CA Spectrum® Infrastructure Manager

SpectroSERVER Performance Administration Guide

r9.2.1

CA Technologies Product References

This guide references CA Spectrum® Infrastructure Manager (CA Spectrum).

Contact CA Technologies

Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information you need for your Home Office, Small Business, and Enterprise CA Technologies products. At http://ca.com/support, you can access the following:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Provide Feedback

If you have comments or questions about CA Technologies product documentation, you can send a message to techpubs@ca.com.

If you would like to provide feedback about CA Technologies product documentation, complete our short customer survey, which is available on the CA Support website at http://ca.com/docs.

Contents

Chapter 1: Performance and Tuning Concepts

The OneClick client application used by network administrators to monitor and troubleshoot a CA Spectrum-managed network provides a Performance tab for most network and device models. Network administrators can use this tab to examine a specific device's CPU and memory utilization. However, for the VNM model that represents a SpectroSERVER, there exists a separate, more robust application called Performance View. Performance View lets CA Spectrum administrators with Performance Monitor privileges monitor the performance and system resource utilization of a SpectroSERVER, thereby identifying performance problems and possible corrective actions.

This section contains the following topics:

System Components Monitoring

Each computer system is comprised of four major components: disk, network, memory, and CPU. To help ensure the successful operation of your CA Spectrum system, you may need to tune one or more of these components to eliminate bottlenecks.

Performance View includes the following two features to help you detect and locate bottlenecks:

- A series of tabs that provide information on the system components and SpectroSERVER activities that can affect performance. View the Main tab for overall system and network activity at a glance.
- A health report feature that allows you to run a report on the SpectroSERVER resources over a 24-hour period.

Note: In a distributed SpectroSERVER environment, you can switch the focus of Performance View from one SpectroSERVER to another.

More information:

Overview of the User Interface (see page 14)
Running Health Reports (see page 49)
Connect to a Different SpectroSERVER (see page 25)

Performance Data Analysis

In many cases, the information provided in Performance View can give a direct indication of the source of a performance problem. For instance, the CPU tab lists the 10 processes that are currently using the highest percentages of the CPU. However, exactly what constitutes a problem or bottleneck depends on the specific configuration of your CA Spectrum system as well as your network management priorities.

For guidelines on identifying performance problems and corrective actions, see Evaluating the Performance of a SpectroSERVER (see page 31). You can also run a health report as described in Running Health Reports (see page 49). A health report includes its own analysis of the data that it collects.

More information:

Evaluating the Peformance of a SpectroSERVER (see page 31)
Running Health Reports (see page 49)

Performance Optimization

In general, you can resolve SpectroSERVER performance problems by tuning the SpectroSERVER to increase its performance, adding more SpectroSERVERs to distribute the entire network load, or both.

SpectroSERVER Tuning

Once you have identified the reasons for degraded performance, you can tune the SpectroSERVER to optimize its performance. Tuning can include any of the following measures:

- Modifying the polling interval and poll to log ratio of essential device models and application models, and disabling polling of non-essential models. This reduces the network traffic and the resulting latency that degrades performance.
- Increasing the capacity of the system by increasing memory, CPU speed, or disks.
- Reducing the number of traps that are mapped to CA Spectrum events.
- Reducing the amount of data requested by customized watches and displayed attributes in order to reduce the amount of data requested from the SpectroSERVER and devices.

- Adjusting usage of features such as Live Pipes, Discovery, and automatic device configuration.
- Adjusting client interactions with the SpectroSERVER. For example, reports generated using SPECTRUM Report Console can exert a punctuated or prolonged performance burden on the server depending on what is being reported and how often the reports are run. Command Line Interface (CLI) scripts, manual discoveries, and other manually initiated tasks can also impact SpectroSERVER performance.

More information:

Tuning a SpectroSERVER (see page 61)

Additional SpectroSERVERs

To determine if increasing the number of SpectroSERVERs, rather than tuning, is the best means of achieving desired performance improvements, you can request that CA Support perform a sizing of your CA Spectrum environment. The sizing tool uses information about your network configuration to estimate the following:

- The additional amount of network management traffic generated as a result of CA Spectrum.
- The number and configuration of additional SpectroSERVERs required to efficiently manage the number of models in your CA Spectrum environment.

Chapter 2: Getting Started with Performance View

This section contains the following topics:

Start Performance View

Start Performance View by doing either of the following:

- In the CA Spectrum Control Panel, select SpectroSERVER Performance from the Control menu.

  This connects Performance View to the SpectroSERVER to which the Control Panel is connected.

- Access a command prompt, navigate to the <$SPECROOT>/PView folder, type **pview**, and press Enter.

  If you have set the Show Server List At Startup preference, you are then prompted to select the SpectroSERVER to which to connect.

  Note: If your CA Spectrum environment is distributed, you can specify a SpectroSERVER when you start the application by entering the following:

  ```
  pview -vnm landscape_name
  ```

Note: If you are running CA Spectrum on Windows, the user running the SpectroSERVER process must belong to either the Windows Performance Monitor Users group or the Administrators group to launch Performance View.

More information:

Set Preferences (see page 26)

Overview of the User Interface

Performance View has a single main window from which you access all performance information. The following image identifies the major user interface elements in the main window:

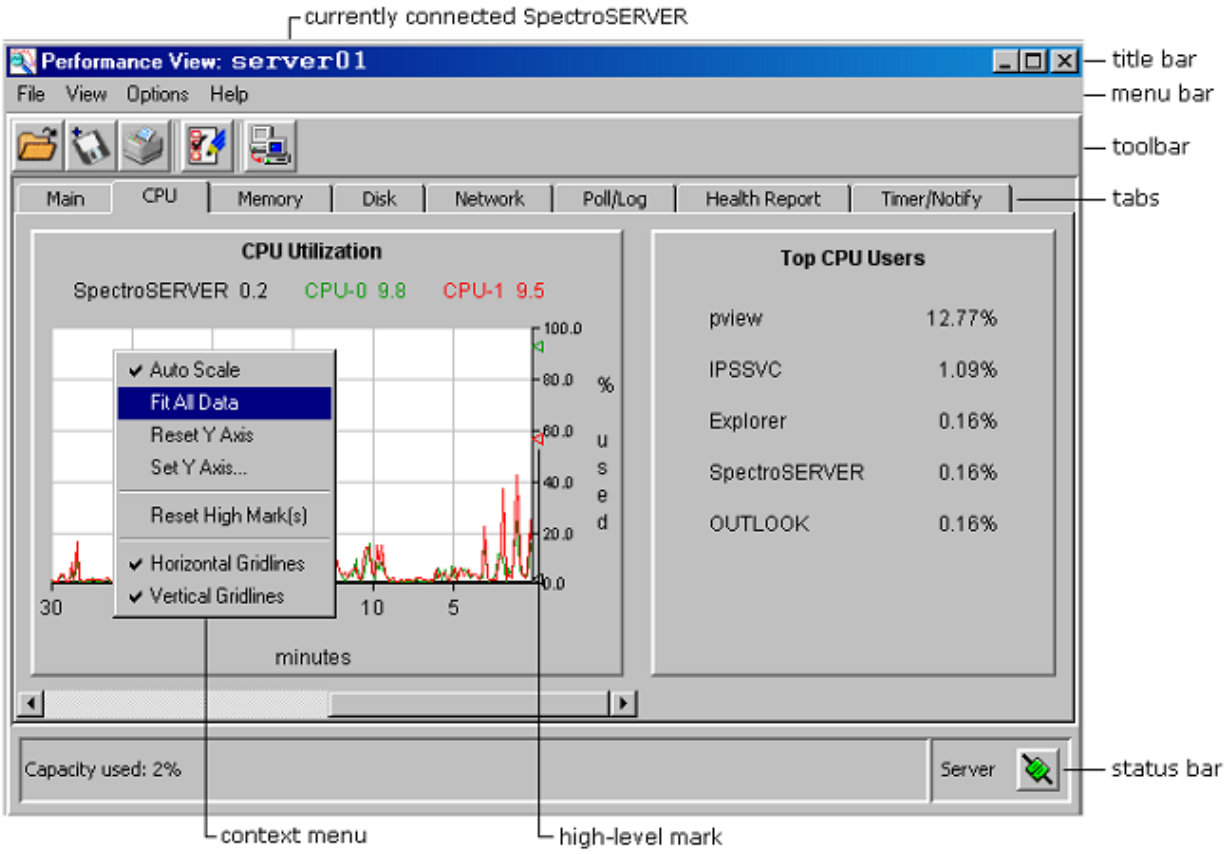

The status bar displays the current state of Performance View and, when applicable, the status of the health report for which data is currently being collected. The color of the connection status icon on the status bar indicates the status of the connection to the SpectroSERVER as follows:

Green

Normal

Yellow

Using backup SpectroSERVER

Red

Contact lost

The tabs in the main window provide detailed information about various aspects of the SpectroSERVER's performance. The information is presented in bar graphs, line graphs, and text.

Bar graphs are used on the Main tab. These graphs use data collected in ten minute running averages. All other attributes and graphs use data collected every 10 seconds.

Line graphs are used on the CPU, Memory, Disk, Network, Poll/Log, and Timer/Notify tabs. These graphs display data collected over 60 minutes. If data is collected for more than 60 minutes, only the most recent 60 minutes worth of data is displayed.

For single-line graphs, instantaneous values are shown in a text box in the top, right corner of the graph. For multiple-line graphs, instantaneous values are shown in color coded labels, where each label represents the instantaneous value of its associated line.

Line graphs also provide high-level marks and associated tooltips for each line in the graph. High-level marks represent the highest data value collected since data collection was last started. High-level marks are located to the right of the line with which they are associated. You can reset the high-level marks and make other changes to the axes of line graphs.

Multiple CPUs and disk drives can be graphed in the same graph. In a bar graph, labels identify which CPU or disk drive is being graphed. In a line graph, colored lines are used to identify which CPU or disk drive is being graphed.

More information:

Adding SpectroSERVERs (see page 69)
Change the Display of Graph Axes (see page 25)
Change the Colors of Graph Lines (see page 28)

Main Tab

The Main tab contains a Physical Memory Utilization graph, a CPU Utilization graph, a Disk Utilization graph, and general data.

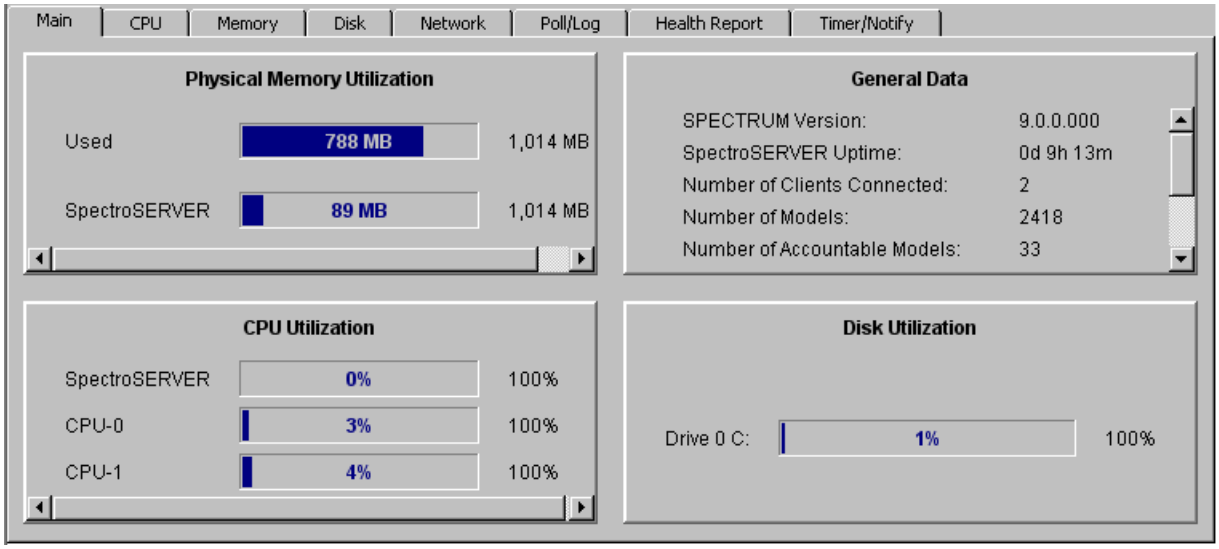

The graphs provide the following information:

Physical Memory Utilization

Displays the amount of physical memory being used by the SpectroSERVER and by other processes running on the machine to which Performance View is connected.

CPU Utilization

Displays the percentage of total CPU processing power being used by the SpectroSERVER and by other processes running on the machine to which Performance View is connected.
Note: Scroll bars are added to the graph area to accommodate multiple CPU bar graphs.

Disk Utilization

Displays how much of a disk's read/write access capacity is being used.

The General Data panel provides the following information:

CA Spectrum Version

Specifies the version of CA Spectrum installed on the server to which Performance View is connected.

