



Monitoring with the SL1 Agent

SL1 version 8.12.0

Table of Contents

Introduction to the SL1 Agent	4
What is an SL1 Agent?	5
What Kind of Data Can an Agent Collect?	5
Supported Operating Systems	6
Agent Architecture	6
Installing an SL1 Agent	8
Installing the Linux Agent	9
Installing the Windows Agent	10
Upgrading an Agent	10
Stopping an Agent	11
Uninstalling an Agent	11
Uninstalling a Linux Agent	11
Uninstalling a Windows Agent	12
Using Agents with SELinux	12
Viewing the Discovered Device	12
Device Classes for Agent-Only Devices	12
Configuring an SL1 Agent	14
Using the Agent Investigator	15
The Config Tab	15
The Polled Data Tab	16
Creating a New Source for Polled Data	17
Configuring Nagios Plug-ins as Polled Data Sources	18
Editing or Deleting Polled Data Sources	19
The Log Sources Tab	19
Creating a New Log Source	20
Editing or Deleting Log Sources	21
The Watched Files Tab	21
Adding Watched Files	21
Editing or Deleting Watched Files	22
Changing the Target Message Collector for the Agent	22
Monitoring Ports Using an Agent	24
What is a Port?	25
Port Security	25
Port Availability	25
Viewing a List of All Open Ports on All Devices	26
Viewing a List of All Open Ports on a Single Device	27
System Settings for Monitoring Port Availability	29
Viewing the TCP/IP Port Monitoring Policies	30
Defining a Monitoring Policy for Port Availability	31
Example Policy for TCP/IP Port Availability	33
Editing a Monitoring Policy for a TCP/IP Port	33
Executing a TCP-IP Port Monitoring Policy	34
Deleting a TCP/IP Port Monitoring Policy	35
Monitoring Processes Using an Agent	36
What is a Process?	37
Viewing the List of Device Processes	37
Generating a Report on Multiple System Processes	39
Generating an Exclusion Report for a Single System Process	41
Viewing the System Process Monitoring Policies	43
Defining a Monitoring Policy for a System Process	44


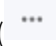
Example System Process Monitoring Policy	47
Editing a System Process Monitoring Policy	48
Executing a System Process Monitoring Policy	49
Deleting a System Process Monitoring Policy	49
Monitoring Logs Using an Agent	51
Using a Log File Monitoring Policy	52
Viewing the List of Log File Monitoring Policies	52
Filtering the List of Log File Monitoring Policies	53
Creating a Log File Monitoring Policy	54
Editing a Log File Monitoring Policy	56
Deleting Log File Monitoring Policies	56
Viewing the List of Log File Monitoring Policies and Aligned Devices	57
Filtering the List of Log File Monitoring Policies and Aligned Devices	58
Aligning a Log File Monitoring Policy to Devices	58
Unaligning Log File Monitoring Policies from Devices	61
Creating an Event Policy for Agent Logs	61
Monitoring Vitals Using an Agent	65
Viewing System Availability Reports for a Device	66
Changing the Method for Measuring Device Availability	69
Viewing CPU and Memory Utilization for a Device	70
Viewing CPU Utilization	70
Viewing Memory Utilization	72
Changing the Dynamic Application Precedence Settings for CPU and Memory Utilization	74
Troubleshooting SL1 Agents	77
Determine if the Agent Process is Running	79
Determine if the Agent Configuration is Valid	80
Determine if the Agent is Able to Upload Data	80
Check the Agent Upload Directory	80
Run the Agent in Debug Mode (Linux)	81
Determine if SL1 is Receiving Agent Data	82
Determine if SL1 Cannot Process Agent Data	82
Determine if the Number of Processes is Inconsistent with Other Applications	83
Troubleshooting Examples	83
Example /var/log/insight/streamer.log for successful discovery	83
Example /var/log/uwsgi/streamer.log for successful discovery in streamer_prime	84
Save incoming data for a specific device ID (streamer_prime)	84
Save incoming data for a specific device ID (Converged Platform or SL1)	84
Additional Troubleshooting Situations and Best Practices	84

Introduction to the SL1 Agent

Overview

This chapter describes SL1 agents and provides instructions for viewing device and interface data collected by agents.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon ().
- To view a page containing all of the menu options, click the Advanced menu icon (.

This chapter includes the following topics:

<i>What is an SL1 Agent?</i>	5
<i>What Kind of Data Can an Agent Collect?</i>	5
<i>Supported Operating Systems</i>	6
<i>Agent Architecture</i>	6

What is an SL1 Agent?

An **SL1 agent** is a program that runs on a device or element monitored by SL1. An agent collects data from the device, interface, or other element and pushes that data back to SL1. You can install and use multiple agents, as needed.

Because an agent is always running on a device, an agent can collect more granular data than can be collected by polling the device periodically. You can monitor devices using agents or by SL1 polling the device, or you can use both methods.

What Kind of Data Can an Agent Collect?

An SL1 agent collects the following data:

- **Device Availability.** SL1 can determine the availability state of a device (available or unavailable) and generate trended availability graphs based on uptime data collected by the agent.
- **Host Performance Metrics.** Using a Dynamic Application, SL1 translates data provided by an SL1 agent to trend the following metrics:
 - Overall CPU Utilization
 - Per-Processor CPU Utilization
 - Disk Average Queue Length
 - Disk Utilization
 - Memory Utilization
 - Network Bytes Read
 - Network Bytes Written

You can view these metrics on the **Device Investigator** page and the **[Performance]** tab of the **Device Reports** panel for a specific device.

- **Host Configuration.** Using a Dynamic Application, SL1 collects the following configuration data based on data provided by the agent:
 - The number and speed of the installed CPUs
 - The overall and per-disk storage size
 - The amount of installed memory

You can view the collected configuration data on the **[Configs]** tab of the **Device Investigator** page and the **Device Reports** panel.

- **System Processes.** The agent collects a list of all processes running on the device. You can view the list of processes on the **[Processes]** tab of the **Device Reports** panel and the **[Processes]** tab of the **Device Investigator** page. Monitoring policies can be configured to trend and alert on process availability, process CPU usage, and process memory usage.

- **Open Ports.** The agent collects a list of open TCP and UDP ports on the device. You can view the list of open ports on the **[TCP/UDP Ports]** tab of the **Device Reports** panel and the **[Ports]** tab of the **Device Investigator** page . Monitoring policies can be configured to trend and alert on port availability.
- **Logs.** The agent can be configured to push logs that match specific criteria from a log file or the Windows Event Log to SL1 . You can view logs collected by the agent on the **Device Investigator** page and the **Device Logs** page for a device and can be configured to trigger events.

Supported Operating Systems

You can install agents on the following operating systems:

- Debian 8 or later
- Ubuntu 14.04.5 or later
- Red Hat 6.10 or later
- CentOS 6.10 or later
- Oracle Linux 6.10 or later
- Windows Server 2016, Windows Server 2016 Core
- Windows Server 2012 R2, Windows Server 2012
- Windows Server 2008 R2
- Windows 10
- Windows 8.1
- Windows 8
- Windows 7
- BusyBox Linux (container guests only)
- Alpine Linux (container guests only)

NOTE: The agent runs on 64-bit Windows and Linux operating systems only.

Agent Architecture

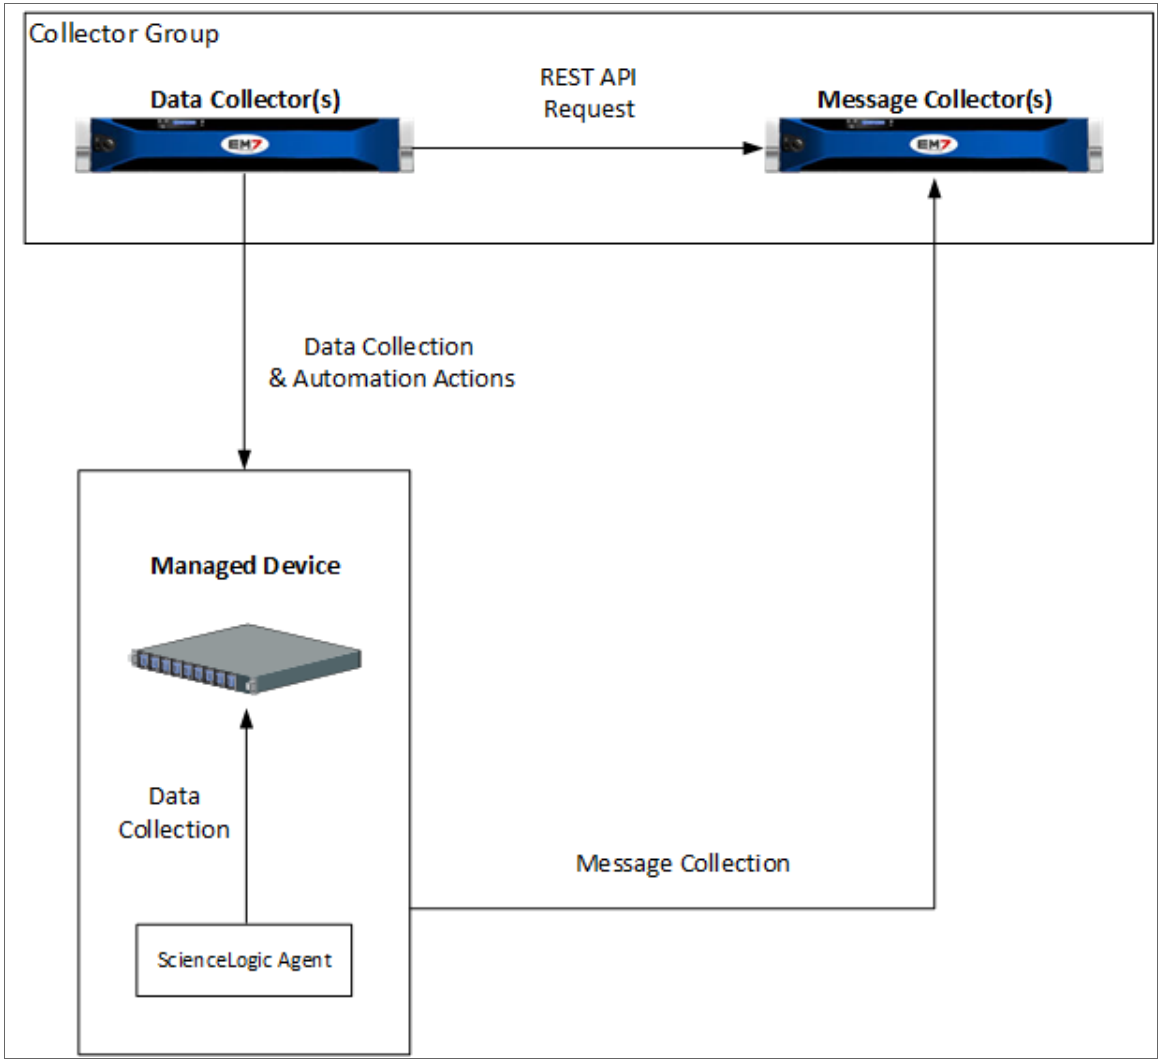
An SL1 agent collects data from the device on which it is installed and transfers that data to a Message Collector in an SL1 system using the HTTPS protocol. In a distributed system, the Data Collector on which the Dynamic Applications and collection processes run then poll the Message Collector using the HTTPS protocol to transfer data to SL1 .

TCP port 443 must be open between the device on which an agent is installed and the Message Collector.

An SL1 agent requires a Message Collector for a distributed architecture. The Message Collector does not need to be dedicated to the agent.

In a distributed architecture, an agent collects data from the device on which it is installed, and then sends messages to a Message Collector.

The diagram below shows the collection layer of a distributed system containing both Data Collectors and Message Collectors in which an agent is installed on a managed device.



Installing an SL1 Agent

Overview

This chapter describes how to install, upgrade, and uninstall SL1 agents for Windows and Linux operating systems.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon (☰).
- To view a page containing all of the menu options, click the Advanced menu icon (⋮).

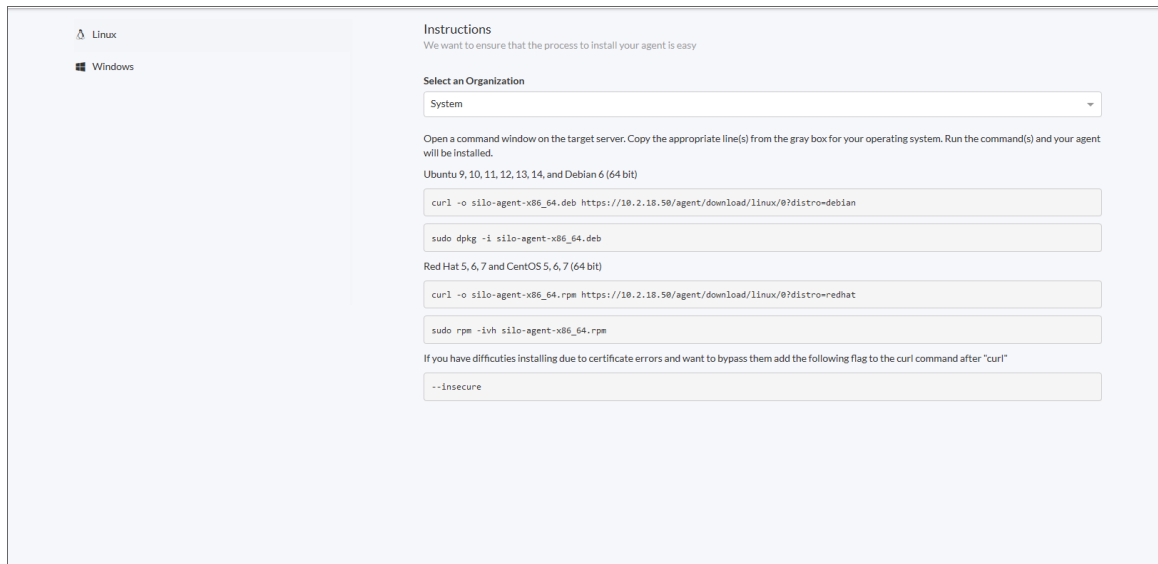
This chapter includes the following topics:

<i>Installing the Linux Agent</i>	9
<i>Installing the Windows Agent</i>	10
<i>Upgrading an Agent</i>	10
<i>Stopping an Agent</i>	11
<i>Uninstalling an Agent</i>	11
<i>Uninstalling a Linux Agent</i>	11
<i>Uninstalling a Windows Agent</i>	12
<i>Using Agents with SELinux</i>	12
<i>Viewing the Discovered Device</i>	12
<i>Device Classes for Agent-Only Devices</i>	12

Installing the Linux Agent

To download and install a Linux agent:

1. On the **Agents** page (Devices > Agents), click the **[New Agent]** button. The **Agent Installation** page appears:

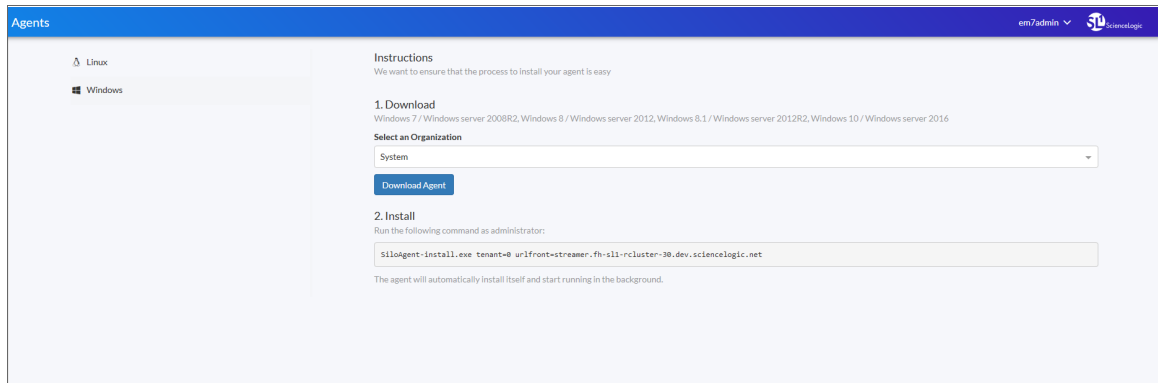


2. Click the **[Linux]** tab.
3. From the **Select an Organization** drop-down list, select an organization for the new agent.
4. Follow the additional installation instructions on the **[Linux]** tab, which includes copying the relevant commands for your operating system. After installation, the agent starts running in the background.

