



CHAPTER 9

InfiniBand Performance Management Tasks

These topics describe the InfiniBand menu tasks for Element Manager that relate to performance management:

- [Using the InfiniBand Menu, page 9-1](#)
- [Enabling and Disabling InfiniBand Port Performance Management, page 9-2](#)
- [Enabling and Managing Port Monitoring, page 9-2](#)
- [Monitoring Connections, page 9-4](#)
- [Monitoring Connections, page 9-4](#)
- [Viewing InfiniBand Port Counters, page 9-12](#)



Note

See [Appendix A, “InfiniBand Concepts”](#) to familiarize yourself with the InfiniBand technology. For hardware-specific information, consult the relevant hardware documentation.

Using the InfiniBand Menu

The InfiniBand menu has two choices for performing InfiniBand performance management tasks:

- Performance Management
- Performance Management (tabular format)

This section describes how to use the Performance Management menu option. Most of the tasks can also be performed by choosing the Performance Management (tabular format) menu option, which presents information and configurable options in tables, but is a less user-friendly way to perform your InfiniBand performance management tasks.

Enabling and Disabling InfiniBand Port Performance Management

Use performance management to view InfiniBand port counters, test connectivity between InfiniBand ports, and monitor InfiniBand ports for errors.

These topics describe how to enable and disable InfiniBand port performance management:

- [Enabling Performance Management, page 9-2](#)
- [Disabling Performance Management, page 9-2](#)

Enabling Performance Management

To enable InfiniBand port performance management, follow these steps:

-
- | | |
|---------------|--|
| Step 1 | From the InfiniBand menu, choose Performance Management .
The Performance Management window opens. |
| Step 2 | Click the subnet of the ports that you want to manage (for instance, fe:80:00:00:00:00:00).
The Port Counter Configuration display appears in the right pane of the window. |
| Step 3 | Click the Enable radio button. |
-

Disabling Performance Management

To disable performance management, follow these steps:

-
- | | |
|---------------|--|
| Step 1 | From the InfiniBand menu, choose Performance Management .
The Performance Management window opens. |
| Step 2 | Click the subnet of the ports that you want to manage (for instance, fe:80:00:00:00:00:00).
The Port Counter Configuration display appears in the right pane of the window. |
| Step 3 | Click the Disable radio button. |
-

Enabling and Managing Port Monitoring

These topics describe how to enable and manage port monitoring:

- [Enabling Port Monitoring, page 9-3](#)
- [Configuring Port Monitoring, page 9-3](#)
- [Configuring Port Monitoring Thresholds, page 9-3](#)
- [Viewing Port Monitoring Errors, page 9-4](#)

Enabling Port Monitoring

To enable port monitoring, follow these steps:

-
- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.
 - Step 2** Expand the subnet of the connections that you want to monitor.
 - Step 3** Select the **Port Monitor** branch.
 - Step 4** Click the **General** tab.
 - Step 5** From the State drop-down menu, choose **Enable**.



Note **Enable** enables port monitoring only for the ports that are configured in the Monitor Port Config table; **enableAll** enables port monitoring for all ports regardless of whether the port is configured in the Monitor Port Config table or not.

- Step 6** Click **Apply**.
-

Configuring Port Monitoring

To configure port monitoring, follow these steps:

-
- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.
 - Step 2** Expand the subnet of the connections that you want to monitor.
The navigation tree expands.
 - Step 3** Select the **Port Monitor** branch.
 - Step 4** Click the **General** tab.
 - Step 5** In the Polling Period field, enter an integer value between 1 and 600 to configure the number of seconds between polls.
 - Step 6** In the Start Delay field, enter an integer value between 1 and 600 to configure the delay between startup and polling.
-

Configuring Port Monitoring Thresholds

To configure port monitoring thresholds, follow these steps:

-
- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.

