Monitoring and Debugging

This chapter offers a comprehensive guide to BGP labeled unicast, detailing its various implementations, including over RSVP-TE for traffic engineering and resiliency. It also covers the exclusion of label allocation for non-advertised routes and steering BGP control-plane traffic over IP-only paths to optimize MPLS network operations.

- BGP-RIB feedback mechanisms for update generation, on page 1
- Enhanced monitoring of NSR statistics, on page 3
- Store and analyze changes in the prefixes received from BGP peer, on page 6
- Monitor BGP memory statistics, on page 9
- Enhanced monitoring of BGP keepalive messages, on page 11
- Enhanced monitoring of BGP memory utilization, on page 16
- Verify enhanced next hop monitoring, on page 18
- Enhanced monitoring of version-rate statistics, on page 21

BGP-RIB feedback mechanisms for update generation

The BGP-RIB feedback mechanisms for update generation are BGP route control mechanisms that

- ensure routes are installed locally before advertising them to neighbors
- track which route versions the forwarding information base (FIB) has consumed, and
- send updates only for routes confirmed installed in the FIB to prevent premature advertisements.

These mechanisms help you avoid traffic loss caused by premature route advertisements after events like router reloads, line card online insertion and removal (LC OIR), or link flaps when alternate paths become available

You configure BGP to wait for routing information base (RIB) feedback before sending updates by using the **update wait-install** command in the `router address-family IPv4` or `router address-family VPNv4` configuration mode. This command ensures that BGP sends updates only after routes are confirmed installed in the forwarding information base (FIB), preventing premature route advertisements.

You can verify this configuration using the following commands:

- **show bgp**
- **show bgp neighbors**

- **show bgp process performance-statistics**

This configuration helps you avoid traffic loss caused by premature route advertisements after events such as router reloads, line card online insertion and removal (LC OIR), or link flaps when alternate paths become available.

Guidelines for BGP-RIB feedback mechanisms

To prevent traffic loss and ensure reliable routing, always advertise BGP routes only after they are confirmed as installed in the Forwarding Information Base (FIB) via the BGP-RIB feedback mechanism.

- ensure that BGP installs routes in the Routing Information Base (RIB) and waits for feedback from the RIB about installation in the FIB.
- confirm that the RIB tracks which route versions are in the FIB using the BCDL feedback mechanism.
- send BGP update messages only for routes confirmed as installed in the FIB, preventing premature advertisements that could cause packet loss or blackholing.

Configure BGP to wait for RIB feedback before sending updates

Enable BGP to delay advertising updates until routes are confirmed as installed in the FIB, preventing premature updates and possible traffic loss.

Use this configuration to enhance BGP routing reliability by ensuring updates are sent only after successful RIB-to-FIB installation confirmation.

Before you begin

- Verify you have administrator access to the router.
- Identify the AS number and desired address family (e.g., IPv4 unicast, VPNv4).

Follow these steps to configure BGP to wait for RIB feedback:

Procedure

Step 1 Enter router configuration mode for the desired address family.

Example:

```
Router# configure
Router(config)# router bgp 1
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# update wait-install
```

Step 2 Save the configuration.

Example:

```
Router#(config-bgp-af)# commit
```

Step 3 Use the **show bgp process** command to view the delay of the BGP process update since the last router reload.

Example:

```
Router# show bgp process
Wed Aug 24 00:40:48.649 PDT

BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System number format: ASPLAIN
Autonomous System: 100
Router ID: 192.168.0.2 (manually configured)
Default Cluster ID: 192.168.0.2
Active Cluster IDs:  192.168.0.2
------------------------------
------------------------------
Update wait-install enabled:
  ack request 2, ack rcvd 2, slow ack 0
  startup delay 10 secs
```

BGP on the router now delays advertising route updates until RIB confirms that routes are installed in the FIB, ensuring reliable route propagation and preventing traffic loss.

What to do next

Monitor BGP updates and verify network stability after applying the configuration.

Enhanced monitoring of NSR statistics

The Enhanced monitoring of non-static routing (NSR) statistics is a network monitoring feature that

- provides detailed packet processing metrics during critical network operations,
- tracks BGP update processing speed, traffic volume, and packet sequence number integrity, and
- enables early detection of packet loss or sequencing inconsistencies to maintain routing reliability and network stability.

This feature helps you ensure continuous and reliable BGP routing by synchronizing state information between primary and standby routing engines during software upgrades or failover events. It gives you real-time insights into NSR activities so you can proactively address issues before they impact network performance.

NSR synchronization and monitoring for network stability and failover

- NSR synchronizes routing and forwarding data between primary and standby routing engines to minimize disruption during failover.
- Monitoring NSR statistics enables proactive identification and resolution of potential BGP routing issues.
- The feature supports comprehensive data analysis during pivotal network events, enhancing overall network stability and reliability.

Table 1: Feature History Table

Feature Name	Release Name	Description
Enhanced Monitoring of NSR Statistics	Release 24.2.1	You can maintain uninterrupted network functionality during upgrades or failovers with Non-Stop Routing (NSR), ensuring consistent data across primary and standby engines. The Enhanced Monitoring of NSR Statistics feature offers metrics on NSR packet handling, providing processing times, counts, and sequence numbers in real-time. If no new packets are received, the last known statistics persist, keeping the displayed data current. CLI: The feature modifies the output of the show command given below: • **show bgp nsr** YANG Data Model: • `Cisco-IOS-XR-ipv4-bgp-oper` (see GitHub, YANG Data Models Navigator)

View enhanced NSR statistics

Obtain and analyze key metrics to assess the efficiency and robustness of BGP operations with NSR functionality.

Use this task to monitor BGP update processing speed, packet volume over intervals, and sequence number integrity to prevent packet loss and maintain network stability.

Before you begin

Ensure you have access to the device CLI and necessary permissions to run show commands.

Follow these steps to view enhanced monitoring of NSR statistics:

Procedure

Verify the key metrics to evaluate BGP operations with NSR. Monitor BGP update processing speed, packet volume over intervals, and sequence number integrity to prevent packet loss and ensure network stability.

