



Managing Faults with Cisco MGC Node Manager

This chapter contains the following sections:

- [Overview of Fault Management Features, page 6-1](#)
- [Managing Faults with Cisco MGC Node Manager, page 6-3](#)
- [How Cisco MGC Node Manager Handles Events, page 6-17](#)
- [Commissioning, Decommissioning, and Rediscovering Devices, page 6-23](#)
- [Forwarding Traps to Other Systems, page 6-24](#)
- [Specifying the Length of Time Alarms Are Stored, page 6-27](#)
- [, page 6-27](#)

See [Appendix A, “Alarm Message Reference”](#) for information about alarm events.

Overview of Fault Management Features

One of the most important aspects of network management is the ability to identify problems on the system and to take action to resolve them quickly and efficiently. For example, a power supply failure in a chassis may be critical to the running of the network and need prompt attention.

In network management, these problems are typically called *faults*, and the message (or lack of message) that comes from the network device is called an *alarm* or *alarm event*.

Cisco MGC Node Manager provides fault management of the devices in the Cisco MGC Node Manager node. When the Cisco MGC host detects a problem with one of its logical connections, it generates a trap. Cisco MGC Node Manager receives the trap and delegates it to the graphical object that represents that logical connection. For example, if Cisco MGC Node Manager receives a trap that the link to a Cisco SLT is down, it delegates the trap to the object that represents that link.

To facilitate monitoring the network and identifying potential problems, Cisco MGC Node Manager propagates the alarm state of network elements upward through each object view. When an object receives an alarm, the object changes color to reflect its new state, and all parent objects also change color to reflect the most severe alarm on any of the child objects.



Note

If there are multiple alarm events propagated in an object tree, only the most severe is displayed at the highest level.

Cisco MGC Node Manager periodically polls managed devices to make sure that each device is still reachable using SMNP (status polling). If the device is not reachable, an annotation appears on the display in the Map Viewer, an alarm is generated, and the object is placed in an error state. Cisco MGC Node Manager continues to poll the device until connectivity is re-established. At that point, the alarm is cleared, the annotation on the display is removed, and the object is returned to its normal state. Duplicate alarms are filtered out.

For example, when a C7 IP Link goes out of service, a major alarm is immediately raised and propagated up to the Cisco MGC host object. If the IP connection to a Cisco MGC node is lost, a critical alarm is raised. A failover causes a major alarm.

In addition to managing alarms sent by SNMP traps, Cisco MGC Node Manager monitors system resources on the Cisco MGC host, HSI server, and BAMS and raises alarms for events such as an application being down or file usage above a specified percentage.

To investigate an alarm that is displayed in the Cisco MGC Node Manager Map Viewer, you typically perform the following steps:

- Drill down through the tree view to the object that has raised the alarm,
- From the object, open the Event Browser.
- In the Event Browser, identify the alarm details and take appropriate action to resolve the problem.

In addition to using the Event Browser to check on alarms flagged in the Map Viewer, you can use the Query Editor included in the Event Browser to filter alarms on any desired criteria. Diagnostic services can be invoked on events so that faults can be managed from the window that shows the event.

With Cisco MGC Node Manager, you can also forward alarms to any configured remote host and continuously export alarm events, as they are raised, to a text file.

The “[Managing Faults with Cisco MGC Node Manager](#)” section on page 6-3, describes the main tasks and procedures for managing faults. If you are interested in some of the principles Cisco MGC Node Manager applies in its fault management, see the “[How Cisco MGC Node Manager Handles Events](#)” section on page 6-17.

What Is Managed

Cisco MGC Node Manager performs fault management on the Cisco MGC node devices including the Cisco MGC host connectivity network. This includes the logical connections from the active Cisco MGC host to the:

- Interfaces (Ethernet, TDM)
- STPs
- Point codes (SS7 routes)
- Remote MGCs
- TCAP nodes
- Cisco Media Gateways

The logical connections from the active Cisco MGC host are shown as subnodes under the common Cisco MGC host object. If the standby Cisco MGC host is not processing calls, only the network connectivity of the active Cisco MGC host is shown.

For a reference describing the alarm events for specific devices, see [Appendix A, “Alarm Message Reference.”](#)

In addition to accepting SNMP traps from managed devices, Cisco MGC Node Manager also generates alarm events based on its own internal traps.

Managing Faults with Cisco MGC Node Manager

Managing faults with Cisco MGC Node Manager includes the following tasks:

- Task 1. Adjusting status polling settings if needed.
- Task 2. Customizing event management by setting up scoreboards and threshold crossing alerts.
- Task 3. Monitoring the network for alarm events, using the Map Viewer alarm display and the customized alerts you set up in task 2.
- Task 4. Using the Event Browser to view, filter, and clear events and to start diagnostic services on a problem device.
- Task 5. Using troubleshooting tools, opened directly from the Event Browser.

The following sections describe each of these tasks.

Task 1. Adjusting Status Polling Settings If Needed

Status polling checks the device status, such as up, down, active, or standby, as shown on the Status tab of the Properties dialog box for the device. You can specify a different status polling frequency for each device type, such as Cisco SLTs or Cisco LAN switches. All devices of that type and all components of such devices will have the same polling frequency. For example, if you set a 5 minute polling frequency for one Cisco SLT in the network, the frequency is applied to all Cisco SLTs in the network and to all monitored elements on the Cisco SLT, such as the TDM interfaces.

For the Cisco SLT and Cisco LAN switch, the polling frequency is set at the device level. All components of the device have the same polling frequency as the parent device. For example, setting a frequency of every 10 minutes for a Cisco SLT also polls its TDM interfaces every 10 minutes.

Default status polling frequency is every 2 minutes.

**Note**

To stop status polling when there is a known problem with a device or the device is taken out of service, decommission it. See the [“Commissioning, Decommissioning, and Rediscovering Devices”](#) section on page 6-23.

Use the following procedure to set or change a status polling frequency:

Step 1

In the Map Viewer, right-click the desired device, and choose **States**.

**Note**

For the MGC host, including its MGC node signaling or trunking components, select the MGC Host object.

**Note**

If the **States** option is not available, select the parent device.

The States dialog box appears.

- Step 2** Set the desired Status polling frequency. To change from minutes to hours, select from the pull-down menu. For all devices and types of polling, the minimum frequency is 1 minute and the maximum is 24 hours.
- Step 3** Click the **Save** tool to save the changes.
- Step 4** Close the dialog box.