- Step 2** Expand the subnet of the connections that you want to monitor.
The navigation tree expands.
- Step 3** Select the **Port Monitor** branch.
- Step 4** Click the **Threshold** tab.
- Step 5** Enter an integer value in the fields where you want to apply a threshold. Enter **none** in the fields to which you do not want to apply a threshold.
- Step 6** Click **Apply**.
-

Viewing Port Monitoring Errors

To view port monitoring errors, follow these steps:

-
- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.
- Step 2** Expand the subnet of the connections that you want to monitor.
The navigation tree expands.
- Step 3** Select the **Port Monitor** branch.
- Step 4** Click the **Port Errors** tab.
Port errors are displayed.
-

Monitoring Connections

To monitor connections, you complete tasks such as:

- [Creating a Connection to Monitor, page 9-4](#)
- [Viewing Monitored Connections, page 9-5](#)
- [Viewing Connection Counters, page 9-6](#)
- [Viewing Connection Monitor Counters, page 9-7](#)
- [Testing Connections, page 9-8](#)
- [Viewing Port Counters of Connections, page 9-8](#)

Creating a Connection to Monitor

To create a connection to monitor, follow these steps:

-
- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.
- Step 2** Expand the subnet of the connections that you want to monitor.

Step 3 Choose **Connection Counters**.

The Monitored Connection tab appears in the right pane of the window.

Step 4 Click **Add**.

The Add Connection window opens.

Step 5 In the Source LID field, enter a source LID.**Note**

To view available source and destination LIDs, return to the main Element Manager display, click the **InfiniBand** menu, choose **Subnet Management**, and then click the **SwitchRoute** tab. For more information, see the [“Viewing and Managing InfiniBand Routes” section on page 8-33](#).

Step 6 In the Destination LID field, enter a destination LID.**Step 7** Check the **Enable Connection Monitoring** check box.**Note**

If this check box is not selected, you can view only counter information and cannot view monitoring information.

Step 8 Click **Add**.

The connection entry appears under the Monitored Connections tab.

Viewing Monitored Connections

These instructions assume that you have already defined connections to monitor. To view monitored connections, follow these steps:

Step 1 From the InfiniBand menu, choose **Performance Management**.

The Performance Management window opens.

Step 2 Expand the subnet of the connections that you want to monitor.

The navigation tree expands.

Step 3 Select the **Connection Counters** branch.

The Monitored Connection tab appears in the right pane of the window. [Table 9-1](#) describes the fields in this pane.

Table 9-1 *Monitored Connections Field Descriptions Pane*

Field	Description
Subnet Prefix	Subnet prefix of the monitored connection.
Source LID	16-bit source Local ID of the connection.
Destination LID	16-bit destination Local ID of the connection.

Table 9-1 *Monitored Connections Field Descriptions Pane (continued)*

Field	Description
Error Status	Displays unknown, exceeded, or notExceeded to indicate if the error value has exceeded the threshold that you configured. To configure thresholds, see the “Configuring Port Monitoring Thresholds” section on page 9-3 .
Util Status	Displays unknown, exceeded, or notExceeded to indicate if the utilization value has exceeded the threshold that you configured. To configure thresholds, see the “Configuring Port Monitoring Thresholds” section on page 9-3 .

Viewing Connection Counters

Each hop in the display is a port on a node. When connections move through nodes, they enter the node in one hop (GUID A, port a), and exit in another hop (GUID A, port b). Though the GUIDs of subsequent hops may match, the ports do not match. To view connection counters, follow these steps:

- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.
- Step 2** Expand the subnet of the connections that you want to monitor.
- Step 3** Expand the **Connection Counters** branch.
- Step 4** Select the connection with counters that you want to view.
- Step 5** Click the **Connection Counters** tab.

[Table 9-2](#) describes the fields in the display.

Table 9-2 *Connection Counters Field Descriptions*

Field	Description
Subnet Prefix	Subnet prefix of the subnet on which each hop resides.
Node Guid	Global unique ID of the node (switch chip, HCA, or TCA) of the next-hop port.
Port Number	Port number (on the appropriate node) of the hop.
Chassis Guid	Global unique ID (GUID) of the chassis.
Slot Number	Slot of the port.
Ext Port Number	External port number of the port.
Data Is Valid	Displays true or false.
Symbol Errors	Number of symbol errors on the port.
Link Recovery Errors	Number of link recovery errors on the port.
Link Downs	Number of link-down errors on the port.
Received Errors	Number of received errors that the port experienced.
Received Remote Physical Errors	Number of physical errors that the port experienced.

