Live!

Thank you


cisco Live!
ALL IN