Note For information on Performance and Configuration polling, see [Chapter 7, “Managing the Performance of Cisco MGC Node Manager Devices.”](#) For information on Auto-Discovery polling, see [Chapter 5, “Deploying Your Network in Cisco MGC Node Manager,” “Synchronizing the MNM with Device Changes”](#) section on page 5-20.

Task 2. Customizing Event Management

The Cisco EMF Event Manager provides three tools that can be used together to customize how you manage events:

- Thresholding regimes can be used to set up criteria for raising alarms on groups of devices based on selected performance measurements that cross a specified threshold. The thresholding regime also specifies what notification profile is used when the threshold is crossed and the alarm is raised.
- Notification profiles define how you want to be notified of the threshold-crossing alert, such as with a pop-up message window or a sound. You can also have a script run when the threshold is crossed.
- Event groups let you group events according to your own criteria, such as event severity or device type.

After an overview of each of these Event Manager tools, this section gives an example of how to create and use a scoreboard and how to set threshold-crossing alerts.

About Thresholding Regimes

Thresholding is the ability to configure the management system to actively monitor the network and notify the operator when some aspect of the network performance deviates from preset criteria.

Typically, you apply a standard set of criteria to an entire set of objects as part of a management policy. An example policy might be:

```

Poll all routers every 15 minutes and check if their CPU utilization is higher than 80%.
If it is higher than this, raise a warning alarm on the routers that breach this
condition.

```

A thresholding regime has a set of trigger conditions and the set of object groups to which these trigger conditions are to be applied.

Each trigger condition is made up of the following components:

- Expression to be checked (for example, CPU > 80%)
- Frequency that the expression should be checked; for example, every 15 minutes
- Notifications profile to run when the expression is satisfied

Setting up a thresholding regime allows you to apply or change the management policy of all 5000 routers at once rather than having to apply it to each one individually. You can change the central regime to apply the new policy to all objects within a group.

Once a threshold has been crossed, you can have the system notify you or carry out a sequence of actions. The specification of the actions to carry out is called a *notification profile*. Notification profiles are described in the [“About Notification Profiles” section on page 6-5](#).

About Notification Profiles

Notification profiles consist of a specified series of notifications that should be carried out as a result of the profile being triggered by a thresholding regime. Thresholding regimes are described in the [“About Thresholding Regimes” section on page 6-4](#).

Notification types available are:

- **Beep Once**—Produces a single beep
- **Raise Window**—Brings all windows that contain the controller object icon to the front of the window stack
- **Flash Icon**—Causes the controlling object icon to flash in active windows
- **Beep Continuously**—Produces a continuous beep
- **Popup Dialog**—Opens a window that contains a user defined message
- **Play Sound**—Plays a user-defined sound
- **Run Script**—Causes a user-defined script to run
- **Raise Event**—Generates a Cisco EMF event

All notifications can be given a time delay, allowing a simple form of escalation process to be implemented. For example:

- When a notification profile is triggered, raise a minor event; if the notification profile has not been reset within 30 minutes, raise a major alarm.

Once a notification profile is triggered, a running instance of this profile is created. This is a copy of the profile that is used to keep track of the current status of active notifications. Notification profiles can be viewed as templates that are used at trigger time to create an active running version. You can view the state of any notification profiles currently running on an object.

About Event Groups

A typical telecommunications network can generate a large volume of events. Only a small proportion of these events may affect service or require immediate attention. Others may still be of interest but are not urgent. For effective network management, you must be able to quickly identify the critical issues from the less than critical events.

You may also want to categorize the handling of certain events based on geographical location or the technical knowledge of certain users.

Event groups allow you to easily divide events into manageable groups based on user defined filtering criteria, such as:

- Event severity
- Event state
- Type of network element affected by the event

For display purposes, you can arrange these event groups on *scoreboards*. Each scoreboard shows a summary box for each group, allowing you to see the state of a group at a glance.

Having multiple scoreboards allows multiple users to keep track of different sets of events easily without being distracted by events that are of no interest to them.

In a similar way to thresholding regimes, event groups can also be configured to run notification profiles that carry out a series of actions when certain trigger conditions are satisfied. Event groups have three possible trigger conditions:

- When the first event enters the group, invoke notification profiles
- When the first event on an object enters the group, invoke notification profiles
- When *any* event enters the group, invoke notification profiles

For a description of tools used with event groups, see the [“About Thresholding Regimes”](#) section on page 6-4 and the [“About Notification Profiles”](#) section on page 6-5.

Creating and Using Scoreboards

You can create a scoreboard to display the alarms you are interested in. For example, you might create a single scoreboard to display the critical, major, and minor alarms received for your entire network, as well as alarms site by site.

The major tasks involved are:

- Create a scoreboard using the Notification Profiles application
- Using the event group application, create an event group for the alarm criteria you are interested in
- Add the event group to the scoreboard

Two examples are provided in the next section.

Scoreboard Example 1

Use a scoreboard to monitor all alarms on a network using the following procedure:

-
- Step 1** In Notification Profiles, create a notification to display a pop-up window.
- Step 2** Create a new event group:
- a. On the Launchpad, click **Event Groups**.
 - b. Click **Edit > Create > Event Group**.
 - c. Fill in a name and description for the group.
 - d. Click **Edit Query** to modify the default query.
 - e. For Severity, select **Critical/Major/Minor**.
 - f. Click **File > Exit**. You are prompted to save the query. Click **Yes**.
 - g. Click **Forward**.
 - h. From the list of trigger conditions, choose **Trigger every time an event enters the Event Group**.
 - i. Click **Edit**.
 - j. From the list of notification profiles, select the pop-up window.
 - k. Click **Finish**.
 - l. Click **Forward**.

- m. Click **Finish**.
-

Example 2

Use a scoreboard to monitor alarms at a particular site using the following procedure

- Step 1** Create a new event group as described in Example 1.
 - Step 2** For Severity, select **Critical/Major/Minor**.
 - Step 3** For Event Status, keep the default, **Active Only**.
 - Step 4** For Object Scope, select all objects for the desired site.
 - Step 5** Click **Finish**.
 - Step 6** In Notification Profiles, select the scoreboard used in Example 1 and add the Event Group.
-

Setting Threshold Crossing Alerts

You can trigger a threshold crossing alert (TCA) when a particular performance indicator crosses a specific threshold. In the example here, a TCA is created to alert you when the CPU utilization of a Cisco MGC host crosses a specified threshold.