Table 9-2 *Connection Counters Field Descriptions (continued)*

Field	Description
Received Switch Relay Errors	Number of switch relay errors that the port experienced.
Transmitted Discards	Number of transmitted discards that occurred on the port.
Transmitted Constraint Errors	Number of Transmitted Constraint errors that the port experienced.
Received Constraint Errors	Number of Received Constraint errors that the port experienced.
Local Link Integrity Errors	Number of logical link integrity errors on the port.
Excessive Buffer Overrun Errors	Number of excessive buffer overrun errors on the port.
VL15 Dropped	Number of VL15 drops on the port.
Transmitted Data	Volume of transmitted data on the port.
Received Data	Volume of received data on the port.
Transmitted Packets	Volume of transmitted packets on the port.
Received Packets	Volume of received packets on the port.

Viewing Connection Monitor Counters

To view connection monitor counters, follow these steps:

- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.
- Step 2** Expand the subnet of the connections that you want to monitor.
- Step 3** Expand the **Connection Counters** branch.
- Step 4** Select the connection with counters that you want to view.
- Step 5** Click the **Connection Monitor Counters** tab.

[Table 9-3](#) describes the fields in the tab.

Table 9-3 *Connection Monitor Counters Field Descriptions*

Field	Description
Node Guid	Global unique ID of the InfiniBand node of the hop port.
Port Number	Port number of the hop.
Chassis Guid	GUID of the chassis that includes the connection.
Slot Number	Slot number of the port(s) in the connection.
Ext Port Number	External port number of the connection port.
Error Type	Type of error that occurred.

Testing Connections

To test connections, follow these steps:

-
- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.
 - Step 2** Expand the subnet of the connections that you want to monitor.
 - Step 3** Expand the **Connection Counters** branch.
 - Step 4** Select the connection with counters that you want to view.
 - Step 5** Click the **Test Connection** tab.
 - Step 6** Click **Test**.
-

Viewing Port Counters of Connections

To view port counters, follow these steps:

-
- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.
 - Step 2** Expand the subnet of the connections that you want to monitor.
 - Step 3** Expand the **Connection Counters** branch.
 - Step 4** Expand the connection with port counters that you want to view.
 - Step 5** Select the port (in GUID port number format) with counters that you want to view.

[Table 9-4](#) describes the fields in this display.

Table 9-4 Port Counters Field Descriptions

Field	Description
Subnet Prefix	Subnet prefix of the subnet on which each hop resides.
Node Guid	Global unique ID of the node (switch chip, HCA, or TCA) of the next-hop port.
Port Number	Port number (on the appropriate node) of the hop.
Chassis Guid	GUID of the chassis that includes the connection.
Slot Number	Slot number of the port(s) in the connection.
Ext Port Number	External port number of the connection port.
Symbol Errors	Total number of symbol errors detected on one or more lanes.
Link Recovery Errors	Total number of times the port training state machine has successfully completed the link error recovery process.
Link Downs	Total number of times that the port training state machine has failed the link error recovery process and downed the link.

Table 9-4 Port Counters Field Descriptions (continued)

