Monitoring the System

This chapter provides information for monitoring system status and performance using the **show** commands found in the Command Line Interface (CLI). These command have many related keywords that allow them to provide useful information on all aspects of the system ranging from current software configuration through call activity and status.

The selection of keywords described in this chapter is intended to provide the most useful and in-depth information for monitoring the system. For additional information on these and other **show** command keywords, refer to the *Exec Mode show Commands* chapter of the *Command Line Interface Reference*.

Note A VPC-DI or VPC-SI virtual machine (VM) has no knowledge of the hypervisor under which it is running or the commercial off-the-shelf (COTS) server. To monitor the status of the hypervisor and COTS server, refer to the user documentation supplied with these components of this system.

Important In Release 21.1 and forward, use the **do show** command to run all Exec Mode **show** commands while in Global Configuration Mode. It is not necessary to exit the Config mode to run a **show** command. The pipe character | is only available if the command is valid in the Exec mode.

- SNMP Notifications, page 1
- Monitoring System Status and Performance, page 2
- Monitoring the DI Network, page 3
- Monitoring the SF, page 7
- Clearing Statistics and Counters, page 11

SNMP Notifications

In addition to the CLI, the system supports Simple Network Management Protocol (SNMP) notifications that indicate status and alarm conditions. Refer to the *SNMP MIB Reference* for a detailed listing of these notifications.

Monitoring System Status and Performance

This section contains commands used to monitor the status of tasks, managers, applications and other software components in the system. Output descriptions for most of the commands are located in the *Statistics and Counters Reference*.

Table 1: System Status and Performance Monitoring Commands

To do this:	Enter this command:
View Administrative Information	
Display Current Administrative User Access	
View a list of all administrative users currently logged on the system	**show administrators**
View the context in which the administrative user is working, the IP address from which the administrative user is accessing the CLI, and a system generated ID number	**show administrators session id**
View information pertaining to local-user administrative accounts configured for the system	**show local-user verbose**
View statistics for local-user administrative accounts	**show local-user statistics verbose**
View information pertaining to your CLI session	**show cli**
Determining System Uptime	
View system uptime (time since last reboot)	**show system uptime**
View NTP Server Status	
View NTP servers status	**show ntp status**
View System Resources	
View all system resources such as CPU resources and number of managers created	**show resources [cpu]**
View System Alarms	
View information about all currently outstanding alarms	**show alarm outstanding all verbose**
View system alarm statistics	**show alarm statistics**
View Congestion-Control Statistics	
View Congestion-Control Statistics	**show congestion-control statistics**
View Remote Management Statistics	
View SNMP notification statistics	**show snmp notifies**
View SNMP access statistics	**show snmp accesses**
View SNMP trap history	**show snmp trap history**

To do this:	Enter this command:
View SNMP Trap Statistics	**show snmp trap statistics**
View Port Counters	
View datalink counters for a specific port	**show port datalink counters** *slot#/port#*
View Port Network Processor Unit (NPU) counters for a specific port	**show port npu counters** *slot#/port#*
View System Information and Network Interfaces	
View information about system components, storage devices and network interfaces	**show hardware**
View Card Information and Statistics	
View diagnostics for all cards or for a card in a specific slot/port; (for VPC, slot = VM)	**show card diag** *slot/port*
View detailed information for all cards or a card in a specific slot/port (for VPC, slot = VM)	**show card info** *slot/port*
View operating status for all cards or VMs	**show card table**
View the contents of the boot configuration (param.cfg) file [VPC-DI]	**show cloud configuration**
View information about installed hardware and whether it is optimal or not for a specific card or all cards in the system [VPC-DI]	**show cloud hardware**
View monitored statistics about the VPC-DI network relative to a specific card [VPC-DI]	**show cloud monitor di-network**

Important The commands or keywords/variables that are available are dependent on platform type, product version, and installed license(s).

Important Some commands have different outputs depending on the platform type.

Monitoring the DI Network

The DI network is the private L2 network that interconnects the VMs. The DI network transports user traffic from the received VM to the serving Session Manager on a different VM, and also transports CF to SF

communications such as CLI commands, health checks, status changes. If the link is compromised unexpected things can happen (such as slow response to CLI commands), potentially resulting in service interruption.