Task 1

Use the following procedure to create an object group for a Cisco MGC host processor:

- Step 1** On the Launchpad, click **Object Group**.
 - Step 2** Right-click **Object Group**, and choose **Create New Group**.
 - Step 3** Fill in a name and description for the group.
 - Step 4** Click **Query Setup**, and click **Add Object(s)**.
 - Step 5** In the Host View, select the Cisco MGC host, and choose the Processor-1 object.
 - Step 6** Click **Apply**, to add this object to the group. You can add similar objects from the other deployed MGC hosts if you have multiple hosts deployed.
 - Step 7** Click **File > Close**, and save the query when prompted.
 - Step 8** Click **File > Close** again, and save object group changes when prompted.
-

Task 2

Use the following procedure to create a trigger condition:

- Step 1** On the Launchpad, click **Thresholds**.
- Step 2** Click **Edit > Create New Thresholding Regime**.
- Step 3** Give the regime a name and a description.

- Step 4** From the object group, and choose the object group created above.
 - Step 5** Click **Forward**.
 - Step 6** Click **Add** to create a new threshold.
 - Step 7** From the list of attributes, choose **mgcProcessor**.
 - Step 8** Under that object, choose **HOST-RESOURCES-MIB.hrProcessorTablemgcProcessor**.
 - Step 9** Under that object, choose **hrProcessorTable**, and then choose **Utilization**.
 - Step 10** From the list of operators, choose **>**.
 - Step 11** Specify a value, such as 70, and click **Add**. You have now created a trigger condition.
 - Step 12** Click **Forward**, and choose whether to use the default reset condition.
 - Step 13** Click **Forward**, and specify how often the trigger or reset condition should be checked.
 - Step 14** Click **Forward**, and choose the notification profile to associate with this thresholding profile.
 - Step 15** Click **Finish**.
 - Step 16** Click **Forward** to activate this thresholding regime.
 - Step 17** Click **Forward** and **Finish** to save this thresholding regime.
-

Task 3. Monitoring the Network for Alarm Events

You can monitor the network for alarm events in two ways:

- Using the Map Viewer Node View, you can see color-coded alarm indicators displayed on problem objects. In the Node View, alarms are propagated up from child elements to parent devices, so by watching just the main network devices, you can see when alarm events have occurred in any of their subcomponents. By drilling down, you can find the affected network element and then open the Event Browser to inspect the problem.

For details on using the Map Viewer and understanding its display, see [Chapter 3, “Getting Started with Cisco MGC Node Manager,” “Using the Map Viewer” section on page 3-10](#).

- Using customized event management tools such as scoreboards and threshold-crossing alerts, you can have Cisco MGC Node Manager notify you of selected problems. For information on these tools, see the [“Task 2. Customizing Event Management” section on page 6-4](#).

Task 4. Using the Event Browser

In Cisco MGC Node Manager, an event represents a notification from a managed entity that a certain condition has just occurred. These events usually represent error conditions on managed elements.

Each event is associated with the object for which it provides notification. Therefore, an object can have a number of events at any one time.

The Event Browser provides a tool to manage the network efficiently; you can list, query, and sort all or some events according to how you want to manage the network. The Event Browser can be started from:

- Map Viewer—to check on events for one or more selected devices
- LaunchPad—to run a query for particular events

You can have more than one Event Browser session open at a time, and each session can have different queries specified. All users can see any event. When an event is received, it is shown as active and unacknowledged (the Clear and Acknowledge column indicators on the event browser window are shown as grey). At this stage, no one has taken responsibility to deal with it. In the Event Browser window, you can acknowledge that a particular event is one that you are going to deal with, and all other users then see that the event is being handled. When the event is cleared, it is shown in the Event Browser window, so other users know that the event requires no further attention.

Some events are cleared automatically according to predefined clear correlation rules. These rules are described in the [“Automatic Alarm Clearing” section on page 6-19](#).

**Note**

The BAMS File Rename Failure alarm (POL115) must be manually cleared, not only in Cisco MGC Node Manager but also on the BAMS server before new alarms of that type will be generated.

Opening the Event Browser for One or More Selected Devices

Use this procedure when you have identified an alarm event for a particular device or devices in the Map Viewer.

Step 1 In the Map Viewer, select the device or devices you are interested in.

Step 2 Right-click the device, and choose **Tools > Event Browser**.

The Event Browser window opens, displaying events for the selected devices. Go to the [“Using the Event Browser to Manage Events” section on page 6-9](#).

Opening the Event Browser to Run a Query

Use this procedure when you want to check the network for alarm events of a particular type.

On the Launchpad, click the **Events** icon.

The Event Browser opens to the Query Editor for you to define a query to display events that match the query criteria. For more information, see the [“Filtering Events Using Queries” section on page 6-14](#). Once you have created a query, go to the [“Using the Event Browser to Manage Events” section on page 6-9](#).

Using the Event Browser to Manage Events

You can open the Event Browser from the Map Viewer to check events for specific devices or directly from the Launchpad to run queries.

Use the Event Browser to:

- Get details on events.
- View event history.
- Acknowledge an event, which shows that you have taken responsibility for managing that event. If you cannot continue to manage an event, it can be unacknowledged and then becomes available to other users.
- When the fault has been corrected and the event requires no further attention, clear the event. It is then removed from the Event Browser.

- Start diagnostic or other services to troubleshoot the event.

