Chapter 14

Verify the Routing Engine CPU Memory

This chapter describes how to verify the Routing Engine CPU memory on your Juniper Networks router. (See Table 35.)

Table 35: Checklist for Verifying the Routing Engine CPU Memory

Verify the Routing Engine CPU Memory Tasks	Command or Action
Check the Routing CPU Memory Usage on page 180	
1. Check Overall CPU and Memory Usage on page 180	show system process extensive
2. Check Routing Protocol Process (rpd) Memory Usage on page 183	show route summary show task memory detail
3. Display Tasks on page 185	show task show task memory show task *task-name*

Check the Routing CPU Memory Usage

Purpose Software processes on the router can consume a considerable amount of CPU and memory. The routing protocol process (rpd) can consume enormous amounts of memory to store information needed for the operation of routing and related protocols, such as Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), Intermediate System-to-Intermediate System (ISIS), Resource Reservation Protocol (RSVP), Label Distribution Protocol (LDP), and Multiprotocol Label Switching (MPLS).

Steps To Take To verify the traffic passing through the router and check memory utilization, follow these steps:

1. Check Overall CPU and Memory Usage on page 180
2. Check Routing Protocol Process (rpd) Memory Usage on page 183
3. Display Tasks on page 185

Step 1: Check Overall CPU and Memory Usage

Purpose You can display exhaustive system process information about software processes that are running on the router and have controlling terminals. This command is equivalent to the UNIX top command. However, the UNIX top command shows real-time memory usage, with the memory values constantly changing, while the show system processes extensive command provides a snapshot of memory usage in a given moment.

Action To check overall CPU and memory usage, enter the following JUNOS command-line interface (CLI) command:

```
user@host> show system processes extensive
```

Sample Output

```
user@R1> show system processes extensive
last pid:  5251;  load averages:  0.00,  0.00,  0.00  up 4+20:22:16    10:44:41
58 processes:  1 running, 57 sleeping

Mem: 57M Active, 54M Inact, 17M Wired, 184K Cache, 35M Buf, 118M Free
Swap: 512M Total, 512M Free


  PID USERNAME PRI NICE  SIZE    RES STATE    TIME   WCPU    CPU COMMAND
 4480 root       2   0  3728K  1908K select 231:17  2.34%  2.34% chassisd
 4500 root       2   0  1896K   952K select   0:36  0.00%  0.00% fud
 4505 root       2   0  1380K   736K select   0:35  0.00%  0.00% irsd
 4481 root       2   0  1864K   872K select   0:32  0.00%  0.00% alarmd
 4488 root       2   0  8464K  4600K kqread   0:28  0.00%  0.00% rpd
 4501 root       2 -15  1560K   968K select   0:21  0.00%  0.00% ppmd
 4510 root       2   0  1372K   812K select   0:13  0.00%  0.00% bfdd
    5 root      18   0     0K     0K syncer   0:09  0.00%  0.00% syncer
 4485 root       2   0  3056K  1776K select   0:07  0.00%  0.00% snmpd
 4499 root       2   0  3688K  1676K select   0:05  0.00%  0.00% kmd
 4486 root       2   0  3760K  1748K select   0:05  0.00%  0.00% mib2d
 4493 root       2   0  1872K   928K select   0:03  0.00%  0.00% pfed
 4507 root       2   0  1984K  1052K select   0:02  0.00%  0.00% fsad
 4518 root       2   0  3780K  2400K select   0:02  0.00%  0.00% dcd
    8 root     -18   0     0K     0K psleep   0:02  0.00%  0.00% vmuncachedaemo
```

```
   4 root     -18   0     0K     0K psleep   0:02  0.00%  0.00% bufdaemon
4690 root       2   0     0K     0K peer_s   0:01  0.00%  0.00% peer proxy
4504 root       2   0  1836K   968K select   0:01  0.00%  0.00% dfwd
4477 root       2   0   992K   320K select   0:01  0.00%  0.00% watchdog
4354 root       2   0  1116K   604K select   0:01  0.00%  0.00% syslogd
4492 root      10   0  1004K   400K nanslp   0:01  0.00%  0.00% tnp.sntpd
4446 root      10   0  1108K   616K nanslp   0:01  0.00%  0.00% cron
4484 root       2   0 15716K  7468K select   0:01  0.00%  0.00% mgd
4494 root       2  15  2936K  2036K select   0:01  0.00%  0.00% sampled
5245 remote     2   0  8340K  3472K select   0:01  0.00%  0.00% cli
   2 root     -18   0     0K     0K psleep   0:00  0.00%  0.00% pagedaemon
4512 root       2   0  2840K  1400K select   0:00  0.00%  0.00% l2tpd
   1 root      10   0   852K   580K wait     0:00  0.00%  0.00% init
5244 root       2   0  1376K   784K select   0:00  0.00%  0.00% telnetd
4509 root      10   0  1060K   528K nanslp   0:00  0.00%  0.00% eccd
4508 root       2   0  2264K  1108K select   0:00  0.00%  0.00% spd
2339 root      10   0   514M 17260K mfsidl   0:00  0.00%  0.00% newfs
4497 root       2   0  2432K  1152K select   0:00  0.00%  0.00% cosd
4490 root       2 -15  2356K  1020K select   0:00  0.00%  0.00% apsd
4496 root       2   0  2428K  1108K select   0:00  0.00%  0.00% rmopd
4491 root       2   0  2436K  1104K select   0:00  0.00%  0.00% vrrpd
4487 root       2   0 15756K  7648K sbwait   0:00  0.00%  0.00% mgd
5246 root       2   0 15776K  8336K select   0:00  0.00%  0.00% mgd
   0 root     -18   0     0K     0K sched    0:00  0.00%  0.00% swapper
5251 root      30   0 21732K   840K RUN      0:00  0.00%  0.00% top
4511 root       2   0  1964K   908K select   0:00  0.00%  0.00% pgmd
4502 root       2   0  1960K   956K select   0:00  0.00%  0.00% lmpd
4495 root       2   0  1884K   876K select   0:00  0.00%  0.00% ilmid
4482 root       2   0  1772K   776K select   0:00  0.00%  0.00% craftd
4503 root      10   0  1040K   492K nanslp   0:00  0.00%  0.00% smartd
   6 root      28   0     0K     0K sleep    0:00  0.00%  0.00% netdaemon
4498 root       2   0  1736K   932K select   0:00  0.00%  0.00% nasd
4506 root       2   0  1348K   672K select   0:00  0.00%  0.00% rtspd
4489 root       2   0  1160K   668K select   0:00  0.00%  0.00% inetd
4478 root       2   0  1108K   608K select   0:00  0.00%  0.00% tnetd
4483 root       2   0  1296K   540K select   0:00  0.00%  0.00% ntpd
4514 root       3   0  1080K   540K ttyin    0:00  0.00%  0.00% getty
4331 root       2   0   416K   232K select   0:00  0.00%  0.00% pccardd
   7 root       2   0     0K     0K pfeacc   0:00  0.00%  0.00% if_pfe_listen
  11 root       2   0     0K     0K picacc   0:00  0.00%  0.00% if_pic_listen
   3 root      18   0     0K     0K psleep   0:00  0.00%  0.00% vmdaemon
   9 root       2   0     0K     0K scs_ho   0:00  0.00%  0.00% scs_housekeepi
  10 root       2   0     0K     0K cb-pol   0:00  0.00%  0.00% cb_poll
```