The available monitoring capabilities to verify the health of the DI network are detailed here:

Inter-SF DI Network Tests

Each SF periodically sends non-blocking UDP test packets to each of other active and standby SFs, and keeps track of the responses to calculate latency and packet loss. Test packets are sent once per second. Both jumbo and non-jumbo test packets are sent alternately. A non-jumbo UDP test packet has a payload size of 200 bytes, and a jumbo test packet has a payload size of 4000 bytes. These statistics are recorded:

- Dropped packet counts—On receiving a test packet from another SF, the receiving SF sends back a reply. If an SF does not receive a test packet reply within one second, it marks the packet as dropped.
- Dropped jumbo packets—Same calculation as dropped packet counts, but only counts jumbo test packets.
- Number of packets with long latency—If the SF receives a test packet reply after 200 milliseconds, it marks the packet as having long latency.

Note Counters are cleared after an SF reboot.

The reporting interval starts at 15 seconds and can range to 3600 seconds. If there is no error detected during an interval, no warning log is generated and the reporting interval doubles until the interval is 3600 seconds. When an error is detected during an interval, a warning log is generated and the reporting interval is reduced in half until there are no more packets dropped.

If there are any packets lost or long latency counts, a WARNING event is generated. An example warning is shown here:

```
2016-Jan-10+22:00:01.477 [hat 3081 warning] [5/0/5146 <hatcpu:50> hatcpu.c:1307] [software

internal system syslog] Over the past 15 seconds, tests from card 5 to 4 had 1 total drops,

0 jumbo drops, 0 long latency.
```

Use the command **show heartbeat stats card** *cardnumber* **cpu** *cpunumber* to view the statistics collected regarding inter-SF communications.

DI network monitoring is enabled by default. Use the command **debug heartbeat test** to stop and start SF packet tests on specific SFs, or to clear test packet counters on a specific SF.

You can also use the command **show cloud monitor di-network** to display the DI network monitoring statistics. Sample output from the **show cloud monitor di-network summary** command is shown here for Card number 3:

```
Card 3 Test Results:

ToCard   Health   5MinLoss   60MinLoss
   1      Good       0.0%       0.0%
   2      Good       0.0%       0.0%
   4       Bad       6.32%      5.36%
   5      Good       0.0%       0.0%
   6      Good       0.0%       0.0%
```

The display shows the test packet loss rate for the past five minutes and past 60 minutes. If the rate is larger than 1%, the health status is marked as "Bad".

SF to Standby CF DI Network Tests

During an SF boot up, each SF sends both non-jumbo and jumbo ping packets to the standby CF to ensure that the standby CF is reachable.

During SF normal operation, the SF periodically sends non-blocking UDP test packets to the standby CF, and keeps track of the responses to calculate latency and packet loss. This functionality is the same as described for the *Inter-SF DI Network Tests*.

SF Secondary IP Address DI Network Tests

During an SF boot up, each SF sends both non-jumbo and jumbo ping packets to the active CF using the SF primary IP address. In addition, each SF also sends non-jumbo ping packets to active CF using each of its secondary IP addresses. If any of these pings fails, the SF notifies the active CF and the SF reboots.

Standby CF to Active CF DI Network Tests

During the standby CF boot up, the standby CF sends both non-jumbo and jumbo ping packets to the active CF.

DI-Network Bulk Statistics

The **mon-di-net** schema provides the following bulk statistics for monitoring the health of the DI-network on a VPC-DI platform. This information is similar to that provided in the output of the **show cloud monitor di-network summary** Exec mode command.

- src-card – Source card slot number on which monitoring has been performed.
- dest-card – Destination card slot number to which traffic was routed.
- total-pkts-5mins – Total number of packets sent over the past 5 minutes.
- total-drops-5mins – Total number of packets that were dropped over the past 5 minutes.
- total-pkts-60mins – Total number of packets sent over the past 60 minutes.
- total-drops-60mins – Total number of packets that were dropped over the past 60 minutes.
- total-pkts – Total number of all packets sent.
- total-pkts-jumbo – Total number of jumbo packets sent.
- total-drops – Total number of jumbo and non-jumbo test packets that were dropped.
- total-drops-jumbo – Number of jumbo test packets that were dropped.
- latency-warnings – Total number of times the latency has exceeded the threshold.
- long-rtt – Longest Round Trip Time (RTT) in milliseconds.
- average-rtt – Average Round Trip Time (RTT) in milliseconds.

The **mon-di-net** BulkStats Mode command configures the collection of statistics for the Mon-DI-Net schema.

See the *Bulk Statistics* chapter for information about configuring bulk statistic collection.

DI-Network Heartbeat Thresholds

This feature adds the capability to define thresholds for the internal DI-network for percentage heartbeat loss in order to monitor the card-to-card network health in a VPC-DI deployment.

