

Troubleshoot BGP Virtual Memory (RLIMIT) Issue on IOS XR

Contents

[Introduction](#)

[Background Information](#)

[Issue Summary](#)

[Limitation](#)

[Possible Workaround/Solution](#)

Introduction

This document describes the BGP virtual memory (RLIMIT) issue on Cisco routers and outlines steps to take when encountering this issue.

Background Information

Rlimit defines the Resource Limit for a process in XR and varies depending on each process memory requirements. These limits can differ between releases as they can be adjusted based on new needs and discoveries. Rlimit is determined by fixed memory allocations for components such as shared memory, kernel, and dllmgr, making it non-configurable through CLI.

Issue Summary

Memory usage spiked to 90% after the BGP peer connection was established. This could also cause the BGP process to crash.

```
RP/0/RSP0/CPU0:Jul 15 01:04:24.815 GMT: bgp[1087]: %HA-HA_WD_LIB-4-RLIMIT :wd_handle_sigxfsz: Reached 9
RP/0/RSP0/CPU0:Jul 15 01:04:24.815 GMT: bgp[1087]: %ROUTING-BGP-4-VIRTUAL_MEMORY_LIMIT_THRESHOLD_REACHE
```

This command shows the maximum amount of memory that any process can access.

```
RP/0/RSP0/CPU0:ASR#show bgp process performance-statistics | i RLIMIT
Platform RLIMIT max: 2281701376 bytes
```

This command shows the dynamic limit in the heap:

```
RP/0/RSP0/CPU0:ASR#show bgp instance all scale
BGP instance 0: 'default'
=====
```

```

VRF: default
Neighbors Configured: 2      Established: 2
Address-Family  Prefixes Paths  PathElem  Prefix  Path  PathElem
                Memory  Memory  Memory
  IPv4 Unicast  112649  225065  112649   9.88MB  13.74MB  6.77MB
  IPv6 Unicast   6358   12581   6358   645.73KB  786.31KB  391.17KB
-----
Total          119007  237646  119007  10.51MB  14.50MB  7.15MB
node:          node0_RSP0_CPU0

```

```

-----
JID  Text      Data      Stack      Dynamic  Dyn-Limit  Shm-Tot  Phy-Tot  Process
-----
1067      1M      10M      572K      2001M    2175M      145M     2012M    bgp
343       8K      12K      128K      421M     1024M      30M      422M     mibd_infra
1141     22M      5M      1012K     374M     2048M      95M      380M     netconf
Total text: 22893 pages
      data: 24102 pages
      stack: 6765 pages
      malloced: 21257 pages

```

Limitation

The RLIMIT restriction is a critical factor on cXR 32-bit systems, where a memory ceiling is enforced. This limitation directly impacts the memory available for BGP processes.

However, on eXR 64-bit systems, the RLIMIT is significantly increased. This enhancement multiplies the available memory for BGP processes, providing a more robust environment for handling larger routing tables and more peers.

Please find the comparison of Memory Allocation:

Device with RSP880-LT-TR and eXR has the RLIMIT for BGP as 7.4GB

```

RP/0/RSP0/CPU0:ASR#show processes memory detail 10523
JID  Text      Data      Stack      Dynamic  Dyn-Limit  Shm-Tot  Phy-Tot  Process
=====
1087      2M      1030M    136K      41M     7447M      131M     183M     bgp

```

Device having RSP880-LT-TR and cXR has the RLIMIT for BGP as 2.5GB

```

RP/0/RSP0/CPU0:ASR#show processes memory detail 1087
JID  Text      Data      Stack      Dynamic  Dyn-Limit  Shm-Tot  Phy-Tot  Process
-----
1087      1M      10M      356K      31M     2574M      35M      41M      bgp

```

Possible Workaround/Solution

To address the memory issue with BGP, these steps can be considered.

- Upgrade to 64-bit System
 - BGP benefits from a larger memory allocation on a 64-bit system, approximately 8GB as defined by RLIMIT. This upgrade can help manage the increased memory demands of BGP.
- Change ASR9k Profile
 - Switch the ASR9k profile from the default setting to the L3XL profile. This adjustment increases the memory allocation for BGP, which can help alleviate memory pressure.
 - Note that changing to the L3XL profile reduces the memory available for other processes. Therefore, it is essential to evaluate the impact on the overall system performance.
 - Before implementing the L3XL profile, thoroughly review the platform documentation to understand its implications and ensure compatibility with your system requirements.
- Evaluate "soft-reconfiguration inbound always" knob
 - The use of the 'soft-reconfiguration inbound always' knob is highly memory-intensive, especially if additional paths are present.
 - Check BGP peers that lack route refresh capability and ensure this knob is only enabled for those specific peers.
 - Remove this knob from peers that do support route refresh to reclaim memory.
- Implement Route-Policy to Deny Some Prefixes
 - Create a route-policy to deny certain prefixes, which can help reduce the memory usage by limiting the number of routes that need to be processed and stored.
- Reduce the Number of BGP Peers
 - Decrease the number of BGP peers on the router to lower the overall memory consumption. This step is particularly useful if you have a large number of peers that contribute to the high memory usage.
- Restart BGP Process or Reload Router
 - Manually restarting the BGP process or reloading the router can help free up memory. This is a temporary solution but can be effective in mitigating immediate memory issues.
- Evaluate Memory-Intensive Features
 - Be aware that certain features like Non-Stop Routing (NSR), additional-paths, and maximum-path can contribute to increased memory usage.
 - Assess the necessity of these features and consider disabling or optimizing them if they are not critical to your network operations.

These steps can better manage memory usage and ensure the stability and performance of your BGP processes.

If the issue is still not recovered, collect logs and reach out to Cisco TAC:

```
show tech-support
show tech-support routing bgp
show processes memory detail <job id> location 0/rsp0/cpu0
show processes memory detail <job id> location 0/rsp1/cpu0
show memory summary location all
show memory heap <job id> location 0/rsp0/cpu0
show memory heap <job id> location 0/rsp1/cpu0
show memory heap dllname <job id>
show bgp scale
show bgp scale standby
show bgp all all process performance-statistics
show bgp all all process performance-statistics detail
```