

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, on page 1
- Monitoring System Status and Performance, on page 2
- Monitoring the DI Network, on page 3
- Monitoring the SF, on page 14
- Clearing Statistics and Counters, on page 19

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**
View SNMP Trap Statistics	**show snmp trap statistics**
View Port Counters	

To do this:	Enter this command:
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".

Note Users might observe higher internal DI-network latency when the node is idle or handling low traffic. The increased latency results from upgrades in multiple components, such as drivers, kernel, and libraries, to support new operating systems. This behavior is expected and does not cause any service impact.

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 12 for instructions to enable this feature.

Monitor VPC-DI Network

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	All
Applicable Platform(s)	VPC - DI
Feature Default	Enabled - Always-on
Related Changes in This Release	Not Applicable
Related Documentation	• *Command Line Interface Reference* • *VPC-DI System Administration Guide* • *Statistics and Counters Reference*

Revision History

Revision Details	Release
First introduced.	Pre 21.8

Feature Description

In a DI-network, packet loss occurs when the DI-network ports are saturated or the underlying network infrastructure is unreliable. The Monitor VPC-DI network feature enables the identification and quantification of Control Plane and Data Plane packet loss on the VPC-DI system.

VPC-DI collects and aggregates the Control Plane and Data Plane monitor data for use in CLI reports and threshold alarms.

The feature also provides the ability to set the criteria for the VPC-DI to declare a card fault. Currently, a card fault is raised when a fixed number of High Availability Task (HAT) Control Plane heartbeats between the Active CF and an SF cards are consecutively bounced. The number of consecutive misses can be configured using this feature. This feature adds a secondary Data Plane configuration parameter that can be used to effectively discriminate between the DI-network packet loss and packet processing failure scenarios.

How It Works

The Control Plane and Data Plane monitors generate two fundamental DI-network traffic types on a fixed or recurring basis and tracks losses. The tracking data is meant to provide a view into DI-network communication loss or disruption.

Control Plane packets are typically unicast bi-directional UDP/TCP streams between cards; essentially request and response pairs between StarOS proclets.

Data Plane traffic consists of unicast IP protocol 254 packets transferred between cards. This traffic is service port ingress or egress that the StarOs internally transfers to the appropriate application instance (ingress) or service port interface (egress) and is not acknowledged (that is, there are no response packets). For example, an ingress packet arriving on an SF3 port that the Session Manager instance services on SF5, traverses the DI-network from SF3 to SF5.

All operational cards (that is, CFs and SFs with an Active or Standby operational state) transmit and receive monitor packets. The monitor traffic is fully meshed wherein all cards transmit monitor packets to all other cards and receive monitor packets from all other cards.

Data Plane packets are generated at a rate of 10 per second. Control Plane monitor packets are generated at a rate of 5 per second. The packet headers for both are marked with default priority.

StarOS collects and aggregates the monitor transmit, receive, and drop data for all card connections. The **show cloud monitor controlplane** and **show cloud monitor dataplane** CLI commands display current 15 second, 5 minute, and 60 minute data. The 5 minute and 60 minute loss percentages are available as variables in the bulkstats mon-di-net schema. The 5 minute and 60 minute loss percentages are also accessible as threshold alarms/traps.

Note that low or non-zero drop percentages are normal. Because measurements involve correlation across card pairs that are not perfectly synchronized, a response can arrive in the interval adjacent to the one in which the request was generated. This is reflected as a drop in the request interval.

When seen on a sustained basis, higher drop or loss percentages may indicate DI-network configuration or operational issues, traffic overload, or VM or Host issues. The cloud monitor provides the ability to see and characterize DI-network traffic loss; further investigation is typically required to identify the root cause.

Limitations

The Monitor VPC-DI Network feature has the following limitations.

- Only supported on the VPC-DI platform.

- Not license-controlled.

Configuring the Monitor VPC-DI Network Feature

The following section provides information about the CLI commands available to enable or disable the feature.

Configuring Card Fault Detection

Use the following commands to configure secondary card fault detection criteria. This command is configured in the Global Configuration mode.

```
configure
  high-availability fault-detection card dp-outage seconds
  end
```

NOTES:

- **default**: Restores the default dp-outage value. The default value is 2 seconds.
- Note that the dp-outage deferral is limited. If the consecutive heartbeat bounces are 5 greater than the configured hb-loss parameter, then card failure is declared regardless of the dp-outage configuration.
- The **dp-outage** parameter is restricted to Administrator access on the VPC-DI platform.
- If this CLI is not configured, then the default dp-outage value is 2 seconds.