What It Means The sample output shows the amount of virtual memory used by the Routing Engine and software processes. For example, 118 MB of physical memory is free and 512 MB of the swap file is free, indicating that the router is not short of memory. The processes field shows that most of the 58 processes are in the sleeping state, with 1 in the running state. The process or command that is running is the top command.

The commands column lists the processes that are currently running. For example, the chassis process (chassisd) has a process identifier (PID) of 4480, with a current priority (PRI) of 2. A lower priority number indicates a higher priority.

The processes are listed according to level of activity, with the most active process at the top of the output. For example, the chassis (chassisd) process is consuming the largest amount of CPU resource at 2.34 percent.

The memory field (Mem) shows the virtual memory managed by the Routing Engine and used by processes. The value in the memory field is in KB and MB, and is broken down as follows:

- Active—Memory that is allocated and actually in use by programs.
- Inact—Memory that is either allocated but not recently used or memory that was freed by programs. Inactive memory is still mapped in the address space of one or more processes and, therefore, counts toward the resident set size of those processes.
- Wired—Memory that is not eligible to be swapped, and is usually used for Routing Engine memory structures or memory physically locked by a process.
- Cache—Memory that is not associated with any program and does not need to be swapped before being reused.
- Buf—The size of the memory buffer used to hold data recently called from disk.
- Free—Memory that is not associated with any programs. Memory freed by a process can become Inactive, Cache, or Free, depending on the method used by the process to free the memory.

When the system is under memory pressure, the pageout process reuses memory from the free, cache, inactive and, if necessary, active pages.

The Swap field shows the total swap space available and how much is unused. In the example, the output shows 512 MB of total swap space and 512 MB of free swap space.