Example:

```
Router# show bgp nsr

Fri Jan 30 10:18:48.171 PST PDT

BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System: 100
Router ID: 10.1.0.1 (manually configured)
Default Cluster ID: 10.1.0.1
Active Cluster IDs:  10.1.0.1
Fast external fallover enabled
Neighbor logging is not enabled
Enforce first AS enabled
AS Path ignore is enabled
AS Path multipath-relax is enabled
Default local preference: 100
Default keepalive: 60
Graceful restart enabled
Restart time: 180
Stale path timeout time: 360
RIB purge timeout time: 600
Non-stop routing is enabled
Update delay: 120
Generic scan interval: 60

Address family: IPv4 Unicast
Dampening is not enabled
Client reflection is enabled in global config
Scan interval: 60
Main Table Version: 7034
IGP notification: IGPs notified
RIB has converged: version 1

========== Post Failover Summary for Active instance ==========

Node                Process            Read      Write    Inbound


node0_0_CPU0        Speaker          146.75      18.90       3.46

  Entered mode  Standby Ready              : Jan 30 10:00:39
  Entered mode  TCP NSR Setup              : Jan 30 10:00:39
  Entered mode  TCP NSR Setup Done         : Jan 30 10:00:39
  Entered mode  TCP Initial Sync           : Jan 30 10:00:39
  Entered mode  TCP Initial Sync Done      : Jan 30 10:00:44
  Entered mode  FPBSN processing done      : Jan 30 10:00:44
  Entered mode  Update processing done     : Jan 30 10:00:44
  Entered mode  BGP Initial Sync           : Jan 30 10:00:44
  Entered mode  BGP Initial Sync done      : Jan 30 10:00:44
  Entered mode  NSR Ready                  : Jan 30 10:00:44

Current BGP NSR state - NSR Ready achieved at: Jan 30 10:00:44
NSR State READY notified to Redcon at: Jan 30 10:16:58

NSR Post Failover Summary:

NPL Statistics:

Messages Sent: 384985 .           ACKS Received: 384985 :8
Messages Sent: 8                  ACKS Sent: 8
Send failures: 11541 .            Send ACK Failures:0
Suspends:      11541              Resumes: 11407
Messages Processed: 8             Out of sequence drops: 8
Messages Send Drops: 0            Messages Recv Drops: 0
```

```
Sync Send Timeouts: 8

NPL Packet Processing Statistics:

  Interval       End-Time        Avg Proc  Num of        seq num
   (sec)                         time(us)   pkts     [start - end]

    30     Aug 22 23:08:11.142     140         2     [     74 - 75     ]
    60     Aug 22 23:08:11.142     233         4     [     72 - 75     ]
   180     Aug 22 23:08:11.142     426        22     [     54 - 75     ]


QAD Statistics:

  Messages Sent       : 512           ACKs Received       : 512
  Messages Received   : 8             ACKs Sent           : 8
  Send Failures       : 1             Send ACK Failures   : 0
  Suspends            : 1             Resumes             : 1
  Messages Processed  : 8             Out of sequence drops: 0
Postit Summary:

  Total pending postit messages:  0
  Neighbors with pending postits: 0

Conv  Bestpath   TunnelUpd  Import     RIBUpd     Label      ReadWrite  LastUpd
Process: Speaker

Yes   120        ---        ---        120        120        120        87531

Rib Trigger: enabled
Last RIB down event  Jan 29 09:50:03.069 received
Last RIB convergence Jan 29 09:50:03.069 last ack received.

Address Family IPv4 Unicast converged in 87531 seconds
```

Review the output to conform the detailed metrics on BGP NSR packet handling over specific time intervals.

You obtain detailed NSR health data to prevent packet loss and maintain reliable network operations during critical events.

Store and analyze changes in the prefixes received from BGP peer

The store and analyze changes in the prefixes received from BGP peer is a serviceability feature that

- allow routers to handle analysis of millions of BGP paths across IPv4 and IPv6 address families,
- enables operators to actively monitor and analyze changes, acceptances, and rejections of prefixes received from BGP peers by providing detailed statistics, and
- allows you to store all original copies of routes received from peers, including those not selected as the best path.

This feature helps you improve serviceability by monitoring BGP operations and facilitating debugging in both production and lab environments.

How you use this feature

- Soft reconfiguration stores incoming prefixes before applying policies if the peer does not support route refresh; using the **always** keyword forces storage even when route refresh is supported.
- The **soft reconfiguration inbound always** command enables storage of all updates from a specified neighbor.
- The **soft reconfiguration inbound** command stores the original unmodified route alongside modified or filtered routes, allowing you to perform a "soft clear" after inbound policy changes.
- Prefixes fall into three categories: accepted and unmodified, accepted and modified, or denied.
- Operators can monitor the impact of inbound policies without network disruption, helping to maintain stability and facilitate troubleshooting.

Verify the BGP prefix statistics

Ensure that you verify the soft reconfiguration statistics and evaluate the impact of inbound policies on IPv4 unicast BGP sessions to maintain proper routing behavior and policy compliance.

Before you begin

- Confirm that BGP is configured and that soft reconfiguration is enabled on the relevant neighbors.
- Verify that you know the IP addresses of BGP neighbors and the inbound policies applied.
- Be aware that soft reconfiguration is memory intensive; ensure the router has sufficient resources.
- Understand that the `soft-reconfiguration inbound` command stores updates from neighbors, enabling soft resets even if the neighbor does not support route refresh.

Procedure

Step 1 Use the **show bgp ipv4 unicast summary soft-reconfig-stats**command to verify the soft reconfiguration statistics for IPv4 unicast BGP sessions.

Example:

```
Router# show bgp ipv4 unicast summary soft-reconfig-stats

....
.....

Neighbor     Spk  AS  MsgRcvd MsgSent TblVer InQ  OutQ  Up/Down     St/PfxRcd  SoftChgd  Denied
10.10.10.4    0   3    15       12      0    0     0    00:46:06 2           0       0
Total                                                             2           0       0

Legend:

Total PfxRcd: Sum of accepted unmodified and modifed paths
Total SoftChgd: Sum of accepted modified paths
Total Denied: Sum of Denied paths
```

This output displays the soft reconfiguration statistics for IPv4 unicast BGP sessions given below:

- Total PfxRcd represents the sum of accepted unmodified and modified paths.

