CA Spectrum®

SpectroSERVER Performance Administration Guide

Release 9.4

This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the “Documentation”) is for your informational purposes only and is subject to change or withdrawal by CA at any time.

This Documentation may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This Documentation is confidential and proprietary information of CA and may not be disclosed by you or used for any purpose other than as may be permitted in (i) a separate agreement between you and CA governing your use of the CA software to which the Documentation relates; or (ii) a separate confidentiality agreement between you and CA.

Notwithstanding the foregoing, if you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all CA copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to CA that all copies and partial copies of the Documentation have been returned to CA or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS DOCUMENTATION “AS IS” WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL CA BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF CA IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice.

The manufacturer of this Documentation is CA.

Provided with “Restricted Rights.” Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

Copyright © 2014 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

CA Technologies Product References

This guide references CA Spectrum® (CA Spectrum).

Contact CA Technologies

Contact CA Support

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

- 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

Providing Feedback About Product Documentation

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

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: Overview

This section contains the following topics:

Performance and Tuning Options

In highly dynamic IT environments, monitoring the capacity of your systems and optimizing your CA Spectrum deployment cannot be a one-time task. Regular, periodic reviews are required to keep CA Spectrum operating optimally in larger environments (more than 1,000 monitored devices). For a well-developed, up-to-date summary of the performance and tuning best practices that we have derived from years of testing and supporting CA Spectrum, be sure to read the *Deployment Capacity and Optimization Best Practices Guide*, which is available on the Documentation Bookshelf.

CA Spectrum also provides self-monitoring. The OneClick client application lets network administrators monitor and troubleshoot a CA Spectrum-managed network. The Performance tab in the OneClick interface supports most network and device models. You can use this tab to analyze the CPU and memory utilization of a specific device.

For the VNM model that represents a SpectroSERVER, we also include a robust application to monitor performance. Performance View lets CA Spectrum administrators monitor the performance and system resource utilization of a SpectroSERVER. Use Performance View to identify performance problems and to determine the appropriate corrective actions to take. Your user account requires Performance Monitor privileges to deploy Performance View.

System Component Monitoring

Each computer system comprises four major components: disk, network, memory, and CPU. To ensure the successful operation of your CA Spectrum system, you can 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 about the system components and SpectroSERVER activities that can affect performance. Check the Main tab for overall system and network activity at a glance.
- A health report feature that lets you 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:

Running Health Reports (see page 45)
Overview of the User Interface (see page 14)
Connect to a Different SpectroSERVER (see page 22)

Performance Data Analysis

In many cases, the information in Performance View identifies the source of a performance problem. For instance, the CPU tab lists the 10 processes that are currently using the highest percentages of CPU time. 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 27). You can also run a health report, which includes analysis of the performance data. For more information, see Running Health Reports (see page 45).

More information:

Running Health Reports (see page 45)
Evaluating the Peformance of a SpectroSERVER (see page 27)

Performance Optimization

You can typically resolve SpectroSERVER performance problems by tuning the SpectroSERVER to improve server performance. You can also add SpectroSERVERs to distribute the network load. The following topics discuss these options.

SpectroSERVER Tuning

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

- Modifying the polling interval and the poll-to-log ratio of essential device models and application models, and disabling polling of non-essential models. Increasing the polling interval reduces the network traffic. As a result, latency, which degrades performance, can be reduced.
- 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 that is requested by customized watches and displayed attributes. Performance improves if you can reduce the amount of data that is requested from the SpectroSERVER and devices.
- Adjusting the usage of features such as Live Pipes, Discovery, and automatic device configuration.
- Adjusting client interactions with the SpectroSERVER. For example, reports that are generated using SPECTRUM Report Console can exert a punctuated or prolonged performance burden on the server. The performance impact depends on the data that is reported and on reporting frequency. Command Line Interface (CLI) scripts, manual discoveries, and other manually initiated tasks can also affect SpectroSERVER performance.

More information:

Tuning a SpectroSERVER (see page 55)

Additional SpectroSERVERs

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

- The additional network management traffic that CA Spectrum generates.
- The number and configuration of additional SpectroSERVERs that are required to efficiently manage the number of models in your 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 that owns the Control Panel.

- From a command prompt, navigate to the <*$SPECROOT*>/PView folder and enter **pview**.

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

  Note: In a distributed CA Spectrum environment, you can specify a SpectroSERVER when you start the application by entering the following command:

  `pview -vnm landscape_name`

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

More information:

Set Preferences (see page 22)

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:

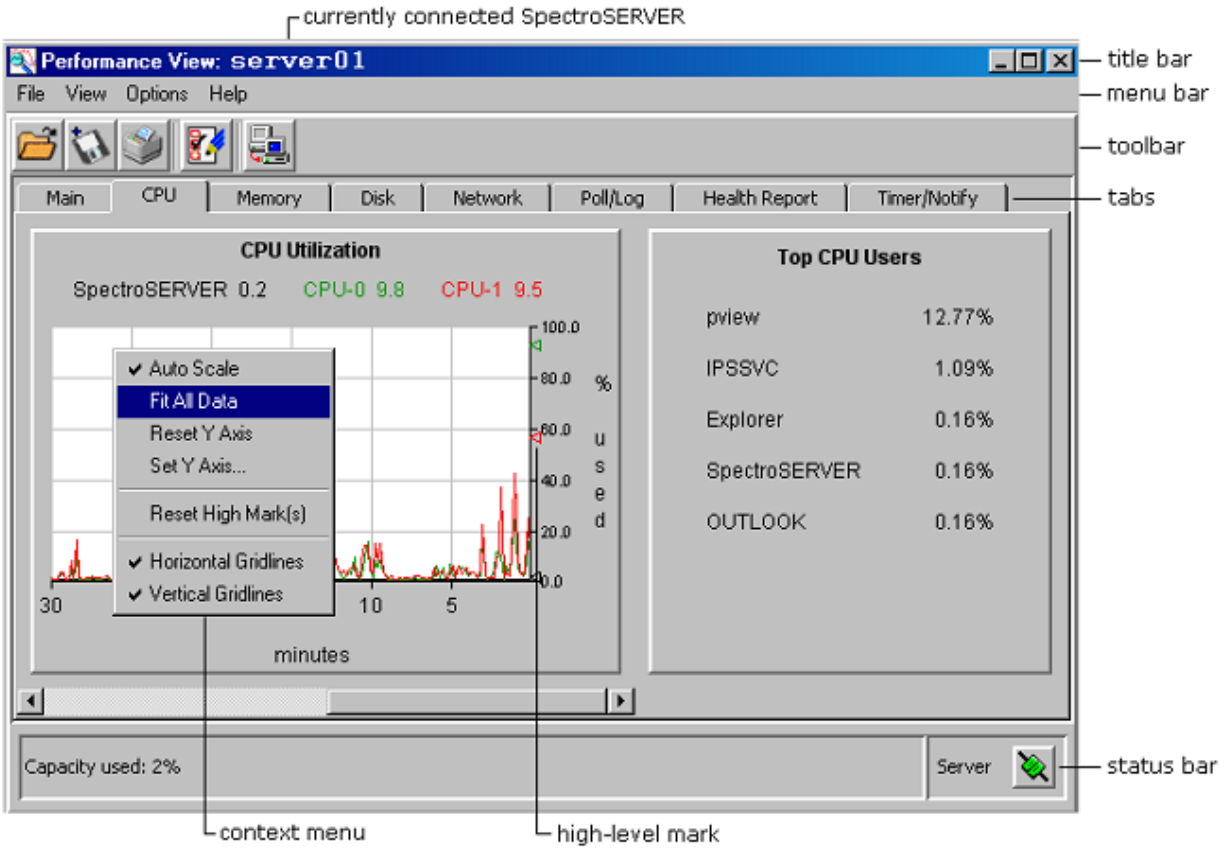

The status bar displays the current state of Performance View and, when applicable, the status of the current health report. The connection status icon is color-coded, indicating the status of the connection to the SpectroSERVER:

Green

Normal

Yellow

Using backup SpectroSERVER

Red

Contact lost

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

Bar graphs appear on the Main tab. These graphs use data that is collected in 10-minute running averages. All other attributes and graphs use data collected every 10 seconds.