Finally, the memory usage of each process is listed. The SIZE field indicates the size of the virtual address space, and the RES field indicates the amount of the program in physical memory, which is also known as RSS or Resident Set Size. In the sample output, the chassis (chassisd) process has 3728 KB of virtual address space and 1908 KB of physical memory.

For additional information about the show system processes extensive command, see “Stop and Start JUNOS Software” on page 37.

Step 2: Check Routing Protocol Process (rpd) Memory Usage

Purpose When you notice a lot of memory usage, you can obtain detailed information about the memory utilization of routing tasks to get an idea of what is going on. The routing process (rpd) is the main task that uses Routing Engine memory.

Action To check routing process memory usage, enter the following JUNOS CLI operational mode commands:

```
user@host> show route summary
user@host> show task memory detail
```

Sample Output

```
user@host> show route summary
Autonomous system number: 209
Router ID: 205.175.0.170

inet.0: 179783 destinations, 898393 routes (179771 active, 146 holddown, 157
hidden)
              Direct:     17 routes,     17 active
               Local:     18 routes,     18 active
                 BGP: 896632 routes, 178010 active
              Static:     32 routes,     31 active
               IS-IS:   1694 routes,   1694 active

inet.2: 8766 destinations, 22700 routes (8766 active, 124 holddown, 73 hidden)
              Direct:     17 routes,     17 active
               Local:     18 routes,     18 active
                 BGP:  20939 routes,   7006 active
              Static:     32 routes,     31 active
               IS-IS:   1694 routes,   1694 active

inet.3: 1614 destinations, 1719 routes (1614 active, 0 holddown, 0 hidden)
               IS-IS:   1613 routes,   1551 active
                RSVP:     45 routes,     45 active
                 LDP:     61 routes,     18 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
              Direct:      1 routes,      1 active

mpls.0: 371 destinations, 371 routes (371 active, 0 holddown, 0 hidden)
               MPLS:      3 routes,      3 active
               RSVP:    303 routes,    303 active
                LDP:     65 routes,     65 active

user@R1> show task memory detail
  ------------------------ Overall Memory Report -------------------------
  Size TP    Allocs   Mallocs  AllocBytes  MaxAllocs    MaxBytes  FreeBytes
    12         8140    186959     2341188     200824     2409888      54972
    16         4061       182       67888       4586       73376       5840
    16 T          -         -           -     393571     6297136          -
    20       688588        51    13772780     713704    14274080     423956
[...Output truncated...]
  8192  P        91         -      745472        195     1597440          -
 12288  P         -         -           -          1       12288          -
 block             5         -      137200         14      137732       6160
  pool            50         -         896        100        1792       3200
 alloc             -         8      383744         10      397365       9472
  -------------------------------------------------------------------------
                                 389169664              578341705   72977920
```

```
  ------------------------ Allocator Memory Report ------------------------
  Name                  Size Alloc DTP    Alloc     Alloc MaxAlloc  MaxAlloc
                             Size       Blocks     Bytes  Blocks     Bytes
  patricia_root_struct    8    12         7741     92892    8130     97560
  sockaddr_un.i802        8    12            2        24       2        24
  sockaddr_un.tag         8    12          371      4452     995     11940
  if_addr_entry           8    12            -         -       1        12
  gw_entry_list           8    12            1        12       1        12
  isis_proto_list         8    12           25       300      30       360
  struct krt_scb         12    16            4        64       6        96
  ldp_rt_data            12    16           61       976     133      2128
  config_list            12    16         2353     37648    2353     37648
  TED NodeInfo           12    16          845     13520     907     14512
  isis_area_addr         12    16          544      8704     612      9792
  isis_nh_list           12    16          237      3792     922     14752
  isis_tsi               12    16           17       272      19       304
  bgp_use_block          12    16            -         -     112      1792
  isis_route_walk_cont   12    16  T         -         -       1        16
  bgpg_rtinfo_entry      12    16  T         -         -  393571   6297136
  task_floating_socket   16    20            1        20       1        20
[...Output truncated...]
  rt_parse_memory      4092  4096  TP        -         -       1      4096
  noblock_buffer_blk   4092  4096  TP        5     20480     811   3321856
  bgp_buffer           4100  8192   P       91    745472     100    819200
  bgp_outbuf           4104  8192   P        -         -      94    770048
  ldp_buffer           4108  8192   P        -         -       7     57344
  RPD SNMP             8268 12288   P        -         -       1     12288
  LDP config            various              1       896       1       896
  ------------------------------------------------------------------------
                                                349037508          543172620

  -------------------------- Malloc Usage Report --------------------------
  Name                       Allocs     Bytes MaxAllocs  MaxBytes  FuncCalls
  MGMT.local                      1        12         1        12          1
  RSVP                            -         -         1      2048     156084
  BGP_Group_Tweak-RTClien         2        24         2        24          2
  [...Output truncated...]
  LDP                             2        24         2        24          2
  KRT Request                     -         -         1        16     446888
  BGP_Group_Packet-Design         2        24         2        24         38
  [...Output truncated...]
  MPLS                        22272   1221656     22274   1221784     228522
  BGP.0.0.0.0+179            186419   2237028    192292   2307504  282141191
  IS-IS I/O./var/run/ppmd         1     66536        43    103916  695536231
  IS-IS                        2407    361372      5887    446076  889294754
  BGP RT Background               3     66556         3     66556          3
  SNMP Subagent./var/run/         -        24         1      9144    3677022
  KRT                             2    205616         3    207900         10
  ASPaths                     13901   1581544     18023   2067605  293868769
  RT                             27       556        28       580       2815
  Scheduler                     194      2604       199      2684      41382
  --Anonymous--          4294944918 4293764616 4294967294 4294967292   45560848
  --System--                  38565  35474324     38684  35487048  235115763
  ------------------------------------------------------------------------
                                     40015436            41923181

    Dynamically allocated memory:  485789696      Maximum:  541736960
         Program data+BSS memory:    2101248      Maximum:    2101248
              Page data overhead:    3039232      Maximum:    3039232
            Page directory size:     512000      Maximum:     512000
                                   ----------
               Total bytes in use:  491442176 (70% of available memory)
```