Use the following procedure to manage events in the Event Browser:

-
- Step 1** Open the Event Browser window as described in the [“Opening the Event Browser for One or More Selected Devices”](#) section on page 6-9 or the [“Opening the Event Browser to Run a Query”](#) section on page 6-9.
- Step 2** (Optional) Change the view options:
- To change sort order, choose **Edit > Sorting Options**, and select the desired fields to sort on.
 - To change how the severity column is color-coded, choose **View > Set Color Coding**.
- Step 3** (Optional) Turn automatic updating off or on:
- Choose **View > Enable Auto Update** to toggle between automatic and manual updating.
- Auto Update is the default state and allows you to view incoming events that are automatically updated in the window. If you are using manual updating, click **Refresh** periodically to see new events.
- Step 4** (Optional) View event history to see any events from the last seven days that match the current query but have had their status changed by being acknowledged, cleared, or unacknowledged:
- Choose **View > Event History**.
- Step 5** Select one or more events by clicking event severity, name, time, or description.
- Step 6** (Optional) View a full description of an event, including acknowledge and clearing details:
- Double-click the event. The Full Event Description window appears. For more information, refer to the [“About the Full Event Description Window”](#) section on page 6-12.
- Step 7** Do one of the following to change the event state, as appropriate:
- To acknowledge that you are handling the event, click **ACK**. The indicator changes to the color of the severity of the event.
 - To acknowledge the event, you can also right click on the event and select **Event State > Acknowledge** from the popup menu.
 - To unacknowledge an event, right click the event and select **Event State > Unacknowledge** from the popup menu.
 - To clear the event when it has been resolved, select the event, and click **Clear Events**. This displays the Events Clearing window. Enter the reason for clearing the event, and click **Apply**. The indicator changes to the new color of the severity of the event.
 - (If you acknowledged the event or are the administrator) To unacknowledge an event that is not resolved but you are not handling, click **ACK**.
- Step 8** (Optional) Click **Print**, to save the contents of all or part of the Browser to a file or to print a paper copy.
- Step 9** (If automatic updating is off) As desired, click **Refresh** to view the new events that meet the current criteria.
- Step 10** When done, close the window.

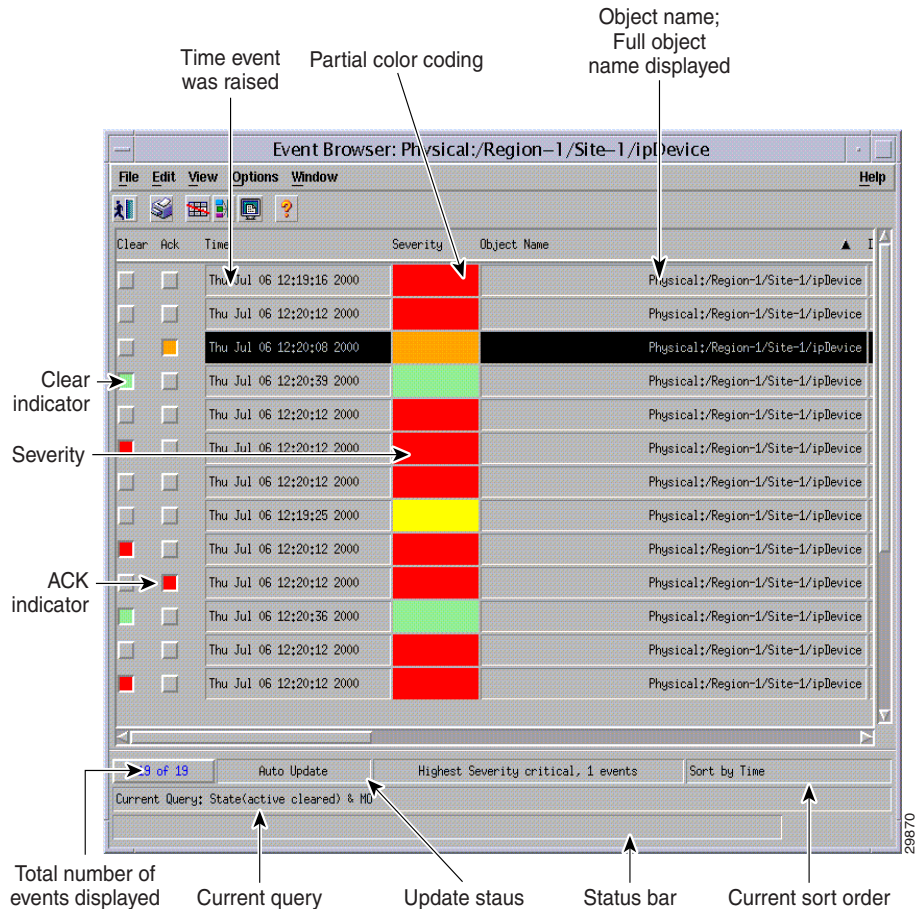


Note Query criteria are discarded when you close the window.

About the Event Browser Window

Use the Event Browser window to view and manage events, either on devices selected in the Map Viewer or selected through a query. The window is shown in [Figure 6-1](#).

Figure 6-1 Event Browser Window



Main Panel

The main panel in the Event Browser window, shown in [Figure 6-1](#), displays information about events including:

- Object name
- Time the event was raised
- Severity of the event (color-coded)
- Description of the event

Two indicators, color-coded to the severity of the event, are available to the left of the object name:

- Clear—An indicator to show if an event is active or cleared
- Ack—An indicator to show if an event is acknowledged or unacknowledged







**Note**

The option to unacknowledge an event is available only to an administrator or to the user who acknowledged the event initially.

Event Severity Color-Coding

Each event has a severity that indicates the importance of the event and is identified with a corresponding color as shown in [Table 6-1](#).

Table 6-1 Colors Used to Indicate Severity

	Color	Severity of Event
	Red	Critical
	Orange	Major
	Yellow	Minor
	Cyan	Warning
	Green	Normal
	White	Informational

Status Bar Information

The Event Browser window also displays information in the status bar:

- Progress bar (indicates that events are being added to the display)
- Current Update status (this can be auto or manual)
- Current query
- Current sort order, for example, sort by time
- Total number of events displayed (This number is shown in blue until it is acknowledged by the user by clicking the number.)


**Note**

The Event Browser can display a maximum of 10,000 entries. The status bar indicates whether there are more events on the system.

About the Full Event Description Window

Double-clicking an event in the Event Browser displays the Full Event Description window (see [Figure 6-2](#)). This window provides details of the event, including acknowledge and clearing details.

Figure 6-2 Full Event Description Window



Note

If the event has not been cleared, the Event State is Active and the Clearing Method, User Responsible for Clearing, and Clearing Time and Date sections are disabled. You cannot alter the information displayed.

If an event has been cleared, you can view the method used to clear it by clicking **Clearing Event**.

The Full Event description window displays the following information:

- Object name—Name of the Cisco EMF managed object the event was reported against.
- Time and Date—The time and date the event was reported.
- Severity—The severity of the reported event.
- Source Domain—The communications domain that reported the event.
- Management Domain—The Management domain that reported the event.
- Event Description—A brief description of the reported event.
- Event State—Whether the event is active or cleared. If the event has been cleared, the Clearing Method, User Responsible for Clearing, and Clearing Time and Date sections become active.