Line graphs on the CPU, Memory, Disk, Network, Poll/Log, and Timer/Notify tabs display data that is collected over 60 minutes. Only the most recent 60 minutes of data is displayed.

For single-line graphs, instantaneous values are shown in a text box. 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 tooltips. High-level marks represent the highest data value that was collected since data collection began. You can reset the high-level marks and can make other changes to the axes of line graphs. Multiple CPUs and disk drives can be graphed in the same chart.

More information:

Change the Colors of Graph Lines (see page 24)
Adding SpectroSERVERs (see page 61)
Change the Display of Graph Axes (see page 25)

Main Tab

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

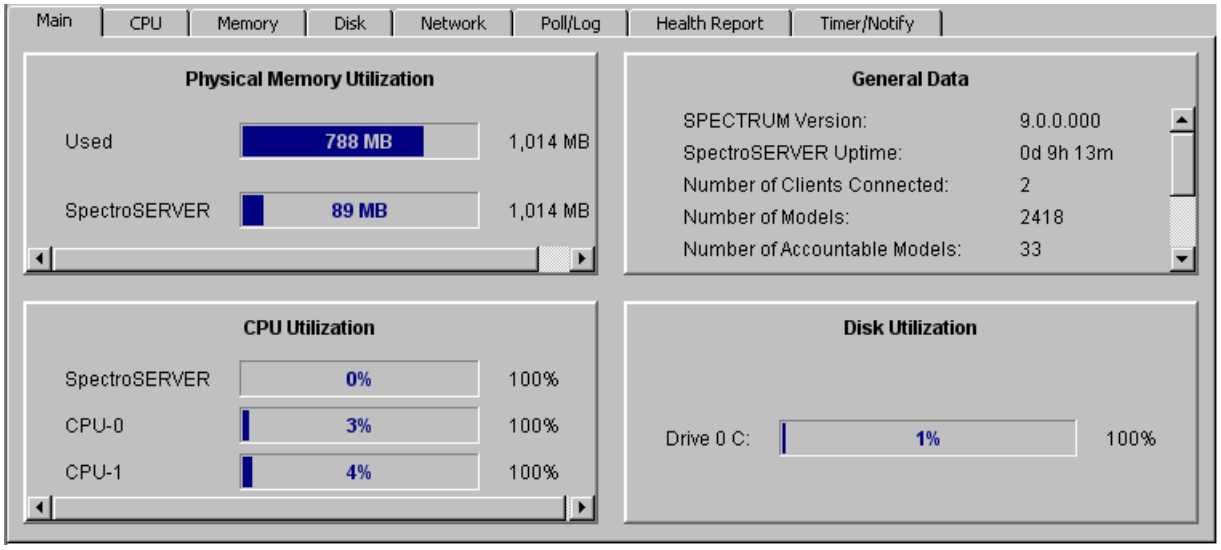

Check the Main tab for the following information:

Physical Memory Utilization

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

CPU Utilization

Displays the percentage of total CPU processing power that is used by the SpectroSERVER and other processes that are running on the server.

Disk Utilization

Displays the amount of disk read/write access capacity that is used on the server.

Check the General Data panel for the following information:

CA Spectrum Version

Specifies the version of CA Spectrum that is installed on the server.

SpectroSERVER Uptime

Specifies the amount of time that the server has been running. The format for Uptime is <days>d <hours>h <minutes>m.

Number of Clients Connected

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

Number of Models

Specifies the total number of models, including device models and other 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 that are modeled with model types derived from the Device model type.

Number of Polled Models

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

Number of Polled Attributes

Specifies the current number of attributes that the SpectroSERVER is polling.

Number of Logged Attributes

Specifies the current number of attributes that the SpectroSERVER is logging.

Note: Polled attributes are used to determine whether a device is up or down (for fault isolation). Logged attributes are used to gather statistical information. Logging of polled attributes is optional, and logged attributes are not always polled.

More information:

Physical Memory Utilization – Main Tab (see page 37)

CPU Tab

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

CPU Utilization

Displays the SpectroSERVER usage as a percentage of the total CPU capacity. Each system CPU is included in the line graph.

Top CPU Users

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

More information:

CPU Tab (see page 41)

Memory Tab

The Memory tab provides the following information:

SS Memory Utilization

Displays the amount of memory that the SpectroSERVER is using.

Paging Activity

Displays total paging activity.

Top Memory Users

Displays the top 10 memory users (from highest to lowest) and the memory that each process is using.

More information:

Paging Activity Graph - Memory Tab (see page 40)
Memory Tab (see page 41)

Disk Tab

The Disk Utilization graph on the Disk tab displays the percentage of disk read/write capacity that is being used.

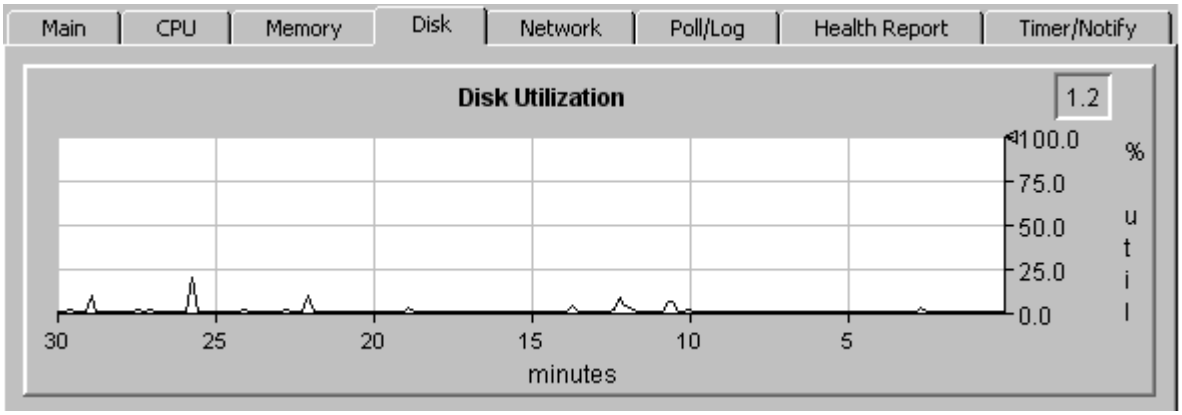

More information:

Disk Utilization Graph – Disk Tab (see page 38)
Network I/O Graph – Network Tab (see page 39)

Network Tab

The Network tab contains a Network I/O graph and a Traps Received graph. Check this tab for 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, that the VNM receives.

More information:

Network I/O Graph – Network Tab (see page 39)

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 expected to complete the actual completion time.

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 expected to complete and the actual completion time.

Log Threads in Use

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

More information:

Poll/Log Tab (see page 31)

Health Report Tab

The Health Report tab displays status information about the current health report. After the data has been collected, the report and the average percentage of SpectroSERVER capacity during the reporting period are displayed on the tab.

More information:

Running Health Reports (see page 45)

Timer/Notify Tab

The Timer/Notify tab lets you monitor the performance of timer and notification threads.