Field	Description
Received Errors	Total number of packets containing an error that was received on the port. These errors are as follows: <ul style="list-style-type: none"> Local physical errors (ICRC, VCRC, FCCRC, and all physical errors that cause entry into the bad state) Malformed data packet errors (Lver, length, VL) Malformed link packet errors (operand, length, VL) Packets discarded due to buffer overrun
Received Remote Physical Errors	Total number of packets marked with the EBP delimiter received on the port.
Received Switch Relay Errors	Total number of packets received on the port that were discarded because they could be forwarded by the switch relay. Reasons for this are as follows: <ul style="list-style-type: none"> DLID mapping. VL mapping. Looping (output port = input port).
Transmitted Discards	Total number of outbound packets discarded by the port because the port is down or congested. Reasons for this are as follows: <ul style="list-style-type: none"> Output port is in the inactive state. Packet length has exceeded neighbor MTU. Switch lifetime limit has been exceeded. Switch HoQ limit has been exceeded.
Transmitted Constraint Errors	Total number of packets not transmitted from the port for the following reasons: <ul style="list-style-type: none"> FilterRawOutbound is true and packet is raw. PartitionEnforcementOutbound is true and packet fails the partition key check, the IP version check, or the transport header version check.
Received Constraint Errors	Total number of packets received on the port that are discarded for the following reasons: <ul style="list-style-type: none"> FilterRawInbound is true, and packet is raw. PartitionEnforcementInbound is true and the packet fails the partition key check, the IP version check, or the transport header version check.
Local Link Integrity Errors	Number of times that the frequency of packets containing local physical errors exceeded local_phy_errors.
Excessive Buffer Overrun Errors	Number of times that overrun errors' consecutive flow control update periods occurred with at least one overrun error in each period.
VL15 Dropped	Number of incoming VL15 packets dropped due to resource limitations on port selected by PortSelect.

Table 9-4 Port Counters Field Descriptions (continued)

Field	Description
Transmitted Data	<p>(Optional) Will be zero if not implemented. Total number of data octets, divided by 4, transmitted on all VLs from the port selected by PortSelect. This includes all octets between (and not including) the start of packet delimiter and VCRC. It excludes all link packets.</p> <p>You may choose to count data octets in groups larger than four but are encouraged to choose the smallest group possible. Results are still reported as a multiple of four octets.</p>
Received Data	<p>(Optional) Will be zero if not implemented. Total number of data octets, divided by 4, received on all VLs from the port selected by PortSelect. This includes all octets between (and not including) the start of packet delimiter and VCRC. It excludes all link packets.</p> <p>You may choose to count data octets in groups larger than four but are encouraged to choose the smallest group possible. Results are still reported as a multiple of four octets.</p>
Transmitted Packets	<p>(Optional) Will be zero if not implemented. Total number of data packets, excluding link packets, transmitted on all VLs from the port selected by PortSelect.</p>
Received Packets	<p>(Optional) Will be zero if not implemented. Total number of data packets, excluding link packets, received on all VLs from the port selected by PortSelect.</p>

Resetting Counters

You can reset counters for the following:

- [Resetting Counters on a Hop, page 9-10](#)
- [Resetting Counters on All Ports on a Node, page 9-11](#)
- [Resetting Counters on All Ports in a Connection, page 9-11](#)
- [Resetting All Counters in a Subnet, page 9-11](#)

Resetting Counters on a Hop

To reset counters on a hop, follow these steps:

-
- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.
- Step 2** Expand the subnet of the connections that you want to monitor.
- Step 3** Expand the **Connection Counters** branch.

- Step 4** Expand the connection that includes the hop that you want to clear.
 - Step 5** Right-click the hop with counters you want to clear, and choose **Clear counters on this Hop**.
-

Resetting Counters on All Ports on a Node

To reset counters on all ports of a node, follow these steps:

-
- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.
 - Step 2** Expand the subnet of the connections that you want to monitor.
 - Step 3** Expand the **Connection Counters** branch.
 - Step 4** Expand the connection that includes the node that you want to clear.
 - Step 5** Right-click the node with counters you want to clear, and choose **Clear counters on this Node**.
-

Resetting Counters on All Ports in a Connection

To reset counters on all ports in a connection, follow these steps:

-
- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.
 - Step 2** Expand the subnet of the connections that you want to monitor.
 - Step 3** Expand the **Connection Counters** branch.
 - Step 4** Right-click the connection with counters you want to clear, and choose **Clear counters on this Connection**.
-

Resetting All Counters in a Subnet

To reset all counters in a subnet, follow these steps:

-
- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.
 - Step 2** Expand the subnet of the connections that you want to monitor.
 - Step 3** Right-click the **Connection Counters** branch, and choose **Clear Counters for All Connections**.
-

Viewing InfiniBand Port Counters

These topics describe how to view InfiniBand port counters:

- [Viewing Port Counters, page 9-12](#)
- [Enabling or Disabling Monitoring a Port, page 9-14](#)
- [Enabling or Disabling Monitoring a Node, page 9-15](#)
- [Viewing Cumulative Port Counters, page 9-15](#)

Viewing Port Counters

To view port counters, follow these steps:

-
- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.
- Step 2** Expand the subnet of the connections that you want to monitor.
- Step 3** Expand the **Port Counters** branch.
- Step 4** View port counters using one of the following methods:
- Click the GUID with port counters that you want to view; all available port counters appear.
 - Expand the GUID of the node with port counters that you want to view, and then select the port with counters that you want to view.

Counters appear for that individual port. [Table 9-5](#) describes the fields in the port counters display.

Table 9-5 *Port Counters Field Descriptions*

Field	Description
Subnet Prefix	Subnet prefix of the subnet on which each hop resides.
Node Guid	Global unique ID of the node (switch chip, HCA, or TCA) of the next-hop port.
Port Number	Port number (on the appropriate node) of the hop.
Chassis Guid	GUID of the chassis that includes the connection.
Slot Number	Slot number of the port(s) in the connection.
Ext Port Number	External port number of the connection port.
Symbol Errors	Total number of symbol errors detected on one or more lanes.
Link Recovery Errors	Total number of times the port training state machine has successfully completed the link error recovery process.
Link Downs	Total number of times the port training state machine has failed the link error recovery process and downed the link.

Table 9-5 Port Counters Field Descriptions (continued)

Field	Description
Received Errors	Total number of packets containing an error that were received on the port. These errors are as follows: <ul style="list-style-type: none"> Local physical errors (ICRC, VCRC, FCCRC, and all physical errors that cause entry into the “bad” state) Malformed data packet errors (Lver, length, VL) Malformed link packet errors (operand, length, VL) Packets discarded due to buffer overrun
Received Remote Physical Errors	Total number of packets marked with the EBP delimiter received on the port.
Received Switch Relay Errors	Total number of packets received on the port that were discarded because they could be forwarded by the switch relay. Reasons for this are as follows: <ul style="list-style-type: none"> DLID mapping. VL mapping. Looping (output port = input port).
Transmitted Discards	Total number of outbound packets discarded by the port because the port is down or congested. Reasons for this are as follows: <ul style="list-style-type: none"> Output port is in the inactive state. Packet length has exceeded neighbor MTU. Switch lifetime limit has been exceeded. Switch HoQ limit has been exceeded.
Transmitted Constraint Errors	Total number of packets not transmitted from the port for the following reasons: <ul style="list-style-type: none"> FilterRawOutbound is true, and packet is raw. PartitionEnforcementOutbound is true and the packet fails the partition key check, the IP version check, or the transport header version check.
Received Constraint Errors	Total number of packets received on the port that are discarded for the following reasons: <ul style="list-style-type: none"> FilterRawInbound is true, and packet is raw. PartitionEnforcementInbound is true and packet fails partition key check, IP version check, or transport header version check.
Logical Link Integrity Errors	Number of times that the frequency of packets containing local physical errors exceeded local_phy_errors.
Excessive Buffer Overrun Errors	Number of times that overrun errors consecutive flow control update periods occurred with at least one overrun error in each period.
VL15 Dropped	Number of incoming VL15 packets dropped due to resource limitations on port selected by PortSelect.