Configuring Packet Loss Threshold on Control Plane

Use the following commands to measure percentage packet loss over the corresponding time interval on the Control plane. The threshold alarm and SNMP trap are raised for any card to card connection that exceeds the configured loss percentage over the indicated time period. This command is configured in the Global Configuration mode.

```
configure
  [ default ] threshold cp-monitor-5min-loss pct [ clear pct ]
  end
```

```
[ default ] threshold poll cp-monitor-5min-loss interval duration
```

```
configure
  [default] threshold cp-monitor-60min-loss pct [ clear pct ]
  end
```

```
[default] threshold poll cp-monitor-60min-loss interval duration
```

NOTES:

- **default**: Clears the configured thresholds for the Control Plane.
- **clear** *pct* : Clears the configured percentage of packet loss.
- **interval** *duration*: Specifies the amount of time (in seconds) that comprises the polling interval. *duration* must be an integer from 60 through 60000. The default is 300 seconds.
- This command is disabled by default.

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:

- ThreshControlPlaneMonitor5MinsLoss / ThreshClearControlPlaneMonitor5MinsLoss
- ThreshControlPlaneMonitor60MinsLoss / ThreshControlPlaneMonitor60MinsLoss

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

Configuring Packet Loss Threshold on Data Plane

Use the following commands to measure percentage packet loss over the corresponding time interval on the Data plane. The threshold alarm and SNMP trap are raised for any card to card connection that exceeds the configured loss percentage over the indicated time period. This command is configured in the Global Configuration mode.

```
configure
   [ default ] threshold dp-monitor-5min-loss pct [ clear pct ]
   end
```

```
[ default ] threshold poll dp-monitor-5min-loss interval duration
```

```
configure
   [default] threshold dp-monitor-60min-loss pct [ clear pct ]
   end
```

```
[ default ] threshold poll dp-monitor-60min-loss interval duration
```

NOTES:

- **default**: Disables the configured thresholds for the Data Plane.
- **clear** *pct* : Clears the configured packet loss.
- **interval** *duration*: Specifies the amount of time (in seconds) that comprises the polling interval. *duration* must be an integer from 60 through 60000. The default is 300 seconds.
- This command is disabled by default.

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:

- ThreshDataPlaneMonitor5MinsLoss / ThreshClearDataPlaneMonitor5MinsLoss
- ThreshDataPlaneMonitor60MinsLoss / ThreshDataPlaneMonitor60MinsLoss

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

Monitoring and Troubleshooting

This section provides information regarding CLI commands available in support of monitoring and troubleshooting the feature.

Show Command(s) and/or Outputs

This section provides information regarding show commands and/or their outputs in support of this feature.

show cloud monitor controlplane

This new show command is introduced to display the following output for the most recent Control Plane monitor information.

show cloud monitor controlplane

```
  Cards      15 Second Interval     5 Minute Interval    60 Minute Interval
 Src  Dst    Xmit   Recv  Miss%    Xmit   Recv  Miss%    Xmit   Recv  Miss%
 ---  ---   ------ ------ ------  ------ ------ ------  ------ ------ ------
  01   02      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%
  01   03      75     75   0.0%    1500   1500   0.0%   18000  17996   0.0%
  01   04      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%
  01   05      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%
  01   06      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%

  02   01      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%
  02   03      75     75   0.0%    1500   1500   0.0%   18000  17997   0.0%
  02   04      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%
  02   05      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%
  02   06      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%

  03   01      75     75   0.0%    1500   1500   0.0%      -incomplete-
  03   02      75     75   0.0%    1500   1500   0.0%      -incomplete-
  03   04      75     75   0.0%    1500   1500   0.0%      -incomplete-
  03   05      75     75   0.0%    1500   1500   0.0%      -incomplete-
  03   06      75     75   0.0%    1500   1500   0.0%      -incomplete-

  04   01      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%
  04   02      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%
  04   03      75     75   0.0%    1500   1500   0.0%   18000  17996   0.0%
  04   05      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%
  04   06      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%

  05   01      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%
  05   02      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%
  05   03      75     75   0.0%    1500   1500   0.0%   18000  17996   0.0%
  05   04      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%
  05   06      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%

  06   01      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%
  06   02      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%
  06   03      75     75   0.0%    1500   1500   0.0%   18000  17997   0.0%
  06   04      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%
  06   05      75     75   0.0%    1500   1500   0.0%   18000  18000   0.0%
```

show cloud monitor dataplane

This new show command is introduced to display the following output for the most recent Data Plane monitor information.