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

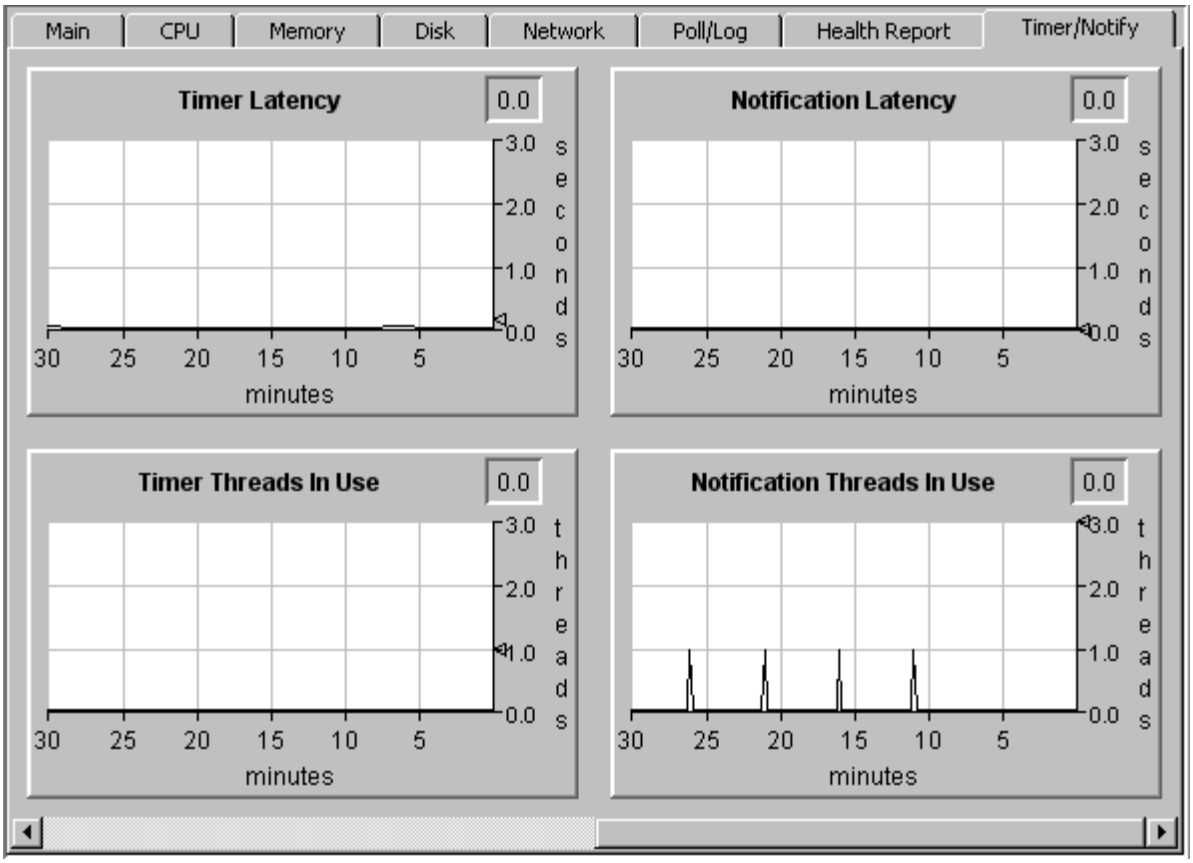

The graphs provide the following information:

Timer Latency

Displays the average timer latency in seconds. Time latency is the interval between when a scheduled timer thread is expected to complete and the actual time of completion.

Timer Threads in Use

Displays the number of timer threads in use.

Notification Latency

Displays the average notification latency in seconds. Notification latency is the interval in seconds between when a scheduled notification thread is expected to complete and the actual time of completion.

Notification Threads in Use

Displays the number of notification threads in use.

More information:

Timer/Notify Tab (see page 33)

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 (). The icon appears in the bottom right corner of the Performance View user interface.

The color of the icon indicates the status:

Green

Normal

Yellow

Using backup SpectroSERVER

Red

Contact lost

To view additional connection details for the SpectroSERVER, click the connection status icon. A dialog provides a server connection log.

Connect to a Different SpectroSERVER

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

Follow these steps:

1. Click  (Change SpectroSERVER) on the toolbar.

   The Select SpectroSERVER dialog opens.

2. (Optional) Filter the list of servers by taking one of the following steps:

   - To exclude servers whose names do not contain a specific text string, select Filter from the list. Enter the string in the text box.
   - To parse the list for servers whose names contain a specific text string, select Search from the list. Enter the string in the text box. The first server that matches the criteria is highlighted. Click Next to cycle through the servers that match.

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

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

Set Preferences

Select Preferences to customize the appearance and behavior of the Performance View user interface.

Follow these steps:

1. Click  (Set user preferences).

   The Preferences dialog opens.

2. Configure preference settings (see page 23) as desired.

3. Click OK.

   The settings that you selected are applied to Performance View.

More information:

Change the Colors of Graph Lines (see page 24)

Configure Preferences

You can configure the following preferences from the Preferences dialog:

Save Settings At Exit

Saves all of the settings that you change during a Performance View session. Saved settings are applied to subsequent sessions. The following settings are available:

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

  Note: You can show or hide tabs using the View menu. To modify the order of tabs, hide them and then display them in the desired order.

- Display settings for the status bar, the toolbar, and tooltips
- The last server that 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 that are 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 to show all pending warnings (for example, the warning when a health report has been created but not saved) before Performance View closes.

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

Sends the health report automatically to the addresses that are specified in the Email Report To field.

Email Report To

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

Title Font

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

Label Font

Specifies 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

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

All graphs use this same color palette. Changing the color for a graph line changes that color in all graphs.

Change the Colors of Graph Lines

You can customize graph settings and change the color of a line in a line graph.

Follow these steps:

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 taking one of the following steps:

   - 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.

View Menu

Use the View menu to customize the Performance View user interface. Select menu items to display or hide interface elements such as the Status Bar, tooltips, or tabs.

Change the Display of Graph Axes

You can customize the axes of individual line graphs in several ways. Right-click in a graph to see a context menu. A check mark indicates that the option is enabled. The menu contains the following display options for graphs:

Auto Scale

Sizes the Y axis value of a graph to the highest value in the collected data. For instance, if the current vertical access value is 20, and data is collected for a value of 300, the Y axis is 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 and the number of divisions (equal intervals) between the maximum and minimum values.

Reset High Marks

Returns all high-level marks in a graph to zero. The data that is collected from the current point forward determines the new high-level marks. High-level marks represent the highest data values that were collected since data collection started.

Horizontal Gridlines

Shows or hides the horizontal gridlines.

Vertical Gridlines

Shows or hides the vertical gridlines.

Chapter 3: Evaluating the Peformance of a SpectroSERVER

This section contains the following topics:

Examining Thread Latency

The topics in this section describe how to determine whether thread latency indicates CA Spectrum performance issues. The types of threads and indicators of latency are also discussed.

Threads and Thread Latency

CA Spectrum performance partially depends on timely thread allocation. A *thread* is a set of commands that perform a function or a set of functions. Each thread can run independently from other threads.

SpectroSERVER is a single-threaded application regarding the CPU, but a multi-threaded application internally. Within its own process, the SpectroSERVER creates and manages multiple threads that run simultaneously for tasks such as polling, logging, notifications, timers, and more.

Note: The Archive Manager runs in its own thread. As a result, you can use multiple CPUs on a server: one for the SpectroSERVER and another for the Archive Manager. However, deploying three or more CPUs can degrade performance.

In a multi-threaded context, while some threads are waiting (for user input, responses from devices, or data retrieval, for example), other threads can be running. A thread runs to log data, respond to traps, and to connect to SSAPI applications, for example. 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.

A SpectroSERVER maintains a pool of threads that are shared by the CA Spectrum processes that perform polling, logging, client requests, inference handler timer, inference handler notification, model activation, and model destruction. A SpectroSERVER subsystem uses threads from the pool–up to their individual limits–during periods of increased processing activity. These limits prevent any one SpectroSERVER subsystem from dominating resources and consuming all of the available threads.

When the common pool of threads is exhausted, new threads are created. The pool grows to meet the needs of the increased activity. Threads that a process no longer requires are returned to the common pool for later use. When a thread remains unused for a specified period, 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 with high thread latencies, delays occur in device polling, logging, and other tasks. As a result, delays contribute to CA Spectrum response time. For example, if a critical network device became inoperable, a delay would occur 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 can affect the performance of the SpectroSERVER. Performance View provides usage and latency statistics for these threads.