- Total SoftChgd represents the sum of accepted modified paths.
- Total Denied represents the sum of denied paths.

Step 2 Use the **show bgp ipv4 unicast neighbors dryrun-policy pass** command to verify the dry run policy impact on inbound policy.

Example:

```
Router# show bgp ipv4 unicast neighbors 10.10.10.1 dryrun-policy pass

Sat Oct 14 01:22:02.946 EDT

Policy Statistics
-----------------
     AFI:                          IPv4 Unicast
     Direction:                    Inbound
     In-use Policy:                pass
     Dry-run Policy:               pass
     Remote-as:                    300
     Total Networks walked:        257
     Total Paths walked:           72257
     Dry Run elapsed time(ms):     8
-------------------------------------------------------------------------------------------
                                       Dry-run-Policy      In-use-Policy      Delta
-------------------------------------------------------------------------------------------
Neighbor: 10.10.10.1
     Accepted Unmodified:              257                 257                0
     Accepted Modified:                0                   0                  0
     Pre-inbound policy copy:          0                   0                  0
     Denied:                           0                   0                  0
     Estimated Total Paths Memory:     26.10KB             26.10KB            0.00
-------------------------------------------------------------------------------------------
    /* The values in the table provides information confirms that the BGP session with the
       neighbor 10.10.10.1 is passing the dry run policy.
       The values indicates that the BGP updates from the neighbor comply with the specified policies

      without actually applying the policies. The values in the table provides insight into potential

       routing changes without committing the policies. */
```

The table values confirm that the BGP session with neighbor 10.10.10.1 is successfully passing the dry run policy. This indicates that BGP updates from the neighbor comply with the specified policies without actually applying them. These values provide insight into potential routing changes without committing the policies, enabling you to evaluate policy impact safely.

Step 3 Use the **show bgp scale detail** command to verify the statistics on configured and established neighbors, and address-family prefixes, paths, and memory usage.

Example:

```
Router# show bgp scale detail
Fri Feb 2 12:49:38.349 EST

VRF: default
Neighbors Configured: 2 Established: 2

Address-Family  Prefixes  Paths  PathElem  Prefix       Path      PathElem
                                           Memory       Memory    Memory

IPv4 Unicast    3         5      3         564.00       520.00    369.00

SoftReconfig Changed  1  104.00 ---> This field shows that soft reconfiguration has been enabled. It
```

```
 also displays the number of prefixes that were accepted and modified, and the amount of memory
consumed by the prefix.
------------------------------------------------------------------------------
Total           3          5        3      564.00       520.00    369.00

Total VRFs Configured: 0
```

The field SoftReconfig Changed 1 104.00 indicates that soft reconfiguration is enabled. It also shows the number of prefixes that were accepted and modified, along with the memory consumed by these prefixes, providing insight into resource usage related to soft reconfiguration.

Monitor BGP memory statistics

The BGP memory statistics is a serviceability feature that:

- periodically monitors memory usage by the BGP process every 60 seconds,
- logs any memory changes exceeding 1% of the configured resource limit (rlimit), and
- triggers syslog notifications at 85%, 90%, and 95% of the rlimit to alert administrators before critical memory exhaustion.
- The **show bgp memory history** command is a diagnostic tool that displays memory usage trends and variations over time, allowing administrators to analyze historical memory consumption patterns.
- Syslog notifications are automated alerts generated when memory usage reaches predefined thresholds, enabling proactive management of memory resources before critical limits are reached.
- A serviceability feature is a functionality designed to maintain system health by providing monitoring, alerting, and diagnostic capabilities that support troubleshooting and capacity planning.

The **show bgp memory history** command displays detailed historical memory usage information, including total memory in megabytes, percent of rlimit used, memory differences between records, and counts of networks, paths, path elements, and attributes for the default VRF.

Monitor BGP Memory Statistics

This task instructs you to check the historical memory usage of the BGP process to monitor memory trends and manage resources proactively.

Procedure

Step 1 Use the **show bgp memory history** command on the active router to verify the BGP memory history:

Example:

```
Router# show bgp memory history

History of memory changes recorded for a threshold greater than 1.0% of rlimit.
Last shown record displays current values.
Network information for default VRF.
```

```
Time              Memory(MB)  Rlimit(%)  Memory diff(MB)  Networks  Paths     PathElems  Attributes

Oct  2 16:30:37    152         1          152              400       400       400        9
Oct  2 16:31:37    343         4          191              396952    396869    396952     725
Oct  2 16:32:37    425         5          81               524567    513979    524567     8408
Oct  2 16:42:38    741         9          316              1178605   1241533   1178604    10753
Oct  2 16:43:38    985         12         243              1778234   1859254   1778234    11214
Oct  2 19:42:39    901         11         -84              1800688   678607    1800688    10911
Oct  2 19:45:39    766         9          -136             1332259   688784    1332259    10943
```

The show command output displays memory change history for thresholds above 1.0% of the resource limit, with the latest record reflecting current values on the active router. It includes network details for the default VRF such as timestamps, memory in MB, rlimit percentage, memory differences, and counts of networks, paths, path elements, and attributes.

Step 2 Use the **show bgp memory history standby** command on the standby router to verify the BGP memory history:

Example:

```
Router# show bgp memory history standby
Sat Mar  2 00:26:46.874 UTC

History of memory changes recorded for a threshold greater than 1.0% of rlimit.
Last shown record displays current values.
Network information for default VRF.

Time              Memory(MB)  Rlimit(%)  Memory diff(MB)  Networks  Paths     PathElems  Attributes

Feb  9 03:39:04    98          1          98               0         0         0          0

Feb  9 03:42:04    2913        35         2814             2372674   14546789  2372674    170613

Feb  9 03:43:04    3129        38         216              2466877   16016399  2466877    181072

Feb  9 03:44:04    3310        40         180              2510788   17302274  2510788    190648

Feb  9 03:45:04    3601        43         291              2579305   19470841  2579305    210759

Feb  9 03:46:04    3825        46         224              2657361   20952659  2657361    240920

Feb  9 03:47:04    4063        49         238              2747506   22538284  2747506    262756

Feb  9 03:48:04    4298        52         234              2830363   24126386  2830363    284014

Feb  9 03:49:04    4530        55         231              2909578   25734085  2909578    304881

Feb  9 03:50:04    4753        58         222              2984782   27302279  2984782    324646

Feb  9 03:51:04    4961        60         208              3057329   28792696  3057329    342571

Feb  9 03:52:05    5177        63         215              3135909   30322183  3135909    360386

Feb  9 03:53:05    5393        65         216              3223111   31851898  3223111    377234

Feb  9 03:54:05    5550        67         156              3229253   33250926  3229253    382132

Feb  9 03:55:05    5694        69         143              3229253   34599173  3229253    385339

Feb  9 03:56:05    5832        71         138              3229253   35912290  3229253    387534

Feb  9 03:57:05    5987        73         155              3229257   37416025  3229257    389403

Feb  9 03:58:05    6133        74         145              3229257   38817868  3229257    390404
```

```
Mar  2 00:26:46    6248        76         114              3229257   39991732  3229257    390551
```

The show command output displays memory change history for thresholds above 1.0% of the resource limit, with the latest record reflecting current values on the standby router. It includes network details for the default VRF such as timestamps, memory in MB, rlimit percentage, memory differences, and counts of networks, paths, path elements, and attributes.

Enhanced monitoring of BGP keepalive messages

An enhanced monitoring of BGP keepalive messages is a network monitoring feature that

- manages per-neighbor input and output queues,
- triggers throttling when packet volume exceeds thresholds or TCP buffers are full, and
- monitors keepalive intervals and hold timers to track session stability.

This feature logs precise times when neighbors' queues enter and exit throttled states, records the maximum duration of throttling events, and maintains historical data including a circular buffer of the last 10 throttling states per neighbor that persists across resets. It also tracks parameters such as maximum hold time elapsed since the last keepalive message, timestamps of these events, and counts of hold time threshold crossings, which indicate session stability or potential issues.

- This monitoring helps you maintain network stability and fairness by controlling message processing rates.
- You can view detailed session stability and throttling statistics by running the **show bgp neighbor detail** command.

View enhanced monitoring of BGP keepalive messages

To view detailed information about BGP session stability and message handling to monitor keepalive message timing and throttling behavior.

The enhanced monitoring of BGP keepalive messages feature enables adminstrators to ensure network stability and fairness by managing BGP message processing rates. It provides key insights into session stability, throttling, and timer metrics essential for diagnosing and troubleshooting delayed or lost keepalive messages that affect BGP session continuity.

Before you begin

Before you begin, gather the necessary details to use the Enhanced Monitoring of BGP Keepalive Messages feature:

- Access to the router CLI with permissions to run commands, such as **show bgp neighbor detail**
- Configured and active BGP neighbors to monitor
- Knowledge of keepalive interval and hold timer settings for interpreting monitoring data

- Readiness to analyze throttling statistics, hold time events, and queue metrics for session stability and troubleshooting

Procedure

Use the **show bgp neighbor detail** command to view the data on BGP neighbors, including critical timers and queue metrics.

Example:

```
....
Max Hold Time elapsed was 6001 msec at Sep 12 17:02:36.954, crossed 40%: 2, 70%: 0

  Max Hold Time elapsed before reset was 9001 msec at Sep 12 17:01:53.397, crossed 40%: 7, 70%: 2

  First message received at Sep 12 16:45:00.973, sent at Sep 12 16:45:00.975

  First message before reset received at Sep 12 16:42:16.573, sent at Sep 12 16:42:16.574

  Max read throttled duration was 6769 msec starting at Sep 12 16:45:01.487, max InQ 1000 processed
 930

    Most recent read throttle periods (in msec):

    Start Time             Duration     Max InQ    Messages

    Sep 12 17:00:16.937          14         104          45

    Sep 12 17:00:16.954           9         136          74

    Sep 12 17:00:47.358          11         125         135

    Sep 12 17:01:02.658           2          83           0

    Sep 12 17:01:02.693           7         110           0

    Sep 12 17:01:02.705          13         139          74

    Sep 12 17:01:17.856           5          92          60

    Sep 12 17:01:17.877           3          91          30

    Sep 12 17:01:17.891          10         135          74

    Sep 12 17:01:33.128          21         132         193

  Max read throttled duration before reset was 5013 msec starting at Sep 12 16:42:17.079, max InQ 76
 processed 0

    Most recent read throttle periods before reset (in msec):

    Start Time             Duration     Max InQ    Messages

    Sep 12 16:42:17.079        5013          76           0

  Max write throttled duration was 685 msec starting at Sep 12 16:45:08.486, max OutQ 1501 queued 57


    Most recent write throttle periods (in msec):
```

```
    Start Time             Duration    Max OutQ    Messages

    Sep 12 17:01:38.799          46         398          23

    Sep 12 17:01:38.846         202         342          57

    Sep 12 17:01:39.049          47         320          23

    Sep 12 17:01:39.097         202         264          57

    Sep 12 17:01:39.299          46         242          22

    Sep 12 17:01:39.346         204         185          58

    Sep 12 17:01:39.551          45         164          23

    Sep 12 17:01:39.597         202         108          57

    Sep 12 17:01:39.799          46          86          23

    Sep 12 17:01:39.847         202          30           8

  Max write throttled duration before reset was 205 msec starting at Sep 12 16:42:21.849, max OutQ
1003 queued 1

    Most recent write throttle periods before reset (in msec):

    Start Time             Duration    Max OutQ    Messages

    Sep 12 16:42:21.849         205        1003           1

    Sep 12 16:42:22.055          20         925          23

    Sep 12 16:42:22.075          16         869          56
....
```

This table shows key fields from the show bgp neighbor detail output related to enhanced BGP keepalive monitoring, including timers, throttling, queue sizes, and message stats for assessing session stability.

This table describes the significant fields shown in the display.

Table 2: Fields pertaining to enhanced monitoring of BGP keepalive messages in the output of show bgp neighbor details command

Field	Description
Max Hold Time elapsed	Maximum amount of time that has passed since the last BGP keepalive message was received from a neighbor before a BGP session is considered to be down.