Acknowledge Details

- Acknowledgement User—Identifies the user who acknowledged the event

- Acknowledgement Time and Date—Identifies when the event was acknowledged

Clearing Details

- Clearing Method—Indicates if the event was cleared by the network or by a user.
- User Responsible for Clearing—Displays the user name responsible for clearing the event.
- Clearing Time and Date—Indicates the time and date the event was cleared.
- Reason for clearing—The information that was entered in the Events Clearing window, which is completed when the Clear indicator is selected.

Filtering Events Using Queries


The Event Browser monitors all events on all devices. To work efficiently, you might want to specify the objects on the network with which you are concerned. The Event Browser gives you the option to do this through queries that can be configured to match your requirements. With queries, you can choose to include or exclude devices or criteria. For example, you could choose to monitor a particular device, specify a time period, and look only at events that are warnings or are critical. You define a query so that the Event Browser displays only the events that meet the criteria you defined.



Note

A query applies to the current Event Browser session only; stored queries are not supported. You can modify a current query, but once you close the Event Browser, the query is discarded.

Use the following steps to define a query:

-
- Step 1** Do one of the following to open the Query Editor:
- On the Launchpad, click the Events icon.
 - If the Event Browser is already open, choose **Edit > Query Setup** or click the **Query Filter** tool:
- 
- The Query Editor window, similar to [Figure 6-3](#), is displayed.
- Step 2** Set filtering (query) criteria:
- To add a value to the query, select it in the Available Values list and click >> to place the value in the Selected Value list. To remove a value, select the value in the Selected Values list and click <<. See the “[About the Query Editor Window](#)” section on page 6-15 for details.
 - To activate selected values on a given tab, click the **Activate** box. A dark gray tab is active (On); its query is used in the Event Browser. A light gray tab is inactive (Off); its query is not used.
- See the “[About the Query Editor Window](#)” section on page 6-15 for details.
- Step 3** Click **Apply**, and close the Query Editor.
- You see the following message:
- Save Query Changes?
- Step 4** Click **Yes**.
- The Event Browser begins collecting the data using the criteria you selected and displays it in the Event Browser window.

**Note**

Query changes are saved for the current session only. When you close the Event Browser, the query criteria are reset to the default.

Modifying a Query

Use the following steps to modify a query:

Step 1 Choose **Edit > Query Setup** or click the Query Filter tool:



The Query Editor window ([Figure 6-3](#)) is displayed with the current settings.

Step 2 Modify filtering (query) criteria:

- To add a value to the query, select the value in the Available Values list and click >> to place it in the Selected Value list. To remove a value, select it in the Selected Values list and click <<. See the [“About the Query Editor Window”](#) section on page 6-15 for details.
- To activate selected values on a given tab, click the **Activate** box. A dark gray tab is active (On); its query is used in the Event Browser. A light gray tab is inactive (Off); its query is not used.

Step 3 Click **Apply**, and close the Query Editor.

You see the following message:

Save Query Changes?

Step 4 Click **Yes**.

The Event Browser begins collecting the data using the criteria you selected and displays the data in the Event Browser window.

Query changes are saved for the current session only. When you close the Event Browser, the query criteria are reset to the default.

About the Query Editor Window

The Query Editor is shown in [Figure 6-3](#).

Figure 6-3 Query Editor Window

The criteria that can be used to specify a query are grouped on tabs. After selecting criteria from the Available Values list on a tab, click the **Activate** box to activate those criteria. A dark gray tab is active (On); its query is used in the Event Browser. A light gray tab is inactive (Off); its query is not used.

The Query Editor includes these tabs:

- Severity—Critical, major, minor, warning, normal, or informational.
- Time—Time range for which you want to view events, specified with time of day, day of the week, and date.
- Event Status—Acknowledged or unacknowledged, active or cleared.
- Source Domain—Where the event was generated: SNMP, the managed network, internal, or generated by Cisco MGC Node Manager.
- Mgmt Domain—The management domain of the SNMP trap information. The SNMP Management Information Base (MIB) information typically defines the equipment type generating a trap.
- User—Name of the user associated with an acknowledged or cleared event.
- Event Class—Type of event.
- Object Scope—Use to select all the events of a node and its children. You select from an object tree, specifying the number of levels to view for a selected node. To specify scope:
 - On the Object Scope tab, click **Add Scope**. The View Scope selector appears.
 - Select the desired node.
 - In the Number of Levels field, type the number of levels to view.
 - Select **Descendants**.
 - Click **Apply**.
- Object Class—Type of object, such as managed, container, network, site.
- Object Attribute Presence—For various object types, attributes to query.
- Object Attribute Value—For specified object types and attributes, values to query for.

Task 5. Using Troubleshooting Tools

Once an alarm has been identified, you can use Cisco MGC Node Manager to launch a variety of diagnostic and troubleshooting tools, discussed in detail in [Chapter 8, “Other Network Management Tasks,” “Using Diagnostic Tools” section on page 8-51](#).

In the Event Browser, you can right-click a device and open troubleshooting tools such as:

- Opening a Diagnostics dialog box on most devices. The Diagnostics dialog box provides shortcuts for common diagnostics that normally require using UNIX or MML commands. You can ping the device for connectivity, use Traceroute, check the alarm log, check the status of running processes, display the BAMS system log, and audit the BAMS trunk groups, cross-checking them with the Cisco MGC host configuration, and retrieve state information on various network elements.



Note

The alarm log for the Cisco PGW, BAMS, HSI, SLT, and Catalyst is the file traplog.log. Cisco EMF messages go to a separate file, mgcTrapProcessor.log.

- Opening the Cisco MGC Toolbar (also known as the Cisco MGC toolkit), which contains a suite of diagnostic and troubleshooting tools. For details, see [Chapter 8, “Other Network Management Tasks,” “Using the Cisco MGC Toolbar” section on page 8-54](#).
- Open CiscoView to troubleshoot problems on the Cisco SLT or Cisco LAN Switch. You can also use Cisco MGC Node Manager to Telnet to a device or to launch an X terminal window. For details, see [Chapter 8, “Other Network Management Tasks,” “Using Cisco MGC Node Manager To Launch Device Configuration” section on page 8-4](#).