Note: By default, the Poll_Log_Ratio attribute on device models is set to 0, which effectively disables native CA Spectrum logging. To log device, attribute, and port statistics, we recommend using SSLogger instead of the native method, which writes the information to the Archive Manager database. SSLogger is a CA Spectrum command-line application that logs statistics directly to ASCII files. This type of logging reduces the load on the Archive Manager database and eliminates the need to export the data. SSLogger also provides increased control over the type and frequency of data that is logged.

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

Poll threads

Poll devices on the network. CA Spectrum uses polling to manage the operation and performance of the network. The SpectroSERVER code that manages poll threads is named the Poll Manager.

Log threads

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

Timer threads

Notify inference handlers that have registered timers, also known as wake-up calls.

Notification threads

Notify inference handlers about changes in an attribute for which the inference handlers have registered.

More information:

Notification Threads (see page 34)
Timer Threads (see page 34)

Access the Thread Information View

Access the Thread Information view to see information about thread performance and status.

Follow these steps:

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. This subview shows the threads in use, threads available, and peak value for the threads within the SpectroSERVER process. While it shows a list of important threads that are used in CA Spectrum, the list is not exhaustive. Some of the threads available in this list include the following types:

Poll Threads

Used to read the polled attributes for a model on the Polling_Interval of that model.

Log Threads

Used to read and log the logged attributes for a model on the 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.

You can use the Thread Information view to change the available value for each thread type. However, the values are typically left at their defaults. If one of the thread types is consistently running at the limit, and CPU cycles are available, increasing a limit can reduce the latency. However, if CPU utilization is already above 80%, increasing thread limits does not increase throughput. By contrast, the combination of high CPU utilization and increased thread limits can reduce throughput by increasing 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 indicates whether the allocation of polling threads and logging threads is affecting CA Spectrum performance. The tab provides four graphs to report on 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 CA Spectrum threads. The number of poll threads that the system uses at any one time is proportional to the number of models that are polled on the network and the polling rate for each model.

If the number of required poll threads exceeds the number of available threads, pending poll thread requests are queued until a poll thread becomes available. In such a case, the number of poll threads may be insufficient for the current state of the network. Contact CA Support for assistance.

Log Latency Graph

Log latency is the interval between when a scheduled logging thread is expected to complete and the actual completion time. 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 logging is set to occur every 60 seconds, the data is actually logged every 70 seconds.

Running for prolonged periods with high log latency can result in delayed logging and other serious performance problems. An average value in the Log Latency graph that is equal to or greater than 30 seconds can indicate that SpectroSERVER performance is degraded. Options for improving CA Spectrum performance include tuning the system, off-loading system demand, or upgrading system speed or capacity.

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 the system uses at any one time is proportional to the amount of data that is logged.

If the number of required threads exceeds the number of available log threads, pending log thread requests are queued until a log thread becomes available. In such a case, the number of log threads may be insufficient for the current state of the network.

By default, the Poll_Log_Ratio attribute on device models is set to 0, which effectively disables native CA Spectrum logging. To log device, attribute, and port statistics, we recommend using SSLogger instead of the native method, which writes the information to the Archive Manager database. SSLogger is a CA Spectrum command-line application that logs statistics directly to ASCII files. This type of logging reduces the load on the Archive Manager database and eliminates the need to export the data. SSLogger also provides increased control over the type and frequency of data that is logged.

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 expected to complete and the actual completion time. 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 the polling interval for a model is every 60 seconds, the model is actually polled every 70 seconds.

Running for prolonged periods with high poll latency can result in delayed device polling. 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.

An average value in the Poll Latency graph that is equal to or greater than three seconds and that is sustained for a considerable period can indicate degraded SpectroSERVER performance. 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

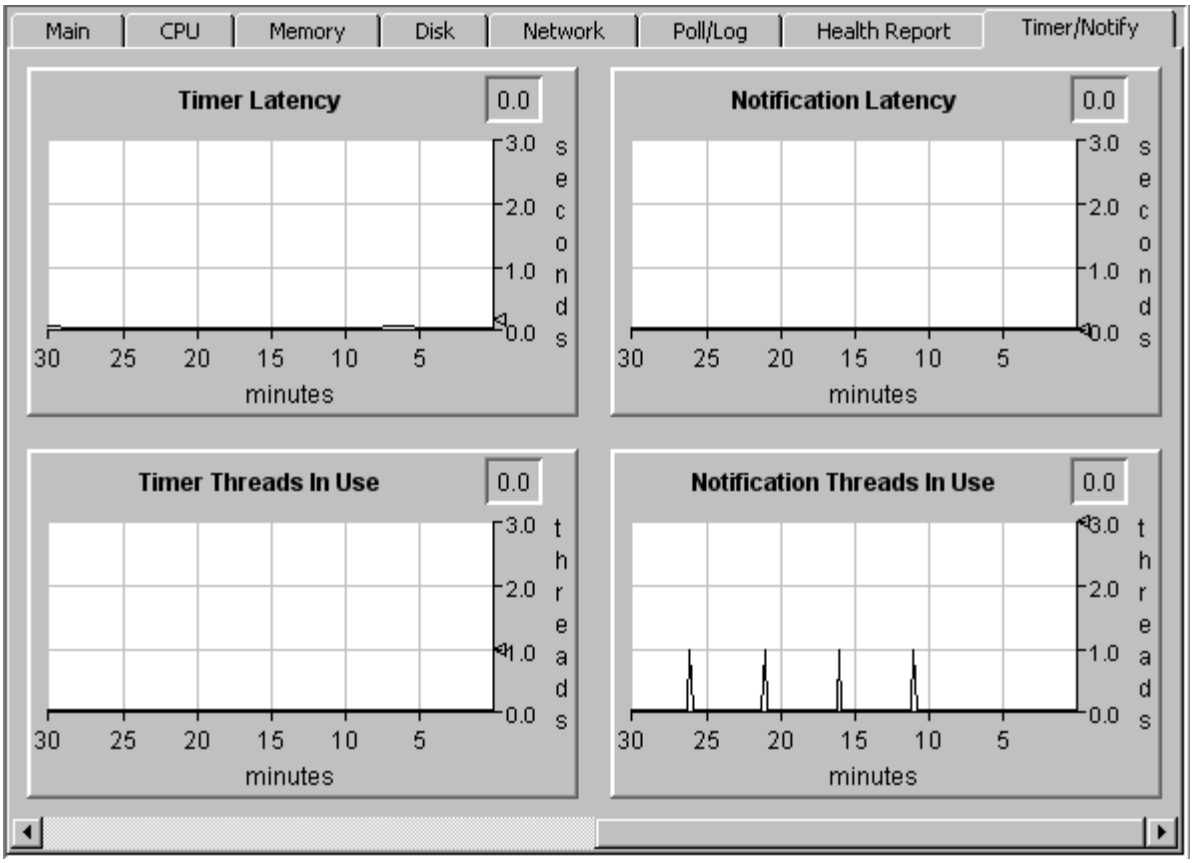

Theory of Operations

The SpectroSERVER is a poll-driven and event-driven system. Each SpectroSERVER actively polls managed elements for state changes, generates events for these 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 that are associated with a model type to define the behavior of the model type. Inference handlers execute on behalf of instantiated models of the model type. Either notification threads or timer threads can trigger these handlers.

Notification Threads

A SpectroSERVER process named the Notification Manager reads attribute changes received from polled devices and runs inference handlers to process the data. Another SpectroSERVER process, the Poll Manager, detects changes in attributes whose poll flags are set.

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

Notification threads are the mechanism that is used to notify inference handlers of a change in an attribute. These 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 that 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 (or "wake-up calls"). When an inference handler registers a “wake-up call” with the SpectroSERVER, the SpectroSERVER uses a timer thread to run the trigger method for that inference handler after a specified 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 that the primary address of a router becomes non-operational, and a secondary address must be used. An inference handler registers with the SpectroSERVER for a timer so that it can determine when the primary address is again operational. The inference handler then runs the trig_timer method.