Field	Description
Max Hold Time elapsed was 6001 msec at Sep 12 17:02:36.954, crossed 40%: 2, 70%: 0	Maximum amount of time that has passed since the last BGP keepalive message was received from a neighbor before a BGP session is considered to be down. In this specific output, the fields indicate the following: Max Hold Time elapsed was 6001 msec: indicates that the maximum time interval between receiving keepalive messages from the neighbor was 6001 milliseconds or approximately 6 seconds. at Sep 12 17:02:36.954: Timestamp when this maximum hold time was observed. crossed 40%: 2, 70%: 0: Number of times the hold time crossed certain thresholds. The hold time crossed the 40% threshold twice and the 70% threshold zero times, suggesting that the hold time reached a significant portion of its configured value but did not exceed it by a large margin.
Max Hold Time elapsed before reset was 9001 msec at Sep 12 17:01:53.397, crossed 40%: 7, 70%: 2	Maximum duration between receiving BGP (Border Gateway Protocol) keepalive messages from a neighbor before the BGP session was reset. In this specific output, the fields indicate the following: Max Hold Time elapsed before reset was 9001 msec: Maximum time interval between receiving keepalive messages from the neighbor before the BGP session reset was 9001 milliseconds or approximately 9 seconds. at Sep 12 17:01:53.397: Timestamp when this maximum hold time before reset was observed. crossed 40%: 7, 70%: 2: Number of times the hold time crossed certain thresholds. The hold time crossed the 40% threshold seven times and the 70% threshold two times, suggesting that the hold time frequently approached significant portions of its configured maximum value before the BGP session reset
First message received at Sep 12 16:45:00.973, sent at Sep 12 16:45:00.975	Timestamp when the first message from a BGP neighbor was received by the local router. In this specific output, the fields indicate the following: First message received at Sep 12 16:45:00.973: First message from the BGP neighbor was received at 16:45:00 on September 12th sent at Sep 12 16:45:00.975: Timestamp when the corresponding message was sent by the BGP neighbor, which was nearly simultaneously, just 0.002 seconds later.
First message before reset received at Sep 12 16:42:16.573, sent at Sep 12 16:42:16.574	Timestamp when the first message from a BGP neighbor was received by the local router before a reset occurred. In this specific output, the fields indicate the following: First message before reset received at Sep 12 16:42:16.573: first message from the BGP neighbor was received at 16:42:16 on September 12th, before a reset occurred. sent at Sep 12 16:42:16.574: Timestamp when the corresponding message was sent by the BGP neighbor, which was nearly simultaneous, just 0.001 seconds later.

Field	Description
`Max read throttled duration was 6769 msec starting at Sep 12 16:45:01.487, max InQ 1000 processed 930`	Maximum duration during which the read process was throttled, indicating a restriction or limitation on the rate of reading data. In this specific output, the fields indicate the following: Max read throttled duration was 6769 msec: Maximum duration of throttling for reading data was 6769 milliseconds (approximately 6.769 seconds). starting at Sep 12 16:45:01.487: Timestamp when this maximum throttling duration started, which was at 16:45:01 on September 12th. max InQ 1000 processed 930: Maximum input queue (InQ) size was 1000, and during the throttled duration, 930 items were processed.
Start Time	Timestamp when the read throttle period started.
Dry Run elapsed time(ms)	Time taken for the dry run in milliseconds.
Duration	Duration of the throttle period in milliseconds, indicating how long the read process was restricted or limited.
Max InQ	Maximum size of the input queue during the throttle period. The input queue typically holds incoming data packets waiting to be processed.
Messages	Number of messages or data packets processed during the throttle period.
Max read throttled duration before reset was 5013 msec starting at Sep 12 16:42:17.079, max InQ 76 processed 0	Maximum duration of a read throttle period on the network device. In this specific output, the fields indicate the following: Max read throttled duration before reset: Maximum duration of the read throttle period, which was 5013 milliseconds or approximately 5.013 seconds. Starting at Sep 12 16:42:17.079: Timestamp when the read throttle period started, which was at 16:42:17 on September 12th *Max InQ 76 processed 0*: The segment *Max InQ 76* indicates that the maximum size of the input queue during the throttle period was 76. The segment *processed 0* indicates that no messages or data packets were processed during this throttle period.
Max write throttled duration was 685 msec starting at Sep 12 16:45:08.486, max OutQ 1501 queued 57	Maximum duration of the write throttle period, which was 685 milliseconds. In this specific output, the fields indicate the following: Max write throttled duration: Maximum duration of the write throttle period, which was 685 milliseconds. Starting at Sep 12 16:45:08.486: Timestamp when the write throttle period started, which was September 12th at 16:45:08.486. Max OutQ: Maximum size of the output queue during the throttle period. In this case, it was 1501, which typically holds data packets waiting to be transmitted. Queued: Number of items queued in the output queue during the throttle period. In this case, it was 57.

Field	Description
Max write throttled duration before reset was 205 msec starting at Sep 12 16:42:21.849, max OutQ 1003 queued 1	Maximum duration of a write throttle period on a network device before a reset occurred. In this specific output, the fields indicate the following: Max write throttled duration before reset: Maximum duration of the write throttle period before a reset occurred, which was 205 milliseconds. Starting at Sep 12 16:42:21.849: Timestamp when the write throttle period started, which was on September 12th at 16:42:21.849. Max OutQ: Maximum size of the output queue during the throttle period. In this case, it was 1003, indicating the maximum number of items that were waiting to be transmitted. Queued: Number of items queued in the output queue during the throttle period. In this case, it was 1.
Start Time:	Timestamp when the write throttle period started.
Duration	Duration of the write throttle period in milliseconds.
Max OutQ	Maximum size of the output queue during the throttle period. The output queue typically holds data packets waiting to be transmitted.
Messages	Number of messages or data packets transmitted during the throttle period.

Enhanced monitoring of BGP memory utilization

Enhanced monitoring of BGP memory utilization is a network monitoring feature that

- tracks memory usage by the BGP process within a device,
- provides periodic memory state checks and logs changes, and
- triggers notifications when memory usage approaches critical thresholds.