Installing the Windows Agent

To download and install a Windows agent:

1. On the **Agents** page (Devices > Agents), click the **[New Agent]** button. The **Agent Installation** page appears.
2. Click the **Windows** tab:



3. From the **Select an Organization** drop-down list, select an organization for the new agent.
4. Click the **[Download Agent]** button.
5. Save the SiloAgent-install.exe file for installing the agent.
6. To install the agent, run the following command on your SL1 system as an Administrator:

```
SiloAgent-install.exe tenant=0 urlfront=<URL_for_your_SL1_system>
```
7. To verify that the installation was successful, open the Windows Task Manager or enter the TASKLIST command to view running processes. The SiloAgent process will be running on the Windows machine.

Upgrading an Agent

When you have the latest version of an agent, a check mark icon (✓) appears in the **Newest Version** column for that agent. To upgrade to the latest version of an agent:


1. On the **Agents** page (Devices > Agents), locate the agent you want to upgrade.
2. Click the **[Upgrade]** button. The agent starts the upgrade process.

Stopping an Agent

You can delete any agent that you have installed, as well as any other agent on the **Agents** page. When you use the Delete option, SL1 deletes the data gathered by that agent, and that agent no longer appears on the **Agents** page.

NOTE: Using the delete option for an agent stops the agent on the device, but does *not* actually remove the agent from the device. As a best practice, use the delete process to delete the data gathered by the agent (the uninstallation process does not delete this data), and then uninstall that agent, if needed. For uninstallation details, see [Uninstalling an Agent](#).

To delete an agent and its data:

1. On the **Agents** page (Devices > Agents), locate the agent you want to delete.
2. Click the **[Actions]** button () for that agent and select *Delete*. SL1 stops the agent from collecting data.

Uninstalling an Agent

When you uninstall an agent, you remove that agent completely from the new user interface, but you do not lose the data collected by that agent.

Uninstalling a Linux Agent

To uninstall an agent on a Linux system:

1. Log in to the Linux system via the console or SSH as a user that has sudo administrator permissions.
2. Do one of the following:
 - For Red Hat, CentOS, and Oracle, execute the following command:

```
rpm -e scilogd
```
 - For Debian and Ubuntu, execute the following command:

```
dpkg -r scilogd
```
3. Optionally, you can remove the agent configuration directory from the Linux system. The configuration directory can be found at:

```
/etc/scilog (rm -rf /etc/scilog)
```

Uninstalling a Windows Agent

To uninstall an agent on a Windows system:

1. On the Windows system, open the **Control Panel**.
2. Go to the **Programs and Features** page (Control Panel > Programs > Uninstall a program).
3. Select the SiloAgent program from the list, and then click **[Uninstall]**.
4. When the uninstallation process is complete, remove the agent configuration directory from the Windows system. The configuration directory can be found at:

```
Program Files\ScienceLogic\SiloAgent\conf
```

Using Agents with SELinux

When an agent starts, it checks to see if SELinux is in *enforcing* mode, which means SELinux is running and enforcing SELinux policy. If SELinux is in *enforcing* mode, the agent stops with a warning message.

In this situation, you can either disable SELinux or put SELinux into *permissive* mode.

1. Navigate to `/etc/sysconfig/selinux` and choose one of the following options:
 - If you want to use the "brute force" approach to making an agent work with SELinux, change the SELINUX option to `disabled`.
 - If you want SELinux to run and generate logs, but not control permissions, change the SELINUX option to `permissive`.
2. Reboot the server to start the agent again.

Viewing the Discovered Device

If the installation is successful and the agent can communicate with the specified Message Collector over TCP port 443, one of the following automatically happens:

- If the primary IP address of the device is not currently monitored by SL1, then SL1 creates a device record for the device and populates the device record with data provided by the agent. The device record is assigned a device class based on data reported by the agent.
- If the primary IP address of the device is currently monitored by SL1, the device record for the existing device is updated with data provided by the agent.

Device Classes for Agent-Only Devices

During initial discovery, the agent returns operating system type and version information to SL1.

Based on this information, SL1 assigns one of the following device classes to a device monitored only by an agent:

- Microsoft Windows Workstation
- Microsoft Windows Cluster Point
- Microsoft Windows Server 2008 R2
- Microsoft Windows Server 2012
- Microsoft Windows Server 2012 Domain Controller
- Microsoft Windows Server 2008 R2 Domain Controller
- Microsoft Windows 8.1 Workstation
- Microsoft Windows 8 Workstation
- Microsoft Windows Server 2012 R2
- Microsoft Windows 7 Workstation
- Microsoft Windows Server 2012 R2 Domain Controller
- Microsoft Windows 10 Workstation
- Linux Ubuntu 16.04
- Linux Ubuntu 14.04
- Linux Ubuntu 12.04
- Linux Debian 8
- Linux Debian 7
- Linux Debian 6
- Linux Red Hat Enterprise Linux 7
- Linux Red Hat Enterprise Linux 6
- Linux Red Hat Enterprise Linux 5
- Linux Oracle Linux 7
- Linux Oracle Linux 6
- Linux Oracle Linux 5
- Linux CentOS 7
- Linux CentOS 6


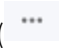
NOTE: If a device is monitored by an agent and via SNMP, the device class assigned by SNMP discovery will take precedence.

Configuring an SL1 Agent

Overview

This chapter describes how to configure the settings on the Message Collector with which the agent communicates. This chapter also covers how to use the **Agent Investigator** page on the **Agents** page (Devices > Agents), which provides access to all of the data associated with an agent, and on that page you can configure the agent settings.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon ().
- To view a page containing all of the menu options, click the Advanced menu icon ().

This chapter includes the following topics:

Using the Agent Investigator	15
<i>The Config Tab</i>	15
<i>The Polled Data Tab</i>	16
<i>The Log Sources Tab</i>	19
<i>The Watched Files Tab</i>	21
Changing the Target Message Collector for the Agent	22

Using the Agent Investigator

The **Agent Investigator** page appears when you click the name of an agent on the **Agents** tab. The **Agent Investigator** page provides access to all of the data associated with an agent, using the following tabs:

- **Config**. Displays the agent name, install agent, and aligned device. On this tab, you can configure the disk space, excludes, includes, and other metrics related to this agent.
- **Polled Data**. Displays the scripts that you execute to gather data by this agent over time. On this tab, you can configure new sources for polled data, including scripts, URLs, JMX data, Windows performance counters, and events.
- **Log Sources**. Displays the log files you are monitoring for the agent.
- **Watched Files**. Displays any watched files you have defined with regular expressions for the agent.

The Config Tab

The **[Config]** tab of the **Agent Investigator** page displays the agent name, agent nickname, install agent, and aligned device. On this tab, you can configure the disk space, excludes, includes, and other settings related to this agent. To view the **Device Investigator** page for the device monitored by this agent, click the device name in the **Device** field.

To configure the agent:

1. Click the **[Config]** tab of the **Agent Investigator** page:

The screenshot shows the configuration page for agent 12r2ent-dc3 on a Windows 2012 R2 system. The page has four tabs: Config (selected), Polled Data (1), Log Sources (1), and Watched Files (0). The main content area is divided into several sections:

- Description:** Name: 12r2ent-dc3, Nickname: 12r2ent-dc3.
- Install Agent:** Windows 2012 R2.
- Device:** 12r2ent-dc3 - 10.40.5.37, Healthy (indicated by a green dot).
- Collect Frequency:** 20 Seconds.
- Disk Space (MB):** 512.
- Excludes:** A text input field for listing processes and directories to be excluded.
- Includes:** A text input field for listing processes and directories to be included.
- Collection Options:**
 - Collect File Information
 - Collect Named Pipe Information
 - Collect Socket Information
 - Collect Thread Information
 - Collect Non-Intercepted Processes
- Without Sockets:** A dropdown menu currently set to "Without Sockets".

On the right side, there are explanatory text blocks for the Excludes, Includes, and Without Sockets settings. A blue "Save" button is located at the bottom center of the configuration area.

2. Complete in the following fields, as needed:
 - **Disk Space.** Specify the amount of disk space that the agent can use to store data. If the agent loses connectivity to the new user interface, this disk space will be used to store collected data until the connection to the new user interface is restored. When connectivity is re-established, the agent uploads all of its stored data.
 - **Excludes.** Type a list of processes and directories, separated by semicolons, that you do not want the agent to monitor.
 - **Includes.** Type a list of processes and directories, separated by semicolons, that you want the agent to monitor. This field ensures that specific processes are monitored.

NOTE: If a process or directory is included in both the **Excludes** field and the **Includes** field, that process or directory *will* be monitored by the agent.

3. Select the following configuration options as needed:
 - **Collect File Information:** Select this option if you want the agent to report the names of files accessed by each monitored process.
 - **Collect Named Pipe Information:** Select this option if you want the agent to collect named pipe information.
 - **Collect Socket Information:** Select this option if you want the agent to collect socket information.
 - **Collect Thread Information:** Select this option if you want the agent to collect thread information.
 - **Collect Non-Intercepted Processes:** Select this option if you want the agent to collect limited information for processes that do not have the agent library in them.
4. In the **All** drop-down list, tell the agent how to aggregate short-lived processes. Your options include the following:
 - **All:** Aggregate every short-lived process into its parent.
 - **None:** Do not aggregate any short-lived process.
 - **Without Sockets:** Aggregate short-lived processes unless those processes have sockets.
5. Click the **[Save]** button to save your configuration settings for the agent.

The Polled Data Tab

The **[Polled Data]** tab of the **Agent Investigator** page lets you configure a new source for polled data and manage existing sources for polled data.

Creating a New Source for Polled Data

To create a new source for polled data:

1. On the **[Polled Data]** tab of the **Agent Investigator** page, click the **[New]** button. A set of new fields appear in the right-hand section:

The screenshot shows the 'Agent Investigator' interface for host '12r2ent-dc3' (Windows 2012 R2). The 'Polled Data (1)' tab is active, displaying a table with one entry: '_SLIPD_topo' with an interval of 45, data type 'Script', and 'ENABLED?' checked. A 'New' button is visible above the table. To the right, the 'New Polled Data' configuration form is open. It includes fields for 'Name (required)', 'Data Type' (set to 'Script'), 'Type Fields (valid JSON required)' (with an example '{"script": "required"}'), 'Polling Interval (required)' (set to 60), 'Shell', 'Username', and an 'Enable' checkbox. A 'Save' button is at the bottom right. The form also contains help text for each field.

2. Complete the following fields, as needed:

- **Name**. Type the name of the service or script you are configuring. This field is required.
- **Data Type** and **Type Fields**. Select a data type for this service or script in the **Data Type** drop-down list, and then type code in the **Type Fields** field that is relevant for that data type. Select any option in the **Data Type** drop-down list to see a code example of that type. Your options include:
 - *Script*. Run a simple JSON script, such as `{"script": "echo \"hello world\"}"`
 - *URL*. Download a script (not a binary) from a URL and execute it, such as `{"url": "https://my_webserver /hello_world.sh"}`. Both http and https are supported.

- *JMX*. Query a bean for a specific process. You must use JConsole to get the list of available beans. If you cannot use JConsole to get the beans from a process, you will not be able to get Java Management Extensions (JMX) data from the agent. This process can only get beans from Java 6, 7 and 8, and only supports the Oracle and OpenJDK runtime environments. The **Username** field on this tab should match the username that is running the process from which you want to get the list of beans.

The following is an example of code for a JMX data type:

```
{ "jmx": {
  "process": "HelloWorld",
  "object": "java.lang:type=Memory",
  "attribute": "HeapMemoryUsage",
  "subattribute": "used",
  "warnthreshold": "3200000",
  "critthreshold": "4500000"
}}
```

- *Windows Performance Counter*. Use the specified Windows performance counter for polled data.
- *Event*. Use the specified event for polled data.
- **Polling Interval**. Specify the delay, in seconds, between samples. If you type "0" in this field, the command is run only once. This field is required.
- **Shell**. If needed, specify the shell for the script or URL to run. By default, the command uses the default Linux shell or Windows command prompt, depending on the agent type. You can also use Bash, Python, and PowerShell shells. For example, you can use the following as a Windows Nagios shell:

```
{ "script": "C:\\nagios\\plugins\\check_winprocess-1.6\\check_winprocess" }
```

- **Username**. For Linux, specify the username under which to run the command; root is the default. For Windows, the user must exist and be logged in at the time the poll is run; the default is to run the command as the system account.

3. To make the polled data data source active, select the **Enabled** checkbox.
4. Click the **[Save]** button.

Configuring Nagios Plug-ins as Polled Data Sources

Because the agent does not have a plug-ins directory, you need to create a polled data script that points to the script you want to run.

1. Download the Nagios plug-ins you want to use.
2. In the new user interface, navigate to the **Agent Investigator** page (Inventory > Agents > agent record) for the agent you want to use with the Nagios plug-ins.

3. Create a new polled data source with a **Data Type** of *Script*. Some examples of JSON code you might add to the **Type Fields** field include:


```
{"script": "/usr/lib64/nagios/plugins/check_procs"}
```

```
{"script": "/usr/lib64/nagios/plugins/check_mysql "}
```

```
{"script": "/usr/lib64/nagios/plugins/check_disk -w 10% -c 5% -p /tmp -p /var  
-C -w 100000 -c 50000 -p / "}
```

Editing or Deleting Polled Data Sources

To edit or delete existing polled data sources:

1. From the list of polled data on the **[Polled Data]** tab, select the source you want to edit or delete.
2. Click the **[Actions]** button () and select an option:
 - *Edit*. When you select this option, the fields for this polled data source appear in the right-hand panel, and you can update the fields as needed.
 - *Delete*. When you select this option, the polled data source is immediately deleted.

The Log Sources Tab

The **[Log Sources]** tab of the **Agent Investigator** page lets you configure log sources that the agent will monitor. These sources include syslogs, event logs, and files.

Creating a New Log Source

To create a new log source:

1. On the **[Log Sources]** tab of the **Agent Investigator** page, click the **[New]** button. A set of new fields appear in the right-hand section:

The screenshot shows the 'Agent Investigator' interface for host '12r2ent-dc3' (Windows 2012 R2). The 'Log Sources (1)' tab is active, displaying a table with one entry: 'Application' with a limit of 2000 and source type 'EVENTLOG'. A 'New Log Source' form is open on the right, with the following fields and options:

- Source:** A text input field with a placeholder 'specify a single log file'.
- Source Type:** A dropdown menu.
- Limit:** A text input field with the value '1000' and a clear button.
- Log Filter:** A text input field.
- Template:** A dropdown menu with a placeholder 'Select a log template that best matches the format of text in the log file to improve indexing for log-search. Or select "agent_time_log" to just index on the time the log file is read.'

A 'Save' button is located at the bottom right of the form.

2. Complete the following fields, as needed:


- **Source Type.** Select the log source type. The **Source** field depends on your choice in this field, so select the **Source Type** first. Your options include:
 - *SysLog.* The agent will monitor a syslog on the device. If you select this option, use the **Source** field to specify the UDP port number to listen on.
 - *Event Log.* The agent will monitor the Windows logs on the device. If you select this option, specify an Event Log category in the **Source** field (Application, Security, System).
 - *File.* The agent will monitor a file on the file system of the device. If you select this option, type the full path of the file to monitor in the **Source** field.
- **Source.** Specify the relevant source information based on what you selected in the **Source Type** field. This field is required.
- **Limit.** Specify the maximum number of lines in the log source. Optional.
- **Log Filter.** Specify a regular expression that will evaluate the log messages in the specified syslog, file, or Windows log. If a log message matches this regular expression, the agent sends that log message to SL1. Optional.

- **Template.** Select a log template that best matches the text format in the log file you want to monitor. Using a template in this way improves indexing for log searching. If you want to index at the time the log file is read, select `agent_time_log` from the drop-down list.

3. Click the **[Save]** button.

Editing or Deleting Log Sources

To edit or delete existing log sources:

1. From the list of log sources on the **[Log Sources]** tab, select the source you want to edit or delete.
2. Click the **[Actions]** button () and select an option:
 - *Edit.* When you select this option, the fields for this log source appear in the right-hand panel, and you can update the fields as needed.
 - *Delete.* When you select this option, the log source is immediately deleted.