show cloud monitor dataplane

```
  Cards      15 Second Interval     5 Minute Interval    60 Minute Interval
```

```
 Src  Dst    Miss     Hit     Pct     Miss     Hit     Pct     Miss     Hit     Pct
 ---  ---  ------ ------ ------  ------ ------ ------  ------ ------ ------
  02   01       0    150   0.0%       0   3000   0.0%       0  36000   0.0%
  03   01       0    150   0.0%       0   3000   0.0%      -incomplete-
  04   01       0    151   0.0%       0   3000   0.0%       0  36000   0.0%
  05   01       0    151   0.0%       0   3000   0.0%       0  36001   0.0%
  06   01       0    150   0.0%       0   3000   0.0%       2  35998   0.0%

  01   02       0    150   0.0%       0   3000   0.0%       0  36000   0.0%
  03   02       0    150   0.0%       0   3000   0.0%      -incomplete-
  04   02       0    150   0.0%       0   3000   0.0%       0  36000   0.0%
  05   02       0    150   0.0%       0   3000   0.0%       0  36000   0.0%
  06   02       0    151   0.0%       0   3001   0.0%       1  35999   0.0%

  01   03       0    151   0.0%       0   3000   0.0%      -incomplete-
  02   03       0    151   0.0%       0   3000   0.0%      -incomplete-
  04   03       0    150   0.0%       0   3000   0.0%      -incomplete-
  05   03       0    150   0.0%       0   3000   0.0%      -incomplete-
  06   03       0    151   0.0%       0   3000   0.0%      -incomplete-

  01   04       0    150   0.0%       0   3001   0.0%       0  36001   0.0%
  02   04       0    150   0.0%       0   3000   0.0%       0  36000   0.0%
  03   04       0    150   0.0%       0   3000   0.0%      -incomplete-
  05   04       1    149   0.7%       1   2999   0.0%       0  36001   0.0%
  06   04       0    150   0.0%       0   3000   0.0%       2  35998   0.0%

  01   05       1    149   0.7%       1   2999   0.0%       0  36000   0.0%
  02   05       0    150   0.0%       0   3000   0.0%       0  36000   0.0%
  03   05       0    150   0.0%       0   3000   0.0%      -incomplete-
  04   05       0    150   0.0%       1   2999   0.0%       1  35999   0.0%
  06   05       0    150   0.0%       0   3000   0.0%       2  35998   0.0%

  01   06       0    150   0.0%       0   3001   0.0%       0  36001   0.0%
  02   06       0    151   0.0%       0   3000   0.0%       1  35999   0.0%
  03   06       0    150   0.0%       0   3001   0.0%      -incomplete-
  04   06       0    150   0.0%       0   3000   0.0%       0  36000   0.0%
  05   06       0    150   0.0%       0   3000   0.0%       0  36000   0.0%
```

Bulk Statistics

The following statistics are included in support of this feature.

mon-di-net Schema

The following bulk statistics are added in the mon-di-net schema in support of the Monitor the VPC-DI Network feature.

Bulk Statistics	Description
cp-loss-5minave	Indicates the average Control Plane loss in prior 5 minutes.
cp-loss-60minave	Indicates the average Control Plane loss in prior 60 minutes.
dp-loss-5minave	Indicates the average Data Plane loss in prior 5 minutes.
dp-loss-60minave	Indicates the average Data Plane loss in prior 60 minutes.

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.

Configuration Support for Heartbeat Value

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	All
Applicable Platform(s)	• ASR 5500 • VPC - DI
Feature Default	Disabled - Configuration Required

Related Changes in This Release	Not Applicable
Related Documentation	• *ASR 5500 System Administration Guide* • *Command Line Interface Reference* • *VPC-DI System Administration Guide* • *Statistics and Counters Reference*

Revision History

Revision Details	Release
In this release, the default heartbeat value between the management and data processing card can be modified to prevent the management card from incorrectly detecting and reporting the packet processing card as failed.	21.8
First introduced.	Pre 21.2

Feature Changes

In certain deployment scenarios, the management card reports the packet processing card as failed when it is unable to detect a heartbeat for about two seconds. This assumed failure is observed when the hearbeat is delayed or lost due to congestion in the internal DI network.

This release addresses this issue.

Previous Behavior: The management card reports the packet processing card as failed due to its inability to detect the heartbeat within the default value of two seconds, thereby causing an unplanned switchover.

New Behavior: To prevent the management card from incorrectly detecting and reporting the packet processing card as failed, the default heartbeat value between the management and data processing card can now be modified.

Customer Impact: Prevents the management card from wrongful reporting of the data processing card and unplanned switchover.

Command Changes

high-availability fault-detection

The above CLI command is enhanced to include the **card hb-loss** *value* keyword, which is used to configure the heartbeat value between the management and packet processing cards. This command is configured in the Global Configuration Mode.

```
configure
  [default] high-availability fault-detection card hb-loss value
  end
```

NOTES:

- **default**: Restores the heartbeat value to the default value of 2 heartbeats.

- **card**: Specifies the packet processing card.
- **hb-loss** *value*: Configures the heartbeat loss value. The default value is 2 heartbeats.
- The heartbeat value between a management to management card is set to the default value of 2 heartbeats.
- This command modifies the heartbeat value only between the management and packet processing cards.
- By default, this CLI is disabled.

Monitoring and Troubleshooting

This section provides information regarding show commands and/or their outputs in support of this feature.

show heartbeat stats hb-loss all

This show command now includes the value for the following new fields for the all packet processing cards.

- Max Bounces
- Total HB Miss
- Total HB Card Failure
  - Card/Cpu
  - Total
  - Age/Intf/Seqno/TimeStamp
  - AFD(oldest first)

show heartbeat stats hb-loss card *card-number*

This show command now includes the value for the following new fields for the specified packet processing card.

- Max Bounces
- Total HB Miss
- Total HB Card Failure
  - Card/Cpu
  - Total
  - Age/Intf/Seqno/TimeStamp
  - AFD(oldest first)

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.

Field	Description
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
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.