This feature serves as an early warning system to help you maintain network stability and performance by proactively managing BGP memory consumption.

Key features of enhanced monitoring of BGP memory utilizations:

- Periodic Memory State Check: The system automatically monitors memory usage every 60 seconds to detect issues promptly while minimizing overhead.
- Logging Mechanism: Memory changes are recorded in both the BGP trace log and a circular buffer, ensuring reliable data for troubleshooting.
- Significant Change Detection: Any increase in memory usage exceeding 1% of the configured resource limit (rlimit) since the last report is logged to identify trends or spikes.

- Resource Limit (rlimit): This predefined threshold limits the maximum memory BGP can use to prevent system instability caused by excessive memory consumption.
- Syslog Notifications: Notifications are sent to syslog at 85%, 90%, and 95% of the resource limit, escalating in frequency as memory usage nears the limit to alert administrators effectively.

Monitor BGP memory utilization

Track BGP memory usage changes when utilization exceeds 1% of the configured resource limit (rlimit) to maintain network stability.

Use the **show bgp memory history** command to review BGP memory usage changes exceeding 1% of the resource limit for proactive monitoring.

Before you begin

Make sure you have appropriate access rights to run the **show bgp memory history** command on the device.

Follow these steps to monitor BGP memory utilization:

Procedure

Step 1 Use the **show bgp memory history** command to view the data on memory utilization on the active router.

Example:

```
Router# show bgp memory history
History of memory changes recorded for a threshold greater than 1.0% of rlimit.
Last shown record displays current values.
Network information for default VRF.

Time               Memory(MB)  Rlimit(%)  Memory diff(MB)  Networks  Paths     PathElems  Attributes

Oct  2 16:30:37    152         1          152              400       400       400        9
Oct  2 16:31:37    343         4          191              396952    396869    396952     725
Oct  2 16:32:37    425         5          81               524567    513979    524567     8408
Oct  2 16:42:38    741         9          316              1178605   1241533   1178604    10753
Oct  2 16:43:38    985         12         243              1778234   1859254   1778234    11214
Oct  2 19:42:39    901         11         -84              1800688   678607    1800688    10911
Oct  2 19:45:39    766         9          -136             1332259   688784    1332259    10943
```

This command displays memory changes exceeding 1.0% of rlimit, with the last record showing current values and network information for the default VRF on the active router.

Step 2 Use the **show bgp memory history standby** command to view the data on memory utilization on a standby router.

Example:

```
Router# show bgp memory history standby
History of memory changes recorded for a threshold greater than 1.0% of rlimit.
Last shown record displays current values.
Network information for default VRF.

Time               Memory(MB)  Rlimit(%)  Memory diff(MB)  Networks  Paths     PathElems  Attributes

Feb  9 03:39:04    98          1          98               0         0         0          0
```

```
Feb  9 03:42:04    2913        35         2814            2372674   14546789  2372674    170613
Feb  9 03:56:05    5832        71         138             3229253   35912290  3229253    387534
Feb  9 03:57:05    5987        73         155             3229257   37416025  3229257    389403
Feb  9 03:58:05    6133        74         145             3229257   38817868  3229257    390404
Mar  2 00:26:46    6248        76         114             3229257   39991732  3229257    390551
```

This command displays memory changes exceeding 1.0% of rlimit, with the last record showing current values and network information for the default VRF on the standby router.

This table shows key fields from the show bgp memory history details output related to enhanced BGP memory utilization.

Table 3: Fields pertaining to enhanced monitoring of BGP memory utilization in the output of show bgp memory history details command

Field	Description
Time	Timestamp when the rmeasurement was taken.
Memory (MB)	Total memory in megabytes (MB) used by the routing process at the specified time.
Rlimit (%)	Percentage of the memory resource limit that is being used.
Memory diff (MB)	Quantity of memory usage in megabytes (MB) that has increased or decreased since the last report.
Networks	Number of network prefixes known to the router.
Paths	Number of distinct paths to various destinations.
PathElems	Number of path elements (such as AS numbers) involved in routing.
Attributes	Number of unique BGP attributes in use, such as local preference, and MED.

You obtain a detailed historical view of BGP memory usage for proactive monitoring of resource consumption.

Verify enhanced next hop monitoring

Use these commands to monitor BGP next-hop reachability, metric changes, and event counters to analyze routing dynamics and network stability.

This feature helps you promptly identify and analyze routing path changes, which is essential for maintaining network stability and performance, especially in large-scale environments. By using specific show commands, you can obtain detailed insights into next-hop status, event counters, and historical event data, enabling effective troubleshooting and optimization of routing decisions.

Before you begin

Before you begin, make sure that you have proper access rights and privileges to run BGP-related show commands on the router.

Follow these steps to monitor BGP enhanced next hop monitoring:

Procedure

Step 1 Use the **show bgp nexthops** command to verify the details of nexthop reachability and metric change counters.

Example:

```
Router# show bgp nexthops
....
Next Hop        Reachable   Unreachable   MetricIncrease   MetricDecrease

0.0.0.0
10.10.10.1              1             0                0                0
203.0.113.1             2             1                0                0
192.168.0.3             1             0                1                2
192.168.0.5             1             0                0                0

.....
```

Step 2 Use the **show bgp nexthops wide** command to verify detailed information about BGP next-hop processing times, status codes, event counters, and metrics for each gateway address family.

Example:

```
Router# show bgp nexthops wide
....
Next Hop             Status      Metric    Tbl-ID        Notf      LastRIBEvent    RefCount     R
  U    MI    MD

0.0.0.0                                                                            25/3
10.10.10.1             [R][C][NL]       0      e0000000         1/0    00:14:52 (Cri)    17/20
 1     0     0     0
203.0.113.1            [R][NC][NL]      2      e0000000         0/0    00:02:06 (Reg)     5/7
 1     0     0     0
192.168.0.3            [R][NC][NL]      3      e0000000         0/0    00:02:06 (Reg)    12/246
 1     0     0     0
192.168.0.5            [R][NC][NL]      2      e0000000         1/0    00:14:17 (Cri)    16/270
 1     0     0     0

.....
```

Step 3 Use the **show bgp nexthops** *ipaddress* command to verify detailed BGP next-hop information for the IP address 10.10.10.1, including VRF, nexthop ID, flags, advertising neighbors, RIB details, event history, and reference counts.