When heartbeat loss (on any of the cards) crosses a set limit of threshold, this feature raises alarms/SNMP trap to indicate the loss.

The internal High Availability Task (HAT) tracks the percentage heartbeat loss over the past 5 minutes and past 60 minutes between cards and can generate SNMP alarms if a threshold has been crossed or a previous alarm has been cleared.

There can be multiple cards in the system and any card can raise this same trap ID but with different card information.

The scope of this funtionality is across the system. It is not specific to any service and is configured at the Global Configuration mode.

See Configure DI-Network Heartbeat Thresholds, on page 6 for instructions to enable this feature.

Configure DI-Network Heartbeat Thresholds

The following steps describe how to configure threshold levels to generate SNMP alarms if the percentage of heartbeats lost exceeds the configured level.

Note The internal High Availability Task (HAT) is alway monitoring the heartbeats across the VMs on the internal DI-Network. This information can be displayed at any time using the **show cloud monitor di-network summary** Exec mode command.

```
configure
  monitoring hat-5min-loss
  threshold hat-hb-5min-loss  high_thresh [ clear low_thresh ]
default threshold hat-hb-5min-loss
  [ default ] threshold poll hat-hb-5min-loss interval duration
configure
  monitoring hat-60min-loss
  threshold hat-hb-60min-loss  high_thresh [ clear low_thresh ]
default threshold hat-hb-60min-loss
  [ default ] threshold poll hat-hb-5min-loss interval duration
```

Note For supplemental information related to this feature, refer to the *Global Configuration Mode Commands* section of the *Command Line Reference*.

The following alarms/traps are generated when these thresholds are exceeded:

- ThreshHatHb5MinLoss / ThreshClearHatHb5MinLoss.
- ThreshHatHb60MinLoss / ThreshClearHatHb60MinLoss.

See the *SNMP MIB Reference* for more details about these alarms/traps.