SpectroSERVER Uptime

Specifies the amount of time, in days, hours, and minutes, that the server to which Performance View is connected has been running. The format for Uptime is: <days>d <hours>h <minutes>m.

Number of Clients Connected

Specifies the total number of clients connected to the selected SpectroSERVER.

Number of Models

Specifies the total number of models, including device and non-device models.

Number of Accountable Models

Specifies the total number of models that are included in the device count calculation (used in CA Spectrum software licensing).

Number of Device Models

Specifies the total number of physical devices modeled with model types derived from the Device model type.

Number of Polled Models

Specifies the total number of models that have Polling Status set to TRUE and a non-zero polling interval.

Number of Polled Attributes

Specifies the current number of attributes being polled by the SpectroSERVER.

Number of Logged Attributes

Specifies the current number of attributes being logged by the SpectroSERVER.

Note: Polled attributes are used to determine if a device is up or down (fault isolation). Logged attributes are used to gather statistical information. Polled attributes may or may not be logged, and logged tabulates may or may not be polled.

More information:

Physical Memory Utilization – Main Tab (see page 42)

CPU Tab

The CPU tab contains a CPU Utilization graph and information on the top CPU users.

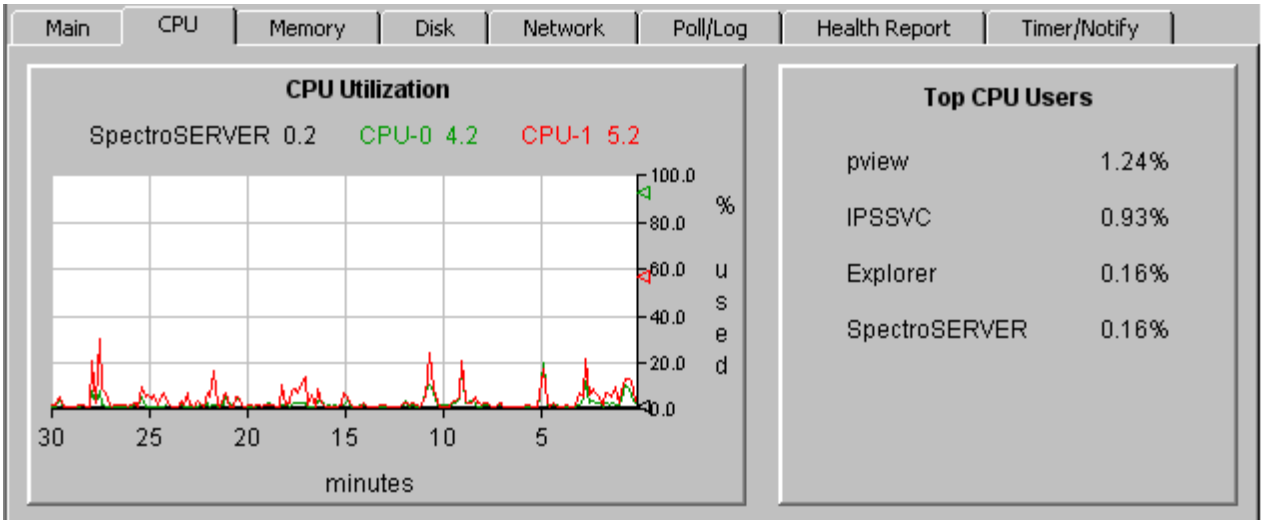

CPU Utilization

Graph displays the SpectroSERVER usage as a percentage of the total CPU capacity. Each system CPU is displayed.

Top CPU Users

Panel displays the top 10 CPU users from highest to lowest and the CPU utilization percentage for each.

More information:

CPU Tab (see page 46)

Memory Tab

The Memory tab provides the following information.

SS Memory Utilization

Graph displays how much memory the SpectroSERVER is using.

Paging Activity

Graph displays total paging activity.

Top Memory Users

Panel displays the top 10 memory users (from highest to lowest) and the memory being used by each process.

More information:

Paging Activity Graph - Memory Tab (see page 45)
Memory Tab (see page 46)

Disk Tab

The Disk tab contains a Disk Utilization graph that displays how much of a disk's read/write access capacity is being used.

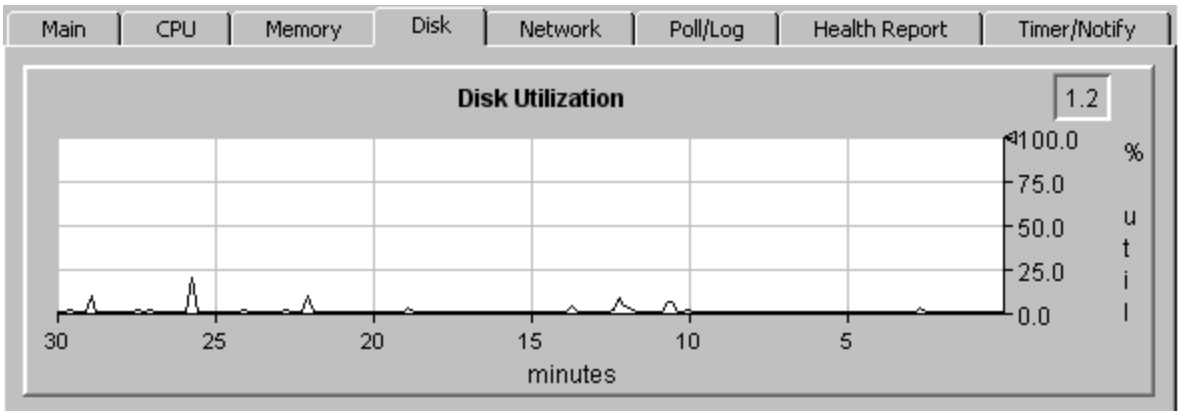

More information:

Disk Utilization Graph – Disk Tab (see page 43)
Network I/O Graph – Network Tab (see page 44)

Network Tab

The Network tab contains a Network I/O graph and a Traps Received graph.

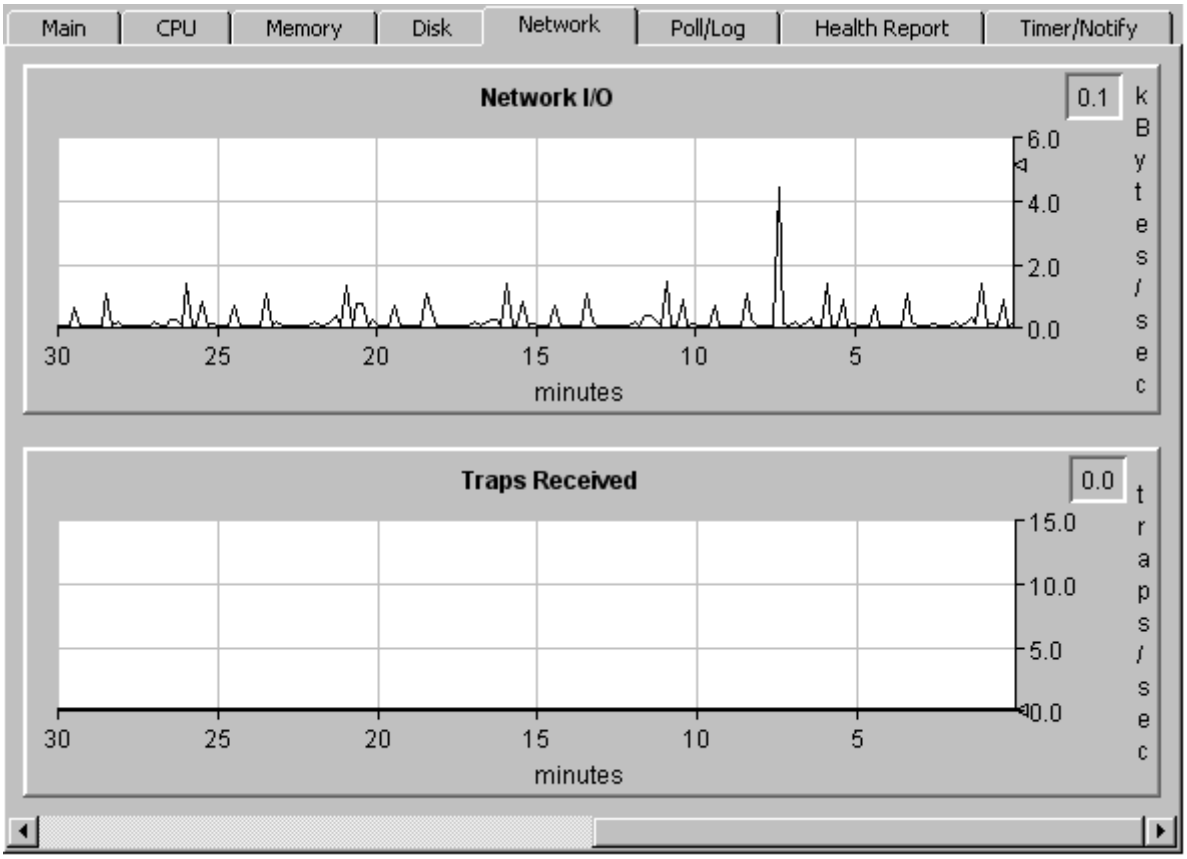

The graphs provide the following information:

Network I/O

Reflects VNM read/write bytes only. This graph does not include traffic from any other sources.

Traps Received

Displays the number of unsolicited messages, such as SNMP traps, received by the VNM.

More information:

Network I/O Graph – Network Tab (see page 44)

Poll/Log Tab

The Poll/Log tab contains the following graphs:

Poll Latency

Displays the average poll latency, which is the interval in seconds between when a scheduled polling thread is supposed to complete and when it actually completes.

Poll Threads in Use

Displays the number of poll threads in use. A poll thread is allocated to every polling operation.

Log Latency

Displays the average log latency, which is the interval in seconds between when a scheduled logging thread is supposed to complete and when it actually completes.

Log Threads in Use

Displays the number of log threads in use. A log thread is allocated to every logging operation.

More information:

Poll/Log Tab (see page 35)

Health Report Tab

The Health Report tab displays status information regarding the health report for which data is currently being collected. After the data has been collected, both the report and the average percentage of the SpectroSERVER's capacity during the reporting period are displayed on the tab.

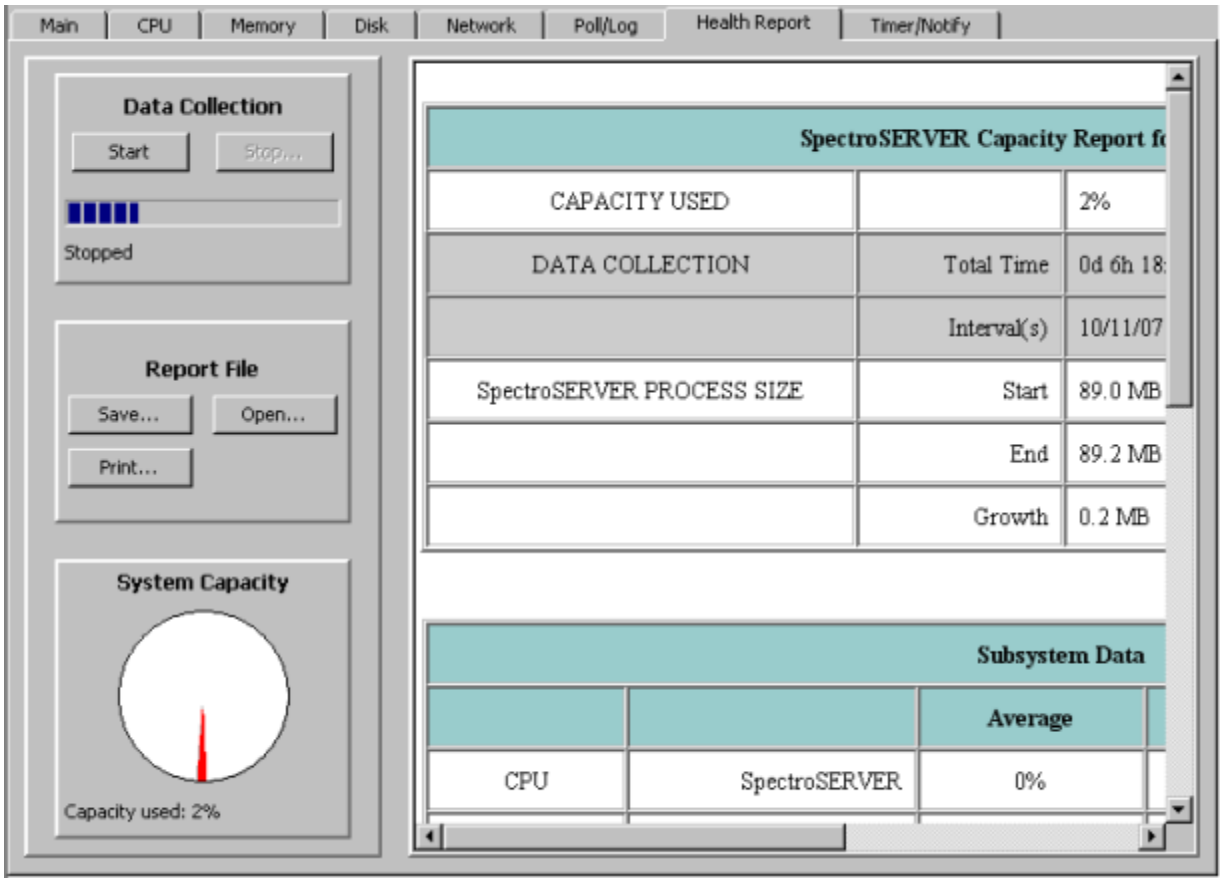

More information:

Running Health Reports (see page 49)

Timer/Notify Tab

The Timer/Notify tab contains a Timer Latency graph, a Timer Threads in Use graph, a Notification Latency graph, and a Notification Threads in Use graph.