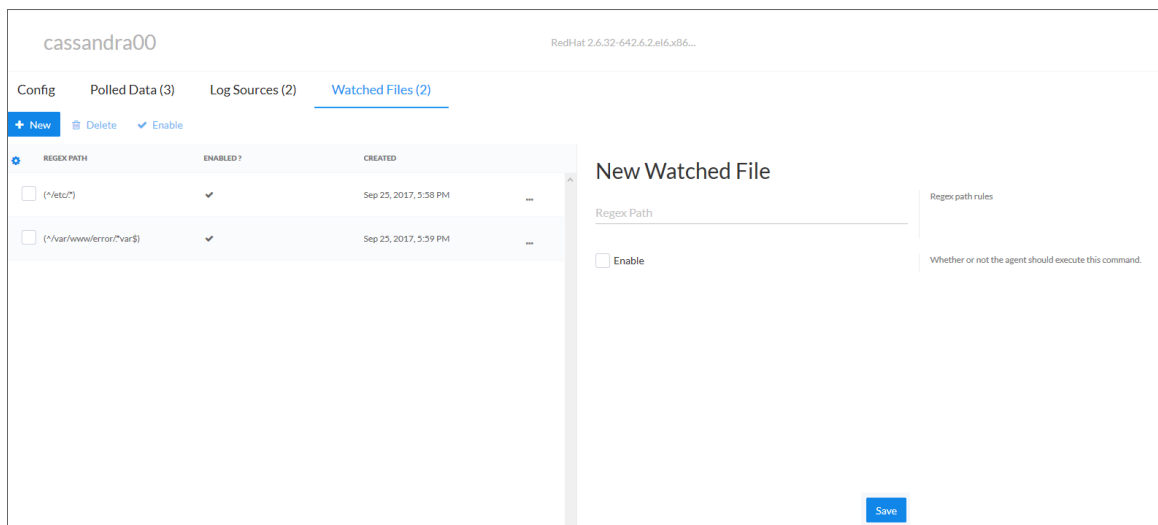
The Watched Files Tab

The **[Watched Files]** tab of the **Agent Investigator** page lets you add regular-expression rules for a set of files you want the agent to watch for you.

Adding Watched Files

To add watched files:

1. On the **[Watched Files]** tab of the **Agent Investigator** page, click the **[New]** button. A set of new fields appears in the right-hand section:



2. In the **Regex Path** field, type the regular expression rules for the file you want to watch, enclosed in parentheses (). Some regex examples include:


```
(^/etc/.*)
```

```
(^/var/www/error/.*var$)
```

3. To make the watched file active, select the **Enabled** checkbox.
4. Click the **[Save]** button.

Editing or Deleting Watched Files

To edit or delete existing watched files:

1. From the list of watched files on the **[Watched Files]** tab, select the file you want to edit or delete.
2. Click the **[Actions]** button () and select an option:
 - *Edit*. When you select this option, the fields for this watched file appear in the right-hand panel, and you can update the fields as needed.
 - *Delete*. When you select this option, the watched file is immediately deleted.

Changing the Target Message Collector for the Agent

You can specify with which Message Collector the agent communicates by editing the main configuration file on your Linux or Windows system.

NOTE: Edit the main configuration file for the purposes of troubleshooting or changing the target Message Collector only. Any other changes made to the main configuration file will be overwritten automatically by the appliance performing message collection.

To reconfigure the agent to communicate with a different Message Collector:

1. Either go to the console of the device where the agent resides or open an SSH session to that device.
2. Using a text editor like "vi", open the main configuration file.

- On a Linux system, the main configuration file is:

```
/etc/scilog/scilog.conf
```

- On a Windows system, the main configuration file is:

```
Program Files\ScienceLogic\SiloAgent\conf\scilog.conf
```

3. Locate the following line and change the IP address to the IP address of the new Message Collector:

```
URL https://<IP address>/SaveData.py/save_data
```

4. Locate the following line and change the IP address to the IP address of the new Message Collector:

```
URLfront <IP address>
```

5. Save and exit the text editor.
6. On a Linux system, restart the scilogd service.

```
sudo /etc/init.d/scilogd restart
```
7. On a Windows system, restart the SiloAgent Service service.


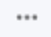
```
net stop "SiloAgent Service"  
net start "SiloAgent Service"
```

Monitoring Ports Using an Agent

Overview

This chapter describes monitoring ports on devices monitored by an agent.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon ().
- To view a page containing all of the menu options, click the Advanced menu icon ().

This chapter includes the following topics:

<i>What is a Port?</i>	25
<i>Port Security</i>	25
<i>Port Availability</i>	25
<i>Viewing a List of All Open Ports on All Devices</i>	26
<i>Viewing a List of All Open Ports on a Single Device</i>	27
<i>System Settings for Monitoring Port Availability</i>	29
<i>Viewing the TCP/IP Port Monitoring Policies</i>	30
<i>Defining a Monitoring Policy for Port Availability</i>	31
<i>Example Policy for TCP/IP Port Availability</i>	33
<i>Editing a Monitoring Policy for a TCP/IP Port</i>	33
<i>Executing a TCP-IP Port Monitoring Policy</i>	34
<i>Deleting a TCP/IP Port Monitoring Policy</i>	35

What is a Port?

Ports are used to route packets on a server to the appropriate application. Ports are like an apartment number in an apartment building; the street address (IP address) gets the message to the right building, and the apartment number (port number) gets the message to the right person. For example, port 80 is the standard port number for HTTP traffic, and port 80 packets are processed by a Web server.

Ports can use the UDP protocol or the TCP protocol. UDP does not include a handshake, does not ensure packets are sent in a particular order, does not return error messages, and will not automatically try to resend or re-receive a packet; TCP will do all these things. Commonly used UDP ports include port 53 for DNS and port 161 for SNMP. Commonly used TCP ports include port 80 for HTTP, port 25 for SMTP, and port 20 for FTP.

Ports 0-1023 are used by common Internet applications such as HTTP, FTP, and SMTP. Ports 1024-49151 can be registered by vendors for proprietary applications.

Port Security

The **Port Security** page (Devices > Device Manager > bar-graph icon > Performance) displays a list of all open ports on a device.

For SNMP and pingable devices, SL1 scans each device's TCP ports using NMAP.

For devices monitored using the SL1 agent, the agent reports open TCP and UDP ports. By default, the list of discovered ports is then automatically updated in SL1 every 5 minutes per agent.

The **Port Security** page displays open port information collected using NMAP and the SL1 agent, where applicable.

For SNMP and pingable devices, SL1 scans all the ports of each managed device every day. If any new ports are opened, SL1 updates the **Port Security** page and creates an event to notify users. You can explicitly ask that a device not be scanned nightly using NMAP, but if you do, SL1 will not notify you of newly opened ports on the device.

Port Availability

SL1 can monitor ports for availability. When a port monitor is created, SL1 monitors the port for availability every five minutes. You can choose whether a policy is executed by SL1 using NMAP or locally on the device by the agent.

During polling, a port has two possible availability values:

- 100%. Port is up and running.
- 0%. Port is not accepting connections and data from the network.

The data gathered by the port monitor is used to create port-availability reports.

If a port is not available, SL1 creates an event with the message "port not responding to connection".

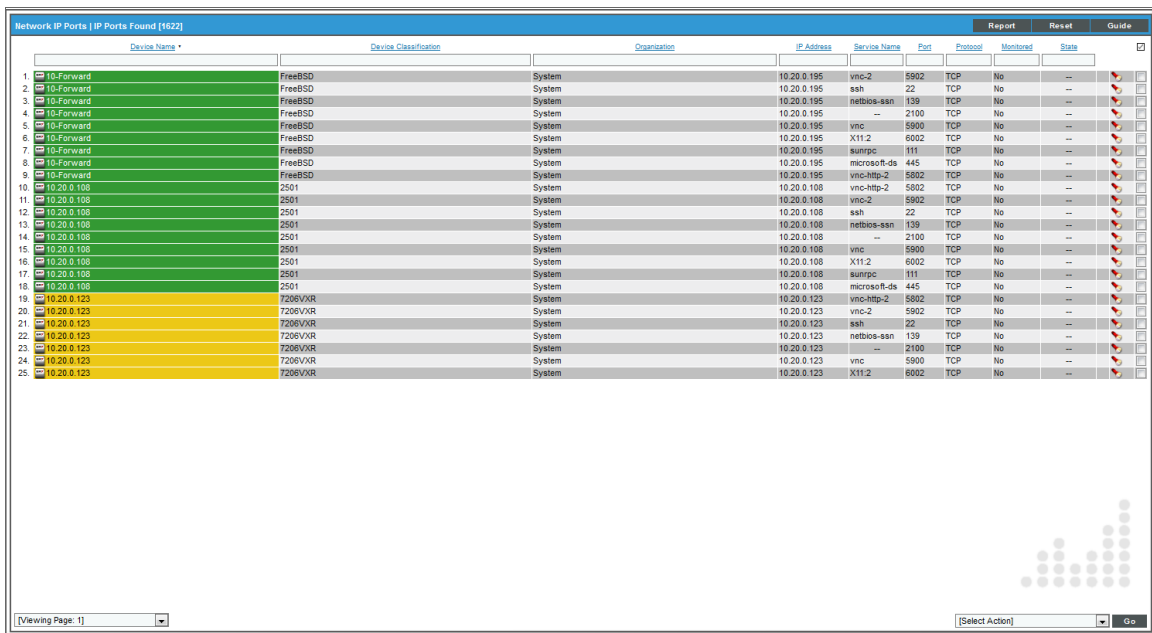
Viewing a List of All Open Ports on All Devices

The **Network IP Ports** page displays a list of all open ports on all devices discovered by SL1 using NMAP and the SL1 agent.

NOTE: Users of type "user" can view only IP ports that are aligned with the same organization(s) to which the user is aligned. This means that the device associated with the port(s) must be aligned with one of the organizations to which the user is aligned. Users of type "administrator" can view all IP ports.

To view the **Network IP Ports** page:

1. Go to the **Network IP Ports** page (Registry > Networks > IP Ports).



The screenshot shows the 'Network IP Ports | IP Ports Found [1622]' page. It features a table with columns: Device Name, Device Classification, Organization, IP Address, Service Name, Port, Protocol, Monitored, and State. The table lists 25 rows of data, including ports like vnc-2, ssh, netbios-ssn, and vnc on various IP addresses. A '4' is visible in a blue box on the right side of the page.

Device Name	Device Classification	Organization	IP Address	Service Name	Port	Protocol	Monitored	State
10-Forward	FreeBSD	System	10.20.0.195	vnc-2	5902	TCP	No	--
10-Forward	FreeBSD	System	10.20.0.195	ssh	22	TCP	No	--
10-Forward	FreeBSD	System	10.20.0.195	netbios-ssn	139	TCP	No	--
10-Forward	FreeBSD	System	10.20.0.195	--	2100	TCP	No	--
10-Forward	FreeBSD	System	10.20.0.195	vnc	5900	TCP	No	--
10-Forward	FreeBSD	System	10.20.0.195	X11-2	6002	TCP	No	--
10-Forward	FreeBSD	System	10.20.0.195	sunrpc	111	TCP	No	--
10-Forward	FreeBSD	System	10.20.0.195	microsoft-ds	445	TCP	No	--
10-Forward	FreeBSD	System	10.20.0.195	vnc-hdp-2	5802	TCP	No	--
10.20.0.108	2501	System	10.20.0.108	vnc-hdp-2	5802	TCP	No	--
10.20.0.108	2501	System	10.20.0.108	vnc-2	5802	TCP	No	--
10.20.0.108	2501	System	10.20.0.108	ssh	22	TCP	No	--
10.20.0.108	2501	System	10.20.0.108	netbios-ssn	139	TCP	No	--
10.20.0.108	2501	System	10.20.0.108	--	2100	TCP	No	--
10.20.0.108	2501	System	10.20.0.108	vnc	5900	TCP	No	--
10.20.0.108	2501	System	10.20.0.108	X11-2	6002	TCP	No	--
10.20.0.108	2501	System	10.20.0.108	sunrpc	111	TCP	No	--
10.20.0.108	2501	System	10.20.0.108	microsoft-ds	445	TCP	No	--
10.20.0.123	7206VXR	System	10.20.0.123	vnc-hdp-2	5802	TCP	No	--
10.20.0.123	7206VXR	System	10.20.0.123	vnc-2	5802	TCP	No	--
10.20.0.123	7206VXR	System	10.20.0.123	ssh	22	TCP	No	--
10.20.0.123	7206VXR	System	10.20.0.123	netbios-ssn	139	TCP	No	--
10.20.0.123	7206VXR	System	10.20.0.123	--	2100	TCP	No	--
10.20.0.123	7206VXR	System	10.20.0.123	vnc	5900	TCP	No	--
10.20.0.123	7206VXR	System	10.20.0.123	X11-2	6002	TCP	No	--

2. The **Network IP Ports** page displays a list of all discovered ports. For each port, the **Network IP Ports** page displays the following:

TIP: To sort the list of ports, click on a column heading. The list will be sorted by the column value, in ascending order. To sort the list by descending order, click the column heading again.

- **Device Name.** Name of the device where the port resides. For devices running SNMP or with DNS entries, the name is discovered automatically. For devices without SNMP or DNS entries, the device's IP address will appear in this field.

- **Device Classification.** The manufacturer (device class) and type of device (sub-class). The Device-Class/Sub-Class is automatically assigned during auto-discovery, at the same time as the Category.
- **Organization.** The Organization associated with the device and port.
- **IP Address.** IP address associated with the open port.
- **Service Name.** The service accessed through the port.
- **Port.** The port number.
- **Protocol.** Either TCP or UDP.
- **Monitored.** Specifies whether SL1 is monitoring this port for availability.
- **State.** This column has a value only if a port-monitoring policy has been defined for the port. This field can have one of two values:
 - *Enabled.* The port-monitoring policy has been activated. SL1 monitors the port and collects availability data about the port.
 - *Disabled.* The port-monitoring policy has not been activated. SL1 will not monitor the port and does not collect availability data about the port.



For more information about filtering the list of IP Ports displayed on the Network IP Ports page, see the **Device Management** manual.

Viewing a List of All Open Ports on a Single Device

NOTE: Users of type "user" can view only IP ports that are aligned with the same organization(s) to which the user is aligned. This means that the device associated with the port(s) must be aligned with one of the organizations to which the user is aligned. Users of type "administrator" can view all IP ports.

The **Port Security** page displays a list of all open ports on a single device.

To view the **Port Security** page for a device:

1. There are two ways to view the **Port Security** page:
 - Go to the **Device Manager** page (Devices > Device Manager). Find the device where you want to view the **Port Security** page. Select the bar graph icon () for that device.
 - Go to the **Network IP Ports** page (Registry > Networks > IP Ports). Find the device for which you want to view the **Port Security** page. Select the flashlight icon () for that device.

2. In the **Device Reports** panel, select the **[TCP/UDP Ports]** tab. The **Port Security** page appears.

Interface IP	Port Number	Service	Protocol	Certificate Issuer	Cert Expiration
0.0.0.0	0		TCP	--	--
::	0		UDP	--	--
::	0		TCP	--	--
0.0.0.0	0		UDP	--	--
10.64.68.20	22	ssh	TCP	--	--
::	22	ssh	TCP	--	--
0.0.0.0	22	ssh	TCP	--	--
10.64.68.20	25	smtp	TCP	--	--
0.0.0.0	25	smtp	TCP	--	--
::	25	smtp	TCP	--	--
127.0.0.1	80	http	TCP	--	--
10.64.68.20	80	http	TCP	--	--
0.0.0.0	80	http	TCP	--	--
0.0.0.0	161	snmp	UDP	--	--
0.0.0.0	162	snmptrap	UDP	--	--
127.0.0.1	199	smux	TCP	--	--
127.0.0.1	323		UDP	--	--
::1	323		UDP	--	--
0.0.0.0	443	https	TCP	--	--
10.64.68.20	443	https	TCP	Silo	2017-06-08 14:50:04
0.0.0.0	514	syslog	UDP	--	--
::	514	shell	TCP	--	--
0.0.0.0	5000	UPnP	TCP	--	--
10.64.68.20	5000	UPnP	TCP	--	--
127.0.0.1	5001	complex-link	TCP	--	--
0.0.0.0	7700		TCP	--	--
::ffff:127.0.0.1	7706		TCP	--	--
::	7706		TCP	--	--

3. For each open port on the device, the **Port Security** page displays the following information:

- **Interface IP.** IP address through which SL1 communicates with the device.
- **Port Number.** The ID number of the port.
- **Service.** The service accessed through the port.
- **Protocol.** Either TCP or UDP.
- **Certificate Issuer.** If the service on this port uses a certificate, this column contains the name of the certificate authority.

NOTE: Certificates are used by secure services like HTTPS, SSL, SSH, and SFTP to verify communication and encrypt message. The certificate issuer (also known as the certificate authority or CA) is an organization that issues digital certificates (digital IDs). These digital IDs (called keys) authenticate the identity of people and organizations over a public system such as the Internet. These keys also allow senders and receivers to encrypt messages and un-encrypt replies.