What It Means The sample output shows summary statistics about the entries in the routing table (show route summary command) and the memory usage breakdown (show task memory detail command) for the routing process (rpd). The two commands provide a comprehensive picture of the memory utilization of the routing protocol process.

The show route summary command shows the number of routes in the various routing tables. In the sample output, the routing tables represented are inet.0, inet.2, inet.3, iso.0, and mpls.0. Within each routing table, all of the active, hold-down, and hidden destinations and routes are summarized for all the protocols from which routes are learned. Routes are in the hold-down state prior to being declared inactive, and hidden routes are not used because of routing policy. Routes in the hold-down and hidden states are still using memory because they appear in the routing table.

In addition, routes are summarized in the following categories: those directly connected to the network (Direct), local routes (Local), and routes learned from configured routing protocols, such as BGP and IS-IS.

The show task memory detail command lists the data structures within the tasks run by the routing protocol process (rpd). Tasks are enabled depending on the router's configuration. For example, isis_area_addr is a data structure resulting from the IS-IS configuration. The AllocBytes field indicates the highest amount of memory used by the data structure. For example, the isis_area_addr data structure has 544 blocks of allocated memory, each block is allocated a value of 16 bytes, resulting in allocated bytes of 8704. The maximum allocated blocks and bytes are high-water marks for a data structure. For more information on displaying task-related information, see "Display Tasks" on page 185.

The Total bytes in use field shows the total amount of memory used by the routing protocol process (rpd).

Step 3: Display Tasks

Purpose You can display information about tasks to further your investigation of a memory problem on the router.

Action To display a list of tasks that are enabled on the router, enter the following JUNOS CLI operational mode commands:

```
user@host> show task
user@host> show task memory
user@host> show task task-name
```