Note: By default, this tab is not shown. To display it, you must select it from the View menu.

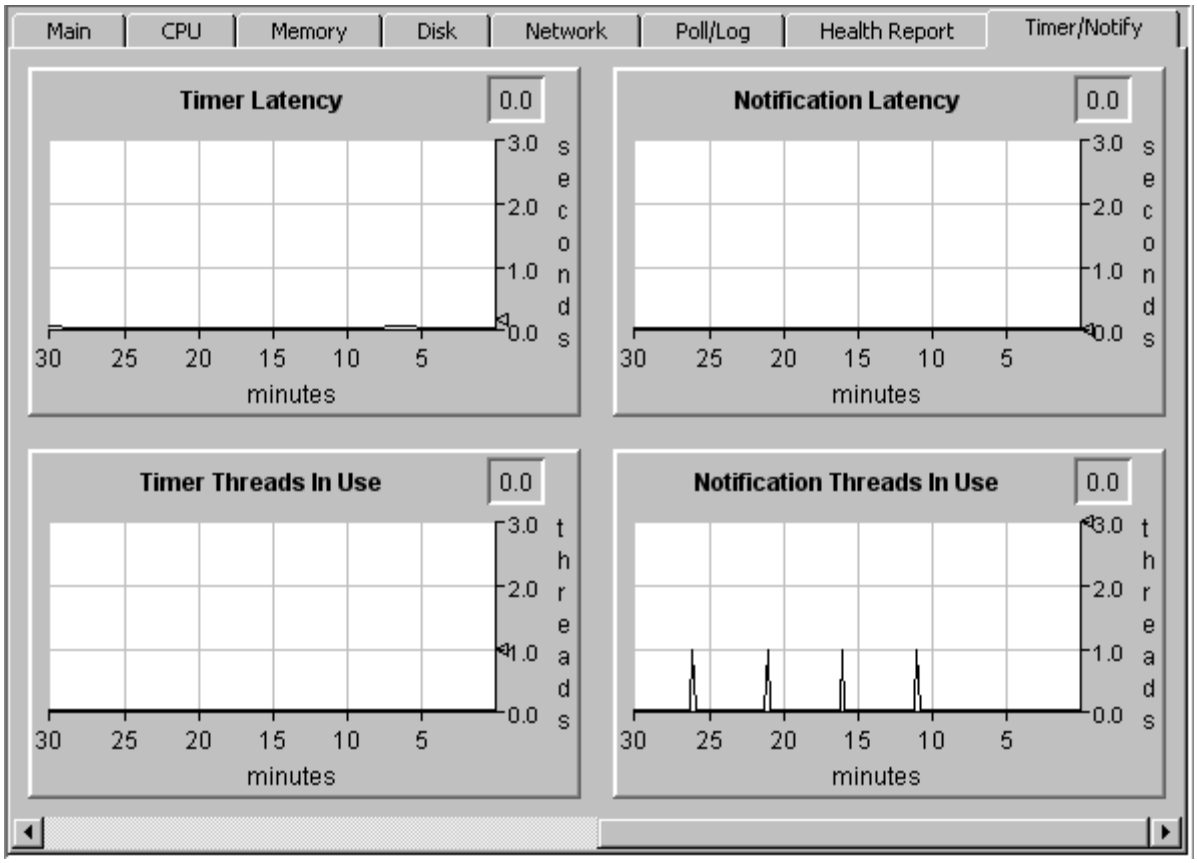

The graphs provide the following information:

Timer Latency

Displays the average timer latency, which is the interval in seconds between when a scheduled timer thread is supposed to complete and when it actually completes.

Timer Threads in Use

Displays the number of timer threads in use.

Notification Latency

Displays the average notification latency, which is the interval in seconds between when a scheduled notification thread is supposed to complete and when it actually completes.

Notification Threads in Use

Displays the number of notification threads in use.

More information:

Timer/Notify Tab (see page 37)

View Connection Details for the Connected SpectroSERVER

To determine the status of the connection between Performance View and the SpectroSERVER, examine the connection status icon () in the bottom, right corner of the Performance View.

The color of the icon indicates the status as follows:

Green

Normal

Yellow

Using backup SpectroSERVER

Red

Contact lost

To view additional connection details for the connected SpectroSERVER, click the connection status icon. The dialog that opens provides a connection log for the server.

Connect to a Different SpectroSERVER

In a distributed SpectroSERVER environment, you sometimes need to change the SpectroSERVER to which Performance View is connected.

To change the SpectroSERVER connected to Performance View

1. Click  (Change SpectroSERVER) on the toolbar.

   The Select SpectroSERVER dialog opens.

2. (Optional) Shorten the list of servers by doing either of the following:

   - To remove servers with names that do not contain a specific text string, select Filter from the drop-down list, and enter the string in the corresponding text box.
   - To parse the list for servers that contain a specific text string, select Search from the drop-down list, and enter the string in the corresponding text box. This highlights the first server that matches the criteria. Click Next as needed to cycle through the servers that match.

3. Select the SpectroSERVER you want to connect Performance View to and click OK.

   The Select SpectroSERVER dialog closes and the Performance View is now connected to the SpectroSERVER you selected.

Change the Display of Graph Axes

You can customize the axes of individual line graphs in several ways. To do so, right-click anywhere in the graph and select from the following menu options on the context menu. A check mark indicates the option is on (active).

Auto Scale

Automatically sizes the Y axis value of a graph to the highest data value collected. For instance, if the current vertical access value is 20, and data is collected for a value of 300, the Y axis is automatically set to 300.

Auto Scale overrides both Fit All Data and Set Y Axis.

Fit All Data

Selects a Y axis scale so that all current data can be displayed on the graph.

Fit All Data overrides Set Y Axis.

Reset Y Axis

Sets the Y axis to the default values.

Set Y Axis

Lets you specify the maximum and minimum values for the Y axis, as well as the number of divisions (equal intervals) between the maximum and minimum values.

Reset High Marks

Each line graph has an associated high-level mark that represents the highest data value collected since data collection was last started. This option sets all high-level marks in a graph back to zero.New high-level marks are then determined by the data that is collected from the current point forward.

Horizontal Gridlines

Lets you show or hide the horizontal gridlines.

Vertical Gridlines

Lets you show or hide the vertical gridlines.

Set Preferences

Preferences let you customize the appearance and behavior of Performance View.

To set preferences

1. Click  (Set user preferences).

   The Preferences dialog opens.
2. Configure the preferences as desired.
3. Click OK.

   The Preferences dialog closes and the preferences are now set.

More information:

Change the Colors of Graph Lines (see page 28)

Configure Preferences

You can configure the following preferences from the Preferences dialog:

Save Settings At Exit

If checked, all of the settings that you change during a Performance View session are saved and applied to subsequent sessions. These settings include the following:

- The Performance View tabs to display and the order in which to display them

  Note: You can show or hide the tabs using the selections on the View menu. To modify their display order, hide them and then re-display them in the desired order.

- The display settings for the status bar, the toolbar, and tooltips
- The last server connected to Performance View
- The directory of the health report that was saved or opened most recently
- The size and on-screen location of the main window
- All other preferences specified in the Preferences dialog

Confirm Exit

Specifies whether you want to be prompted to confirm attempts to exit the Performance View.

Show Warnings At Exit

Specifies whether you want to see all pending warnings (for example, the warning when a health report has been created but not saved) before Performance View is closed.

Show Server List At Startup

Specifies whether you want to be prompted to select the SpectroSERVER to connect to Performance View after you start Performance View.

Note: This setting only applies when you start Performance View from the command line. When you start Performance View from the CA Spectrum Control Panel, you are always connected to the SpectroSERVER to which the Control Panel is connected.

Email Report When Complete

If selected, after a health report is generated, it is automatically emailed to the addresses specified in the Email Report To field.

Email Report To

A comma-separated list of the email addresses to which to automatically send health reports after they are generated.

Title Font

The font, style, and size to use for the graph titles that appear at the top of graphs. To change this preference, click Font, make your selections, and click OK.

Label Font

The font, style, and size to use for graph labels, which are the text elements in a graph other than the graph title. To change this preference, click Font, make your selections, and click OK.

Chart Line Colors

The colors to use for graph lines. The first color button specifies the color for the first attribute in a graph, the second color button specifies the color for the second attribute in a graph, and so on.

This color palette is shared by all graphs so changing the color for a graph line changes that color in all graphs.

Change the Colors of Graph Lines

To change the color of a graph line

1. Click  (Set user preferences).

   The Preferences dialog opens.

2. Beside Chart Line Colors, select the graph (chart) line color to change.

   The Select Color dialog opens.

3. Change the color by doing one of the following:

   - To select a color swatch, click the Swatches tab, and select the swatch.
   - To specify the desired hue, saturation, and brightness of the desired color, click the HSB tab, and specify the values using the slider or the text fields.
   - To specify the red, green, and blue values of the desired color, click the RGB tab. Move the vertical slider along the color spectrum bar to change the hue, and move the small, white circle in the color square to change the saturation and brightness. Alternatively, use the text fields to individually specify the HSB values.

   Note: To exit from the dialog without applying any changes, click Cancel. To return to the color that was active when you opened the dialog, click Reset.

4. Click OK.

   The Set Color dialog closes.

5. Click OK.

   The Preferences dialog closes and your changes are applied.

Displaying and Hiding User Interface Elements

You can display or hide various user interface elements for example, tooltips or a specific tab by clicking the corresponding menu item on the View menu:

- Status Bar
- Toolbar
- Tooltips
- CPU
- Disk
- Health Report
- Main
- Memory
- Network
- Poll/Log
- Timer/Notify

Chapter 3: Evaluating the Peformance of a SpectroSERVER

This section contains the following topics:

Examining Thread Latency

This section describes how to determine if thread latency is symptomatic of CA Spectrum performance problems. The types of threads and indicators of latency are also discussed.

About Threads

Threads and Thread Latency

A *thread* is a set of commands that perform a function or a set of functions. Each thread is able to run independently from other threads. SpectroSERVER is a single-threaded application with respect to the CPU, but a multi-threaded application internally. That is, within its own process, the SpectroSERVER creates and manages multiple threads that run at the same time for tasks such as polling, logging, notifications, timers, and more.

Note: The Archive Manager runs in its own thread, which means you can utilize multiple CPUs on a machine: one for the SpectroSERVER and another for the Archive Manager. However, three or more CPUs can degrade performance.

In a multi-threaded context, while some threads are waiting (for user input, responses from devices, data retrieval, and so on), other threads can be running (for example, logging data, responding to traps, and connecting to SSAPI applications). As each thread runs in the SpectroSERVER process, it takes control of the CPU for a few microseconds and then relinquishes control to allow other threads to run. CA Spectrum performance can be influenced by the timely allocation of threads.

A SpectroSERVER maintains a pool of threads that are shared over time for CA Spectrum's polling, logging, client request, inference handler timer, inference handler notification, model activation, and model destruction processes. A SpectroSERVER sub-system uses threads from the pool-up to their individual limits-during periods of increased processing activity. These maximum limits prevent any one SpectroSERVER sub-system from dominating resources and consuming all of the available threads.

When the common pool of threads is exhausted, new threads are created, and the pool grows to meet the needs of the increased activity. Threads that are no longer needed by a process are returned to the common pool for later use. When a thread remains unused for a specified period of time, it is removed from the pool, and its resources are returned to the system. This process is called aging.

Thread latency is the amount of time between when a thread is supposed to complete and when it actually completes. It can cause problems when the number of outstanding threads accumulates as the threads take increasingly more time to complete. If CA Spectrum runs for a prolonged period of time with high thread latencies, there will be delays in the polling of devices, the logging of data, and so on, and, therefore, there will be delays in CA Spectrum's response time. For example, if a critical network device became inoperable, there would be a delay before CA Spectrum notified a network administrator of the problem.

Note: Thread latency is a symptom, not a cause, of CA Spectrum performance degradation.

Types of Threads that Affect Performance

Poll, log, timer, and notification threads, in addition to any other threads, can affect the performance of the SpectroSERVER.

Poll threads

Poll threads are responsible for polling devices on the network. CA Spectrum uses polling as a way to manage the operation and performance of the network. The SpectroSERVER code responsible for managing poll threads is called the Poll Manager.