Example:

```
Router# show bgp nexthops wide
.....
Reachable Notifications:        2 (last at Sep 11 16:04:56.738)
     Unreachable Notifications:         1 (last at Sep 11 16:04:36.520)
     Metric Increase Notifications:     2
     Metric Decrease Notifications:     1
     Most Recent Events:
      Time                      Event Type      Metric
      Sep 11 16:04:36.520       Unreachable          -
      Sep 11 16:04:56.738       Reachable            2
      Sep 11 16:30:38.402       Reachable           21
      Sep 11 16:31:23.548       Reachable           16
      Sep 11 16:34:59.460       Reachable          101
.....
```

This table shows key fields from the **show bgp nethops** output related to enhanced next hop monitoring, including next-hop IP address, status, metric, notification counts, last RIB event time, reference counts, and event counters for reachability and metric changes, enabling effective tracking and optimization of routing path stability in large-scale networks

Table 4: Fields pertaining to Enhanced Next Hop Monitoring in the output of show bgp nexthops command

Next Hop	The IP address of the next-hop router in the BGP network.
Status	A set of codes indicating the reachability and other status details about the next hop (e.g., Reachable, Unreachable, etc.).
Metric	The metric value used by BGP to determine the best path to the next hop. Lower values are preferred.
Tbl-ID	The unique identifier for the table in which the next-hop information is stored.
Notf	Notifications received/sent related to the next hop, often indicating BGP updates or state changes
LastRIBEvent	The time elapsed since the last Routing Information Base (RIB) event that pertained to this next hop.
RefCount	Reference count, which can indicate how many routes are using this next hop.
R (Reachable)	Event counter for the number of times the next hop has been marked as reachable.
U (Unreachable)	Event counter for the number of times the next hop has been marked as unreachable.
MI (Metric Increased)	Metric value for a particular route that has increased compared to the previous metric value.
MI (Metric Decreased)	Metric value for a particular route that has decreased compared to the previous metric value.
Reachable Notifications	The number of times a route has become reachable, that is a valid route to a destination is available, and the time of the last such notification.
Unreachable Notifications	The number of times a route has become unreachable, that is a previously valid route is no longer available, and the time of the last such notification.
Metric Increase Notifications	The number of times the metric for a route has increased, which typically makes the route less preferred.
Metric Decrease Notifications	The number of times the metric for a route has decreased, which usually makes the route more preferred.

Most Recent Events	List of individual routing events, including the time they occurred, the type of event, and the metric associated with the event. This also indicates the relative desirability of the route, with lower metrics being more preferred.
Unreachable	Indicates a loss of route.
Reachable	Indicates that a route is available.

You have verified BGP next-hop reachability and analyzed event counters, supporting efficient routing stability assessment and troubleshooting.

Enhanced next hop monitoring

Enhanced next hop monitoring is network monitoring feature that

- track changes in BGP next hop metrics and reachability,
- provide detailed event data including counters and timestamps, and
- aggregate statistics globally to assess network health dynamically.

This feature closely integrates BGP with the Routing Information Base (RIB) to receive notifications about next hop status changes, which trigger recalculations of optimal routing paths. This monitoring helps you identify root causes of network variations, especially during high routing activity or instability, which can strain CPU resources.

These key monitoring functions provide detailed insights into next hop status changes and overall network health, enabling effective tracking and analysis:

- Event counter tracking: Counts key events per next hop such as reachable, unreachable, metric increases, and decreases.
- Event history logging: Records the last five events with timestamps for each next hop.
- Recent event tracking: Logs transitions of next hops to reachable or unreachable states with precise timing.
- Global aggregation: Summarizes event counters by Address Family Identifier (AFI) for overall network health insight.
- Temporal counter analysis: Continuously updates event counts and reports changes over 1, 3, and 5-minute intervals, providing a dynamic view of network stability and routing effectiveness.

By using these monitoring capabilities, you can proactively manage BGP next hop changes, improving network reliability and performance.

Enhanced monitoring of version-rate statistics

The enhanced monitoring of version-rate statistics is a BGP monitoring feature that

- identifies sources of BGP churn by tracking and categorizing version bumps,
- calculates version rates over fixed time intervals, and
- generates detailed reports at both Address-Family Identifier (AFI) and neighbor AFI levels.

This feature also maintains a cumulative churn count and classifies version bumps by their origin, such as reachable, unreachable, import, redistribution, or label-related sources.

Additional details of enhanced monitoring of version-rate statistics

The following details provide additional information about how version-rate statistic monitoring features operate:

- Interval mechanism: Version bumps are monitored using fixed time intervals rather than sliding windows. For example, if an interval runs from 12:00 pm to 12:30 pm, data updates occur exactly at 12:30 pm; the next interval starts immediately after.
- NSR synchronization: Both active and standby routers independently record version bumps. Therefore, statistics may differ between active and standby routers.
- Reporting parameters: Reports are generated at both AFI and neighbor AFI levels. These reports include totals and interval-specific rates, categorized into buckets such as reachable, unreachable, import, redistribute, and label. Additional version bumps from other sources are also displayed, based on the main table's version number.

Enhanced monitoring of version-rate statistics

Enable monitoring and analysis of BGP version changes over specified fixed intervals for diagnostic and troubleshooting purposes.

This task helps you analyze BGP version changes over fixed time intervals, categorizing them by neighbor IP, VRF, and address family (AFI). Use these steps to gain visibility into the frequency and source of version changes, which may indicate network events, configuration changes, or route instabilities.

Before you begin

Ensure you have appropriate access rights to run the **show bgp memory history** command on the device.