Monitoring the SF

To view NPU statistics for each active and standby SF, use the **show npu utilization table** command. Statistics are reported for the past five seconds, past five minutes and past 15 minutes. Sample output is shown here:

```
[local]swch91# show npu utilization table

 ******** show npu utilization table card 4 *******
             5-Sec Avg:   lcore00|lcore01|lcore02|lcore03|lcore04|lcore05|lcore06|lcore07|

                      IDLE:           |    99%|       |       |       |       |       |
 |
           QUEUE_PORT_RX:           |     0%|       |       |       |       |       |
 |
           QUEUE_PORT_TX:           |       |       |       |       |       |       |
 |
           QUEUE_VNPU_RX:           |       |       |       |       |       |       |
 |
           QUEUE_VNPU_TX:           |       |       |       |       |       |       |
 |
            QUEUE_KNI_RX:           |       |       |       |       |       |       |
 |
            QUEUE_KNI_TX:           |       |       |       |       |       |       |
 |
         QUEUE_THREAD_KNI:          |       |       |       |       |       |       |
 |
           QUEUE_MCDMA_RX:          |       |       |       |       |       |       |
 |
           QUEUE_MCDMA_TX:          |       |       |       |       |       |       |
 |
       QUEUE_THREAD_MCDMA:          |       |       |       |       |       |       |
 |
        QUEUE_THREAD_VNPU:          |       |       |       |       |       |       |
 |
          QUEUE_CRYPTO_RX:          |       |       |       |       |       |       |
 |
         QUEUE_CRYPTO_IPC:          |       |       |       |       |       |       |
 |
         QUEUE_THREAD_IPC:          |       |       |       |       |       |       |
 |
              MCDMA_FLUSH:          |       |       |       |       |       |       |
 |
    QUEUE_THREAD_TYPE_MAX:          |       |       |       |       |       |       |
 |
             300-Sec Avg:   lcore00|lcore01|lcore02|lcore03|lcore04|lcore05|lcore06|lcore07|

                      IDLE:           |    99%|       |       |       |       |       |
 |
           QUEUE_PORT_RX:           |     0%|       |       |       |       |       |
 |
           QUEUE_PORT_TX:           |       |       |       |       |       |       |
 |
           QUEUE_VNPU_RX:           |       |       |       |       |       |       |
 |
           QUEUE_VNPU_TX:           |       |       |       |       |       |       |
 |
            QUEUE_KNI_RX:           |       |       |       |       |       |       |
 |
            QUEUE_KNI_TX:           |       |       |       |       |       |       |
 |
         QUEUE_THREAD_KNI:          |       |       |       |       |       |       |
 |
           QUEUE_MCDMA_RX:          |       |       |       |       |       |       |
 |
           QUEUE_MCDMA_TX:          |       |       |       |       |       |       |
 |
       QUEUE_THREAD_MCDMA:          |       |       |       |       |       |       |
 |
```

```
      QUEUE_THREAD_VNPU:           |       |       |       |       |       |       |
|
       QUEUE_CRYPTO_RX:            |       |       |       |       |       |       |
|
      QUEUE_CRYPTO_IPC:            |       |       |       |       |       |       |
|
      QUEUE_THREAD_IPC:            |       |       |       |       |       |       |
|
           MCDMA_FLUSH:            |       |       |       |       |       |       |
|
  QUEUE_THREAD_TYPE_MAX:           |       |       |       |       |       |       |
|
          900-Sec Avg:   lcore00|lcore01|lcore02|lcore03|lcore04|lcore05|lcore06|lcore07|

                  IDLE:            |    99%|       |       |       |       |       |
|
         QUEUE_PORT_RX:            |     0%|       |       |       |       |       |
|
         QUEUE_PORT_TX:            |       |       |       |       |       |       |
|
         QUEUE_VNPU_RX:            |       |       |       |       |       |       |
|
         QUEUE_VNPU_TX:            |       |       |       |       |       |       |
|
          QUEUE_KNI_RX:            |       |       |       |       |       |       |
|
          QUEUE_KNI_TX:            |       |       |       |       |       |       |
|
      QUEUE_THREAD_KNI:            |       |       |       |       |       |       |
|
        QUEUE_MCDMA_RX:            |       |       |       |       |       |       |
|
        QUEUE_MCDMA_TX:            |       |       |       |       |       |       |
|
    QUEUE_THREAD_MCDMA:            |       |       |       |       |       |       |
|
     QUEUE_THREAD_VNPU:            |       |       |       |       |       |       |
|
       QUEUE_CRYPTO_RX:            |       |       |       |       |       |       |
|
      QUEUE_CRYPTO_IPC:            |       |       |       |       |       |       |
|
      QUEUE_THREAD_IPC:            |       |       |       |       |       |       |
|
           MCDMA_FLUSH:            |       |       |       |       |       |       |
|
  QUEUE_THREAD_TYPE_MAX:           |       |       |       |       |       |       |
|


thread 1 IDLE                          99.32 %
thread 1 QUEUE_KNI_RX                   0.63 %
thread 1 QUEUE_PORT_RX                  0.05 %
-----------------------------------------------

 ******** show npu utilization table card 5 *******
            5-Sec Avg:   lcore00|lcore01|lcore02|lcore03|lcore04|lcore05|lcore06|lcore07|

                  IDLE:            |    99%|       |       |       |       |       |
|
         QUEUE_PORT_RX:            |       |       |       |       |       |       |
|
         QUEUE_PORT_TX:            |       |       |       |       |       |       |
|
         QUEUE_VNPU_RX:            |       |       |       |       |       |       |
|
         QUEUE_VNPU_TX:            |       |       |       |       |       |       |
|
          QUEUE_KNI_RX:            |     0%|       |       |       |       |       |
|
          QUEUE_KNI_TX:            |       |       |       |       |       |       |
|
```

```
      QUEUE_THREAD_KNI:          |       |       |       |       |       |       |
|
        QUEUE_MCDMA_RX:          |       |       |       |       |       |       |
|
        QUEUE_MCDMA_TX:          |       |       |       |       |       |       |
|
    QUEUE_THREAD_MCDMA:          |       |       |       |       |       |       |
|
     QUEUE_THREAD_VNPU:          |       |       |       |       |       |       |
|
       QUEUE_CRYPTO_RX:          |       |       |       |       |       |       |
|
      QUEUE_CRYPTO_IPC:          |       |       |       |       |       |       |
|
      QUEUE_THREAD_IPC:          |       |       |       |       |       |       |
|
           MCDMA_FLUSH:          |       |       |       |       |       |       |
|
  QUEUE_THREAD_TYPE_MAX:         |       |       |       |       |       |       |
|
         300-Sec Avg:   lcore00|lcore01|lcore02|lcore03|lcore04|lcore05|lcore06|lcore07|
                  IDLE:          |    99%|       |       |       |       |       |
|
         QUEUE_PORT_RX:          |       |       |       |       |       |       |
|
         QUEUE_PORT_TX:          |       |       |       |       |       |       |
|
         QUEUE_VNPU_RX:          |       |       |       |       |       |       |
|
         QUEUE_VNPU_TX:          |       |       |       |       |       |       |
|
          QUEUE_KNI_RX:          |     0%|       |       |       |       |       |
|
          QUEUE_KNI_TX:          |       |       |       |       |       |       |
|
      QUEUE_THREAD_KNI:          |       |       |       |       |       |       |
|
        QUEUE_MCDMA_RX:          |       |       |       |       |       |       |
|
        QUEUE_MCDMA_TX:          |       |       |       |       |       |       |
|
    QUEUE_THREAD_MCDMA:          |       |       |       |       |       |       |
|
     QUEUE_THREAD_VNPU:          |       |       |       |       |       |       |
|
       QUEUE_CRYPTO_RX:          |       |       |       |       |       |       |
|
      QUEUE_CRYPTO_IPC:          |       |       |       |       |       |       |
|
      QUEUE_THREAD_IPC:          |       |       |       |       |       |       |
|
           MCDMA_FLUSH:          |       |       |       |       |       |       |
|
  QUEUE_THREAD_TYPE_MAX:         |       |       |       |       |       |       |
|
         900-Sec Avg:   lcore00|lcore01|lcore02|lcore03|lcore04|lcore05|lcore06|lcore07|
                  IDLE:          |    99%|       |       |       |       |       |
|
         QUEUE_PORT_RX:          |       |       |       |       |       |       |
|
         QUEUE_PORT_TX:          |       |       |       |       |       |       |
|
         QUEUE_VNPU_RX:          |       |       |       |       |       |       |
|
         QUEUE_VNPU_TX:          |       |       |       |       |       |       |
|
          QUEUE_KNI_RX:          |     0%|       |       |       |       |       |
|
          QUEUE_KNI_TX:          |       |       |       |       |       |       |
|
      QUEUE_THREAD_KNI:          |       |       |       |       |       |       |
```

```
|
         QUEUE_MCDMA_RX:           |       |       |       |       |       |       |
|
         QUEUE_MCDMA_TX:           |       |       |       |       |       |       |
|
     QUEUE_THREAD_MCDMA:           |       |       |       |       |       |       |
|
      QUEUE_THREAD_VNPU:           |       |       |       |       |       |       |
|
        QUEUE_CRYPTO_RX:           |       |       |       |       |       |       |
|
       QUEUE_CRYPTO_IPC:           |       |       |       |       |       |       |
|
       QUEUE_THREAD_IPC:           |       |       |       |       |       |       |
|
            MCDMA_FLUSH:           |       |       |       |       |       |       |
|
  QUEUE_THREAD_TYPE_MAX:           |       |       |       |       |       |       |
|

thread 1 IDLE                            99.37 %
thread 1 QUEUE_KNI_RX                     0.55 %
thread 1 QUEUE_PORT_RX                    0.08 %
-----------------------------------------------
```