Log threads

Log threads are responsible for logging data from the network into archive files in the CA Spectrum database. CA Spectrum can use data logging to store information about the operation and performance of the network.

By default, the Poll_Log_Ratio attribute on device models is set to 0, which effectively disables CA Spectrum's native logging method. If you require the logging of device, attribute, and port statistics, it is recommended that you use SSLogger instead of the native method, the latter of which writes the information to the Archive Manager database. SSLogger is a CA Spectrum command-line application that allows you to log statistics directly to ASCII files, which reduces the load on the Archive Manager database and eliminates the need to export the data. Of equal importance, SSLogger gives you greater control over what data to log and how frequently to log it.

Note: For more information about SSLogger, see the *CA Spectrum SS Logger User Guide*.

Timer threads

Timer threads notify inference handlers that have registered timers, also called wake-up calls.

Notification threads

Notification threads notify inference handlers about changes in an attribute for which the inference handlers have registered.

More information:

Timer Threads (see page 39)
Notification Threads (see page 39)

Thread Information View

Access the Thread Information View

To access the Thread Information view

1. Click the Topology tab in the OneClick Console.

   The Topology opens.

2. Click the VNM icon for the SpectroSERVER.

3. Click the Information tab in the Component Details panel.

4. Expand the SpectroSERVER Control, Thread Information subview.

   The Thread Information view opens.

Thread Information View

The Thread Information subview appears within the SpectroSERVER Control subview in the OneClick Console. It shows the threads in use, threads available, and peak value for the various types of threads within the SpectroSERVER process. While it shows a list of important types of threads used in CA Spectrum, be aware that the list is not exhaustive. Some of the threads available in this list include the following:

Poll Threads

Used to read the polled attributes for a model on the model's Polling_Interval.

Log Threads

Used to read and log the logged attributes for a model on the model's Polling_Interval * Poll_Log_Ratio.

Notification Threads

Used to send notifications of attribute changes to inference handlers and CA Spectrum client applications.

IH Timer Threads

Used to trigger timers in inference handlers.

Destroy Threads

Used to send model destruction notifications to inference handlers and client applications.

Model Activate Threads

Used to send model activation notifications to inference handlers and client applications.

Relation Activate Threads

Used to send relation change notifications to inference handlers and client applications.

Client Request Threads(*)

Used to handle client application requests.

Multi Request Threads(*)

Used to handle multi-model requests that originate from inference handlers and client applications.

While you can use the Thread Information view to change the available value for each thread type, generally the values should be left at their defaults. If one of the thread types is consistently running at the limit, and there are available CPU cycles, increasing a limit may reduce the associated latency. However, if CPU utilization is already above 80%, increasing thread limits will not increase throughput and may actually reduce throughput due to the increased thread overhead.

If you find that all available threads are being used for one specific type of thread, contact CA Support for assistance.

Poll/Log Tab

The Poll/Log tab provides an indicator of whether the allocation of polling threads and logging threads is affecting CA Spectrum performance. The tab has four graphs to indicate the amount of thread usage and any associated latency:

- Poll Latency
- Poll Threads in Use
- Log Latency
- Log Threads in Use

Poll Threads In Use Graph

A poll thread is allocated to every polling operation. Poll threads are allocated from a finite number of threads that are resident in CA Spectrum. The number of poll threads that get *used* by the system at any one time is proportional to the number of models to be polled on the network and the number of times each model is polled.

If the number of needed poll threads exceeds the number of available poll threads, pending poll thread requests are queued until the next poll thread becomes available. If this is the case, the number of poll threads may be insufficient for the current state of the network. In this case, call CA Support for assistance.

Log Latency Graph

Log latency is the interval between when a scheduled logging thread is supposed to complete and when it actually completes. The Log Latency graph shows the average latency for the logging process in seconds. For example, if the calculated log latency is 10 seconds, and data is supposed to be logged every 60 seconds, then the data is actually being logged every 70 seconds.

The effect of running for prolonged periods with high log latency is the delayed logging of data and could lead to other serious performance problems. If the average value in the Log Latency graph is equal to or greater than 30 seconds, this indicates that SpectroSERVER performance has been degraded. Options for improving CA Spectrum performance include: tuning the system, off-loading system demand, or upgrading the speed or capacity of the system.

Log Threads In Use Graph

A log thread is allocated to every logging operation. Log threads are allocated from a finite number of threads resident in CA Spectrum. The number of log threads that get used by the system at any one time is proportional to the amount of data being logged.

If the number of needed log threads exceeds the number of available log threads, pending log thread requests are queued until a log thread becomes available. If this is the case, the number of log threads may be insufficient for the current state of the network. In this case, call CA Support for assistance.

By default, the Poll_Log_Ratio attribute on device models is set to 0, which effectively disables CA Spectrum's native logging method. If you require the logging of device, attribute, and port statistics, it is recommended that you use SSLogger instead of the native method, the latter of which writes the information to the Archive Manager database. SSLogger is a CA Spectrum command-line application that allows you to log statistics directly to ASCII files, which reduces the load on the Archive Manager database and eliminates the need to export the data. Of equal importance, SSLogger gives you greater control over what data to log and how frequently to log it.

Note: For more information about SSLogger, see the *CA Spectrum SS Logger User Guide*.

Poll Latency Graph

Poll latency is the interval between when a scheduled polling thread is supposed to complete and when it actually completes. The Poll Latency graph shows the average latency for the polling process in seconds. For example, if the calculated poll latency is 10 seconds, and a model is supposed to be polled every 60 seconds, then the model is actually polled every 70 seconds.

The effect of running for prolonged periods with high poll latency is the delayed polling of devices. Delaying the time to poll devices can increase the response time associated with detecting a network fault. A device could go down, and the system administrator would be unaware.

If the average value in the Poll Latency graph is equal to or greater than three seconds and is sustained for a considerable period of time, this indicates that SpectroSERVER performance has been degraded. Options for improving CA Spectrum performance would include tuning the system, off-loading system demand, or upgrading the speed or capacity of the system.

Timer/Notify Tab

The Timer/Notify tab displays the number of timer threads or notification threads in use and any associated latency. The tab has four graphs:

- Timer Latency
- Timer Threads in Use

- Notification Latency
- Notification Threads in Use

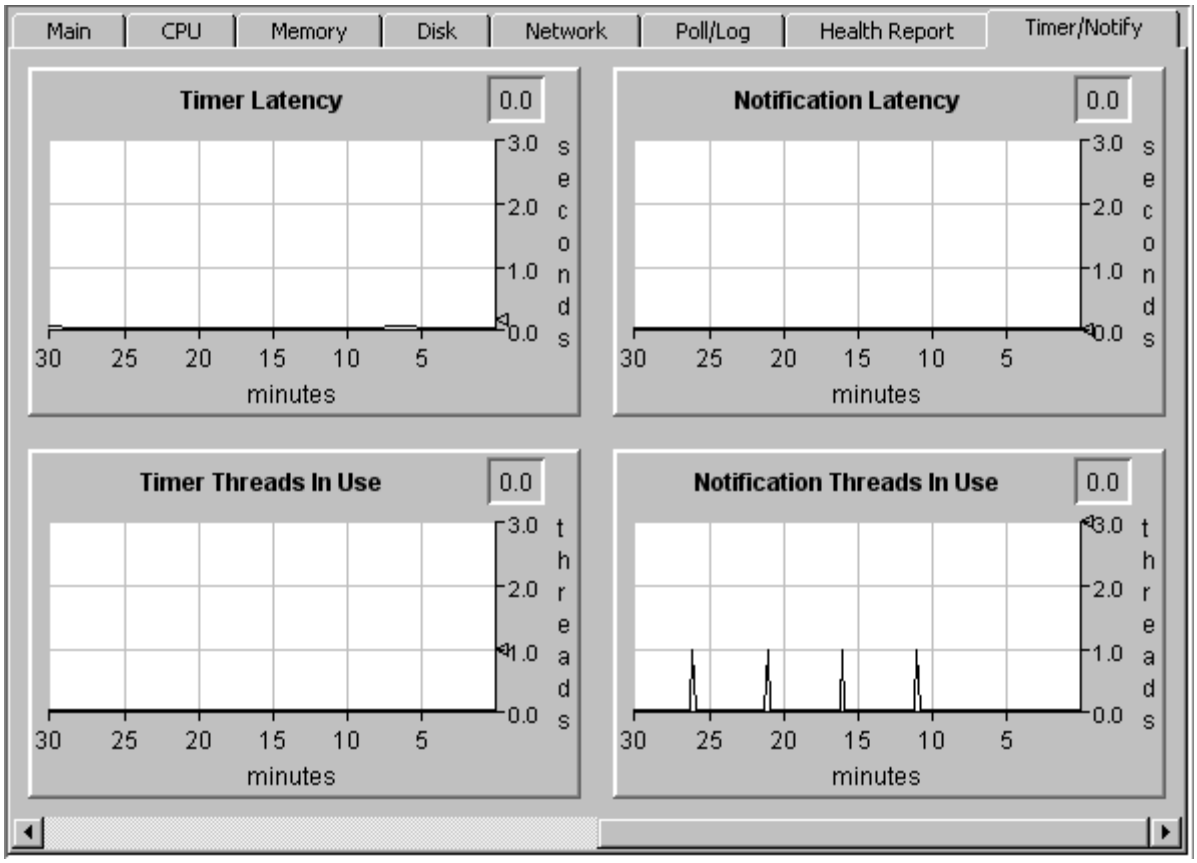

Theory of Operations

The SpectroSERVER is a poll-driven and event-driven system. That is, it actively polls managed elements for state changes, generates events based on the changes, and notifies the inference handlers that have registered for the events. Events for which an inference handler can register include model creation, model destruction, attribute value changes, association creations and destructions, and others.

Inference handlers are code segments associated with a model type that define the behavior of the model type, and that execute on behalf of instantiated models of the model type. They can be triggered by notification threads or timer threads.

Notification Threads

A SpectroSERVER process called the Notification Manager reads attribute changes received from polled devices and then runs inference handlers to process the data. Another SpectroSERVER process called the Poll Manager is responsible for detecting changes in attributes whose poll flags are set.

When a change in an attribute value is detected, the Poll Manager alerts the Notification Manager, and the Notification Manager forwards the events to each registered inference handler.

Notification threads are the mechanism used to notify inference handlers of a change in an attribute. Notification threads are used to run the trig_attr_change and trig_watch_change inference handler methods. Both of these methods are notifications of attribute value changes.

For example, assume the ifInDiscards attribute has changed from a count of 110 to 150. The SpectroSERVER sends a notification thread to the inference handler that has expressed an interest in this attribute. The inference handler then runs the trig_attr_change method.

Timer Threads

Timer threads are used to notify inference handlers that have registered for timers (wake up calls). When an inference handler registers a "wake up call" with the SpectroSERVER, the SpectroSERVER uses a timer thread to run that inference handler's trigger method after a specified time interval. Timer threads are used to run the trig_timer inference handler method in the same way that notification threads are used to run the trig_attr_change and trig_watch_change inference handler methods.

For example, assume a router's primary address becomes non-operational, and a secondary address must be used. An inference handler registers with the SpectroSERVER for a timer so it can determine when the primary address is once again operational. The inference handler then runs the trig_timer method.

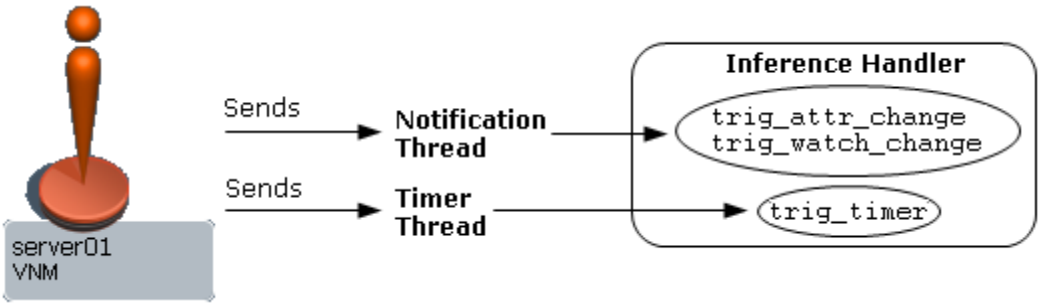

Timer Latency Graph

Timer latency is the interval between when a scheduled timer thread is supposed to complete and when it actually completes. The Timer Latency graph shows the average latency for the timer process in seconds. For example, if the average timer latency is 10 seconds, and an inference handler has registered for a timer thread every 60 seconds, then the inference handler's corresponding trigger method is actually being activated by the timer thread every 70 seconds.

The effect of running for prolonged periods with high timer latency is the delayed activation of inference handler triggers and, therefore, network monitoring. If the Timer Latency graph shows more than three seconds sustained, this indicates performance problems. Call CA Support for assistance.

Timer Threads In Use Graph

The Timer Threads in Use graph displays the number of timer threads in use.

If the number of needed timer threads exceeds the number of available timer threads, pending timer thread requests are queued until a timer thread becomes available. If this is the case, the number of timer threads may be insufficient for the current state of the network. In this case, call CA Support for assistance.