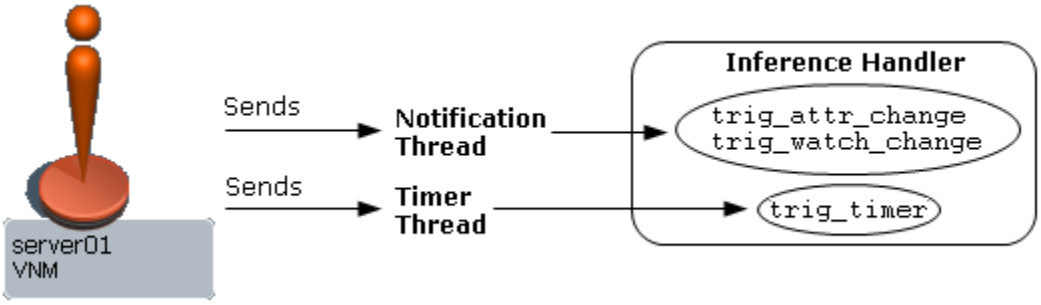

Timer Latency Graph

Timer latency is the interval between when a scheduled timer thread is expected to complete and the actual completion time. 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, the timer thread actually activates the corresponding trigger method for the inference handler every 70 seconds.

Running for prolonged periods with high timer latency can result in the delayed activation of inference handler triggers and, therefore, delayed network monitoring. A Timer Latency graph that shows more than three seconds sustained indicates performance problems. Contact 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 required notification threads exceeds the number of available threads, pending notification thread requests are queued until a notification thread becomes available. In such a case, the number of notification threads may be insufficient for the current state of the network. Contact CA Support for assistance.

Examining Memory Usage

The topics in this section describe how to determine whether swapping or paging activities are causing CA Spectrum performance problems. *Swapping* or *paging* is a processing technique that involves bidirectional transfers of data from a main storage area to an auxiliary storage area. For example, swapping occurs from memory to disk. *Pages* refer to the individual units of data transfer that are used to swap data.

Considerable swapping activity indicates a shortage of system memory. When memory resources are insufficient, data must be temporarily transferred from memory to disk to make room for various processes to run. A large amount of paging activity indicates a high amount of swapping. Because use of the disk is much slower than the use of physical memory, paging can result in performance problems. Paging should therefore be minimized.

The swap space and the physical memory collectively compose the available memory or virtual memory. Virtual memory is frequently a bottleneck for overall system performance. A system typically requires 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 in Performance View can help you determine whether insufficient memory has affected SpectroSERVER performance:

- 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

In addition, you can monitor the memory usage of the SpectroSERVER process in CA Spectrum OneClick. Events are logged and alarms are triggered when defined threshold values are exceeded. These thresholds are described in Using Performance Thresholds (see page 42).

Indicators of Memory Problems

Memory management is important for achieving and maintaining CA Spectrum performance. A 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. A system that lacks sufficient memory can appear to be disk bound because of the high paging and swapping activity occurring on the disk. When memory is the primary bottleneck, you must either increase the amount of memory or reduce the demand for memory to restore acceptable performance. Adding disk capacity or speed has a negligible effect.

Physical Memory Utilization – Main Tab

The Physical Memory Utilization area of the Main tab displays the amount of physical memory that the system is using and the amount of memory that CA Spectrum is using.

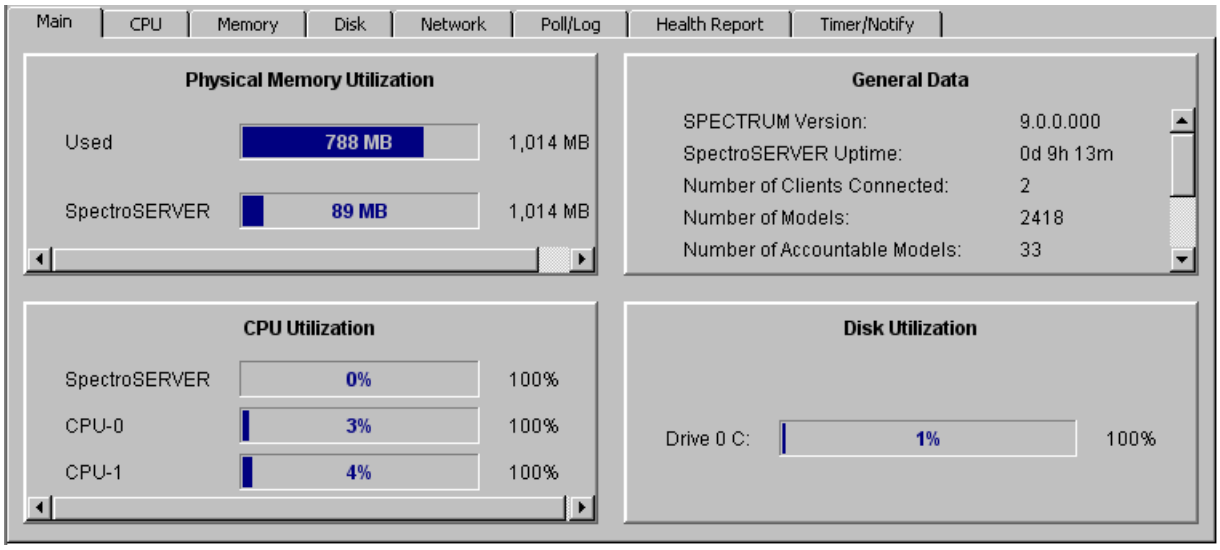

The total physical memory statistic is the actual amount of physical memory that the server to which the Performance View is connected contains. Therefore, this value is system-dependent. If SpectroSERVER uses a large percentage of the virtual memory, consider upgrading the memory or allocating more swap space.

Disk Utilization Graph – Disk Tab

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

If the Disk Utilization graph is high, data logging can 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 61)

Network I/O Graph – Network Tab

The Network I/O graph records network I/O from the Ethernet interfaces on the system. The data includes I/O activity from the SpectroSERVER. Expect to see an increase when models are created in the database or polling intervals are changed.