Sample Output

```
user@R1> show task
Pri Task Name                              Pro  Port So Flags
 10 LMP Client                                 17 <>
 10 IF
 15 INET6
 15 INET
 15 ISO
 15 Memory
 20 RPD Unix Domain Server./var/run/rpd_serv.local           21 <>
 20 RPD Unix Domain Server./var/run/rpd_serv.local           20 <>
 20 RPD Unix Domain Server./var/run/rpd_serv.local           19 <>
 20 RPD Unix Domain Server./var/run/rpd_server_communication    16  <Accept>
 20 RPD Server.0.0.0.0+666                    666 15 <Accept>
```

```
 20 Aggregate
 20 RT
 30 ICMP                                  1
 30 Router-Advertisement
 30 ICMPv6                               58  9 <>
 39 OSPFv2 I/O./var/run/ppmd_control                12 <>
 40 l2vpn global task
 40 BGP RT Background                                <LowPrio>
 40 BGP.::+179                                      179 23 <Accept LowPrio>
 40 BGP.0.0.0.0+179                                 179 22 <Accept LowPrio>
 40 BFD I/O./var/run/bfdd_control                   11 <>
 40 OSPF                                 89
 50 BGP_65001.10.0.0.5+3531                        3531 18 <LowPrio>
 50 BGP_65002.10.1.12.2+1224                       1224 25 <LowPrio>
 50 BGP_Group_internal                               <LowPrio>
 50 BGP_Group_toR2                                   <LowPrio>
 50 TED
 50 ASPaths
 51 Resolve inet.0                                   <LowPrio>
 60 KStat                                          13 <>
 60 KRT Request                                     7 <>
 60 KRT Ifstate                          255  6 <>
 60 KRT                                  255  5 <>
 60 Redirect
 70 MGMT.local                                      24 <>
 70 MGMT_Listen./var/run/rpd_mgmt                   14 <Accept>
 70 SNMP Subagent./var/run/snmpd_stream             10 <>
 80 IF Delete

user@R1> show task memory
Memory                 Size (kB)  %Available  When
  Currently In Use:         3490          1%  now
  Maximum Ever Used:        3535          1%  04/02/04 11:54:46
  Available:              220623        100%  now

user@R1> show task io
Task Name                         Reads  Writes    Rcvd    Sent Dropped
 LMP Client                           1       1       0       0       0
 IF                                   0       0       0       0       0
 INET6                                0       0       0       0       0
 INET                                 0       0       0       0       0
 ISO                                  0       0       0       0       0
 Memory                               0       0       0       0       0
 RPD Unix Domain Server./var/ru       1       0       0       0       0
 RPD Unix Domain Server./var/ru       1       0       0       0       0
 RPD Unix Domain Server./var/ru       0       0       0       0       0
 RPD Unix Domain Server./var/ru       3       0       0       0       0
 RPD Server.0.0.0.0+666               0       0       0       0       0
 Aggregate                            0       0       0       0       0
 RT                                   0       0       0       0       0
 ICMP                                 0       0       0       0       0
 Router-Advertisement                 0       0       0       0       0
 ICMPv6                               0       0       0       0       0
 OSPFv2 I/O./var/run/ppmd_contr   31167       1       0       0       0
 l2vpn global task                    0       0       0       0       0
 BGP RT Background                    0       0       0       0       0
 BGP.::+179                           0       0       0       0       0
 BGP.0.0.0.0+179                      8       0       0       0       0
 BFD I/O./var/run/bfdd_control    30731       1       0       0       0
 OSPF                                 0       0       0       0       0
 BGP_65001.10.0.0.5+3531          20486       0       0       0       0
 BGP_65002.10.1.12.2+1224         20489       6       0       0       0
 BGP_Group_internal                   0       0       0       0       0
```

```
BGP_Group_toR2                    0        0        0        0        0
TED                               0        0        0        0        0
ASPaths                           0        0        0        0        0
Resolve inet.0                    0        0        0        0        0
KStat                             0        0        0        0        0
KRT Request                       0        0       57        0        0
KRT Ifstate                      18        0       16        0        0
KRT                               0        0        2        0        0
Redirect                          0        0        0        0        0
MGMT.local                        0        0        0        0        0
MGMT_Listen./var/run/rpd_mgmt    23        0        0        0        0
SNMP Subagent./var/run/snmpd_s   23        0        0        0        0
IF Delete                         0        0        0        0        0
```

What It Means The sample output shows a list of routing, routing protocol, and interface tasks that are currently running on the router (show task), a summary of memory utilization (show task memory), and the memory utilization of a particular task (show task io). Tasks can be baseline tasks performed regardless of the router configuration, and other tasks that depend on the router configuration. For example, the BGP_Group_internal task is the result of the configuration of BGP on the router, while the INET6 task is a base task associated with the routing process (rpd).

Each task in the show task command output has a priority and a task name. For example, the current priority is 10 for LMP Client and 80 for IF Delete. A lower number indicates a higher priority.

Some tasks have flags attached to them. For example, the BGP.0.0.0.0+179 task has two flags, Accept and LowPrio. The Accept flag indicates that the task is waiting for incoming connections, and the LowPrio flag indicates that the task will be dispatched to read its socket after other, higher priority tasks. Two additional flags are Connect, which indicates that a task is waiting for a connection to complete, and Delete, which indicates that a task has been deleted and is being cleaned up.

The show task io command shows the statistics gathered for each IO operation. The counters show the following:

- Reads—This counter increments when a datagram arrives on a connected socket of the task and the task's read callback is called.
- Writes—This counter increments when a connected socket of a task becomes writable and the tasks's callback is called.
- Rcvd—This counter increments when the task calls the Routing Engine to read a datagram from a socket which may or may not be connected.
- Sent—This counter increments when a task attempts to read or write a datagram on an existing or nonexisting socket.
- Drops—This counter increments when a task attempts to read or write a datagram through the Routing Engine on a prebuilt socket, but the request fails for any reason.