Notification Latency Graph

Notification latency is the interval between when a scheduled notification thread is supposed to complete and when it actually completes. The Notification Latency graph shows the average latency for the notification process in seconds. For example, if the average notification latency is 10 seconds, and an inference handler has registered for a notification thread every 60 seconds, then the inference handler's corresponding trigger method is actually being activated by the notification thread every 70 seconds.

The effect of running for prolonged periods with high notification latency is the delayed activation of inference handler triggers and, therefore, network monitoring.

Notification Threads in Use Graph

The number of notification threads that get used by the system at any one time is proportional to the number of attribute or watch changes occurring on the network. Once a network is up and has achieved stability, the number of notification threads required to monitor the system should remain steady and small.

If the number of needed notification threads exceeds the number of available notification threads, pending notification thread requests are queued until a notification thread becomes available. If this is the case, the number of notification threads may be insufficient for the current state of the network. In this case, call CA Support for assistance.

Examining Memory Usage

This section describes how to determine if swapping or paging activities are causing CA Spectrum performance problems. Swapping or paging is a processing technique that involves transferring data back and forth from a main storage area to an auxiliary storage area for example, from memory to disk. Pages refer to the individual units of data transfer used to swap data.

Considerable swapping activity indicates a shortage of system memory, as data has to be temporarily transferred from memory to disk to make room for various processes to run. A large amount of paging activity is an indication of a high amount of swapping. Because use of the disk is much slower than use of the physical memory, paging can result in performance problems and, therefore, should be kept to a minimum.

The swap space and the physical memory collectively make up the available memory or virtual memory. This is frequently a bottleneck for the overall system performance. As a rule of thumb, the system should have two times the amount of physical memory configured as swap space. If this space is low, reconfigure the system with more memory and swap space.

The following information provided in the Performance View can help you determine if SpectroSERVER performance has degraded due to insufficient memory:

- Physical Memory Utilization graph on the Main tab
- Disk Utilization graph on the Disk tab
- Paging Activity graph on the Memory tab
- SS Memory Utilization graph on the Memory tab

Additionally, memory usage of the SpectroSERVER process can also be monitored in CA Spectrum OneClick. If defined threshold values are exceeded, events are logged and alarms are triggered. These thresholds are described in Using Performance Thresholds (see page 47).

Indicators of Memory Problems

Memory management is important for achieving and maintaining CA Spectrum performance. Memory shortage has an immediate and significant effect on SpectroSERVER fault detection response times.

One direct consequence of running CA Spectrum with insufficient memory is that it appears disk I/O bound. In many ways, a memory-shy system can appear to be disk bound because of the high paging and swapping activity occurring on the disk. When memory is the primary bottleneck, only increasing the amount of memory, or decreasing the demand for memory, returns the system back to acceptable levels of performance. Adding more disk capacity or speed has a negligible effect.

Physical Memory Utilization – Main Tab

The Physical Memory Utilization area of the Main tab displays how much physical memory the system is currently using and how much memory CA Spectrum is currently using.

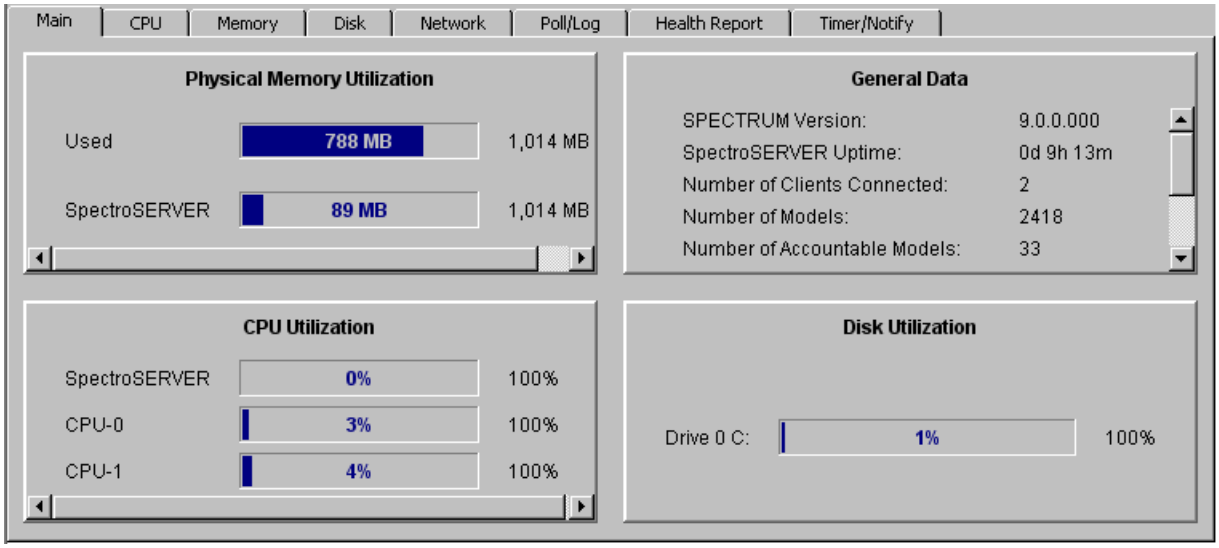

The value for total physical memory is displayed to the right of the Used and SpectroSERVER bar graphs. This is the actual amount of physical memory contained in the machine to which the Performance View is connected. As such, this value is machine-dependent. If SpectroSERVER uses a large percentage of the virtual memory, you should consider upgrading the memory or allocating more swap space.

Disk Utilization Graph – Disk Tab

The Disk Utilization graph records all disk transfers, including all physical disks attached to the system. Disk utilization refers to how busy the disk is that is, the percentage of time it is being used.

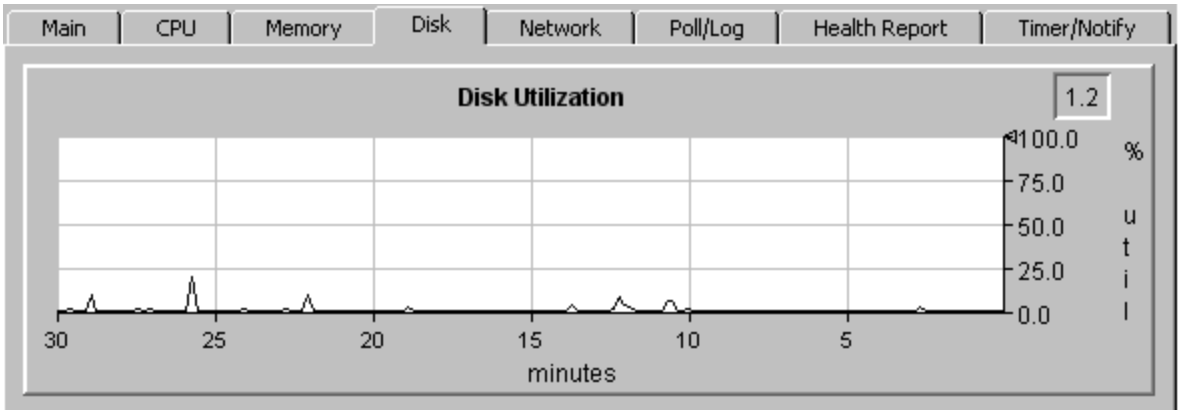

If the Disk Utilization graph is high, data logging might also be high. If this graph shows continuously high numbers, consider changing the logging ratio of some of your models and running the PMCount utility.

More information:

Adding SpectroSERVERs (see page 69)

Network I/O Graph – Network Tab

The Network I/O graph records network I/O from the Ethernet interfaces on the system. This includes I/O activity from the SpectroSERVER. An increase can be expected when either new models are created in the database or the polling intervals are changed.

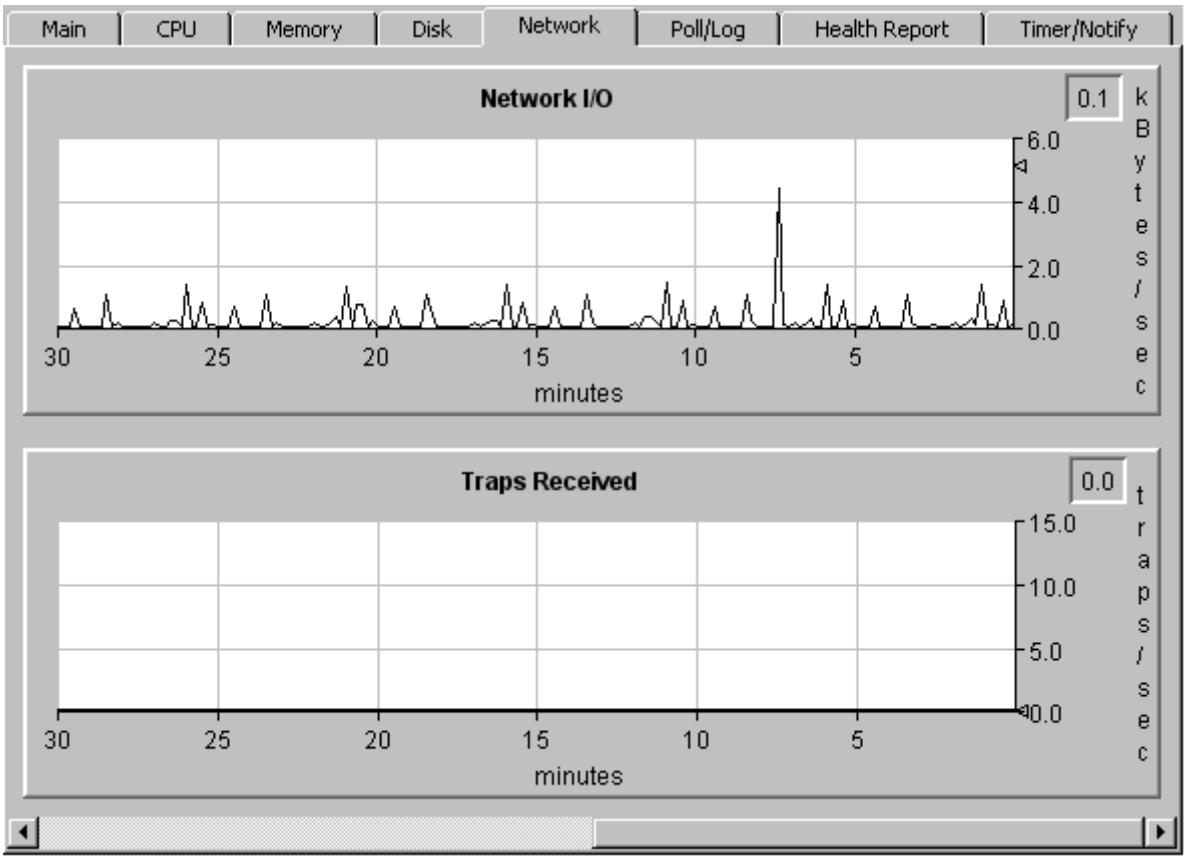

Paging Activity Graph - Memory Tab

The Paging Activity graph displays the number of system pages over time. If the values in this graph are persistently high, there is not enough physical memory. A situation like this can result from having more processes running than available physical memory.

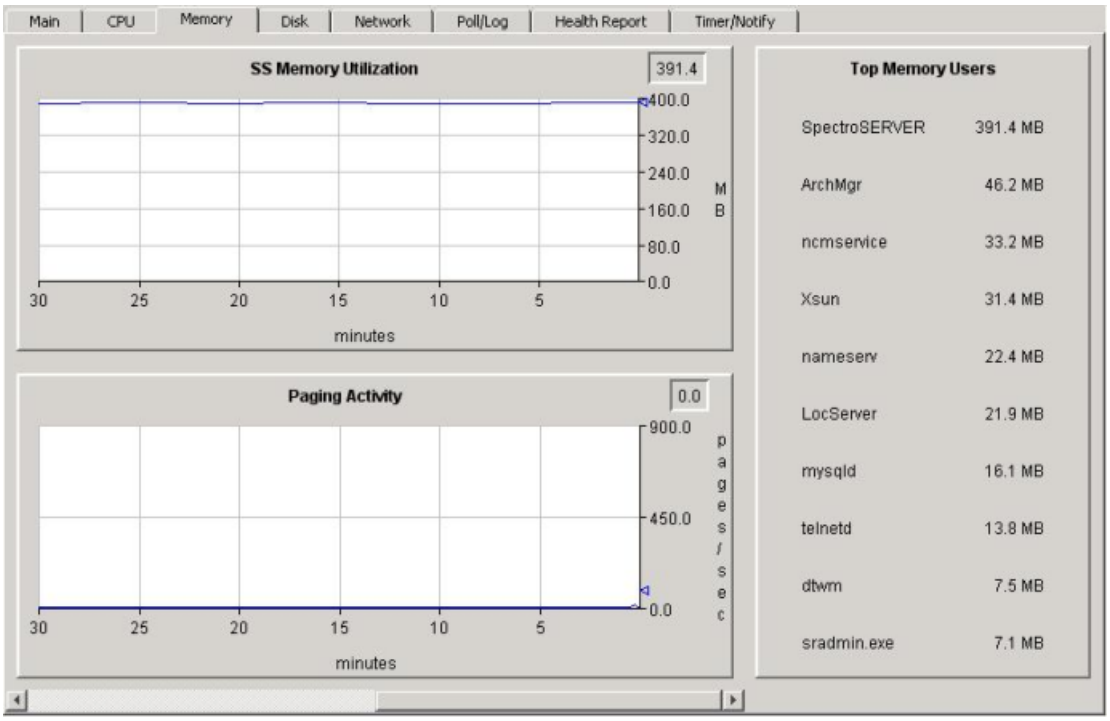

A persistently high value indicates that the system is heavily loaded, and a memory upgrade should be considered. If you see high paging activities, consider reducing the number of non-CA Spectrum processes that are running or increasing the physical memory.