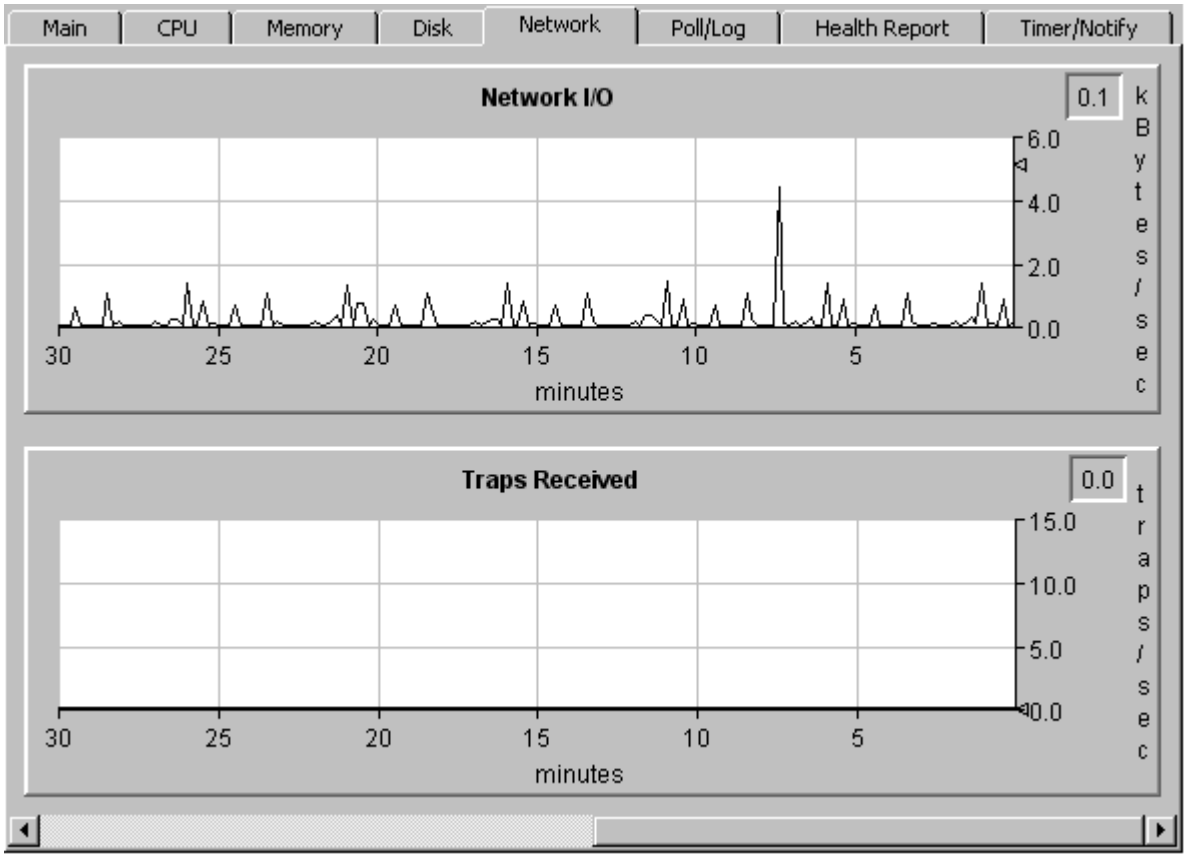

Paging Activity Graph - Memory Tab

The Paging Activity graph displays the number of system pages over time. Values in this graph that are persistently high indicate a lack of adequate physical memory. Such a situation can result from having more processes running than the available physical memory can accommodate.

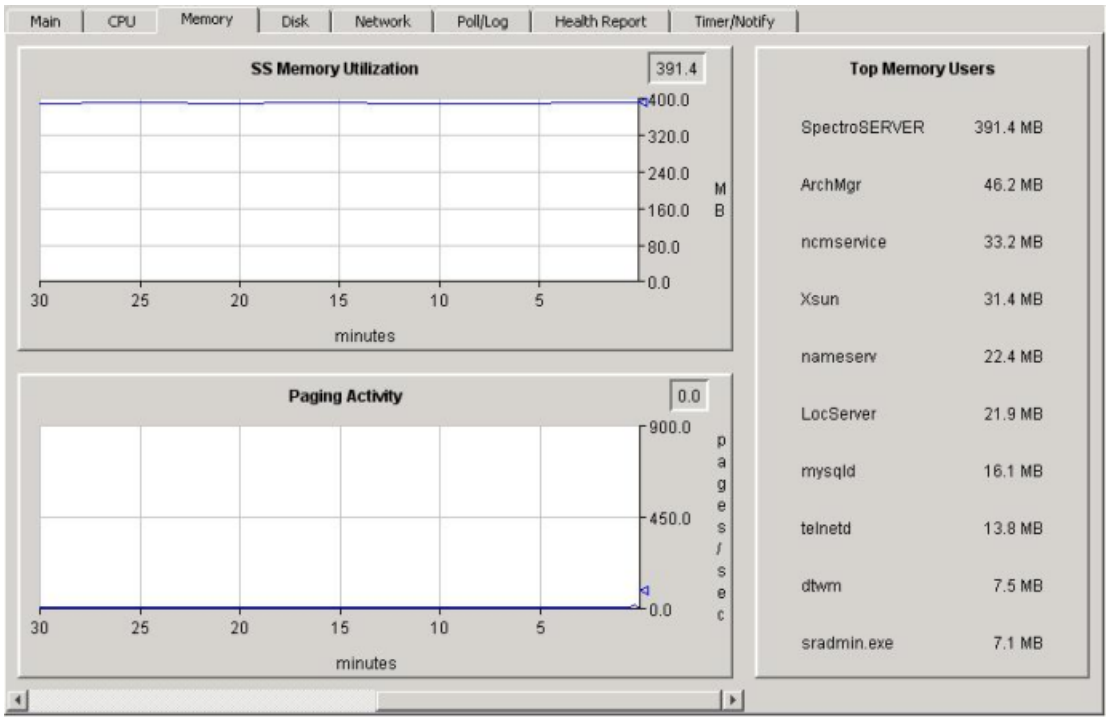

A persistently high value indicates that the system is heavily loaded. Consider a memory upgrade. If you see high paging activity levels, consider reducing the number of non-CA Spectrum processes that are running or increasing the physical memory.

Examining the Application Load

This section describes how to determine whether too many applications are running on the system. These applications include SpectroSERVER and OneClick. The following views can help you determine whether an excessive number of applications that are running on the system are causing performance issues:

- 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.