Follow these steps to monitor BGP version-rate statistics:

Procedure

Step 1 Use the **show bgp sessions version-rate** command to o obtain a detailed analysis of BGP version changes across specified intervals;

Example:

```
Router# show bgp sessions version-rate
Thu Nov  2 11:40:41.107 IST
Interval definition(s):
  Interval   Duration (min)  Start time            End time
  Live                   5  Nov  2 11:37:22.029   Nov  2 11:40:41.413
  1                      5  Nov  2 11:32:22.029   Nov  2 11:37:22.029
  2                     15  Nov  2 11:22:22.029   Nov  2 11:37:22.029
```

```
  3                    30  Nov  2 10:52:22.027   Nov  2 11:22:22.027

Neighbor    VRF        AFI                  Total           Live
                                                          Total  Reach  UnReach
-------------------------------------------------------------------------------
10.10.10.1     default    All                    5    0      0        0
10.10.10.1                IPv4 Unicast           5    0      0        0
192.168.0.5    default    All                  606    0      0        0
192.168.0.5               IPv4 Unicast          63    0      0        0
192.168.0.5               VPNv4 Unicast        240    0      0        0
192.168.0.5               IPv6 Labeled-unicast  63    0      0        0
192.168.0.5               VPNv6 Unicast        240    0      0        0
192.168.0.5               RT Constraint          0    0      0        0
10:10:10::1    default    All                    5    0      0        0
10:10:10::1               IPv6 Unicast           5    0      0        0
10.0.1.1       1          All                    5    0      0        0
10.0.1.1                  IPv4 Unicast           5    0      0        0
10:0:1::1      1          All                    5    0      0        0


/* The output was too wide, so it was segmented; the below segment continues from above.
 */
 Interval 1              Interval 2            Interval 3  Spk AS
Total  Reach UnReach Total Reach UnReach Total Reach  UnReach
-------------------------------------------------------------------
  5      5       0     5     5      0       0     0      0      0   200
  5      5       0     5     5      0       0     0      0
282    282       0   282   282      0       0     0      0      0   100
 31     31       0    31    31      0       0     0      0
110    110       0   110   110      0       0     0      0
 31     31       0    31    31      0       0     0      0
110    110       0   110   110      0       0     0      0
  0      0       0     0     0      0       0     0      0
  5      5       0     5     5      0       0     0      0      0    200
  5      5       0     5     5      0       0     0      0
  0      0       0     0     0      0       0     0      0      0    200
  0      0       0     0     0      0       0     0      0
  0      0       0     0     0      0       0     0      0      0    200

/* The output was too wide, so it was segmented; the below segment continues from above.
 */

InQ OutQ NBRState     NSRState
----------------------------
0   0    Established  NSRReady

0   0    Established  NSRReady

0   0    Established  NSRReady

0   0    Established  NSRReady

0   0    Established  NSRReady
```

The output displays interval definitions, neighbor information (IP, VRF, AFI), and version change statistics categorized by interval (e.g., 5, 15, 30 minutes). It includes detailed breakdowns by traffic type and source.

The show command output categorizes changes by neighbor IP addresses, VRFs, and AFIs, such as IPv4 Unicast and VPNv4 Unicast, facilitating the analysis of traffic types and identification of version number fluctuation origins.

Step 2 Use the **show bgp sessions version-rate live** command to view the real-time BGP session version-rate statistics, capturing changes within the most recent 5-minute interval across all BGP neighbors and address families.

Example:

```
Router# show bgp sessions version-rate live
Thu Nov  2 11:40:48.314 IST
Interval definition(s):
  Interval   Duration (min)  Start time            End time
  Live                   5  Nov  2 11:37:22.029   Nov  2 11:40:48.610

Neighbor   VRF       AFI                   Total        Live           Spk
                                                    Total Reach Unreach
------------------------------------------------------------------------------
10.10.10.1    default   All                    5      0     0     0        0
10.10.10.1              IPv4 Unicast           5      0     0     0
192.168.0.5   default   All                  606      0     0     0        0
192.168.0.5             IPv4 Unicast          63      0     0     0
192.168.0.5             VPNv4 Unicast        240      0     0     0
192.168.0.5             IPv6 Labeled-unicast  63      0     0     0
192.168.0.5             VPNv6 Unicast        240      0     0     0
192.168.0.5             RT Constraint          0      0     0     0
10:10:10::1    default  All                    5      0     0     0        0
10:10:10::1             IPv6 Unicast           5      0     0     0
10.0.1.1       1        All                    5      0     0     0        0
10.0.1.1                IPv4 Unicast           5      0     0     0
10:0:1::1      1        All                    5      0     0     0        0

/* The output was too wide, so it was segmented; the below segment continues from above.
 */
AS  InQ  OutQ  BRState    NSRState
-------------------------------------
200  0   0    Established  NSRReady

100  0   0    Established  NSRReady

200  0   0    Established  NSRReady


200  0   0    Established  NSRReady

200  0   0    Established  NSRReady
```

The command reports current version changes and session metrics, including total version changes, reachability, speaker ID, AS number, queue sizes, and session states.

This provides a concise snapshot of BGP session activity and health.

Step 3 Use the **show bgp sessions version-rate brief** command to view concise information only for the "Live" interval, which is typically the most recent 5-minute window.

Example:

```
Router# show bgp sessions version-rate brief live

Thu Nov  2 11:40:55.743 IST
Interval definition(s):
  Interval   Duration (min)  Start time            End time
  Live                   5  Nov  2 11:37:22.029   Nov  2 11:40:56.059

Neighbor      VRF         Spk    AS   InQ  OutQ  NBRState     NSRState  Total       Live
10.10.10.1    default      0   200     0     0  Established  NSRReady     5          0
192.168.0.5   default      0   100     0     0  Established  NSRReady   606          0
10:10:10::1   default      0   200     0     0  Established  NSRReady     5          0
10.0.1.1        1          0   200     0     0  Established  NSRReady     5          0
10:0:1::1       1          0   200     0
```

The output focuses exclusively on real-time or near real-time BGP version rate statistics, showing only the "Live" interval column next to the "Total," without historical interval details.

You can monitor and analyze BGP version changes in detail using both historical intervals and real-time snapshots, categorized by neighbor, VRF, and AFI. This facilitates early detection of route instability, excessive churn, or underlying network issues.

What to do next

Review the output for any unexpected spikes or patterns in version changes. Investigate specific neighbors or intervals as needed for further troubleshooting or optimization.