Examining the Application Load

This section describes how to determine if there are too many applications running on the system. These applications would also include SpectroSERVER and OneClick. The following views can help you determine if too many applications are running on the system and, therefore, are causing performance problems:

- Memory tab
- CPU tab

Memory Tab

The Memory tab displays the top users of system memory. This view can help you determine which applications are consuming the most system memory.

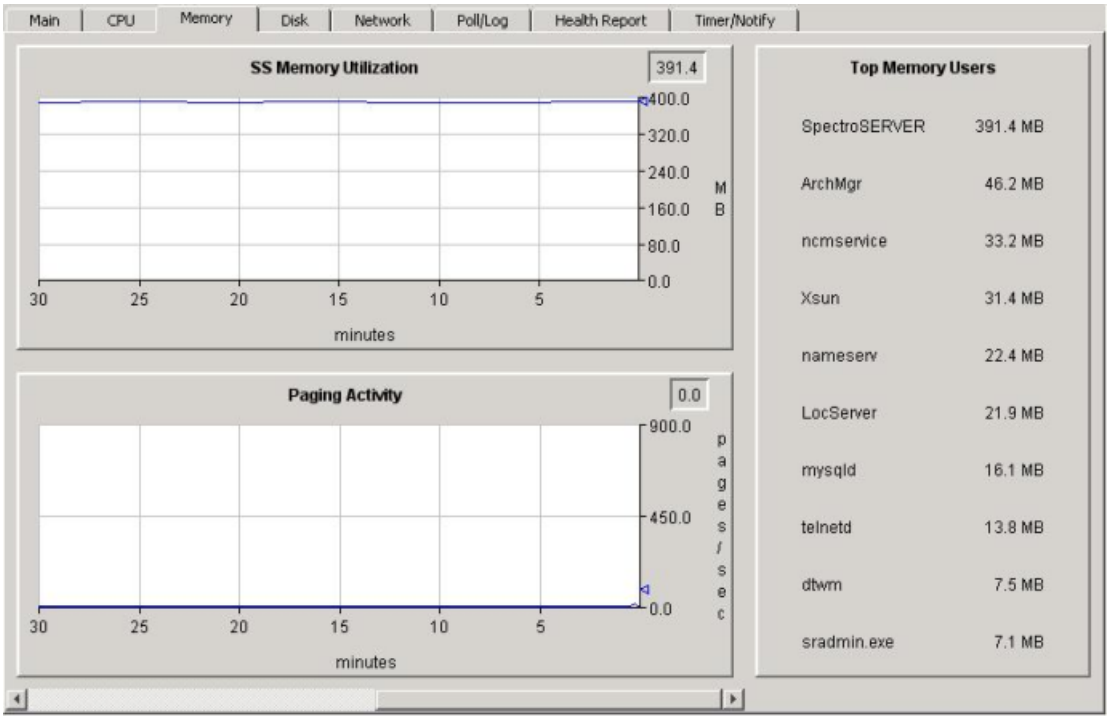

CPU Tab

The CPU tab displays the top users of the system CPU. This view can help determine which applications are consuming the most system CPU.

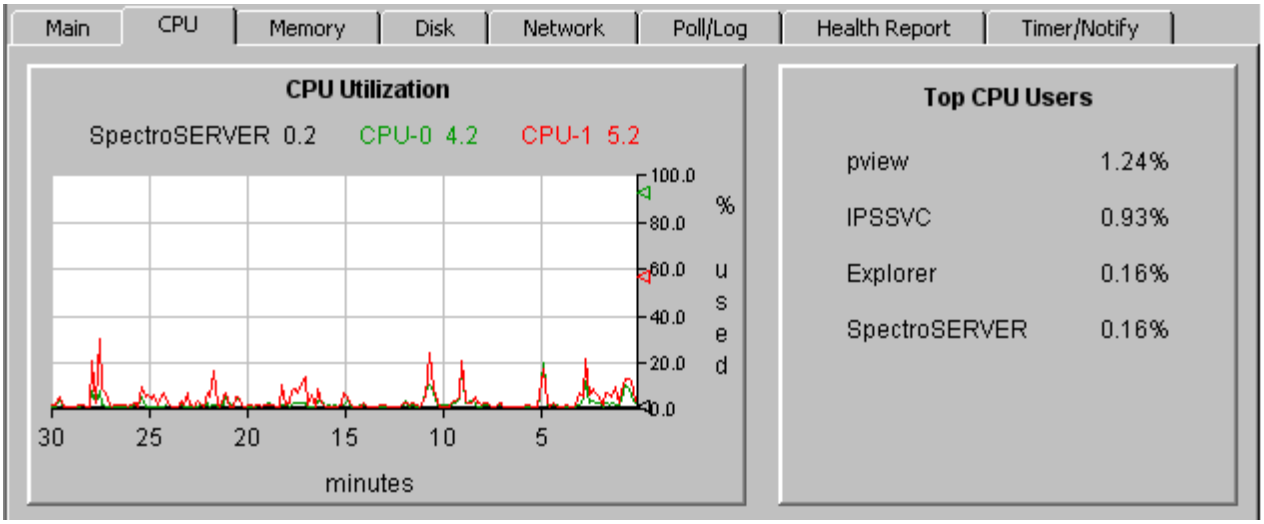

If SpectroSERVER CPU Utilization is low, but CA Spectrum performance is still slow, then other applications may be straining the system.

Examining the Number of Connected Clients

A large number of client connections can place a heavy load on the server. To identify the number of active client connections, click the Main tab, and in the General Data panel, view the value for Number of Clients Connected.

Note: A OneClick web server represents one SpectroSERVER client regardless of how many OneClick clients are running against that web server.

We recommend that you install only CA Spectrum applications except OneClick on the SpectroSERVER machine. Installing OneClick on a single-CPU SpectroSERVER host system can degrade the performance of both the SpectroSERVER and OneClick. Installing OneClick on a separate, dedicated system can maximize the performance of both the SpectroSERVER and OneClick.

Using Performance Thresholds

In addition to performance monitoring capabilities provided in Performance View, the SpectroSERVER process can also be monitored in CA Spectrum OneClick. If defined threshold values are exceeded, events are logged and alarms are triggered. The following performance metrics are monitored:

- SNMP traps (see page 47)
- Memory usage (see page 48)

SNMP Traps

When excessive SNMP trap rates are experienced by the SpectroSERVER process, degradation may occur. If the specified threshold rate is exceeded, corresponding events and alarms will be generated on the SSPerformance and VNM models.

The default trap rate threshold is 100 traps per second. To change it, you must modify the value the EventDisp file, which is located at:

```
$SPECROOT/SS/CsVendor/Cabletron/EventDisp
```

The following line controls the SNMP trap rate value, which is currently set to 100.0 (traps per second):

```
"{v 0x11eca} >= {R 100.0 }",       "0x00010f92 -:-", \
```

Notes:

- The threshold rate must be exceeded for at least 300 seconds for an alarm to trigger. This time period value cannot be changed.
- The trap rate is monitored by attribute vnm_snmp_traps_ps (AttrID = 0x11eca) on the SSPerformance model.

Memory Usage

If the SpectroSERVER process has increased in memory size, the SpectroSERVER can be at risk of termination due to memory exhaustion. If a specified threshold rate for either physical or virtual memory is exceeded for a period of time, corresponding events and alarms will be generated on the SSPerformance and VNM models.

Separate threshold values are used for physical memory and virtual memory. The default memory size for each is 2.5 GB. To change either, you must modify the value the EventDisp file, which is located at:

```
$SPECROOT/SS/CsVendor/Cabletron/EventDisp
```

The following lines control the memory sizes, which are currently set to 2.5 GB:

```
"{v 0x11e8b} >= {R 2500000000.0}", "0x00010f95 -:-", \     <- physical memory
"{v 0x12e62} >= {R 2500000000.0}", "0x00010f98 -:-"        <- virtual memory
```

Note: The threshold rate must be exceeded for at least 300 seconds for an alarm to trigger. This time period value cannot be changed.

Chapter 4: Running Health Reports

This section contains the following topics:

Health Reports

You can use Performance View's Health Report feature to measure and report on the relative health of your SpectroSERVER and the system on which it is installed for any period from 6 to 24 hours.

You can start the reporting process at any time either from the main Performance View window or from the command line. The command line option also lets you take advantage of your host system's native scheduling service to run the report automatically over a specified time period or at regular intervals.

By default, Health Report collects the following data at 10-second intervals for a 24-hour period:

- CPU, disk, and memory usage data
- network I/O and trap data
- poll, log, timer, and notification latency data

The full 24-hour period is recommended because it helps to ensure the collection of a typical day's worth of data. If data representing your typical day's workload can be collected during a shorter period of time, you can use that time period instead. However, be careful not to exclude data collection for jobs that run during non-business hours, such as backups executed at midnight.

Once data collection has occurred for a full 24-hour period, a health report is generated automatically. If you ran the report from the Health Report tab, the report is displayed there, and you can save it to a location of your choice. If you ran the report from the command line, the report is written to a file.

Note: From the Preferences dialog you can configure Performance View to automatically email health reports to a list of addresses after they are generated.

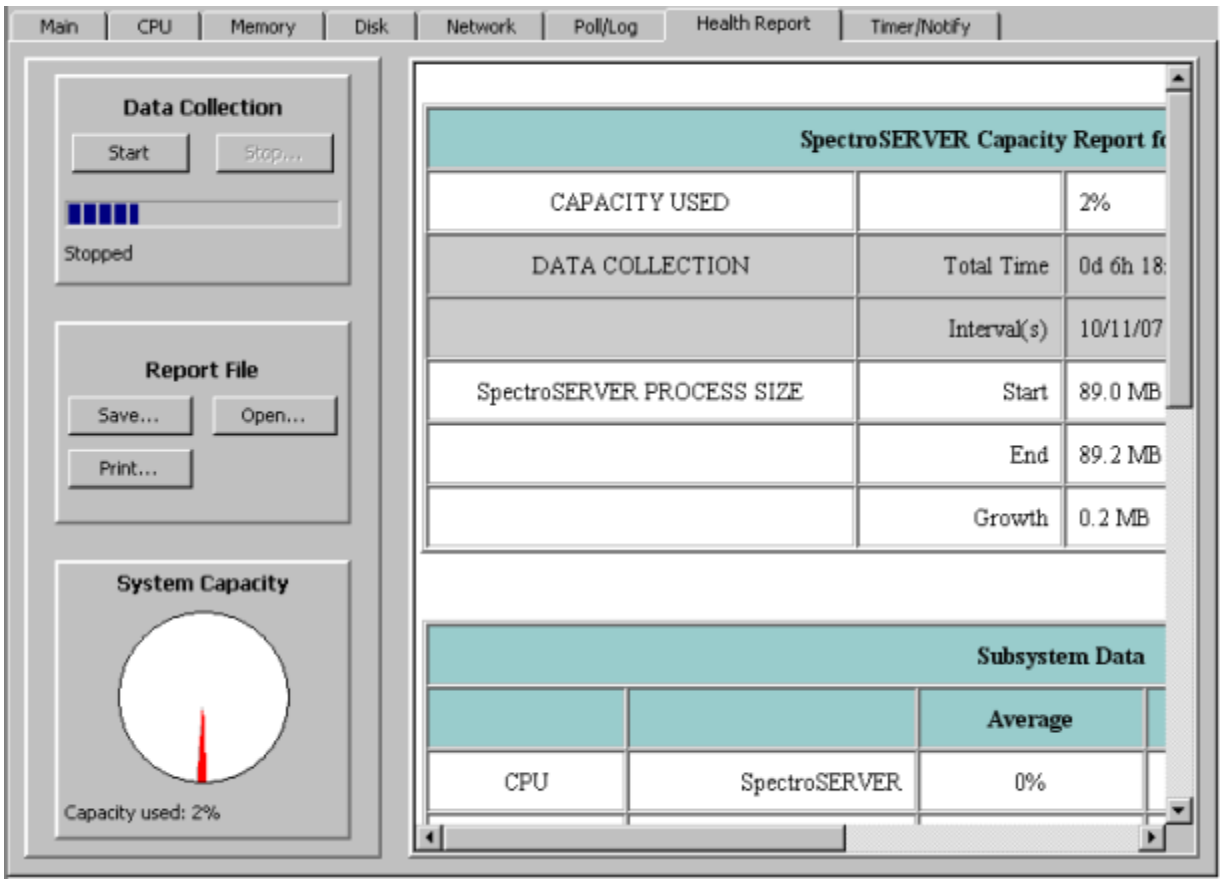

The relative health of each system resource for example, CPU usage is determined by analyzing the Average, Peak, and % Over Critical Value readings collected. Any values that exceed predefined thresholds are flagged (displayed in red) to indicate a potential performance problem.