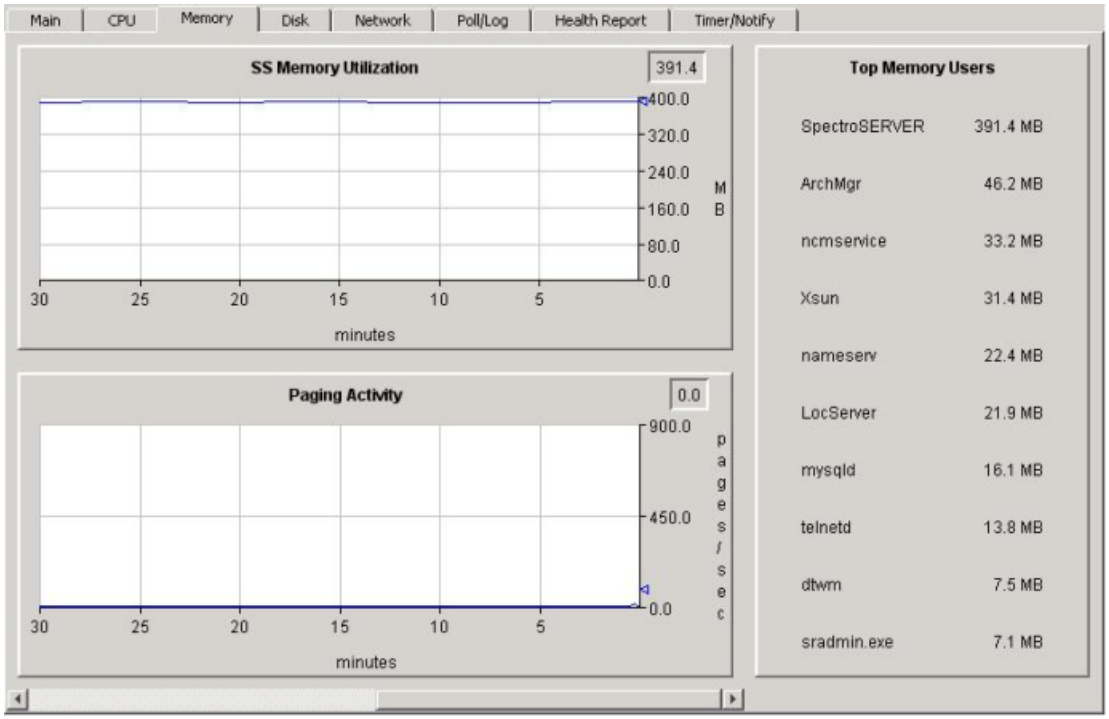

CPU Tab

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

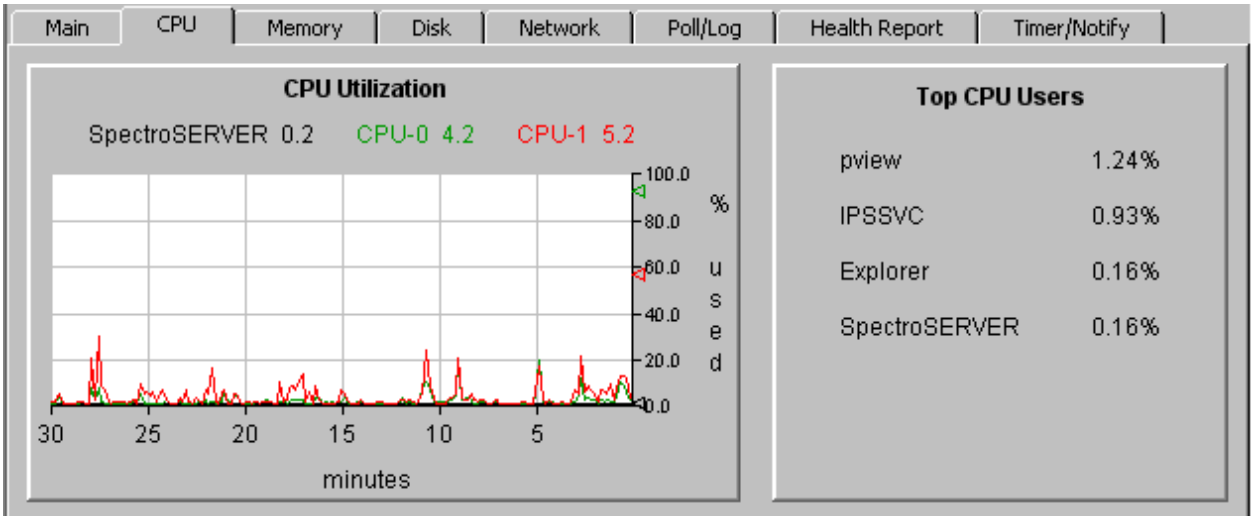

If SpectroSERVER CPU Utilization is low, but CA Spectrum performance is still slow, consider whether other applications are straining the system.

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. In the General Data panel, check 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 installing only CA Spectrum applications on the SpectroSERVER host. But 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 the performance monitoring capabilities of Performance View, CA Spectrum OneClick also provides some system performance statistics. You can also monitor the SpectroSERVER process in OneClick. If defined threshold values are exceeded, events are logged, and alarms are triggered. The following performance metrics are monitored using CA Spectrum performance thresholds:

- SNMP traps (see page 42)
- Memory usage (see page 43)

SNMP Traps

When excessive SNMP trap rates affect the SpectroSERVER process, performance can degrade. If the threshold rate is exceeded, events and alarms are generated on the SSPerformance and VNM models.

The default trap rate threshold is 100 traps per second. To change it, modify the value in the EventDisp file in the following location:

`$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 -:-", \
```

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

The trap rate is monitored by the attribute vnm_snmp_traps_ps (AttrID = 0x11eca) on the SSPerformance model.

Memory Usage

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

Separate threshold values are used for physical memory and virtual memory. The default memory size for each threshold is 2.5 GB. To change either threshold, modify the value in the EventDisp file in the following location:

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

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

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

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

Chapter 4: Running Health Reports

This section contains the following topics:

Health Reports

You can use the Performance View Health Report feature to measure and report on the health of your SpectroSERVER and the system where it is installed. Report options let you select any time period, from 6 to 24 hours.

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 native scheduling service on your host system to run the report automatically 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

We recommend selecting the full 24-hour time period. This interval collects data from a typical day. Or, if data that represents your typical workload is collected during a shorter period of time, select that time period. However, be careful not to exclude data collection for jobs that run during non-business hours, such as backups that are executed late at night.

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. 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: Use the Preferences dialog to configure Performance View to automatically send health reports to a list of email addresses.

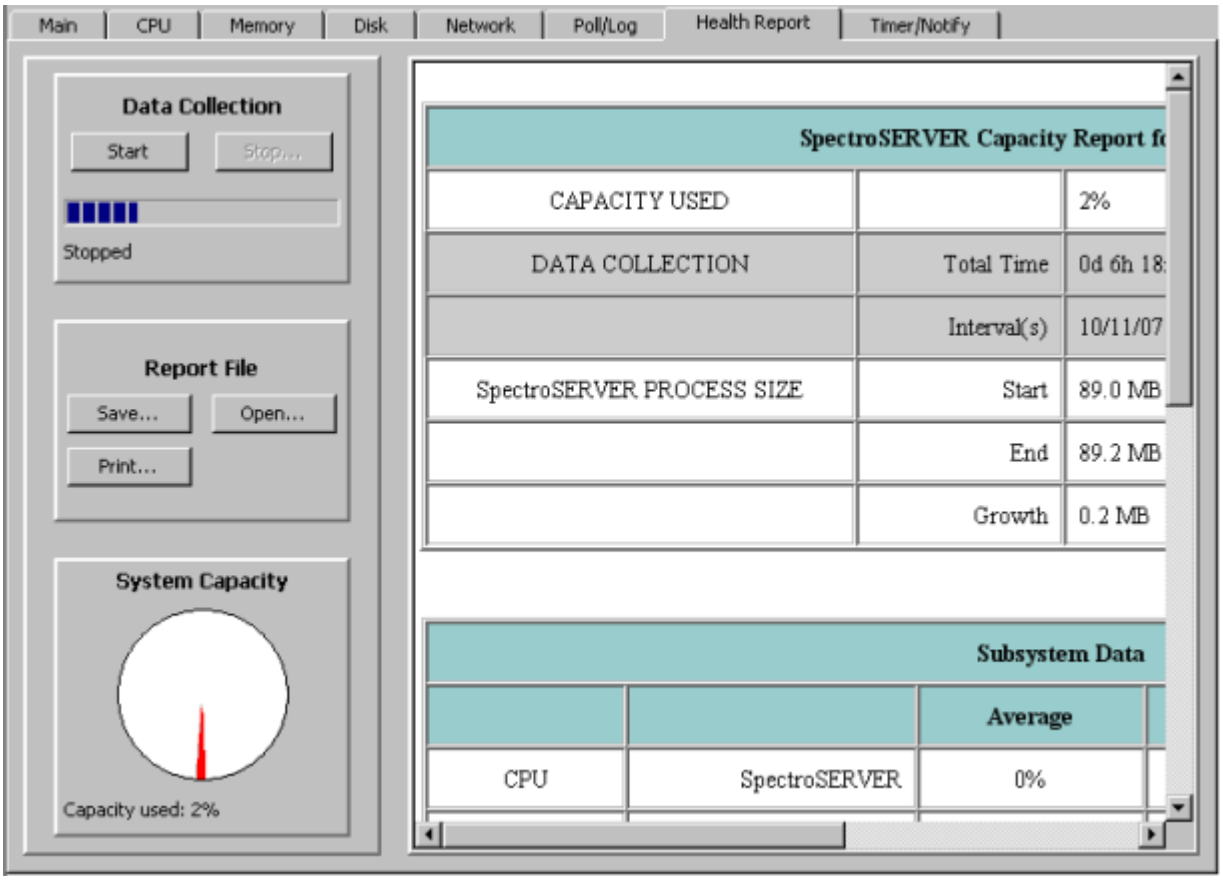

The relative health of each system resource for example, CPU usage is determined by analyzing the Average, Peak, and % Over Critical Value readings collected. Any values that exceed predefined thresholds are flagged (displayed in red) to indicate a potential performance problem.

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 enable health reports, you must start data collection for the Health Report feature.

Follow these steps:

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

   The message area in the panel indicates 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 displays a health report to the right of the Data Collection panel.

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

More information:

Set Preferences (see page 22)
Configure Preferences (see page 23)

Stop Data Collection

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

Follow these steps:

1. Click the Health Report tab.
2. Click Stop in the Data Collection section.
3. (Optional) Select *one* of the following options if you started the data collection less than 24 hours ago:

   Resume Data Collection

   Restarts data collection. For example, if you stopped data collection after one hour, select this option to restart data collection and continue it for 23 more hours (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, to generate a report, you must collect data for at least six hours.

Stop and Analyze Data

Immediately generates a health report from the collected data. This option appears if data has been collected for at least 6 hours. The report remains 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

Ends the data collection process and deletes all collected data. No report is generated.

Save Health Reports

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

Follow these steps:

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, do not include the .htm file extension. It is added automatically.

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

   The health report is saved.

Open Health Reports

Open health reports to view or print them.

Follow these steps:

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, read-only window.

Print Health Reports

You can print the health report that is currently displayed on the Health Report tab.

Follow these steps:

1. Click  (Print the health report).

   The Print dialog opens.

2. Select settings in the Print dialog, and click OK.

   The health report prints.

You can also print a health report that you have saved.

Follow these steps:

1. Open the report (see page 48).

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

2. Click Print.

   The Print dialog opens.

3. Select settings, and click OK.

   The health report prints.

4. Click Close.

   The SpectroSERVER Capacity Report dialog closes.

More information:

Open Health Reports (see page 48)

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 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 time period 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. To specify multiple addresses on Windows systems, enclose the list in quotation marks, for example, “address1,address2,address3”.

Important! 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 fails to send the email notification. It looks for a registry entry under HKEY_CURRENT_USER\\Software\\Microsoft\\Windows NT\\CurrentVersion\\Windows Messaging Subsystem\\Profiles and creates an application event if the entry cannot be found. The local email system can require confirmation steps before sending the email.

Reports that pviewrep generates are saved automatically with the name of the SpectroSERVER host. These files have an .htm extension. Sequential numbers are added to keep subsequent reports from overwriting existing ones. For example, the first report that is generated for a server named “ace” is ace.htm. The second report is ace_1.htm.

By default, reports that are generated with pviewrep are saved to the directory that was used for the most recent health report 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 standard output stream for the program.

Note: Health reports that are distributed automatically by email are in plain text format, not HTML.

Interpreting Health Reports

A health report includes three major sections, which are described in the following topics:

- SpectroSERVER Capacity
- Subsystem Data
- Analysis

SpectroSERVER Capacity

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

Capacity Used

Reports 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

Reports the total amount of time that data was collected. Also reports the start and end times for the individual data collection intervals that make up the total time.

SpectroSERVER Process Size

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

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 the Performance View threshold value for that parameter is flagged (displayed in red). Flagged parameters indicate possible performance problems. These parameters contribute to the message that appears 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 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 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 includes a narrative that describes the results of the analysis that was performed on the collected data. The description that you see depends on whether parameter values exceeded predetermined thresholds and which parameters indicated problems.

Report results fall into three possible categories:

- No parameters are flagged to indicate problems
- 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:

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 overutilized. 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.

A list of problems that can create the conditions that were flagged then follows.

Other narratives are also displayed for various threshold conditions. In some cases, you are advised to contact CA 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, tune the OneClick to improve performance by taking the following steps:

- Modifying the polling interval and poll-to-log ratio of essential device models and application models. Disabling polling of non-essential models. These changes reduce the network traffic and the resulting latency that affect performance.
  - **Polling interval**: The time interval in seconds at which the OneClick reads all device model attributes that are flagged as POLLED.
  - **Poll-to-log ratio**: The number of OneClick device polls that occur prior to logging the attributes that are flagged as LOGGED. The default value is 0 (logging is disabled).

  Polling and logging create the primary workload for OneClick. Changes to polling and logging can have a significant impact on performance. To see the best performance, poll and log only required data.
- Increasing the capacity of the system by increasing memory, CPU speed, or disk space.
- Reducing the number of traps that are mapped to CA Spectrum events.
- Reducing the amount of data that is requested by customized watches and displayed attributes. As a result, less data is requested from the OneClick and devices.

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

Note: This chapter provides information about configuring polling for device and application models. For other suggested 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 as a tuning measure. However, note the following guidelines:

- If you increase the time between polls, less bandwidth is required for management traffic. However, device status is updated less frequently.
- If you decrease the time between polls, device status is updated more frequently. However, more bandwidth is required 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 to poll other, less critical network devices every 180 to 300 seconds. Often, polling and logging for end nodes, such as workstations, can be disabled 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 native CA Spectrum logging. To log device, attribute, and port statistics, we recommend using SSLogger instead of the native method, which writes the information to the Archive Manager database. SSLogger is a CA Spectrum command-line application that logs statistics directly to ASCII files. This type of logging reduces the load on the Archive Manager database and eliminates the need to export the data. SSLogger also provides increased control over the type and frequency of data that is logged.

For more information about SSLogger, 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 fault management. An example is shown in the following illustration:

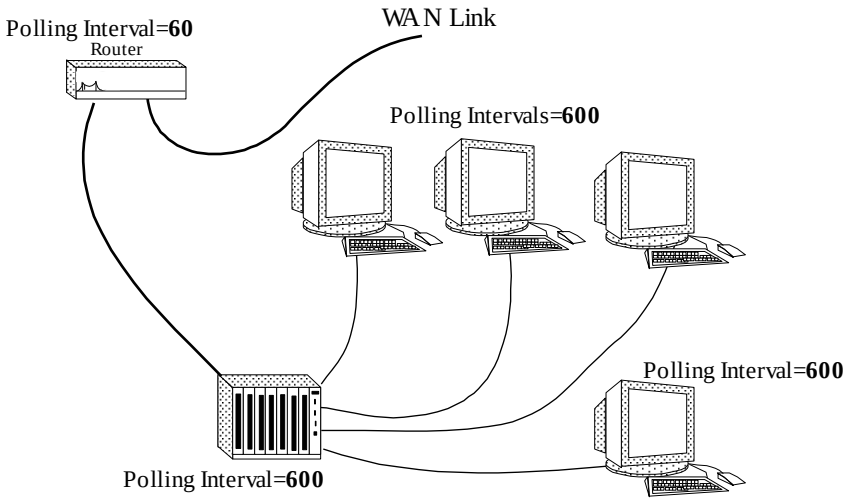

If all of the devices in the example had a default polling interval of 60 seconds, they would all use SpectroSERVER resources every 60 seconds. SpectroSERVER resource utilization is reduced by setting the polling interval to 60 seconds for the router and to 600 seconds for all of the other devices. However, the staggered polling does not impede management capabilities. If any fault occurred on the devices that are 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

Polling and logging create the primary workload for OneClick. Configure the polling interval and the poll-to-log ratio to improve SpectroSERVER performance.

The *polling interval* is the time interval in seconds at which the SpectroSERVER reads all device model attributes that are flagged as POLLED. The *poll to log ratio* is the number of SpectroSERVER device polls that occur prior to logging the attributes that are flagged as LOGGED in the database. The default value is 0 (logging is disabled).

Note: By default, the Poll_Log_Ratio attribute on device models is set to 0, which effectively disables native CA Spectrum logging. To log device, attribute, and port statistics, we recommend using SSLogger instead of the native method, which writes the information to the Archive Manager database. SSLogger is a CA Spectrum command-line application that logs statistics directly to ASCII files. This type of logging reduces the load on the Archive Manager database and eliminates the need to export the data. SSLogger also provides increased control over the type and frequency of data that is logged.

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

Follow these steps:

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

   Note: For more information, see the *Operator Guide*.

   Results are displayed.

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

   The Attribute Editor opens.

3. Use the Attribute Editor to modify the following attributes:

   - Poll Interval
   - Poll to Log Ratio

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

Disable Polling for Multiple Devices

You can disable polling for some device models. For example, you want to disable polling for endpoints, such as workstations, to avoid using bandwidth for network polling traffic. Some administrators do not model endpoints at all because of the alarms that can occur each time the endpoints are powered down.

Follow these steps:

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 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 must now manually add the attribute to the User Defined folder in the tree.

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

Configuring Polling for Multiple Applications

Application model types have different default polling intervals. The default polling interval for some application models is set to zero. To set the polling interval for these application models, use the Attribute Editor. You can quickly retrieve the application models by performing a search of All Application Models on the Locater tab.

In general, we recommend setting the polling interval of application models to 60 seconds.

More information:

About Configuring Polling for Multiple Devices (see page 57)

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. You can set polling intervals for individual devices.

Follow these steps:

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.

Follow these steps:

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 SpectroSERVER, size the network to determine the appropriate number of SpectroSERVERs to add.

Note: For more information, see the *Distributed SpectroSERVER Administrator Guide*.

This section contains the following topics:

Sizing the Network

The CA Spectrum sizing tool determines the number of SpectroSERVERs that are required to efficiently manage your network. The sizing tool can be run by CA Support at your request.

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

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

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

CA Support can help you size your deployment. You can also contact CA Support for detailed information about accessing and running PMCount.

Index

U

W