How Cisco MGC Node Manager Handles Events

Refer to this section if you are interested in the principles applied by Cisco MGC Node Manager in processing and displaying events. It includes:

- [Understanding Event Propagation, page 6-17](#)
- [Understanding Alarm Acknowledgement and Clearing, page 6-18](#)
- [Understanding Status Polling, page 6-20](#)

Understanding Event Propagation

In order to make the identification of potential problems easy, Cisco MGC Node Manager propagates the alarm state of objects upwards through the Physical and Node object views.



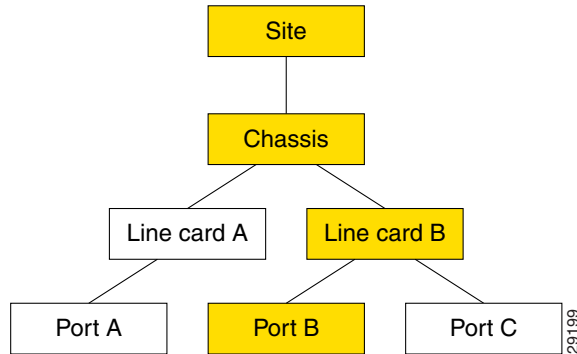
Note

To minimize redundant updating, alarms are propagated only in the Physical and Node views, not in device views. In a device view, a gray dot indicates an alarm somewhere in the tree. Check the relevant device in the Node view to find the alarm.

If an object receives an event, the object change color to reflect its new state, and all parent objects within a view also change color to reflect the most severe alarm on any of the children. The example in the following diagram shows a typical physical view of the network. The line cards are contained within the chassis, the chassis within a bay, and the bay within a site.

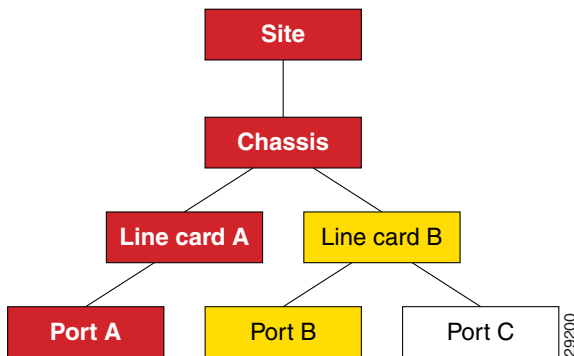
If a minor alarm is received on Port B, then Port B and all of the objects up to the region turn yellow to indicate a potential minor problem, as illustrated in [Figure 6-4](#).

Figure 6-4 Example Minor Event Propagation



If a critical alarm was then received on Port A, that port, and all of the objects up to the region, turn red to indicate a potential critical problem, as illustrated in [Figure 6-5](#).

Figure 6-5 Example Critical Event Propagation



If the critical alarm is cleared, the icons return to yellow.

Cisco MGC Node Manager filters out duplicate traps from a network element. It also filters out traps from network elements that report a problem and reports within a few seconds (up to 6) when the problem is resolved. Cisco MGC automatically clears existing alarms when a network element reports that an alarm condition is no longer present. This reduces the number of unnecessary alarms displayed in the Event Browser. You cannot configure when an alarm should be automatically cleared.

Understanding Alarm Acknowledgement and Clearing

This section describes how event acknowledgement and clearing work in the Event Browser.


When a new event is received, its event state is active and unacknowledged. Acknowledging the event indicates to other users that it is being handled. Once it has been dealt with, you can clear the event. When you cannot clear an event due to an existing problem, the event can be returned to the unacknowledged state and later acknowledged or cleared by another user.

Whether an event is unacknowledged or acknowledged, the event is considered active until it is cleared. The relationship between event states is shown in [Figure 6-6](#).

Some events are cleared automatically when the originating condition is resolved, according to rules described in the “Automatic Alarm Clearing” section on page 6-19.

After events are cleared, they continue to be stored within the system for a configurable amount of time to maintain an event history for an element. These events can be viewed and manipulated in the same way as any other event.

Figure 6-6 State Diagram for Events



Automatic Alarm Clearing

Cisco MGC Node Manager automatically clears alarms based on certain built-in, logical rules. On receipt of an incoming clear alarm, the rules defined in these files indicate which active alarms on a given object should be cleared. For example, a link-up alarm clears a link-down alarm, a process normal alarm clears a process error alarm, and a communication success alarm clears a communication failure alarm.

These rules are maintained in Clear Correlation files. A sample Clear Correlation file is:

```

CLEAR_CORRELATION_RULE
    INCOMING_ALARM_CLASSlinkUpAlarmClass
    ALARM_CLASS_TO_CLEARlinkDownAlarmClass
    END_RULE
  
```

When a clear condition is received, the cleared alarm is automatically removed from the appropriate screens, and the clear alarm is forwarded to northbound systems like any other alarm.

The following tables show the clear conditions for the alarms for each Cisco MGC node device.

Table 6-2 Cisco MGC Host Clear Correlation

Alarm	Clear Condition
processingError	processingNormal
communicationFailure	communicationSuccess
qualityOfServiceError	qualityOfServiceNormal
equipmentError	equipmentNormal
environmentError	environmentNormal

Table 6-3 Cisco SLT Clear Correlation

Alarm	Clear Condition
IF-MIB.linkDown	IF-MIB.linkUp

Table 6-4 Cisco LAN Switch Clear Correlation

Alarm	Clear Condition
IF-MIB.linkDown	IF-MIB.linkUp
CISCO-STACK-MIB.switchModuleDown	CISCO-STACK-MIB.switchModuleUp

Table 6-5 Resource Alarms (Cisco MGC host, HSI server, and BAMS) Clear Correlation

Alarm	Clear Condition
CRITAPP-MIB.critAppDown	CRITAPP-MIB.critAppUp ¹
CRITAPP-MIB.critAppNotAllRunning	CRITAPP-MIB.critAppAllRunning
SIFSMONITOR-MIB.siFsBelowWarningThreshold	SIFSMONITOR-MIB.siFsAboveWarningThreshold ²
SIFSMONITOR-MIB.siFsBelowCriticalThreshold	SIFSMONITOR-MIB.siFsAboveCriticalThreshold ³

1. The varbind criAppName in the trap/clear must match.
2. The varbind siFsMonName in the trap/clear must match.
3. The varbind siFsMonName in the trap/clear must match.

Understanding Status Polling

Cisco MGC Node Manager periodically polls each managed object (the Cisco MGC host, Cisco SLT, LAN switch, HSI server, and BAMS) to ensure that the device is still reachable using SNMP. If the device is not reachable, its state is indicated by annotation on the Map Viewer, and an alarm is generated. In addition, the object is placed into the error state.