The relative health of the SpectroSERVER is determined by applying performance algorithms to the collected data. If it is determined that performance has degraded, the likely causes and recommendations for improving performance are also provided in the report.

Start Data Collection

To start data collection for a health report

1. Click the Health Report tab.
2. Click Start in the Data Collection section.

   The message area in the panel dynamically updates to indicate the number of hours and minutes remaining in the default 24-hour reporting period. After all of the data has been collected, Health Report analyzes the data and generates a health report, which is displayed in the large area to the right of the Data Collection panel.

   Note: As long as the new report remains on display, the average percentage of the SpectroSERVER's capacity used during the reporting period will also be shown in the graph in the System Capacity panel.

More information:

Set Preferences (see page 26)
Configure Preferences (see page 27)

Stop Data Collection

Data collection for a Health Report stops automatically after data has been collected for a total of 24 hours, but you can stop or pause it manually at any time.

To stop data collection for a health report

1. Click the Health Report tab.
2. Click Stop in the Data Collection section.
3. (Optional) Choose *one* of the following options if it has been fewer than 24 hours since you started the data collection:

   Resume Data Collection

   Select this option to restart data collection. For example, if you stopped data collection after one hour, selecting this option restarts data collection and continues it for 23 more hours that is, until the default reporting period of 24 hours has been reached. In other words, the total time of data collection does not have to be contiguous. You can start and stop data collection for the same report as many times as you want. However, you must collect data for at least 6 hours to generate a report.

   Stop and Analyze Data

   This option appears if data has been collected for at least 6 hours. Select this option to immediately generate a health report based on the collected data. The report will remain displayed until you start data collection for a new report or exit Performance View.

   Note: You cannot resume data collection for the same report once you have clicked this button.

   Stop and Delete Data

   Select this option to end the data collection process and delete all collected data. No report is generated.

Save Health Reports

You can save the current health report to preserve the data.

To save a health report

1. Click  (Save the current health report).

   The Choose Directory and Filename for HTML Report dialog opens.

   Note: When specifying a filename for the report, you do not need to include the .htm file extension. It is added automatically.

2. Navigate to the folder in which to save the report, enter a filename for File name, and click Save.

   The health report is saved.

Open Health Reports

You need to open health reports if you want to view or print them.

To open a health report

1. Click  (Open a previously saved health report).

   The Choose Report File to Open dialog opens.

2. Navigate to the report, select it, and click Open.

   The report is displayed in a separate view-only window.

Print Health Reports

You can print a currently displayed health report on the Health Report tab or a saved health report.

To print the currently displayed health report

1. Click  (Print the health report).

   The Print dialog opens.

2. Specify the appropriate settings in the Print dialog and click OK.

   The health report prints.

To print a saved health report

1. Open the report (see page 53).

   The SpectroSERVER Capacity Report dialog opens, displaying the selected report.

2. Click Print.

   The Print dialog opens.

3. Select the appropriate settings and click OK.

   The health report prints.

4. Click Close.

   The SpectroSERVER Capacity Report dialog closes.

More information:

Open Health Reports (see page 53)

Run Health Reports from the Command Line

You can start data collection for a health report at any time by entering the desired parameters from the command line, or you can use your system's native scheduling service to execute the command at a specified time or at regular intervals.

The command line executable is named pviewrep and is located in the *<$SPECROOT>*\PView directory.

The syntax for the pviewrep command is as follows:

```
pviewrep vnm -c collectTime -e addrList
```

vnm

Specifies the name of the SpectroSERVER for which to run a report.

-c *collectTime*

Specifies the number of hours for which to collect data.

Note: The minimum number of hours required for a report is six hours. If you specify fewer than six hours, the report still collects six hours of data before it is generated.

-e *addrList*

Specifies a comma-separated list of email addresses to which to send the completed report. If you are specifying more than one address on Windows systems, the list must be enclosed in quotation marks, for example, "address1,address2,address3".

Important! If you are working in a Windows environment, you must have the Windows Messaging Subsystem or Messaging Application Programming Interface (MAPI) subsystem installed to be able to send messages using the -e option. If the subsystem is not installed, the executable will fail to send the email notification; it will look for a registry entry under HKEY_CURRENT_USER\\Software\\Microsoft\\Windows NT\\CurrentVersion\\Windows Messaging Subsystem\\Profiles and will create an application event if the entry cannot be found. The email system installed might require confirmation steps before sending the email.

Reports generated by pviewrep are saved automatically with the name of the SpectroSERVER host machine and an .htm extension. Sequential numbers are added to keep subsequent reports from overwriting existing ones. For example, the first report generated for a machine named "ace" will be ace.htm, the second report will be ace_1.htm, and so on.

By default, reports generated with pviewrep are saved to the directory used for the most recent health report generated from the Performance View main window. If that directory is not available, the report is saved to the current working directory. If the file cannot be written to the current working directory, it is written to the program's standard output stream.

Note: Health reports distributed automatically via email will be in plain text format, not HTML.

Interpreting Health Reports

A health report includes three major sections:

- SpectroSERVER Capacity
- Subsystem Data
- Analysis

SpectroSERVER Capacity

The SpectroSERVER Capacity section of a health report provides the following information:

Capacity Used

The average percentage of SpectroSERVER capacity being used during the data collection period. This percentage is also shown graphically in the System Capacity panel on the Health Report tab.

Data Collection

The total amount of time that data was collected, as well as the start and end times for the individual data collection intervals that make up the total time.

SpectroSERVER Process Size

The amounts of RAM that the SpectroSERVER used at the beginning and the end of the data collection period, as well as the difference (Growth) between the two amounts.

Subsystem Data

The Subsystem Data section of a health report provides Average, Peak, and % Over Critical Value readings for individual parameters within various performance categories (CPU, LATENCIES, DISK, MEMORY, and NETWORK).

Any Average or % Over Critical Value reading that exceeds Performance View's predetermined threshold value for that parameter is flagged (displayed in red). Flagged parameters indicate possible performance problems and are used to determine the message displayed in the Analysis area of the report.

The Average, Peak, and % Over Critical Value columns are defined as follows:

Average

The sum of all of the values recorded for the parameter during the total data collection period (Total Time) divided by the number of collection points. A collection point occurs every 10 seconds.

Peak

The highest value recorded for a parameter during the total data collection period.

% Over Critical Value

The percentage of the data collection period during which a value for a parameter exceeded the predetermined threshold value for that parameter.

Analysis

The Analysis section of a health report provides a narrative describing the results of the analysis performed on the collected data. The narratives are predetermined for each of the different scenarios.

An example of a narrative is shown in the following image.

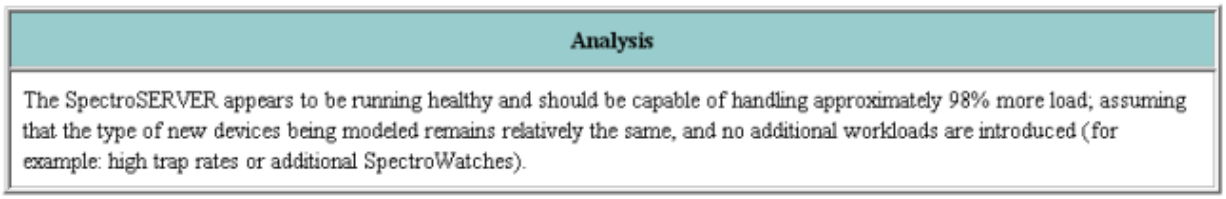

Analysis
The SpectroSERVER appears to be running healthy and should be capable of handling approximately 98% more load; assuming that the type of new devices being modeled remains relatively the same, and no additional workloads are introduced (for example: high trap rates or additional SpectroWatches).

Analysis Narrative

The narrative that appears in the Analysis part of a health report depends on whether parameter values exceeded predetermined thresholds and, if so, which ones.

Report results fall into three possible categories:

- No parameters are flagged
- No Average reading is flagged, but one or more % Over Critical Value readings are flagged
- One or more Average readings are flagged

If all Average numbers and % Over Critical Value numbers are fine (not flagged), the following narrative is displayed:

The SpectroSERVER appears to be running healthy and should be capable of handling approximately (100 - % Capacity) % more load; assuming that the type of new devices being modeled remains relatively the same, and no additional workloads are introduced (for example: high trap rates or additional Watches).

If all Average numbers are fine, but % Over Critical Value number is flagged, the following narrative is displayed along with one of the explanations from the subsequent bulleted list:

On average, the SpectroSERVER is running within an acceptable resource utilization range, however, as indicated by a high "% over threshold" value, there are excessive periods of time where one (or more) of the system resources are over-utilized. This could be an indication that the resource could be close to a premature bottleneck. Based on the calculated values from this data collection period, it appears that the following problems might exist.

If the % Over Critical Value number for CPU is flagged:

There were periods of time where the CPU was too busy to handle the current workload. This could be caused by other applications running on the SpectroSERVER machine periodically consuming available CPU or that the system is approaching its maximum management capacity. Only the bare minimum CA Spectrum applications should be running on the SpectroSERVER machine.

If the % Over Critical Value number for log latency is flagged, but it is not flagged for paging:

There are excessive log latencies, which could be a result of too much/too frequent logging or unexpected workloads, such as high event rates.

If the % Over Critical Value number for paging and log latency are both flagged:

There was excessive memory swapping on the system. This could be caused by other applications running on the SpectroSERVER machine periodically consuming available memory. Only the bare minimum CA Spectrum applications should be running on the SpectroSERVER machine. On Solaris machines, check that priority paging has been enabled. On machines running NT Server, check that the NT page stealing algorithm is set to "balance."

If the % Over Critical Value number for paging is flagged, but it is not flagged for log latency:

There were periods of time of excessive memory swapping. Although the system is running out of memory, it is not affecting system performance yet, but could be an issue shortly. This could be caused by other applications running on the SpectroSERVER machine periodically consuming available memory. Only the bare minimum CA Spectrum applications should be running on the SpectroSERVER machine. On Solaris machines, check that priority paging has been enabled. On machines running NT Server, check that the NT page stealing algorithm is set to "balance."

If the % Over Critical Value number for paging and disk utilization are both flagged, but it is not flagged for log latency:

There were periods of time of excessive memory swapping. Although the system is running out of memory, it is not affecting system performance yet, but could be an issue shortly. This could be caused by other applications running on the SpectroSERVER machine periodically consuming available memory. Only the bare minimum CA Spectrum applications should be running on the SpectroSERVER machine. On Solaris machines, check that priority paging has been enabled. On machines running NT Server, check that the NT page stealing algorithm is set to "balance."

If the % Over Critical Value number for disk utilization is flagged, but it is not flagged for paging:

There were periods of time of excessive disk utilization. The system could be close to a disk bottleneck possibly due to too much/too frequent logging or unexpected workloads, such as high event rates.

If the % Over Critical Value number for notification is flagged:

It appears that there was an excessive amount of notifications being generated. This could be a result of models being created or reconfigured, or an unusually high rate of events had been generated during the data collection period. Try running the collection again.

If the % Over Critical Value number for traps is flagged:

It appears that there was an excessive amount of traps being generated. This could be caused by an unusually high rate of events being generated during the data collection period. Try running the collection again.

If the % Over Critical Value number is flagged for only poll latency:

There were periods of time of excessive polling latency. This could be a result of network problems preventing contact with a specific device(s). Check that there are no abnormal network issues and run the data collection again.

If any other scenario exists where all Average numbers are fine, but at least one % Over Critical Value number is flagged:

There may have been a situation which caused the SpectroSERVER machine to experience abnormal behavior. Try running the data collection again, and if the problem persists, contact support.

If any Average number is flagged, then the following narrative is displayed along with one of the explanations from the subsequent bulleted list:

The SpectroSERVER is currently running at full capacity and no other models or additional work should be introduced. Steps should be taken to minimize the current workload or some of the models/work may have to be moved off to an additional SpectroSERVER. CA Support can perform a sizing to verify that your machine configuration is adequate for your existing database.

If the Average number is flagged for CPU, poll, or timer latency:

It appears that the SpectroSERVER has too much work to process. Increase polling intervals on all devices to the minimum required that will achieve business goals/needs.

If the Average number is flagged for log latency and disk utilization or disk queue length, but it is not flagged for paging:

It appears that the SpectroSERVER is doing too much logging. Either increase the poll-log ratio (so that logging occurs less frequently) or turn off logging on models that are not necessary. A disk array running RAID 0 could help alleviate this problem if one is not already in use.