- **Cert. Expiration.** The expiration date of the certificate.

System Settings for Monitoring Port Availability

Although you are not required to define system settings for port availability, you might find it useful to understand how these settings affect port monitoring.

The **Behavior Settings** page (System > Settings > Behavior) includes the following settings that affect policies for port availability:

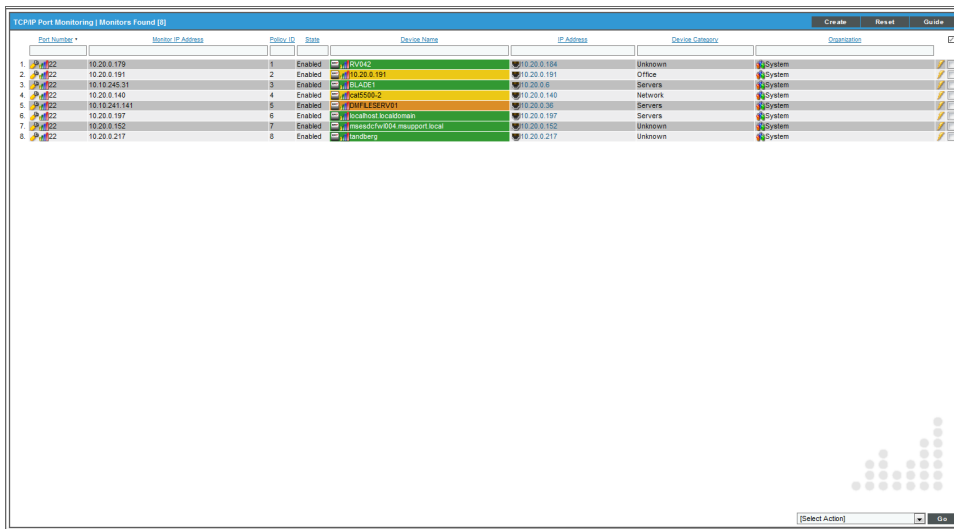
- **Port Polling Type.** Specifies how SL1 should poll ports for availability using NMAP. The choices are:
 - *Half Open.* Uses a faster TCP/IP connection method (a TCP SYN scan, `nmap -sS`) and does not appear on device's logs.
 - *Full Connect.* Uses the standard TCP/IP connection (`TCP connect()` scan, `nmap -sT`) to detect open ports.

Viewing the TCP/IP Port Monitoring Policies

You can view a list of TCP/IP port monitoring policies from the **TCP/IP Port Monitoring** page (Registry > Monitors > TCP-IP Ports).

The **TCP/IP Port Monitoring** page displays the following information for each TCP/IP port monitoring policy:

NOTE: Users of type "user" can view only IP ports that are aligned with the same organization(s) to which the user is aligned. This means that the device associated with the port(s) must be aligned with one of the organizations to which the user is aligned. Users of type "administrator" can view all IP ports.



Port Number	Monitor IP Address	Policy ID	State	Device Name	IP Address	Device Category	Organization	Disabling
1. IP:22	10.20.0.179	1	Enabled	RD42	10.20.0.184	Unknown	System	
2. IP:22	10.20.0.191	2	Enabled	10.20.0.191	10.20.0.191	Office	System	
3. IP:22	10.10.245.31	3	Enabled	BUCKET	10.20.0.6	Servers	System	
4. IP:22	10.20.0.140	4	Enabled	server02	10.20.0.140	Network	System	
5. IP:22	10.10.241.141	5	Enabled	PMLESERVER1	10.20.0.36	Servers	System	
6. IP:22	10.20.0.197	6	Enabled	localhost.ucaboman	10.20.0.197	Servers	System	
7. IP:22	10.20.0.192	7	Enabled	localhost.ucaboman.local	10.20.0.192	Helpdesk	System	
8. IP:22	10.20.0.217	8	Enabled	landberg	10.20.0.217	Unknown	System	


- **TCP/IP Port Number.** Port number of the port to be monitored.
- **Monitor IP Address.** IP address associated with the port to be monitored. For devices with multiple IP addresses, the IP address for the port policy might be different than the IP address used by SL1 to communicate with the device.
- **Policy ID.** Unique, numeric ID, assigned to the policy automatically by SL1.
- **Device Name.** Name of the device associated with the policy.
- **IP Address.** IP address of the device associated with the policy. This is the IP address SL1 uses to communicate with the device.
- **Device Category.** Device category of the device associated with the policy.
- **Organization.** Organization for the device associated with the policy.

Defining a Monitoring Policy for Port Availability

NOTE: Users of type "user" can view only IP ports that are aligned with the same organization(s) to which the user is aligned. This means that the device associated with the port(s) must be aligned with one of the organizations to which the user is aligned. Users of type "administrator" can view all IP ports.

You can define a port monitoring policy in the **TCP/IP Port Policy** modal page. You can access the **TCP/IP Port Policy** page either from the **Device Manager** page (Devices > Device Manager) or from the **TCP/IP Port Monitoring** page (Registry > Monitors > TCP-IP Ports).

To access the **TCP/IP Port Policy** modal page from the **Device Manager** page:

1. Go to the **Device Manager** page (Devices > Device Manager)
2. In the **Device Manager** page, find the device that you want to associate with the monitoring policy. Select wrench icon () for the device.
3. In the **Device Administration** panel for the device, select the **[Monitors]** tab.
4. From the **[Create]** menu in the upper right, select **Create TCP/IP Port Policy**.
5. The **TCP/IP Port Policy** modal page appears.

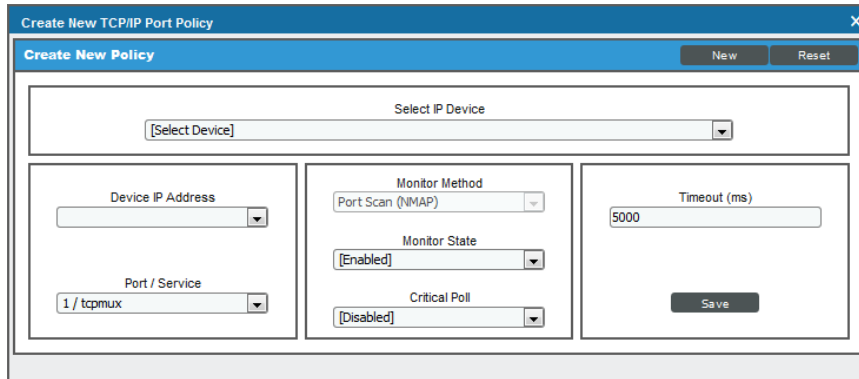
To access the **TCP/IP Port Policy** modal page from the **TCP/IP Port Monitoring** page:

1. Go to the **TCP/IP Port Monitoring** page (Registry > Monitors > TCP-IP Ports).
2. Select the **[Create]** button.
3. The **TCP/IP Port Policy** modal page appears.

To define a port monitoring policy:

1. Navigate to the **TCP/IP Port Policy** modal page. See the procedures above for more information.

2. In the **TCP/IP Port Policy** modal page, supply a value in each of the following fields:



- **Select Device.** Select a device from this drop-down list to align with this policy. By default, the current device is selected in this field.
- **Device IP Address.** IP address through which SL1 communicates with the device.
- **Port/Service.** Port number and the corresponding service running on the port.
- **Monitor Method.** Select whether the policy will be executed using NMAP or using the agent. This option is available only if you selected a device on which the agent is installed.
- **Monitor State.** Specifies whether SL1 should start collecting data specified in this policy from the device. Choices are:
 - *Enabled.* SL1 will collect the data specified in this policy, from the device, at the frequency specified in the **Process Manager** page (System > Settings > Processes) for the **Data Collection: TCP Port Monitor** process.
 - *Disabled.* SL1 will not collect the data specified in this policy, from the device, until the **State** field is set to *Enabled*.
- **Critical Poll.** Frequency with which SL1 should "ping" the device. If the device does not respond, SL1 creates an event. The choices are:
 - *Disabled.* SL1 will not ping the device.
 - *Enabled.* SL1 will ping the device every 15, 30, 60, or 120 seconds, as specified.

NOTE: SL1 uses **Critical Poll** data to create events when mission-critical ports are not available. SL1 does not use this critical poll data to create port-availability reports. SL1 will continue to collect port availability only every five minutes.

3. Click **[Save]**.

Example Policy for TCP/IP Port Availability

The screenshot shows a configuration window for a TCP/IP Port Policy. The window has a blue header with the title 'TCP/IP Port Policy' and a close button. Below the header is a sub-header 'Editing Policy [1]' with 'New' and 'Reset' buttons. The main area contains a 'Select Device' dropdown menu with 'em7ao' selected. Below this are three columns of settings: 'Device IP Address' (10.64.68.20), 'Port / Service' (22 / ssh), 'Monitor Method' (Port Scan (NMAP)), 'Monitor State' (Enabled), and 'Critical Poll' (Disabled). A 'Timeout (ms)' field is set to 5000. 'New' and 'Reset' buttons are at the top right, and a 'Save' button is at the bottom right.

- This policy monitors a TCP/IP port on the device "cisco_10.2.1.29", at IP address 10.1.0.205.
- The policy will monitor port 22 for availability.

Editing a Monitoring Policy for a TCP/IP Port

You can edit a port monitoring policy on the **TCP/IP Port Policy** modal page. You can access the **TCP/IP Port Policy** modal page either from the **Device Manager** page (Devices > Device Manager) or from the **TCP/IP Port Monitoring** page (Registry > Monitors > TCP-IP Ports).

To access the **TCP/IP Port Policy** modal page from the **Device Manager** page:

1. Go to the **Device Manager** page (Devices > Device Manager)
2. In the **Device Manager** page, find the device that you want to associate with the monitoring policy. Select the wrench icon (🔧) for the device.
3. In the **Device Administration** panel, select the **[Monitors]** tab.
4. In the **Monitoring Policies** page, find the port policy you want to edit and select its wrench icon (🔧).
5. The **TCP/IP Port Policy** modal page appears.

To access the **TCP/IP Port Policy** modal page from the **TCP/IP Port Monitoring** page:

1. Go to the **TCP/IP Port Monitoring** page (Registry > Monitors > TCP-IP Ports).
2. Find the device and port for which you want to edit the monitoring policy. Select the wrench icon (🔧) for the port.
3. The **TCP/IP Port Policy** modal page appears.

To edit a port monitoring policy:

1. If you have not done so already, navigate to the **TCP/IP Port Policy** modal page. See the procedures above for more information.
2. In the **TCP/IP Port Policy** modal page, edit the values in one or more of the fields.

The screenshot shows a modal window titled "TCP/IP Port Policy" with a sub-header "Editing Policy [1]". It contains several fields: "Select Device" with a dropdown menu showing "em7ao"; "Device IP Address" with a dropdown menu showing "10.64.68.20"; "Port / Service" with a dropdown menu showing "22 / ssh"; "Monitor Method" with a dropdown menu showing "Port Scan (NMAP)"; "Monitor State" with a dropdown menu showing "Enabled"; "Critical Poll" with a dropdown menu showing "Disabled"; and "Timeout (ms)" with a text input field showing "5000". There are "New" and "Reset" buttons at the top right, and a "Save" button at the bottom right.

3. Click **[Save]** when done.

Executing a TCP-IP Port Monitoring Policy

After creating or editing a TCP-IP port monitoring policy, you can manually execute the policy and view detailed logs of each step during the execution. To do so:

NOTE: After you define a TCP-IP port monitoring policy and enable the policy, SL1 or the SL1 agent will automatically execute the policy every five minutes. However, you can use the steps in this section to execute the policy immediately and see debug information about the execution of the policy.

1. In the **TCP/IP Port Monitoring** page (Registry > Monitors > TCP-IP Ports), find the policy you want to run manually.
2. Select the lightning bolt icon (⚡) to manually execute the policy.
3. While the policy is executing, SL1 spawns a modal page called **Session Logs**. The **Session Logs** page provides detailed descriptions of each step during the execution. This is helpful for diagnosing possible problems with a policy.

You can view reports for executed port monitoring policies. See the **Device Management** manual for more information.

Deleting a TCP/IP Port Monitoring Policy

You can delete a port monitoring policy from the **TCP/IP Port Monitoring** page. You can delete individual, multiple, or all existing port monitoring policies. When you delete a TCP/IP Port Monitoring policy, SL1 no longer uses the policy to collect data from the aligned device.

To delete a port monitoring policy:

1. Go to the **TCP/IP Port Monitoring** page (Registry > Monitors > TCP-IP Ports).
2. In the **TCP/IP Port Monitoring** page, select the checkbox(es) for each port monitoring policy you want to delete. Click the checkmark icon (☑) to select all of the system process policies.
3. In the **[Select Action]** menu in the bottom right of the page, select *Delete Monitors*.

Row Number	Device IP Address	Subnet ID	State	Device Name	IP Address	Device Category	Organization
1	10.20.0.178	1	Enabled	RV942	10.20.0.184	Unknown	System
2	10.20.0.191	2	Enabled	10.20.0.191	10.20.0.191	Office	System
3	10.10.0.2531	3	Enabled	10.10.0.2531	10.10.0.2531	Servers	System
4	10.20.0.140	4	Enabled	10.20.0.140	10.20.0.140	Network	System
5	10.10.241.141	5	Enabled	10.10.241.141	10.10.241.141	Servers	System
6	10.20.0.197	6	Enabled	10.20.0.197	10.20.0.197	Servers	System
7	10.20.0.152	7	Enabled	msesdctw004.msupport.local	10.20.0.152	Unknown	System
8	10.20.0.217	8	Enabled	linborg	10.20.0.217	Unknown	System


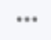
4. Click **[Go]** to delete the port monitoring policy.
5. The policy is deleted from SL1. The associated reports (from the Device Reports > **[Performance]** tab) are also deleted.

Monitoring Processes Using an Agent

Overview

This chapter describes viewing system processes for devices monitored with an agent. It also describes using system process reports and monitoring policies to monitor processes.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon ().
- To view a page containing all of the menu options, click the Advanced menu icon (.

This chapter includes the following topics:

<i>What is a Process?</i>	37
<i>Viewing the List of Device Processes</i>	37
<i>Generating a Report on Multiple System Processes</i>	39
<i>Generating an Exclusion Report for a Single System Process</i>	41
<i>Viewing the System Process Monitoring Policies</i>	43
<i>Defining a Monitoring Policy for a System Process</i>	44
<i>Example System Process Monitoring Policy</i>	47
<i>Editing a System Process Monitoring Policy</i>	48
<i>Executing a System Process Monitoring Policy</i>	49
<i>Deleting a System Process Monitoring Policy</i>	49

What is a Process?

A process is a program that is currently running or has been run in the past and is currently idle. Sometimes a process is called a task.

There are two methods for monitoring processes:

- For devices monitored using SNMP, SL1 automatically collects a list of all processes running every two hours.
- For devices monitored using the SL1 agent, SL1 automatically collects a list of all processes running every five minutes.

SL1 allows you to create policies that monitor system processes every five minutes:

- If a device is not monitored using the SL1 agent, the policy collection is performed using SNMP.
- If a device is monitored using the SL1 agent, the policy collection is performed by the agent.

For each monitored process, you can create a policy that specifies:

- Whether or not to generate an event if the process is running.
- How much memory each instance of a process can use.
- How many instances of a process can run simultaneously.
- If policy collection is performed by the agent, how much memory all instances of a process can use in total.
- If policy collection is performed by the agent, how much CPU all instances of a process can use in total.

Viewing the List of Device Processes

The **Device Processes** page displays a list of all processes discovered by SL1 on all devices.