Table 2: show npu utilization table

Field	Description
IDLE	Idle time in each core
QUEUE_PORT_RX	Time spent processing RX port
QUEUE_PORT_TX	Time spent processing TX port
QUEUE_VNPU_RX	Time spent processing RX vNPU
QUEUE_VNPU_TX	Time spent processing TX vNPU
QUEUE_KNI_RX	Time spent processing RX kernal network interface (KNI). The KNI is the path to the kernal from the IFTASK.
QUEUE_KNI_TX	Time spent processing TX KNI
QUEUE_THREAD_KNI	Thread dedicated to KNI processing
QUEUE_MCDMA_RX	Time spent processing RX multi-channel direct memory acess (DMA) [MCDMA]. The MCDMA is the path from the IFTASK to the SESSMGR.
QUEUE_MCDMA_TX	Time spent processing TX MCDMA.
QUEUE_THREAD_MCDMA	Thread dedicated to MCDMA processing
QUEUE_THREAD_VNPU	Thread dedicated to VNPU processing
QUEUE_CRYPTO_RX	Time spent processing IPSec

Field	Description
QUEUE_CRYPTO_IPC	Time spent processing IPSec inter-process communication (IPC)
MCDMA_FLUSH	Time spent flushing out MCDMA packets
QUEUE_THREAD_TYPE_MAX	Not used

Clearing Statistics and Counters

It may be necessary to periodically clear statistics and counters in order to gather new information. The system provides the ability to clear statistics and counters based on their grouping (PPP, MIPHA, MIPFA, etc.).

Statistics and counters can be cleared using the CLI **clear** command. Refer to the *Exec Mode Commands* chapter of the *Command Line Interface Reference* for detailed information on using this command.