Table 9-5 Port Counters Field Descriptions (continued)

Field	Description
Transmitted Data	<p>(Optional) Will be zero if not implemented. Total number of data octets, divided by 4, transmitted on all VLs from the port selected by PortSelect. This includes all octets between (and not including) the start of the packet delimiter and the VCRC. It excludes all link packets.</p> <p>You may choose to count data octets in groups larger than four but are encouraged to choose the smallest group possible. Results are still reported as a multiple of four octets.</p>
Received Data	<p>(Optional) Will be zero if not implemented. Total number of data octets, divided by 4, received on all VLs from the port selected by PortSelect. This includes all octets between (and not including) the start of the packet delimiter and the VCRC. It excludes all link packets.</p> <p>You may choose to count data octets in groups larger than four but are encouraged to choose the smallest group possible. Results are still reported as a multiple of four octets.</p>
Transmitted Packets	(Optional) Will be zero if not implemented. Total number of data packets, excluding link packets, transmitted on all VLs from the port selected by PortSelect.
Received Packets	(Optional) Will be zero if not implemented. Total number of data packets, excluding link packets, received on all VLs from the port selected by PortSelect.

Enabling or Disabling Monitoring a Port

To enable or disable port monitoring for a specific port, follow these steps:

- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.
- Step 2** Expand the subnet of the connections that you want to monitor.
- Step 3** Expand the **Port Counters** branch.
- Step 4** Expand the GUID of the node with port counters that you want enable or disable.
- Step 5** Right-click the port for which you want to enable or disable monitoring.
- Step 6** From the drop-down menu, select **Enable Port Monitoring** or **Disable Port Monitoring**.

Enabling or Disabling Monitoring a Node

To enable or disable port monitoring for a specific node, follow these steps:

-
- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.
 - Step 2** Expand the subnet of the connections that you want to monitor.
 - Step 3** Expand the **Port Counters** branch.
 - Step 4** Right-click the GUID of the node with the port counters that you want enable or disable.
 - Step 5** From the drop-down menu, select **Enable Port Monitoring on this Node** or **Disable Port Monitoring on this Node**.
-

Viewing Cumulative Port Counters

To view cumulative port counters, follow these steps:

-
- Step 1** From the InfiniBand menu, choose **Performance Management**.
The Performance Management window opens.
 - Step 2** Expand the subnet of the connections that you want to monitor.
 - Step 3** Expand the **Port Counters** branch.
 - Step 4** Expand the node of the port with cumulative counters that you want to view.
 - Step 5** Click the port with navigation counters that you want to view.
 - Step 6** Click the **Port Cumulative Counters** tab.

[Table 9-6](#) describes the fields in the tab.

Table 9-6 Cumulative Port Counters Field Descriptions

Field	Description
Subnet Prefix	Subnet prefix of the subnet on which each hop resides.
Node Guid	Global unique ID (GUID) of the node (switch chip, HCA, or TCA) of the next-hop port.
Port Number	Port number (on the appropriate node) of the hop.
Chassis Guid	GUID of the chassis.
Slot Number	Slot of the port.
Ext Port Number	External port number of the port.
Error Status	Displays true or false.
Util Status	Displays true or false.
Symbol Errors	Number of symbol errors on the port.

Table 9-6 Cumulative Port Counters Field Descriptions (continued)

Field	Description
Link Recovery Errors	Number of link recovery errors on the port.
Link Downs	Number of link-down errors on the port.
Received Errors	Number of received errors that the port experienced.
Received Remote Physical Errors	Number of physical errors that the port experienced.
Received Switch Relay Errors	Number of switch relay errors that the port experienced.
Transmit Discards	Number of transmitted discards that occurred on the port.
Transmit Constraint Errors	Number of Transmit Constraint errors that the port experienced.
Received Constraint Errors	Number of Received Constraint errors that the port experienced.
Logical Link Integrity Errors	Number of logical link integrity errors on the port.
Excessive Buffer Overrun Errors	Number of excessive buffer overrun errors on the port.
VL15 Dropped	Number of VL15 drops on the port.
Transmit Data	Volume of transmitted data on the port.
Received Data	Volume of received data on the port.
Transmit Packets	Volume of transmitted packets on the port.
Received Packets	Volume of received packets on the port.
Transmit Rate	Rate at which the packets are transmitted.
Received Rate	Rate at which the packets are received.