To view the list of all processes running on all discovered devices:

1. Go to the **Device Processes** page (Devices > Processes).

Device Name	Organization	IP Address	Device Class / Sub-Class	Process	PID	Memory	Run State
1. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AKNAGent.exe	2360	8192 KB	Running
2. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	ALV.exe	3400	2004 KB	Running
3. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVCSAD.exe	1568	36060 KB	Running
4. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVCSMgr.exe	3148	15500 KB	Running
5. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVCSChangeFilter.exe	8972	23724 KB	Running
6. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVCSAD.exe	5740	51950 KB	Running
7. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVCSDBCatalog.exe	5864	53300 KB	Running
8. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVCS.exe	4112	36060 KB	Running
9. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVMSProxySrv.exe	8200	45008 KB	Running
10. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVMagStoreMonitorSrv.exe	8200	68420 KB	Running
11. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVContentMgr.exe	8044	59960 KB	Running
12. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVRepOrSrvSrv.exe	8054	20020 KB	Running
13. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVSaverMgrSrv.exe	2480	20224 KB	Running
14. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVSvcChangeNotifier.exe	4140	10320 KB	Running
15. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AV7Svc.exe	2844	21084 KB	Running
16. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVUISyncSrv.exe	7804	47340 KB	Running
17. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVAgent.exe	2832	12704 KB	Running
18. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVjava-mvc.exe	3352	20332 KB	Running
19. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVCastionTSAPolicy.exe	8204	34800 KB	Running
20. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVSvc.exe	3956	5456 KB	Running
21. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVgoramc.exe	2452	4996 KB	Running
22. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVCSMgrConnector.exe	8300	40760 KB	Running
23. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVCSMVC.exe	2508	35748 KB	Running
24. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVSvc.exe	1256	6952 KB	Running
25. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVContentMgr.exe	8116	57320 KB	Running
26. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVContentMgr.exe	4192	4524 KB	Running
27. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVContentMgr.exe	8412	34820 KB	Running
28. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVSvc.exe	2008	11024 KB	Running
29. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVAgent.exe	5312	30436 KB	Running
30. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVAgent.exe	1520	20100 KB	Running
31. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVAgent.exe	3800	23600 KB	Running
32. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVInfo.exe	2644	48204 KB	Running
33. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVSvc.exe	3940	37276 KB	Running
34. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVSvc.exe	3096	3072 KB	Running
35. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVSvc.scr	8272	8420 KB	Running
36. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVSvc.exe	1340	23172 KB	Running
37. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVManagementAgent.exe	3152	7236 KB	Running
38. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVSvc.exe	5384	24320 KB	Running
39. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVSvc.exe	2240	13300 KB	Running
40. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVSvc.exe	5416	16002 KB	Running
41. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVSvc_exchange.exe	3972	13068 KB	Running
42. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVSvc.exe	392	14008 KB	Running
43. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVSvc.exe	1360	14808 KB	Running
44. AZQNTY	System	10.4.1.14	Microsoft Windows 2003 Server	AVSvc.exe	4140	14808 KB	Running

2. The **Device Processes** page displays the following about each process:

TIP: To sort the list of processes, click on a column heading. The list will be sorted by the column value, in ascending order. To sort the list by descending order, click the column-heading again.

- **Device Name.** Name of the device where the process resides. For devices running SNMP or with DNS entries, the name is discovered automatically. For devices without SNMP or DNS entries, the device's IP address will appear in this field.
- **Organization.** Organization associated with the device where the process resides.
- **IP Address.** IP address of the device where the process resides.
- **Device Classification / Sub-Class.** The manufacturer (device class) and type of device (sub-class). The Device-Class/Sub-Class is automatically assigned during auto-discovery.
- **Process.** The name of the process. A single process name can have multiple entries.
- **PID.** A unique ID for the process. The device's operating system assigns this value.
- **Memory.** The amount of memory currently used/reserved for the process.
- **Run State.** The current state of the process:
 - *Runnable.* Process is ready to run as needed.
 - *Running.* Process is currently running.
 - *Not Running.* Process is in a "waiting" state.
 - *Invalid.* Process is part of an operation that failed. Process was not ended gracefully.

NOTE: Run states are defined by a device's operating system and/or installed agents. Run states may differ between devices.

- **Monitored.** Specifies whether or not SL1 monitors the process:
 - Yes. SL1 currently monitors this process.
 - No. SL1 does not currently monitor this process.

For more information about filtering the list of device processes on the Device Processes page or about viewing the system processes on a single device, see the **Device Management** manual.

Generating a Report on Multiple System Processes

From the **Device Processes** page (Devices > Processes) you can generate a report on all, multiple, or a single process in SL1.

The report will contain all the columns displayed in the **Device Processes** page (Devices > Processes).

Device Processes Report generated by banderton on 2015-04-17 03:47:25

	Device Name	Organization	IP Address	Device Class	Sub-Class	Process	PID	Memory	Run State	Monitored
0.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	boinc.exe	2140	4952 kB	Running	No
1.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	boincmgr.exe	2888	5860 kB	Running	No
2.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	conhost.exe	2668	115 kB	Running	No
3.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	csrss.exe	296	680 kB	Running	No
4.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	csrss.exe	348	664 kB	Running	No
5.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	csrss.exe	1220	544 kB	Running	No
6.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	dwm.exe	1040	284 kB	Running	No
7.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	explorer.exe	2648	3200 kB	Running	No
8.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	LogonUI.exe	704	6576 kB	Running	No
9.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	lsass.exe	452	5148 kB	Running	No
10.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	lsm.exe	464	1920 kB	Running	No
11.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	msdtc.exe	2432	156 kB	Running	No
12.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	msmdsrv.exe	1080	6320 kB	Running	No
13.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	rdpclip.exe	2084	352 kB	Running	No
14.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	ReportingServicesService.exe	1140	64212 kB	Running	No
15.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	services.exe	444	4760 kB	Running	No
16.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	smss.exe	216	80 kB	Running	No
17.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	snmp.exe	1460	3624 kB	Running	No
18.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	spoolsv.exe	272	1148 kB	Running	No
19.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	spsvc.exe	2496	2992 kB	Running	No
20.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	sqlservr.exe	1052	36984 kB	Running	No
21.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	sqlwriter.exe	1484	88 kB	Running	No
22.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	svchost.exe	552	3072 kB	Running	No
23.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	svchost.exe	624	3628 kB	Running	No
24.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	svchost.exe	712	6388 kB	Running	No
25.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	svchost.exe	764	1992 kB	Running	No
26.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	svchost.exe	804	5296 kB	Running	No
27.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	svchost.exe	844	1176 kB	Running	No
28.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	svchost.exe	884	6140 kB	Running	No
29.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	svchost.exe	980	3496 kB	Running	No
30.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	svchost.exe	1108	80 kB	Running	No
31.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	svchost.exe	1832	2632 kB	Running	No
32.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	svchost.exe	1864	108 kB	Running	No
33.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	svchost.exe	2248	100 kB	Running	No
34.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	System	4	48 kB	Running	No
35.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	System Idle Process	1	24 kB	Running	No
36.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	taskhost.exe	2704	3304 kB	Running	No
37.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	wininit.exe	356	80 kB	Running	No
38.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	winlogon.exe	384	280 kB	Running	No
39.	ACME - DB MSSQL 2 - WebACME		192.168.32.113	Microsoft	MSSQL Server	winlogon.exe	1664	80 kB	Running	No
40.	ACME - DB-MSSQL - WebACME		192.168.32.112	Microsoft	Windows Server 2008 R2	csrss.exe	296	844 kB	Running	No
41.	ACME - DB-MSSQL - WebACME		192.168.32.112	Microsoft	Windows Server 2008 R2	csrss.exe	348	452 kB	Running	No
42.	ACME - DB-MSSQL - WebACME		192.168.32.112	Microsoft	Windows Server 2008 R2	csrss.exe	1676	564 kB	Running	No
43.	ACME - DB-MSSQL - WebACME		192.168.32.112	Microsoft	Windows Server 2008 R2	dwm.exe	2272	512 kB	Running	No
44.	ACME - DB-MSSQL - WebACME		192.168.32.112	Microsoft	Windows Server 2008 R2	explorer.exe	2340	4080 kB	Running	No
45.	ACME - DB-MSSQL - WebACME		192.168.32.112	Microsoft	Windows Server 2008 R2	LogonUI.exe	704	1592 kB	Running	No
46.	ACME - DB-MSSQL - WebACME		192.168.32.112	Microsoft	Windows Server 2008 R2	lsass.exe	452	6460 kB	Running	No
47.	ACME - DB-MSSQL - WebACME		192.168.32.112	Microsoft	Windows Server 2008 R2	lsm.exe	460	2156 kB	Running	No
48.	ACME - DB-MSSQL - WebACME		192.168.32.112	Microsoft	Windows Server 2008 R2	msdtc.exe	1276	1516 kB	Running	No
49.	ACME - DB-MSSQL - WebACME		192.168.32.112	Microsoft	Windows Server 2008 R2	msmdsrv.exe	1128	7260 kB	Running	No
50.	ACME - DB-MSSQL - WebACME		192.168.32.112	Microsoft	Windows Server 2008 R2	Oobe.exe	2472	17408 kB	Running	No
51.	ACME - DB-MSSQL - WebACME		192.168.32.112	Microsoft	Windows Server 2008 R2	rdpclip.exe	536	560 kB	Running	No
52.	ACME - DB-MSSQL - WebACME		192.168.32.112	Microsoft	Windows Server 2008 R2	services.exe	444	5864 kB	Running	No
53.	ACME - DB-MSSQL - WebACME		192.168.32.112	Microsoft	Windows Server 2008 R2	smss.exe	216	316 kB	Running	No
54.	ACME - DB-MSSQL - WebACME		192.168.32.112	Microsoft	Windows Server 2008 R2	snmp.exe	1408	3916 kB	Running	No

Page 1


To generate a report on all or multiple device processes in SL1:

1. Go to the **Device Processes** page (Devices > Processes).
2. In the **Device Processes** page, select the **[Report]** button.

The screenshot shows the 'Device Processes | Processes Found [2834]' interface. At the top right, there are buttons for 'Report', 'Reset', and 'Guide'. The 'Report' button is highlighted with a red box. Below the buttons is a table with columns: Device Name, Organization, IP Address, Device Class | Sub-Class, Process, PID, Memory, Run State, and Monitored. The table lists 25 processes, all from the organization 'ACME' and IP address '192.168.32.113'. The processes include various Microsoft SQL Server related executables like 'boinc.exe', 'boincmgr.exe', 'conhost.exe', 'csrss.exe', 'dism.exe', 'explorer.exe', 'LogonUI.exe', 'iass.exe', 'iim.exe', 'msdts.exe', 'msmdsrv.exe', 'rdpclip.exe', 'ReportingServicesService.exe', 'services.exe', 'sms.exe', 'snmp.exe', 'spoolsv.exe', 'spssvc.exe', 'sqlservr.exe', 'sqlwtrtr.exe', 'svchost.exe', and 'svchost.exe'.

NOTE: If you want to include only certain processes in the report, use the "search as you type" fields at the top of each column. You can filter the list by one or more column headings. You can then select the **[Report]** button, and only the processes displayed in the **Device Processes** page will appear in the report.

3. The **Export current view as a report** modal page appears.



4. In the **Export current view as a report** modal page, you must select the format in which SL1 will generate the report. Your choices are:
 - Comma-separated values (.csv)
 - Web page (.html)
 - OpenDocument Spreadsheet (.ods)

- Excel spreadsheet (.xlsx)
- Acrobat document (.pdf)

5. Click **[Generate]**. The report will contain all the information displayed in the **Device Processes** page. You can immediately view the report or save it to a file for later viewing.

Generating an Exclusion Report for a Single System Process

From the **Device Processes** page (Devices > Processes), you can generate an exclusion report for a process. SL1 will generate the report in MS Word format. An exclusion report specifies all devices where the selected process is running and all devices where the selected process is not running. SL1 lists only appropriate servers in this report. For example, Linux servers would not appear in a report for Windows-based processes.

EM7 TM Management Systems		Windows Service Exclusion Report April 17, 2015, 3:49 am	
Devices That Have [ReportingServicesService.exe] Service Installed			
Device	IP Address	Device Class / Sub-Class	Service
			Run State
			Report Summary
			Total Devices 0
			Unique Device Categories 0
			Unique Device Classes 0
			0
			[on + off]
			Services Found
			Services Not Found 0
			Report Created By ScienceLogic EM7 TM


















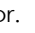







A Process Exclusion Report displays the following:

- Name of the process.
- List of all devices in SL1 where the process is running.
- List of all devices in SL1 where the process is not running. SL1 includes only appropriate servers in this report. For example, Solaris servers would not appear in a report for a Windows 2000 patch.
- The last row in the report displays:
 - Total number of devices in report.
 - Total number of device categories included in the report.
 - Total number of device classes included in the report.
 - Total number of devices where process is running
 - Total number of devices where process is not running.


To generate an exclusion report about a process:

1. Go to the **Device Processes** page (Devices > Processes).

Device Processes | Processes Found [2834] Report Reset Guide

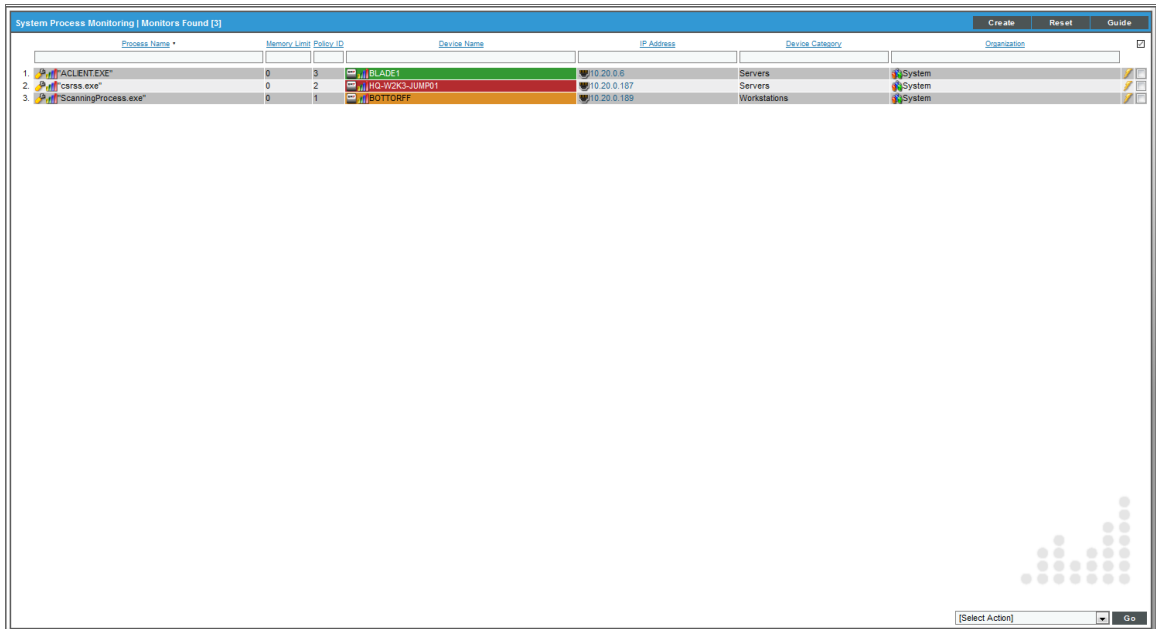
Device Name *	Organization	IP Address	Device Class Sub-Class	Process	PID	Memory	Run State	Monitored
1 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	boinc.exe	2140	4952 kB	Running	No 
2 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	boincmgr.exe	2888	5860 kB	Running	No 
3 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	conhost.exe	2668	116 kB	Running	No 
4 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	csrss.exe	296	680 kB	Running	No 
5 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	csrss.exe	348	664 kB	Running	No 
6 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	csrss.exe	1220	544 kB	Running	No 
7 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	divm.exe	1640	294 kB	Running	No 
8 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	explorer.exe	2648	3200 kB	Running	No 
9 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	LogonUI.exe	704	6576 kB	Running	No 
10 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	lsass.exe	452	5148 kB	Running	No 
11 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	lsim.exe	464	1920 kB	Running	No 
12 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	msdtc.exe	2432	156 kB	Running	No 
13 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	msmsdsv.exe	1080	6320 kB	Running	No 
14 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	rdpclip.exe	2094	352 kB	Running	No 
15 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	ReportingServicesService.exe	1140	64212 kB	Running	No 
16 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	services.exe	444	4760 kB	Running	No 
17 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	smss.exe	216	80 kB	Running	No 
18 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	snmp.exe	1460	3624 kB	Running	No 
19 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	spoolsv.exe	272	1148 kB	Running	No 
20 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	spssvc.exe	2496	2992 kB	Running	No 
21 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	sqlservr.exe	1052	36884 kB	Running	No 
22 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	sqlwiter.exe	1484	88 kB	Running	No 
23 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	svchost.exe	552	3072 kB	Running	No 
24 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	svchost.exe	624	3628 kB	Running	No 
25 ACME - DB MSSQL 2 - WebApp	ACME	192.168.32.113	Microsoft MSSQL Server	svchost.exe	712	6388 kB	Running	No 

[Viewing Page: 1] [Select Action] Go

2. In the **Device Processes** page, find an instance of the process you want to generate an exclusion report for. Select its printer icon ().
3. You will be prompted to save or view the generated report.