If the Average number for traps is flagged:

There are many traps/events being generated that are causing the SpectroSERVER to saturate with work. Check to see if there is a specific device that is causing a trap storm and also re-evaluate business goals/requirements to determine which traps are unnecessary and turn off trap forwarding on those devices, or turn off CA Spectrum's trap handling for those device types.

If the Average number for paging *or* for paging and disk utilization are flagged:

It appears that the system is running low on physical memory. The SpectroSERVER process should never exceed available physical RAM (in other words, the process should not be getting paged in and out as a result of a memory shortage). Any unnecessary applications should remain closed or more memory needs to be added to the SpectroSERVER machine. If there is a local OneClick running on the same system, splitting it off to run remotely would help reduce memory consumption.

If the Average number for notification is flagged:

It appears that there was an excessive amount of notifications being generated. This could be a result of models being created or reconfigured, or an unusually high rate of events had been generated during the data collection period. Try running the collection again.

If any other scenario exists where an Average number is flagged:

There may have been a situation which caused the SpectroSERVER machine to experience abnormal behavior. Try running the data collection again, and if the problem persists, contact support.

Chapter 5: Tuning a SpectroSERVER

This section contains the following topics:

Introduction to Tuning a SpectroSERVER

Once you have determined the reasons for degraded CA Spectrum performance, you can tune the OneClick to improve performance by doing the following:

- Modifying the polling interval and poll to log ratio of essential device models and application models, and disabling polling of non-essential models. This reduces the network traffic and the resulting latency that degrades performance.

  The polling interval is the time interval in seconds at which the OneClick reads all of the attributes of the device models that are flagged as POLLED.

  The poll to log ratio is the number of OneClick polls of a device that occur prior to logging the attributes flagged as LOGGED in the database. The default value is 0; in other words, logging is disabled by default.

  Polling and logging are the primary workload on a OneClick, and changes to polling and logging can have a significant impact on performance. To see the best performance, you should poll and log only what you need.
- Increasing the capacity of the system by increasing memory, CPU speed, or disks.
- Reducing the number of traps that are mapped to CA Spectrum events.
- Reducing the amount of data requested by customized watches and displayed attributes in order to reduce the amount of data requested from the OneClick and devices.

- Adjusting usage of features such as Live Pipes, Discovery, and automatic device configuration.
- Adjusting client interactions with the OneClick. For example, reports generated using Report Manager can exert a punctuated or prolonged performance burden on the server depending on what is being reported and how often the reports are run. Command Line Interface (CLI) scripts, manual discoveries, and other manually-initiated tasks can also impact OneClick performance.

Note: This chapter provides information on configuring polling for device and application models. For assistance in using other measures to improve OneClick performance, contact CA Support.

Polling Intervals

CA Spectrum polls devices to retrieve management information. You can change the polling interval for each device but note the following:

- If you increase the time between polls, you will use less bandwidth for management traffic. However, you will have less frequent updates of device status.
- If you decrease the time between polls, you will have more frequent updates of device status. However, you will use more bandwidth for management traffic.

Default Polling and Logging Intervals

By default, CA Spectrum polls some devices every 60 seconds, polls other devices every 300 seconds, and does not log statistics (the poll to log ratio is set to 0). In many cases, polling of this frequency is unnecessary and slows performance by creating network traffic and resulting latency.

A good rule of thumb is to poll and log critical background devices every 60 seconds and poll other, less critical network devices every 180 to 300 seconds. Often, polling and logging for end nodes, such as workstations, can be turned off to reduce network traffic and the SpectroSERVER workload.

Note: By default, the Poll_Log_Ratio attribute on device models is set to 0, which effectively disables CA Spectrum's native logging method. If you require the logging of device, attribute, and port statistics, it is recommended that you use CA Spectrum SS Logger instead of the native method, the latter of which writes the information to the Archive Manager database. CA Spectrum SS Logger is a CA Spectrum command-line application that allows you to log statistics directly to ASCII files, which reduces the load on the Archive Manager database and eliminates the need to export the data. Of equal importance, CA Spectrum SS Logger gives you greater control over what data to log and how frequently to log it. For more information about CA Spectrum SS Logger, see the *CA Spectrum SS Logger User Guide*.

Staggering Polling Intervals to Reduce SpectroSERVER Workload

You can set staggered polling intervals to reduce network management traffic, spread out the SpectroSERVER workload, and enhance the fault management capabilities at the same time. An example is shown in the following illustration:

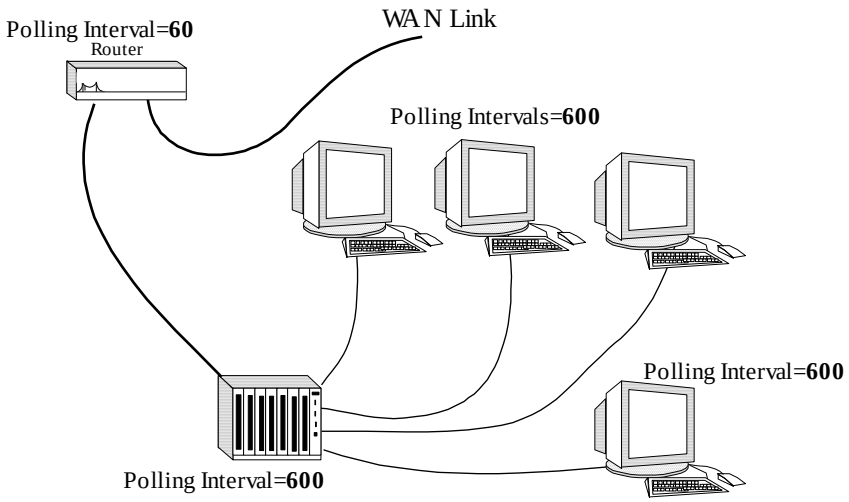

If all of the devices in the example had a default polling interval of 60 seconds, they would all use SpectroSERVER resources every 60 seconds. By setting the polling interval to 60 seconds for the router and to 600 seconds for all of the other devices, the SpectroSERVER resource utilization is reduced. However, management capabilities are not lost because if anything happened to the devices downstream from the router, polling would be interrupted, and an alarm would be generated.

About Configuring Polling for Multiple Devices

To enhance SpectroSERVER performance, you can modify the polling interval, modify the poll-to-log ratio, or disable polling altogether for multiple devices using the Attribute Editor. The Attribute Editor is an advanced OneClick utility that allows you to change one or more attribute values for multiple models at once.

Note: You can also use the Command Line Interface to change the attribute values for multiple models at once. For more information, see the *Command Line Interface User Guide*.

Set the Polling Interval and Poll-to-Log Ratio for Multiple Devices

The *polling interval* is the time interval in seconds at which the SpectroSERVER reads all of the attributes of the device model that are flagged as POLLED. The *poll to log ratio* is the number of SpectroSERVER polls of a device that occur prior to logging the attributes flagged as LOGGED in the database.

By default, the Poll_Log_Ratio attribute on device models is set to 0, which effectively disables CA Spectrum's native logging method. If you require the logging of device, attribute, and port statistics, it is recommended that you use SSLogger instead of the native method, the latter of which writes the information to the Archive Manager database. SSLogger is a CA Spectrum command-line application that allows you to log statistics directly to ASCII files, which reduces the load on the Archive Manager database and eliminates the need to export the data. Of equal importance, SSLogger gives you greater control over what data to log and how frequently to log it.

Note: For more information about SSLogger, see the *CA Spectrum SS Logger User Guide*.

To set the polling interval and poll-to-log ratio for multiple devices

1. Search for the device models that you want to modify from the Locater tab in the OneClick Console.

   Note: For information on using the Locater tab, see the *Operator Guide*.

   Results display in the Results list.

2. Select the models you want to modify, right-click, and select Utilities, Attribute Editor.

   The Attribute Editor opens.

3. Use the Attribute Editor to modify the following attributes, as desired:

   - Poll Interval
   - Poll To Log Ratio

   Note: You can find these attributes under SNMP Communication in the Attributes tree. For more information about using the Attribute Editor, see the *Modeling and Managing Your IT Infrastructure Administrator Guide*.

Disable Polling for Multiple Devices

For some device models, you might want to disable polling altogether. For example, you might want to disable polling for endpoints, such as workstations, so you do not use bandwidth for associated network polling traffic. Some administrators do not model endpoints at all because of the alarms that may occur each time the endpoints are powered down.

To disable polling for multiple devices

1. Search for the device models that you want to modify from the Locater tab.

   Note: For information about searching using the Locater tab, see the *Operator Guide*.

   Search results are displayed in the Results tab.

2. Select the models you want to modify, right-click and select Utilities, Attribute Editor.

   The Attribute Editor opens.

3. Use the Attribute Editor to set the PollingStatus attribute to no (for false) to disable polling.

   You will need to manually add the attribute to the User Defined folder in the tree.

   Note: For more information about using the Attribute Editor, see the *Modeling and Managing Your IT Infrastructure Administrator Guide*.

Configuring Polling for Multiple Applications

Application model types also have polling intervals. Some application models initially have their polling interval set to zero. To set the polling interval for these application models, use the Attribute Editor. To quickly retrieve the application models, use the All Application Models search on the Locater tab.

In general, it is recommended that you set the polling interval of application models to 60 seconds.

More information:

About Configuring Polling for Multiple Devices (see page 64)

Configuring Polling for a Single Device

You can modify the polling interval of a single device, or disable polling for the device, using the Information view of the device model.

Set the Polling Interval for a Single Device

The polling interval is the time interval in seconds at which the SpectroSERVER reads all of the attributes of the device model that are flagged as POLLED.

To set the polling interval for a single device

1. Select the device in the OneClick Console.

   The Component Detail panel displays the information for the selected model in the Information tab.

2. Expand the CA Spectrum Modeling Information subview.

3. Click set in the Poll Interval (sec) field, type the desired polling interval, and press Enter.

   The polling interval is set for this device.

Disabling Polling for a Single Device

You can disable polling for a single device.

To disable polling for a single device

1. Select the device in the Explorer tab or the Topology tab in the OneClick Console.

   The Component Details panel displays the information for the selected model in the Information tab.

2. Expand the CA Spectrum Modeling Information subview.
3. Click set in the Polling field and select Off.

   Polling is disabled for this device.

Chapter 6: Adding SpectroSERVERs

If you have not achieved desired performance levels after tuning your existing SpectroSERVERs, you will need to size the network to determine the appropriate number of SpectroSERVERs to add.

Note: For information on setting up a distributed CA Spectrum environment, see the *Distributed SpectroSERVER Administrator Guide*.

This section contains the following topics:

Sizing the Network

The CA Spectrum sizer determines the number of SpectroSERVERs you need to efficiently manage your network. The sizing tool can be run by CA Support at your request.

Before your network can be sized, you will need to run a utility called PMCount and provide the resulting data about your CA Spectrum environment (number of pollable models in a database, polling intervals, poll-to-log ratios, number of ports, and so on) to CA Support. The sizer uses this raw data to estimate the following:

- The additional amount of network management traffic generated as a result of CA Spectrum
- The number and configuration of additional SpectroSERVERs that are required to efficiently manage the number of models in your CA Spectrum environment

Both the PMCount results and the sizing results can also help you define where polling and logging can be further reduced or disabled, thereby improving performance.

Contact CA Support for detailed information about accessing and running PMCount and having a sizing performed.

Sample Sizing Report

The following is a sample CA Spectrum sizing report. The top two sections are summaries of the supplied data. The bottom two sections contain the results of the sizing.

General Sizing Information

Account Traffic Control Center **Sizing Note** Regional NOC **Percentage of SDM devices** 0% **Number of WebServers** 2 **Number of Alarms** 50 **Percentage of Active Port Monitoring** 100%	**SPECTRUM Version** 8.1rev0 **Percentage of SNMPv3 devices** 0% **Future growth** 0% **Number of Global Collection Elements** 85 **Number of Events/sec** 30

Managed devices

Quantity	Device	Poll Interval (sec)	Poll to Log Ratio	Interfaces	Note
773	Switch/Router (any switch or router)	300	Logging Off	72	
24	Pingable (no SNMP agent)	300	Logging Off		
13	Host (PC, Workstation, Server)	300	Logging Off		
229	GnSNMPDev (other)	300	Logging Off	2	

SpectroSERVER Sizing

# Systems	RAM* (Mb)	% Utilized	Platform	Processor	Speed (MHz)	OS	Other Info
1	1,183	43.8	Unix SUN	Dual UltraSPARC-III	1600	Solaris	

Estimated management traffic load due to polling and logging: 0.25 Mbits/sec.

This includes the management traffic to and from the SpectroSERVER caused by normal steady-state polling and logging. It does not include management traffic due to traps, transient events such as fault isolation, or any real time views that may be open.

Please Note

This initial sizing is based on the raw SpectroSERVER capacity required to manage the network described by the information provided to CA by the client. Actual requirements may vary depending on network architecture and the client's business requirements.

Index

U

W