After the object loses connectivity, Cisco MGC Node Manager continues to poll the object until it can be reached. Once connectivity is re-established, the alarm is cleared and the annotation on Map Viewer is removed. In addition, the object is returned to the normal state.

Cisco MGC Node Manager also displays the status of the Cisco MGC host connectivity network. This includes the logical connections from the active Cisco MGC host to the:

- Interfaces (Ethernet, TDM)
- STPs
- Point codes (SS7 routes)
- Remote MGCs
- TCAP nodes
- Cisco Media Gateways

The logical connections from the active Cisco MGC host are shown as subnodes under the common Cisco MGC host object. If the standby Cisco MGC host is not processing calls, only the network connectivity of the active Cisco MGC host is shown.

Status details are provided below.

Network Interface Status

Cisco MGC Node Manager performs status polling to reflect the state of each network interface. Depending on the operational and administrative status of the interface, the object representing the network interface transitions to another state, as indicated in [Table 6-6](#).

Table 6-6 Transition States

Admin Status	Operational Status	Network Interface State
Up	Up	up
Up	Down	down
Up	In Test	in-test
In Test	N/A	in-test
Down	N/A	off-duty
<not reachable>	N/A	unreachable

Note that the chassis is queried for the state of its interfaces. That is, the status of the interface reported by Cisco MGC Node Manager is identical to the status reported by the chassis on its current management IP address. However, the status of each interface is reported by the chassis via that object's specific IP addresses. In this way, Cisco MGC Node Manager can better reflect the true health of the chassis.

Interface Alarms

When a network interface goes down, the device sends a link-down trap to Cisco MGC Node Manager. When Cisco MGC Node Manager detects this trap, it transitions the object representing that interface to the down state. To handle the case where Cisco MGC Node Manager may have missed a trap, the status polling mechanism raises an alarm if it detects that the interface is down. When the interface comes back up, the device raises a link-up trap. If Cisco MGC Node Manager detects this trap, it transitions the interface back into the normal state. If Cisco MGC Node Manager missed this trap, the next status poll detects that the interface is back up. Internally, Cisco MGC Node Manager transitions the interface back to the normal state and clears the appropriate alarms on the object.

MGC Host Status

Cisco MGC Node Manager periodically checks the status of each MGC node device. The attribute `SNMP:CISCO-TRANSPATH-MIB.tpCompOpStatus` is retrieved and its value is used to determine the required state of the object as indicated in [Table 6-7](#).

Table 6-7 Cisco MGC Host States

Component Status	Network Interface State
ACTIVE	active
STANDBY	standby
OOS	oos

Table 6-7 Cisco MGC Host States (continued)

Component Status	Network Interface State
No answer	not-running
Not reachable	unreachable

BAMS Status

Cisco MGC Node Manager periodically checks the status of each BAMS device. The SNMP:ACECOMM-BAMS-SYSPARM-MIB.sysStatus attribute is retrieved, and its value is used to determine the required state of the object, as indicated in [Table 6-8](#).

Table 6-8 BAMS States

Component Status	Network Interface State
Active	active
Standby	standby
Outage	oos
Other	other
No answer	not-running
Not reachable	unreachable

HSI Status

Cisco MGC Node Manager periodically checks the status of each HSI device. The SNMP:HOST-RESOURCES-MIB.hr.sysStatus attribute is retrieved, and its value is used to determine the required state of the object, as indicated in [Table 6-9](#).

Table 6-9 HSI Status

Component Status	Network Interface Status
Active	active
Other	other
No answer	not-running
Not reachable	unreachable

Trap Receipt Not Guaranteed

Cisco MGC Node Manager does not provide any guarantee that it received a trap from the southbound systems or network elements. Cisco MGC Node Manager does not perform any negotiation with the network elements to detect or recover lost traps.

How Cisco MGC Node Manager Manages Multiple IP Addresses for Status Polling

By default, each Cisco MGC Node Manager object can contain only a single IP address. For example, when the user deploys a Cisco SLT, the user can specify only a single IP address. Cisco MGC Node Manager uses this IP address for all management transactions, including status polling and performance polling. In addition, the IP address is used to map incoming faults to the Cisco MGC Node Manager object. When a trap arrives from the network element, Cisco MGC Node Manager matches the IP address of the trap sender to the IP address of an object in the database.

In reality, a physical device might have more than one IP address. Traps can come from any interface on the device. Since Cisco MGC Node Manager is aware of only a single IP address, traps received from an alternate interface might be dropped.

Any interface on the device might go down (either operationally or administratively). If the management interface goes down, all SNMP-based operations fail. That is, not all SNMP queries are completed nor does status polling or performance polling function. Cisco MGC Node Manager is designed to avoid these situations by using trap proxies and IP address failover.

Commissioning, Decommissioning, and Rediscovering Devices

When a device is administratively off the network, or it has a known problem and you do not want to manage it, you can decommission the device in Cisco MGC Node Manager to stop it from being polled and generating unnecessary alarms.

When a device is decommissioned, no actual changes are made to the device, which still sends traps to Cisco MGC Node Manager. However, the resulting alarm events are not reported and do not initiate any actions or status changes. Status and performance polling are also suspended.

**Note**

When a device is decommissioned, all its subcomponents are also decommissioned.

When the device is back in service, commission it to resume polling. At that point, Cisco MGC Node Manager starts discovery to resolve any component changes that may have occurred while the device was decommissioned.

When a device's subcomponents have changed or you have corrected a problem that interfered with discovery, you can rediscover the device to immediately update the Cisco MGC Node Manager network model with the changes.

**Note**

Rediscovery is necessary only when you want the update to occur before the next auto-discovery polling interval, when any changes are routinely detected.

Use the following procedure to decommission, commission, or rediscover a device:

- Step 1** In the Map Viewer, select the object and right-click.
- Step 2** Choose **States**. You see the States dialog box.
- Step 3** On the States tab, do one of the following:
- Click **Decommission** to stop processing traps from the device.
 - Click **Commission** to resume processing traps after a device had been decommissioned.

- Click **Rediscover** to rediscover a device, updating the network model with any device changes since the last auto-discovery.

You are prompted to confirm the action.

Step 4 Click **Yes**. Cisco MGC Node Manager executes the action. The device state changes to reflect the change.