Viewing the System Process Monitoring Policies

You can view a list of system process monitoring policies from the **System Process Monitoring** page (Registry > Monitors > System Processes). The **System Process Monitoring** page displays the following information about each system process:



The screenshot shows a web interface titled "System Process Monitoring | Monitors Found (3)". It contains a table with the following columns: Process Name, Memory Limit, Policy ID, Device Name, IP Address, Device Category, and Organization. There are three rows of data:

Process Name	Memory Limit	Policy ID	Device Name	IP Address	Device Category	Organization
1. "ACCLIENT.EXE"	0	3	BLADE1	10.20.0.6	Servers	System
2. "csrss.exe"	0	2	10-WORK-JUMP01	10.20.0.187	Servers	System
3. "ScanningProcess.exe"	0	1	BOTTORFF	10.20.0.189	Workstations	System


- **Process Name.** Name of the policy.
- **Memory Limit.** The maximum amount of memory that can be used or reserved by a single instance of the process, as specified in the process policy.
- **Policy ID.** Unique, numeric ID, assigned to the policy automatically by SL1.
- **Device Name.** Name of the device associated with the policy.
- **IP Address.** IP address of the device associated with the policy. This is the IP address SL1 uses to communicate with the device.
- **Device Category.** Device category of the device associated with the policy.
- **Organization.** Organization for the device associated with the policy.

For more information about filtering the list of system process monitoring policies on the **System Process Monitoring** page, see the **Device Management** manual.


Defining a Monitoring Policy for a System Process

You can define a process monitoring policy in the **System Process Policy** modal page. You can access the **System Process Policy** page either from the **Device Manager** page (Devices > Device Manager) or from the **System Process Monitoring** page (Registry > Monitors > System Processes).

To access the **System Process Policy** modal page from the **Device Manager** page:

1. Go to the **Device Manager** page (Devices > Device Manager)
2. In the **Device Manager** page, find the device that you want to associate with the monitoring policy. Select wrench icon () for the device.
3. In the **Device Administration** panel for the device, select the **[Monitors]** tab.
4. From the **[Create]** menu in the upper right, select **Create System Process Policy**.
5. The **System Process Policy** modal page appears.

To access the **System Process Policy** modal page from the **System Process Monitoring** page:

1. Go to the **System Process Monitoring** page (Registry > Monitors > System Processes).
2. Select the **[Create]** button.
3. Click the device icon () for the device you want to align to policy with.
4. The **System Process Policy** modal page appears.

To define a process monitoring policy in the **System Process Policy** modal page:

1. In the **System Process Policy** modal page, supply a value in each of the following fields:

The screenshot shows the 'System Process Policy' modal page. The window title is 'System Process Policy'. Inside, there's a sub-header 'Editing Policy [1]' with 'New' and 'Reset' buttons. The device name is 'SAC-PATCH-DB-9-26'. The form is divided into two main sections. The left section contains: 'Process Name' (dropdown with '[agetty]'), 'Process Argument (regular expression)' (dropdown with '[\\-\\noclear tty1 linux]'), 'Process User' (text input), 'Alert if Restarted' (dropdown with '[No]'), and 'Alert if Found' (dropdown with '[No]'). The right section contains: 'Memory Limit (Kilobytes per instance)' (text input with '9999999'), 'Total Memory Limit (Kilobytes)' (text input), 'Min / Max Instances' (two text inputs), 'Total CPU Utilization Limit (%)' (text input), and 'State' (dropdown with '[Enabled]'). A 'Save' button is at the bottom center.

- **Process Name.** The name of the process. You can either:
 - Select from a list of all processes running on this device.
 - Click on the "+" icon and manually enter the name of a process.
- **Process Argument (regular expression).** The arguments with which the process is invoked. This field includes a drop-down list of all arguments currently in use by the current device for the specified process (specified in the **Process Name** field). If you don't want to use an argument from the drop-down, you can manually enter a valid regular expression in this field. If you want to include special characters in this regular expression, be sure to escape those special characters. The **Create System Process Policy** modal page will display an error message if the regular expression is not valid. SL1 will match the policy to a process if the value in this field appears anywhere in the argument string for that process. For example "win" would match arguments for "windows" and "win2k".
- **Process User.** Search for the following process user or process owner when the process is running. This field is helpful for finding processes running as root or su which should not be.

NOTE: Some hardware includes information about a process user or owner for each process in the SNMP data; some does not. Do not specify a value in the **Process User** field if the device does not include process user or process owner information in its SNMP data. If you specify a process user, and a device does not include process user in its SNMP data, SL1 will not generate an alert, even if it finds this process running

- **Alert if Restarted.** You can use this field to generate an alert in the Device Log if a system process restarts. Your choices are:
 - Yes. Use this setting to check for system processes that have restarted. SL1 checks every 5 minutes to determine if a system process has restarted. If SL1 finds a restarted system process, it will generate an alert in the Device Log.
 - No. Use this setting if you do not want SL1 to check for system processes that have restarted.

NOTE: When a system process has been restarted, it receives a new process ID number. It might take up to 2 hours for this new ID to appear on the **Process Manager** page (System > Settings > Processes).

NOTE: In some cases, this alert might appear if a device is restarted.

- **Alert if Found.** You can use this field in one of two ways: generate an event when a required system process is not running or generate an event when an illicit system process is running. Your choices are:
 - Yes. Use this setting to look for illicit processes.
 - If SL1 finds the illicit process (specified in the **Process Name** field), SL1 will generate an event.
 - If SL1 does not find the illicit process running, SL1 will not generate an event.
 - No. Use this setting to ensure that a required process is running.
 - If SL1 finds the required (specified in the **Process Name** field) running, SL1 does not generate an event.
 - If SL1 does not find the required process running, SL1 generates an event.
- **Memory Limit (Kilobytes per instance).** The amount of memory, in kilobytes, you will allow each instance of the process to use. This is an optional field.
- **Total Memory Limit (Kilobytes).** This setting is available only if the SL1 agent is installed on the selected device. The amount of memory, in kilobytes, you will allow all instances of the process to use in total. This is an optional field.
- **Min Instances.** The minimum number of instances of the process that should be running. If the minimum instances are not running, SL1 generates an event. The event will be of severity "major" and will say "too few processes running."
- **Max Instances.** The maximum number of instances of the process you will allow to run. If the maximum number of instances is exceeded, SL1 generates an event. The event will be of severity "major" and will say "too many processes process running."
- **Total CPU Utilization Limit (%).** This setting is available only if the SL1 agent is installed on the selected device. The amount of overall CPU you will allow all instances of the process to use in total. This is an optional field.
- **State.** Specifies whether SL1 should start collecting data specified in this policy from the device. Choices are:

- *Enabled*. SL1 will collect the data specified in this policy, from the device, at the frequency specified in the **Process Manager** page (System > Settings > Admin Processes) for the **Data Collection: OS Process Check** process.
- *Disabled*. SL1 will not collect the data specified in this policy, from the device, until the **State** field is set to *Enabled*.

2. Click **[Save]**.

NOTE: If you want to change the aligned device, click on the link for **Change Selected Device** before you clicked **[Save]**. After you clicked **[Save]**, you cannot edit the aligned device.

Example System Process Monitoring Policy

The screenshot shows a web-based configuration interface for a 'System Process Policy'. The interface is titled 'System Process Policy' and has a sub-header 'Editing Policy [1] | Click Save to commit changes'. There are 'New' and 'Reset' buttons in the top right. The device being monitored is 'SAC-PATCH-DB-9-26'. The configuration is divided into several sections:

- Process Name:** A dropdown menu with 'crond' selected.
- Process Argument (regular expression):** A dropdown menu with '[\n]' selected.
- Process User:** An empty text input field.
- Alert if Restarted:** A dropdown menu with '[No]' selected.
- Alert if Found:** A dropdown menu with '[No]' selected.
- Memory Limit (Kilobytes per instance):** An empty text input field.
- Total Memory Limit (Kilobytes):** An empty text input field.
- Min / Max Instances:** Two empty text input fields separated by a slash.
- Total CPU Utilization Limit (%):** An empty text input field.
- State:** A dropdown menu with '[Enabled]' selected.

 A 'Save' button is located at the bottom center of the form.

- This policy monitors a system process on the device "em7ao".
- The policy looks for the process "crond".
- If the process is not found running on the device, SL1 generates an event.

Editing a System Process Monitoring Policy

There are two places in SL1 from which you can edit a monitoring policy for a system process:

1. From the **Device Manager** page (Devices > Device Manager):
 - In the **Device Manager** page, find the device that you want to associate with the monitoring policy. Select the wrench icon (🔧) for the device.
 - In the **Device Administration** panel, select the **[Monitors]** tab.
 - In the **Monitoring Policies** page, find the policy you want to edit and select its wrench icon (🔧).

Or:

2. From the **System Process Monitoring** page (Registry > Monitors > System Processes):
 - In the **System Process Monitoring** page, find the policy you want to edit and select its wrench icon (🔧).
3. The **System Process Policy** modal page appears.

The screenshot shows the 'System Process Policy' modal page. The title bar reads 'System Process Policy'. Below the title bar, there is a header bar with the text 'Editing Policy [1] | Click Save to commit changes' and two buttons: 'New' and 'Reset'. The main content area is divided into several sections. At the top, it says 'Device: SAC-PATCH-DB-9-26'. Below this, there are two columns of fields. The left column contains: 'Process Name' (dropdown menu with 'crond' selected), 'Process Argument (regular expression)' (dropdown menu with '[*-n]' selected), 'Process User' (text input field), 'Alert if Restarted' (dropdown menu with '[No]' selected), and 'Alert if Found' (dropdown menu with '[No]' selected). The right column contains: 'Memory Limit (Kilobytes per instance)' (text input field), 'Total Memory Limit (Kilobytes)' (text input field), 'Min / Max Instances' (two text input fields separated by a slash), 'Total CPU Utilization Limit (%)' (text input field), and 'State' (dropdown menu with '[Enabled]' selected). At the bottom center of the modal, there is a 'Save' button.


4. In the **System Process Policy** modal page, you can change the values in one or more of the fields described in the section on [Defining a Monitoring Policy for System Processes](#).
5. To save your changes to the policy, select the **[Save]** button.

Executing a System Process Monitoring Policy

After creating or editing a system process monitoring policy, you can manually execute the policy and view detailed logs of each step during the execution.

NOTE: After you define a system process monitoring policy and enable the policy, SL1 will automatically execute the policy every five minutes. However, you can use the steps in this section to execute the policy immediately and see debug information about the execution of the policy.

To execute a system process monitoring policy:

1. In the **System Process Monitoring** page (Registry > Monitors > System Processes), find the policy you want to run manually.
2. Select the lightning bolt icon () to manually execute the policy.
3. While the policy is executing, SL1 spawns a modal page called **Session Logs**. The **Session Logs** page provides detailed descriptions of each step during the execution. This is very helpful for diagnosing possible problems with a policy.

You can view reports for executed system process monitoring policies. For more information, see the **Device Management** manual.


Deleting a System Process Monitoring Policy

You can delete a system process monitoring policy from the **System Process Monitoring** page. You can delete individual, multiple, or all existing policies. When you delete a system process monitoring policy, SL1 no longer uses the policy to collect data from the aligned device.

To delete a system process policy:

1. Go to the **System Process Monitoring** page (Registry > Monitors > System Processes).
2. In the **System Process Monitoring** page, select the checkbox(es) for each system process policy you want to delete. Click the checkmark icon () to select all of the system process policies.

3. In the **[Select Action]** menu in the bottom right of the page, select *Delete Monitors*.




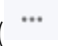
4. Click **[Go]**.
5. The policy is deleted from SL1. The associated reports (from the Device Reports > **[Performance]** tab) are also deleted.

Monitoring Logs Using an Agent

Overview

This chapter describes how to use the agent to monitor logs with Log File Monitoring policies.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon (.
- To view a page containing all of the menu options, click the Advanced menu icon (.

This chapter includes the following topics:

<i>Using a Log File Monitoring Policy</i>	52
<i>Viewing the List of Log File Monitoring Policies</i>	52
<i>Filtering the List of Log File Monitoring Policies</i>	53
<i>Creating a Log File Monitoring Policy</i>	54
<i>Editing a Log File Monitoring Policy</i>	56
<i>Deleting Log File Monitoring Policies</i>	56
<i>Viewing the List of Log File Monitoring Policies and Aligned Devices</i>	57
<i>Filtering the List of Log File Monitoring Policies and Aligned Devices</i>	58
<i>Aligning a Log File Monitoring Policy to Devices</i>	58
<i>Unaligning Log File Monitoring Policies from Devices</i>	61
<i>Creating an Event Policy for Agent Logs</i>	61

Using a Log File Monitoring Policy

A Log File Monitoring policy specifies:

- a file or Windows log on the host device that an agent will monitor
- the logs from the file or Windows log that an agent will send to SL1

You can create, edit, and delete Log File Monitoring policies from the **Log File Monitoring Policies** page. After creating a Log File Monitoring policy, you must align the policy to one or more devices either from the **Log File Monitoring** page or by using a Device Template.

The logs that an agent sends to SL1 are displayed in the **[Logs]** tab in the **Device Administration** and **Device Reports** panels. You can define event policies that specify how logs collected by an agent will trigger events.

Log File Monitoring policies can be included in a PowerPack. For information about including a Log File Monitoring Policy in a PowerPack, see the **PowerPacks** manual.

Viewing the List of Log File Monitoring Policies

The **Log File Monitoring Policies** page (System > Manage > Log File Monitoring Policies) displays a list of all Log File Monitoring policies. From this page, you can also create, edit, and delete Log File Monitoring policies.

Name	Policy ID	Source Type	Source	Filter	Subscribers	Edited By	Last Edited
Kate's test policy	2	Event Log	application	*ERROR	--	em7admin	2016-12-02 20:12:10
Silo	1	File	/var/log/em7/silo.log	*DEBUG.*	3	em7admin	2016-11-16 15:15:42

TIP: To sort the list of Log File Monitoring policies, click on a column heading. The list will be sorted by the column value, in ascending order. To sort by descending order, click the column heading again. The **Last Edited** column sorts by descending order on the first click; to sort by ascending order, click the column heading again.

For each Log File Monitoring Policy, the page displays:

- **Name.** Name of the Log File Monitoring policy.
- **Policy ID.** Unique numeric ID, automatically assigned by SL1 to each Log File Monitoring policy.
- **Source Type.** The source of the logs on the monitored device. Possible values are:
 - *File.* The agent will monitor a file on the file system of the device(s).
 - *Event Log.* The agent will monitor the Windows log on the device(s).
- **Source.** The full path of the log file or the name of the Windows log that the agent will monitor.
- **Filter.** The regular expression that the agent uses to determine whether a log message is sent to SL1.
- **Subscribers.** The number of devices with which the policy is aligned.
- **Edited By.** SL1 user who created or last edited the Log File Monitoring policy.
- **Last Edited.** Date and time the Log File Monitoring policy was created or last edited.

Filtering the List of Log File Monitoring Policies

To filter the list of credentials in the **Log File Monitoring Policies** page, use the search fields at the top of each column. The search fields are find-as-you-type filters; as you type, the page is filtered to match the text in the search field, including partial matches. Text matches are not case-sensitive. Additionally, you can use the following special characters in each filter:

- , (comma). Specifies an "or" operation. For example:
 - "dell, micro" would match all values that contain the string "dell" OR the string "micro".
- & (ampersand). Specifies an "and" operation. For example:
 - "dell & micro" would match all values that contain the string "dell" AND the string "micro".
- ! (exclamation mark). Specifies a "not" operation. For example:
 - "!dell" would match all values that do not contain the string "dell".
- ^ (caret mark). Specifies "starts with." For example:
 - "^ micro" would match all strings that start with "micro", like "microsoft".
 - "^" will include all rows that have a value in the column.
 - "!^" will include all rows that have no value in the column.
- \$ (dollar sign). Specifies "ends with." For example:
 - "\$ware" would match all strings that end with "ware", like "VMware".
 - "\$" will include all rows that have a value in the column.
 - "!\$" will include all rows that have no value in the column.

- min-max. Matches numeric values only. Specifies any value between the minimum value and the maximum value, including the minimum and the maximum. For example:

"1-5" would match 1, 2, 3, 4, and 5.

- - (dash). Matches numeric values only. A "half open" range. Specifies values including the minimum and greater or including the maximum and lesser. For example:

"1-" matches 1 and greater, so it would match 1, 2, 6, 345, etc.

"-5" matches 5 and less, so it would match 5, 3, 1, 0, etc.

- > (greater than). Matches numeric values only. Specifies any value "greater than." For example:

">7" would match all values greater than 7.

- < (less than). Matches numeric values only. Specifies any value "less than." For example:

"<12" would match all values less than 12.

- >= (greater than or equal to). Matches numeric values only. Specifies any value "greater than or equal to." For example:

">=7" would match all values 7 and greater.

- <= (less than or equal to). Matches numeric values only. Specifies any value "less than or equal to." For example:

"<=12" would match all values 12 and less.

- = (equal). Matches numeric values only. For numeric values, allows you to match a negative value. For example:

"=-5" would match "-5" instead of being evaluated as the "half open range" as described above.

Creating a Log File Monitoring Policy

To create a Log File Monitoring policy:

1. Go to the **Log File Monitoring Policies** page (System > Manage > Log File Monitoring Policies).

2. Click **[Create]**. The **Log Monitoring Policy** modal window appears:

Log Monitoring Policy [X]

Create a new policy [Reset]

Name:

Type: [File]

File Path:
The full path of the file from which to capture log messages. You can use a * to match multiple files (eg: /var/log/httpd/*.log).

Limit:
The maximum number of messages to upload per minute from this source. A higher limit means you won't miss log messages, but it may cause a performance hit for uploading the data.

Filter:
Capture log messages which match this string via RegEx. For example, .*ERROR.* to match log entries containing "ERROR"

[Save]

3. Supply values in the following fields:

- **Name**. Enter a name for the policy.
- **Type**. Select the source of the logs on the monitored device. Choices are:
 - *File*. The agent will monitor a file on the file system of the device(s).
 - *Event Log*. The agent will monitor the Windows log on the device(s).
- **File Path**. If you selected *File* in the **Type** field, this field is displayed. Enter the full path of the file to monitor.
- **Source**. If you selected *Event Log* in the **Type** field, this field is displayed. Select the Windows log to monitor. Choices are:
 - *application*
 - *system*
 - *security*
- **Limit**. The maximum log messages the agent sends to SL1 per minute. If the number of matching logs exceeds this value, the agent will stop sending logs to the platform for the remainder of the minute. The limit resets at the beginning of the next minute. For example, suppose you set this field to *10,000*. Suppose the agent monitors a device that has *30,000* log messages. The agent will retrieve *10,000* logs and then wait until the beginning of the next minute. The agent will then retrieve the next *10,000* logs and then wait until the beginning of the next minute. The agent will continue to retrieve *10,000* logs per minute until it has retrieved all the logs from the device.


- **Filter**. Specify a regular expression that will be used to evaluate the log messages in the specified file or Windows log. If a log message matches this regular expression, the agent will send that log message to SL1. If a log message does not match this regular expression, the agent will not send that log message to SL1.

NOTE: For Windows event logs, the SL1 agent adds the Event ID to the value in the *Message* portion of the Windows log before applying the value in the **Filter** field. The agent does not apply the value in the **Filter** field to the *Instance ID* or any other property of a Windows event log entry.

4. Click **[Save]**.

Editing a Log File Monitoring Policy

To edit a Log File Monitoring policy:

1. Go to the Log File Monitoring Policies page (System > Manage > Log File Monitoring Policies).
2. Click the wrench icon () for the Log File Monitoring Policy you want to edit. The **Log Monitoring Policy** modal window appears.
3. Edit the value in one or more fields. For a description of each field, see the [Creating a Log File Monitoring Policy](#) section.
4. Click **[Save]**.

Deleting Log File Monitoring Policies

NOTE: Before you delete a Log File Monitoring Policy, you must un-align that policy from all devices. [See Un-aligning Log File Monitoring Policies](#) for more information.

To delete one or more Log File Monitoring policies:

1. Go to the **Log File Monitoring Policies** page (System > Manage > Log File Monitoring Policies).
2. Select the checkboxes for the **Log File Monitoring Policies** you want to delete.
3. In the **Select Action** drop-down list, select *DELETE Log File Monitoring Policies*.
4. Click **[Go]**.

Viewing the List of Log File Monitoring Policies and Aligned Devices

The **Log File Monitoring** page (Registry > Monitors > Logs) displays a list of existing relationships between devices and Log File Monitoring policies. From the **Log File Monitoring** page, you can also align and unalign devices and Log File Monitoring policies.

Name	Device Name	Policy ID	Source Type	Source	Filter	Limit	Edited By	Last Edited
1. Silo	em7ao	1	File	--	--	20000	em7admin	2016-12-04 19:52:57
2. Silo	win-2012-22	1	File	--	--	--	--	2016-11-16 15:17:27
3. Silo	em7gm	1	File	--	--	--	--	2016-11-16 15:17:27

For each aligned Log File Monitoring policy and device, the page displays:

- **Name**. The name of the Log File Monitoring policy.
- **Device Name**. The name of the device aligned to the Log File Monitoring policy.
- **ID**. The unique numeric ID of the Log File Monitoring policy. The ID is automatically assigned by SL1.
- **Source Type**. The source of the logs in the monitored device. The possible values are:
 - File. The agent monitors a file on the file system of the device. Usually, this is used to monitor Linux log files.
 - Event Log. The agent monitors to Windows log on the device.
- **Source**. The full path of the log file or the name of the Windows log that the agent monitors.
- **Filter**. The regular expression the agent uses to determine if a log should be sent to SL1.

- **Limit.** The maximum log messages the agent sends to SL1 per minute. If the number of matching logs exceeds this value, the agent will stop sending logs to the platform for the remainder of the minute. The limit resets at the beginning of the next minute. For example, suppose you set this field to *10,000*. Suppose the agent monitors a device that has 30,000 log messages. The agent will retrieve 10,000 logs and then wait until the beginning of the next minute. The agent will then retrieve the next 10,000 logs and then wait until the beginning of the next minute. The agent will continue to retrieve 10,000 logs per minute until it has retrieved all the logs from the device.
- **Edited By.** The user who created or last edited the alignment between the device and Log File Monitoring policy.
- **Last Edited.** The date and time the alignment between the device and Log File Monitoring policy was created or last edited.

Filtering the List of Log File Monitoring Policies and Aligned Devices

You can filter the list of Log File Monitoring policies and aligned devices on the **Log File Monitoring** page using the search fields at the top of each column. When you type in each search field, the list of results on the page is automatically updated to match the text, including partial matches.

You can use special characters in each search field to filter. For more information about filtering using special characters, see the [Filtering the List of Log File Monitoring Policies](#) section.

Aligning a Log File Monitoring Policy to Devices

Log File Monitoring policies are aligned to devices either from the **Log File Monitoring** page, or by using a Device Template.

This section describes how to align a Log File Monitoring policy from the **Log File Monitoring** page. It also describes how to use a one-off Device Template to align a Log File Monitoring policy. For more information on Device Templates, including the other methods you can use to create, save, and apply Device Templates, see the **Device Groups and Device Templates** manual.

To align Log File Monitoring policies to one or more devices *from the **Log File Monitoring** page*:

1. Go to the **Log File Monitoring** page (Registry > Monitors > Logs).

2. Click **[Create]**. The Log File Monitor modal page appears.

Log Monitoring Policy [Close]

Create a new Log File monitor [Reset]

Device: ▼
The device on which the log(s) will be monitored

Log Policy: ▼
The log policy definition to be used as the basis for this log monitor

Click on the labels of fields below to enable overrides to the selected policy

File Path:
The full path of the file from which to capture log messages. You can use a * to match multiple files (eg. /var/log/httpd/*.log).

Limit:
The maximum number of messages to upload per minute from this source. A higher limit means you won't miss log messages, but it may cause a performance hit for uploading the data.

Filter:
Capture log messages which match this string via RegEx. For example, *.ERROR.* to match log entires containing "ERROR"

[Save]

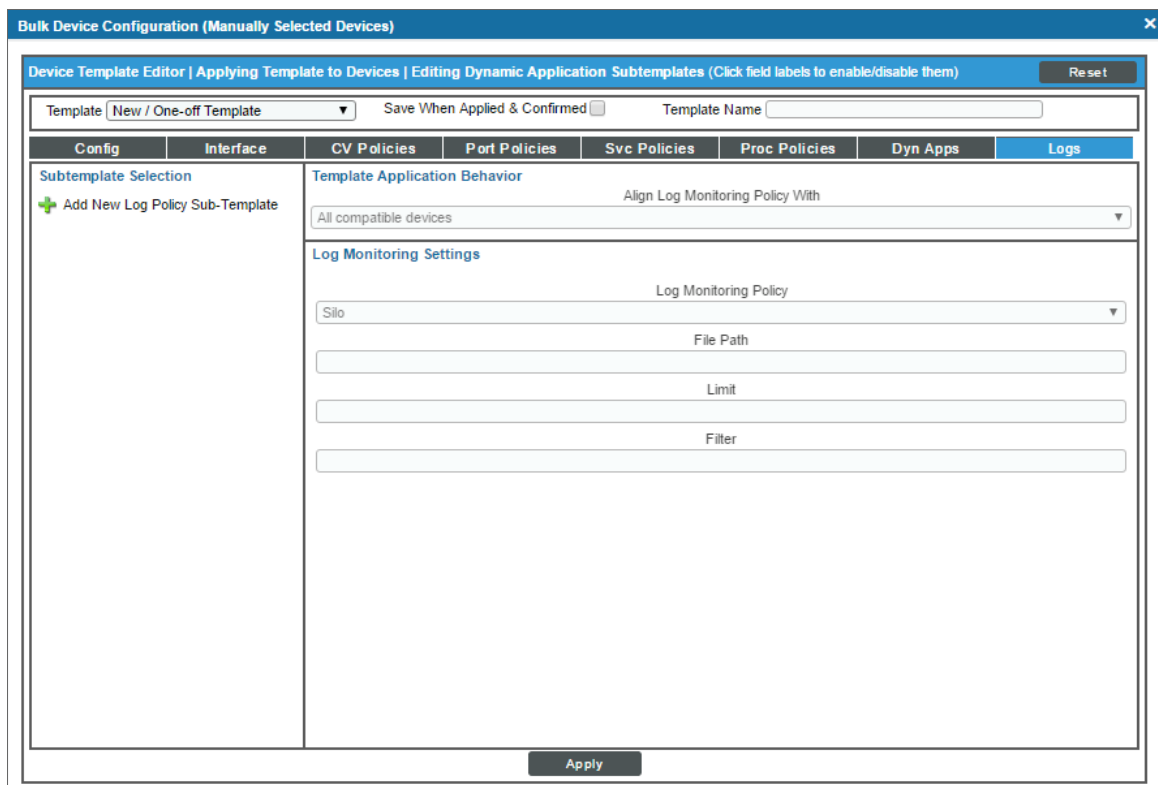
3. In the Log File Monitor modal page, supply values in the following fields:
 - **Device**. Select a device to align with the Log File Monitoring policy.
 - **Log Policy**. Select the Log File Monitoring policy to align with the selected device. Only policies that are appropriate for the selected device will appear. For example, if you chose a Linux device in the **Device** field, the **Log Policy** field will not show policies of the *Event Log* type.
4. If desired, click on the names of the following fields to enable and edit them. These fields allow you to override settings of the policy you selected in the **Log Policy** field for the device selected in the **Device** field:
 - **File Path**. Enter the full file path or the file name to monitor. This field appears only if the type of the policy is *File*.
 - **Limit**. The maximum log messages the agent sends to SL1 per minute. If the number of matching logs exceeds this value, the agent will stop sending logs to the platform for the remainder of the minute. The limit resets at the beginning of the next minute. For example, suppose you set this field to *10,000*. Suppose the agent monitors a device that has 30,000 log messages. The agent will retrieve 10,000 logs and then wait until the beginning of the next minute. The agent will then retrieve the next 10,000 logs and then wait until the beginning of the next minute. The agent will continue to retrieve 10,000 logs per minute until it has retrieved all the logs from the device.

- **File.** Specify a regular expression that will be used to evaluate the log messages in the specified file or Windows log. If and only if a log message matches this regular expression, the agent will send the log message to SL1.

5. Click **[Save]**.

To align Log File Monitoring policies to one or more devices *using a Device Template*:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. Select the checkboxes for the devices with which you want to align Log File Monitoring policies.
3. In the **Select Action** drop-down list, select *MODIFY by Template*.
4. Click **[Go]**. The **Device Template Editor** modal page appears.



5. Click the **[Logs]** tab.

6. Click the Add New Log Policy Sub-Template icon (+).

7. Supply values in the following fields:

- **Align Log Monitoring Policy With.** Select the devices to which the Log File Monitoring policy will be applied.
- **Log Monitoring Policy.** Select the Log File Monitoring policy you want to align with the selected devices.

8. Optionally, you can override one or more settings from the Log File Monitoring policy specifically for the selected devices. To do this, click the field label for each setting you want to override to enable the fields and supply a value in those fields. For a description of each field, see the [Creating a Log File Monitoring Policy](#) section.
9. Repeat steps 6 and 7 for each Log File Monitoring policy you want to align with the devices you selected in step 2.
10. If you want to save this Device Template for future use, select the **Save When Applied & Confirmed** checkbox and enter a name for the Device Template in the **Template Name** field.
11. Click **[Apply]**. The **Setting Confirmation** page is displayed.
12. Click **[Confirm]**. The aligned Log File Monitoring policy will appear on the **Log File Monitoring** page (Registry > Monitors > Logs).

Unaligning Log File Monitoring Policies from Devices

To delete Log File Monitoring Policies, you must first unalign the policy from any devices. You can unalign a Log File Monitoring policy by from the **Log File Monitoring** page.

To unalign devices from a Log File Monitoring policy:

1. Go to the **Log File Monitoring** page (Registry > Monitors > Logs)
2. Select the devices from which the policy must be unaligned.
3. In the **Select Action** drop-down menu, choose *Delete Log File Monitors*.

NOTE: This does not delete the Log File Monitoring policy.

4. Click **[Go]** to unalign the Log File Monitoring policy from the devices.

Creating an Event Policy for Agent Logs

To trigger events in SL1 based on log messages collected by the agent, you must create an event policy that is associated with a Log File Monitoring policy.

To create an event policy that triggers based on log data collected by the agent:

1. Go to **Event Policy Manager** page (Registry > Events > Event Manager).

- In the **Event Policy Manager** page, click **[Create]**. The **Event Policy Editor** page appears:

The screenshot shows the 'Event Policy Editor | Create New Event Policy' window. It has three tabs: 'Policy', 'Advanced', and 'Suppressions'. The 'Policy' tab is active. The form includes the following fields and controls:

- Event Source:** A dropdown menu with 'Syslog' selected.
- Operational State:** A dropdown menu with '[Enabled]' selected.
- Event Severity:** A dropdown menu with '[Major]' selected, and a checkbox for 'Use Modifier'.
- Policy Name:** A text input field.
- Event Message:** A large text area for defining the event message.
- Policy Description:** A rich text editor with a toolbar containing icons for bold, italic, underline, strikethrough, text color, background color, bulleted list, numbered list, link, unlink, and code.

A 'Save' button is located at the bottom center of the form.

- In the **Event Policy Editor** page and set of tabs, you can define a new event. The **Event Policy Editor** page contains three tabs:
 - Policy.** Define basic parameters for the event.
 - Advanced.** Define pattern-matching for the event and also define event roll-ups and suppressions.
 - Suppressions.** Suppress the event on selected devices. When you suppress an event, you are specifying that, in the future, if this event occurs again on a specific device, the event will not appear in the **Event Console** page or the **Viewing Events** page for the device.
- Supply values in the following fields:
 - Event Source.** Select *ScienceLogic Agent*.
 - Policy Name.** The name of the event. Can be any combination of alphanumeric characters, up to 48 characters in length.
 - Operational State.** Specifies whether event is to be operational or not. Choices are *Enabled* or *Disabled*.

- **Event Message.** The message that appears in the **Event Console** page or the **Viewing Events** page when this event occurs. Can be any combination of alphanumeric characters.
 - You can use regular expressions that represent text from the original log message to create the **Event Message**:
 - **%R.** Indicates a regular expression. Surround the regular expression with %R and %/R. For example, %Rfilename: .*? %/R would search for the first instance of the string "filename:" followed by any number of any characters up to the line break. For details on the regular expression syntax allowed by SL1, see <http://www.python.org/doc/howto/>.
 - You can also use the following variables in the **Event Message** field:
 - **%I** ("eye"). This variable contains the value that matches the **Identifier Pattern** field in the **[Advanced]** tab.
 - **%M.** The full text of the log message that triggered the event will be displayed in **Event Message** field.
 - **%T.** Threshold value from the log file will be displayed in **Event Message** field.
- **Event Severity.** Defines the severity of the event. Choices are:
 - **Healthy.** Healthy Events indicate that a device or condition has returned to a healthy state. Frequently, a healthy event is generated after a problem has been fixed.
 - **Notice.** Notice Events indicate a condition that does not affect service but about which users should be aware.
 - **Minor.** Minor Events indicate a condition that does not currently impair service, but the condition needs to be corrected before it becomes more severe.
 - **Major.** Major Events indicate a condition that is service impacting and requires immediate investigation.
 - **Critical.** Critical Events indicate a condition that can seriously impair or curtail service and require immediate attention (i.e. service or system outages).
- **Use Modifier.** If selected, when the event is triggered, SL1 will check to see if the interface associated with this event has a custom severity modifier. If so, the event will appear in the **Event Console** with that custom severity modifier applied to the severity in the **Event Severity** field. For example, if an interface with an **Event Severity Adjust** setting of Sev -1 triggers an event with an **Event Severity** of **Major** and that event has the **Use Modifier** checkbox selected, the event will appear in the **Event Console** with a severity of **Minor**.
- **Policy Description.** Text that explains what the event means and what possible causes are.