Note When a device is rediscovered, if the Event Browser is open, it displays the message, “Discovery is now complete”. With each new discovery, any earlier discovery messages are cleared; only the most recent discovery message appears.

Step 5 Close the dialog box.

Forwarding Traps to Other Systems

You can forward the traps (alarms) collected from managed elements by Cisco MGC Node Manager to other systems. In addition to receiving SNMP traps from node devices, Cisco MGC Node Manager monitors resource usage on the Cisco MGC host, HSI server, and BAMS. Traps are generated, for example, when disk usage exceeds a given threshold or when applications are down. There are two types of trap forwarding:

- [Using cmnmtrapforward](#)
- [Using Northbound Event Interface](#)

Using cmnmtrapforward

Use the **cmnmtrapforward** command to automate the procedure to stop or start forwarding traps to other systems by updating the trapForwardFile file.

To Start Trap Forwarding

Follow these steps to start trap forwarding:

Step 1 From the Cisco EMF base directory, enter the following command:

```
cmnmtrapforward
```

Information similar to the following is displayed:

```
Configure trap forwarding to other hosts? [y/n]: [n]
```

Step 2 Enter Y. Information similar to the following is displayed:

```
Trap Forwarding is configured for the following IPs
```

```
172.16.128.46
```

```
Please enter a Trap Forwarding IP address [?,q]
```

Step 3 Enter the IP address to which traps will be forwarded, and press Enter. Information similar to the following is displayed:

Enter another IP address? [y/n]: [n]

- Step 4** Continue to add IP addresses, or press N when you are finished. The following prompt appears:
Restarting TrapManager...

To Stop Trap Forwarding

Enter the following command from the Cisco EMF base directory to stop trap forwarding:

```
./cmmntrapforward -d <IP address of destination host>
```

The destination IP address is removed from the trapForwardFile file.

Use the **cmmntrapforward -h** command to view more information about this command.

Using Northbound Event Interface

The Northbound Event Interface (NEI) allows for integration with network management systems (NMSs), such as Hewlett Packard-OpenView Element Management Framework (HP-OEMF) and CIC (Cisco Information Center). Using NEI, you can export topological information about managed objects and forwardCisco EMF events to NMSs.

The main purpose of NEI is to convert Cisco EMF events (appearing in the Event Browser) to a particular output for NMSs. Output can be in the form of an SNMP trap, log files, or TCP connections.

NEI has two main functions: exporting and forwarding. To define export and forward filters, you can create a filter file that will contain both types of information.

For further information on NEI click on the below given link

http://www.cisco.com/univercd/cc/td/doc/product/rtrmgmt/cemf/3_2sp7/config/nei.htm



Note

Northbound CORBA flowthrough will not be enhanced in MNM 2.6(1) to support the latest PGW 2200 features.

Sample Filter File in MNM

This section presents a sample filter file, used in MNM, for physical view

```
name="physical-traps.nbf"
exporting
{
    delta
    {
        temp="/tmp/filter.delta.tmp"
        result="/tmp/filter.delta.result"
    }
    dump
    {
        temp="/tmp/filter.dump.tmp"
        result="/tmp/filter.dump.result"
    }
    filter="NbNullExporter"
    origin="Physical:/"
}
```

```

forwarding
{
  filter="NbExtensibleSNMPForwarder"
  snmp-destination="10.20.1.19"
  snmp-port="162"
  enterprise="1.3.6.1.4.1.1469.6"
  added_aqs
  {
    severity="critical"
    severity="major"
    severity="minor"
    status="cleared"
    status="acknowledged"
    status="unacknowledged"
  }
  changed_aqs
  {
    status="acknowledged"
    status="unacknowledged"
  }
  containment-tree="MGC-Node-View"
}

```

Difference between NEI and cmnmtrapforward

This section briefly explains the differences between using NEI and Trap Forwarding. It contains

- [Northbound Event Interface](#)
- [cmnmtrapforward](#)
- [Recommendation](#)

Northbound Event Interface

NEI allows integration with higher fault management systems say HP-OEMF. NEI includes two main areas of Functionality,

- Configurable topological export - to build managed objects as defined within MNM server using CLASS mapping files
- Configurable MNM fault forwarding - forward SNMPv1 and SNMPv2c traps generated from MNM faults to higher fault management systems say HP-OEMF. Customers create filter files with two sections namely "exporting" and "forwarding" .

cmnmtrapforward

Raw traps received by MNM from the devices being managed are forwarded to multiple third party NMSs by adding hostname ipaddress, generic trap id, specific trap id, enterprise oid, details in the trapForwardingFile. This feature does not include the functionality of exporting, forwarding filters as supported by NEI, nor does it forward alarms that are generated by MNM itself, such as communication failures detected during polling of managed devices.

Recommendation

Since trap forwarding has minimal value when compared to NEI, it is recommended to use NEI for forwarding traps to third party NMS.

Specifying the Length of Time Alarms Are Stored

All alarms are automatically stored in the Cisco MGC Node Manager database and purged at regular intervals to make room for new alarms. The Alarm Deleter, built into Cisco EMF, is set up to run at midnight every night. The Alarm Deleter queries the alarm database and deletes alarms that meet the specified criteria. The default is to delete cleared alarms that are seven days old.

If you want to change the frequency with which old alarms are deleted, you can change the values in the alarmDelete.ini file. An example of the file is shown here:

```
[logger]
#include "loggercommon.include"
loggingName = alarmDeleter

[AlarmDeleter]
databaseName = [[OSDBROOT]]/alarm.db
segmentDeletionInterval = 15
ageOfAlarmsInDays= 7
ageOfAlarmsInHours= 0
ageOfAlarmsInMinutes = 0
deleteAllAlarms= 0

[Database]
#include "databaseCommon.include"
```

The variables used in defining the deletion rules are described in [Table 6-10](#).

Table 6-10 Alarm Deleter Attributes

Variable	Description
ageOfAlarmsInDays	The age of the alarm, in days, before it is to be deleted.
ageOfAlarmsInHours	The age of the alarm, in hours, before it is to be deleted.
ageOfAlarmsInMinutes	The age of the alarm, in minutes, before it is to be deleted.
deleteAllAlarms	0 = delete only cleared alarms that match criteria; 1 = delete both active and cleared alarms that match criteria.

■ Specifying the Length of Time Alarms Are Stored