5. Select the **[Advanced]** tab.

6. In the **Log Policy** field, select the Log File Monitoring policy that the agent will use to collect the log message.

7. Enter values in the following fields to specify specific text that must appear in the log message for the event policy to trigger:
 - **First Match String**. A string used to match against the originating log message. To match this event policy, the text of a log message must match the value you enter in this field. Can be any combination of alphanumeric characters. Expression matching in SL1 is case-sensitive.
 - **Second Match String**. A secondary string used to match against the originating log message. To match this event policy, the text of a log message must match the value you enter in this field and the value you entered in the **First Match String** field. This field is optional.

NOTE: The **Match Logic** field specifies whether SL1 should process **First Match String** and **Second Match String** as simple text matches or as regular expressions.

8. Optionally, supply values in the other fields on this page. For more information on the remaining fields, as well as the **[Suppressions]** tab, see the **Events** manual.
9. Click **[Save]**.


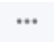
Monitoring Vitals Using an Agent

Overview

This chapter describes using an agent to monitor system vitals, including device availability, CPU utilization, and memory utilization. This chapter also describes how to configure devices to use the agent to collect system vitals.

For more information about monitoring system vitals with SL1, see the **Device Management** manual.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon ().
- To view a page containing all of the menu options, click the Advanced menu icon (.

This chapter includes the following topics:

<i>Viewing System Availability Reports for a Device</i>	66
<i>Changing the Method for Measuring Device Availability</i>	69
<i>Viewing CPU and Memory Utilization for a Device</i>	70
<i>Viewing CPU Utilization</i>	70
<i>Viewing Memory Utilization</i>	72
<i>Changing the Dynamic Application Precedence Settings for CPU and Memory Utilization</i>	74

Viewing System Availability Reports for a Device

The System Availability report displays information about the device's availability. Availability means the device's ability to accept connections and data from the network.

During polling, a device has two possibly availability values:

- 100%. Device is up and running.
- 0%. Device is not accepting connections and data from the network.

By default, the method of discovery determines how the SL1 monitors availability for a device:

- If the agent is installed and creates a device record before the device is discovered as an SNMP or pingable device, availability is measured based on uptime data collected by the agent.
- If the device is discovered as an SNMP or pingable device before the agent is installed, availability is monitored with the method specified in the discovery session (SNMP, ICMP, or TCP).


For devices that SL1 discovers with the discovery tool (Devices > Add Devices button), SL1 determines availability by checking the status of the port specified in the **Availability Port** field in the **Device Properties** page. SL1 collects device-availability data every five minutes, as specified in the process "Data Collection: Availability" (in the **Process Manager** page).

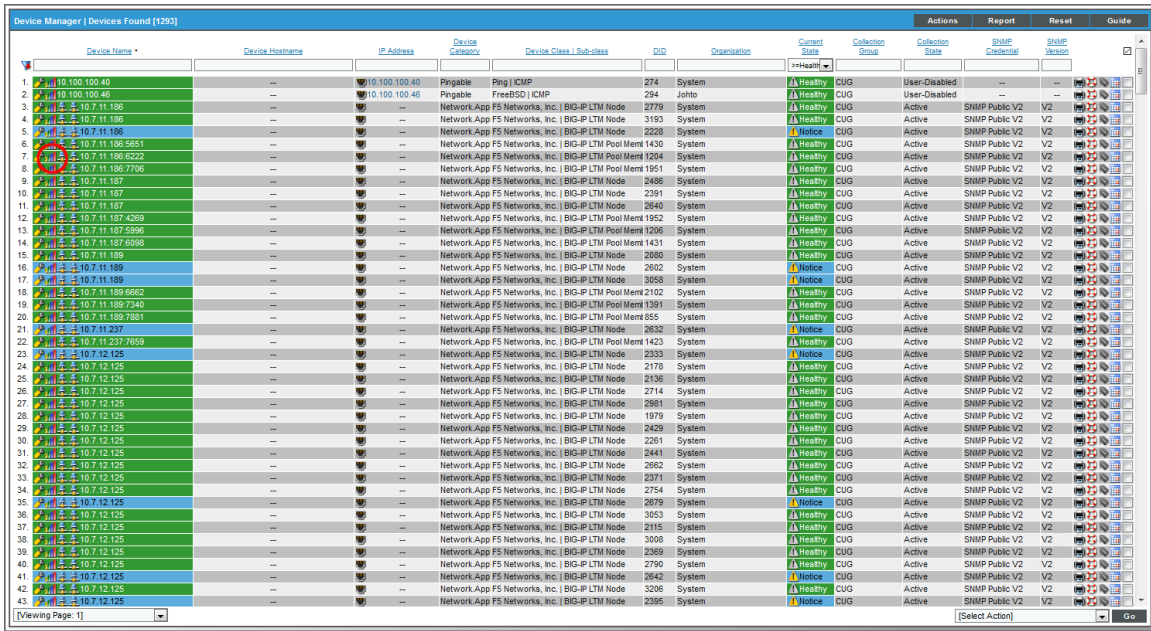
For component devices that SL1 discovers with component mapping Dynamic Applications, SL1 determines availability by checking the status of a collection object.






































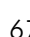



For devices that SL1 discovers with the agent, SL1 collects uptime data from the agent every 5 minutes, and uses this value to determine device availability.

To view the System Availability report for a device:

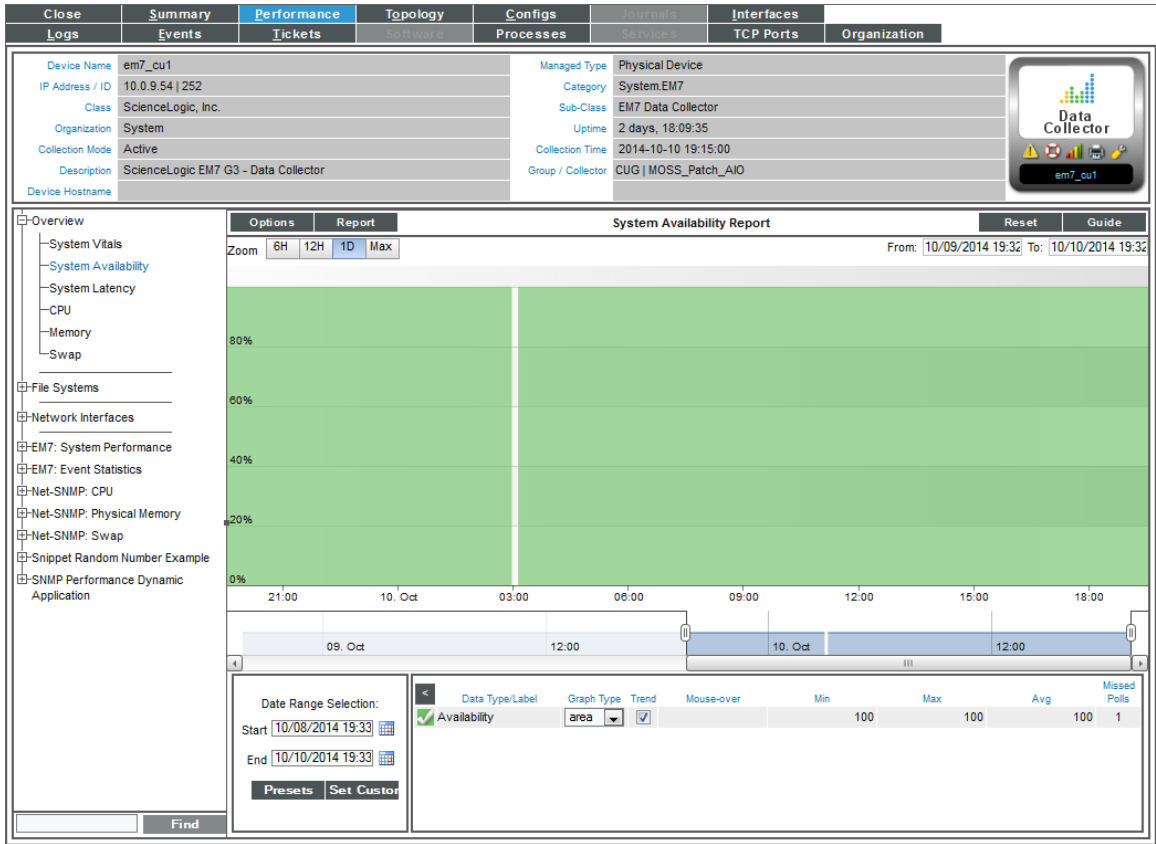
1. Go to the **Device Manager** page (Devices > Device Manager).

2. In the **Device Manager** page, find the device for which you want to view the availability report. Select its bar graph icon () .



Device Name	Device Hostname	IP Address	Device Category	Device Class / Sub-class	D/D	Organization	Current State	Collection Group	Collection Status	SNMP Credentials	SNMP Version	
10.100.100.40	--	10.100.100.40	Pingable	Ping ICMP	274	System	Healthy	CUG	User-Disabled	--	--	
10.100.100.46	--	10.100.100.46	Pingable	FreeBSD ICMP	294	John	Healthy	CUG	User-Disabled	--	--	
10.7.11.186	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2779	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.11.186	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	3193	System	Active	SNMP Public V2	Active	SNMP Public V2	V2	
10.7.11.186	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2228	System	Notice	CUG	Active	SNMP Public V2	V2	
10.7.11.186.5551	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	1430	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.11.186.6222	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Pool Mem	1204	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.11.186.7766	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Pool Mem	1551	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.11.187	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2486	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.11.187	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2391	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.11.187	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2640	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.11.187.4269	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Pool Mem	1952	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.11.187.5996	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Pool Mem	1206	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.11.187.6098	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Pool Mem	1431	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.11.189	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2030	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.11.189	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2002	System	Notice	CUG	Active	SNMP Public V2	V2	
10.7.11.189	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	3058	System	Notice	CUG	Active	SNMP Public V2	V2	
10.7.11.189.6662	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Pool Mem	2102	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.11.189.7340	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Pool Mem	1391	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.11.189.7841	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Pool Mem	655	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.11.192.1	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2632	System	Notice	CUG	Active	SNMP Public V2	V2	
10.7.11.237.7659	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Pool Mem	1423	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2333	System	Notice	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2178	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2136	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2714	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2981	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	1979	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2429	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2261	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2441	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2662	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2271	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2744	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2679	System	Notice	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	3053	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2115	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	3008	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2369	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2790	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2642	System	Notice	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	3206	System	Healthy	CUG	Active	SNMP Public V2	V2	
10.7.12.125	--	--	Network App F5 Networks, Inc.	BIG-IP LTM Node	2395	System	Notice	CUG	Active	SNMP Public V2	V2	

- In the **Device Reports** panel, select the Performance tab.



- In the Performance tab, go to the NavBar (list of links in the left pane), expand the **Overview** link, and select **System Availability**.
- The System Availability report displays system availability for the selected date and time range.
 - The y-axis displays usage, in percent to the left.
 - The x-axis displays time. The increments vary, depending upon the selected data type (from the **[Options]** menu) and the date range (from the **Date Range Selection** pane).
 - Mousing over any point in any line displays (in the **Data Table** pane) the high, low, and average value at the selected time-point.
 - You can use your mouse to scroll the report to the left and right.
 - In a graph of normalized data, clicking on a data point zooms in on that time period and shows the non-normalized data.
- The **[Options]** menu in the upper left of the report displays a menu of options you can apply to data in the current report.
- The **[Reports]** menu in the upper left of the report allows you to export and save the current data and graph as a report. Displays a list of formats for saving the report.

8. The **Data Table** at the bottom of each report allows you to view details about each data point and view information about the entire report. The data table includes the following:
- **Data Type/Label.** For graphs that include multiple types of data on a single graph (for example, availability and latency), each data type has its own row in this table. This column displays the type of data and how it is color coded in the report. Clicking on the check mark toggles on and off the data in the report.
 - **Graph Type.** For selected reports, allows you to specify how you want the data type to be represented in the report. Choices include candlestick, line, stepline, column, area, or stacked. For some reports, the graph type is static and you cannot select a graph type.
 - **Trend.** Toggles on and off a trendline. The trendline shows a bi-directional weighted average, which "smooths" the data for easier consumption. This trending appears as a shaded area superimposed over the graph.
 - **Mouseover.** When you mouseover the graph, this column displays the exact value for each data type at that time point on the graph.
 - **Min.** The column displays the minimum value for the data type in the report.
 - **Max.** This column displays the maximum value for the data type in the report.
 - **Avg.** This column displays the average value for the data type in the report.
 - **Missed Polls.** This column displays the number of times SL1 was unable to collect the data within the time span of the report.

Changing the Method for Measuring Device Availability

By default, discovery determines the method that the SL1 uses to monitor availability of a device:


- If the agent is installed and creates a device record before the device is discovered as an SNMP or pingable device, availability is measured based on whether the agent is reporting data to SL1.
- If the device is discovered as an SNMP or pingable device before the agent is installed, availability is measured based on the method specified in the discovery session (SNMP, ICMP, or TCP).

If a device is monitored using the agent and is discovered as an SNMP or pingable device using the Discovery tool, you can change the method the platform uses to monitor device availability.

To change the method SL1 uses to monitor availability:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).

2. Click the wrench icon (🔧) for the device.

Close	Properties	Thresholds	Collections	Monitors	Schedule					
Logs	Toolbox	Interfaces	Relationships	Tickets	Redirects	Notes	Attributes			
Device Name	em7ao	Managed Type	Physical Device							
IP Address / ID	10.64.68.20 1	Category	System EM7							
Class	ScienceLogic, Inc.	Sub-Class	EM7 All-In-One							
Organization	System	Uptime	0 days, 10:49:54							
Collection Mode	Active	Collection Time	2016-12-02 17:15:09							
Description	ScienceLogic EM7 G3 - All-In-One	Group / Collector	CUG em7ao							
Device Hostname										
Device Properties								Organization	Asset	
								Actions	Reset	Guide
Identification										
Device Name			IP Address			Organization				
em7ao			[10.64.68.20 - verified]			[System]				
Monitoring & Management										
Device Class: ScienceLogic, Inc. EM7 All-In-One										
SNMP Read/Write		[EM7 Default V2]		[None]						
Availability Port		[ScienceLogic Agent]		N/A						
Latency Port		N/A		N/A						
Avail+Latency Alert		[Disable]								
User Maintenance		[Disabled]		[Maintenance Collection Enabled]						
Collection		[Enabled]		[CUG]						
Coll. Type		[Standard]								
Critical Ping		[Disabled]								
Dashboard		None								
Event Mask		[Group in blocks every 10 minutes]								
Save										
Preferences										
Auto-Clear Events <input checked="" type="checkbox"/>										
Accept All Logs <input checked="" type="checkbox"/>										
Daily Port Scans <input checked="" type="checkbox"/>										
Auto-Update <input checked="" type="checkbox"/>										
Scan All IP's <input type="checkbox"/>										
Dynamic Discovery <input checked="" type="checkbox"/>										
Preserve Hostname <input checked="" type="checkbox"/>										
Disable Asset Update <input type="checkbox"/>										
Bypass Interface Inventory <input type="checkbox"/>										

3. In the **Availability Port** field, select the method you want to use to monitor availability:
 - *TCP*. Availability is based on whether the SL1 can connect to the device using the specified TCP port.
 - *ICMP*. Availability is based on whether the device responds to an ICMP ping request from SL1.
 - *SNMP*. Availability is based on whether the device responds to an SNMP get request from SL1.
 - *ScienceLogic Agent*. Availability is based on whether the agent is reporting data to SL1.
4. Click **[Save]**.

Viewing CPU and Memory Utilization for a Device

The agent gathers CPU and memory utilization data for devices.

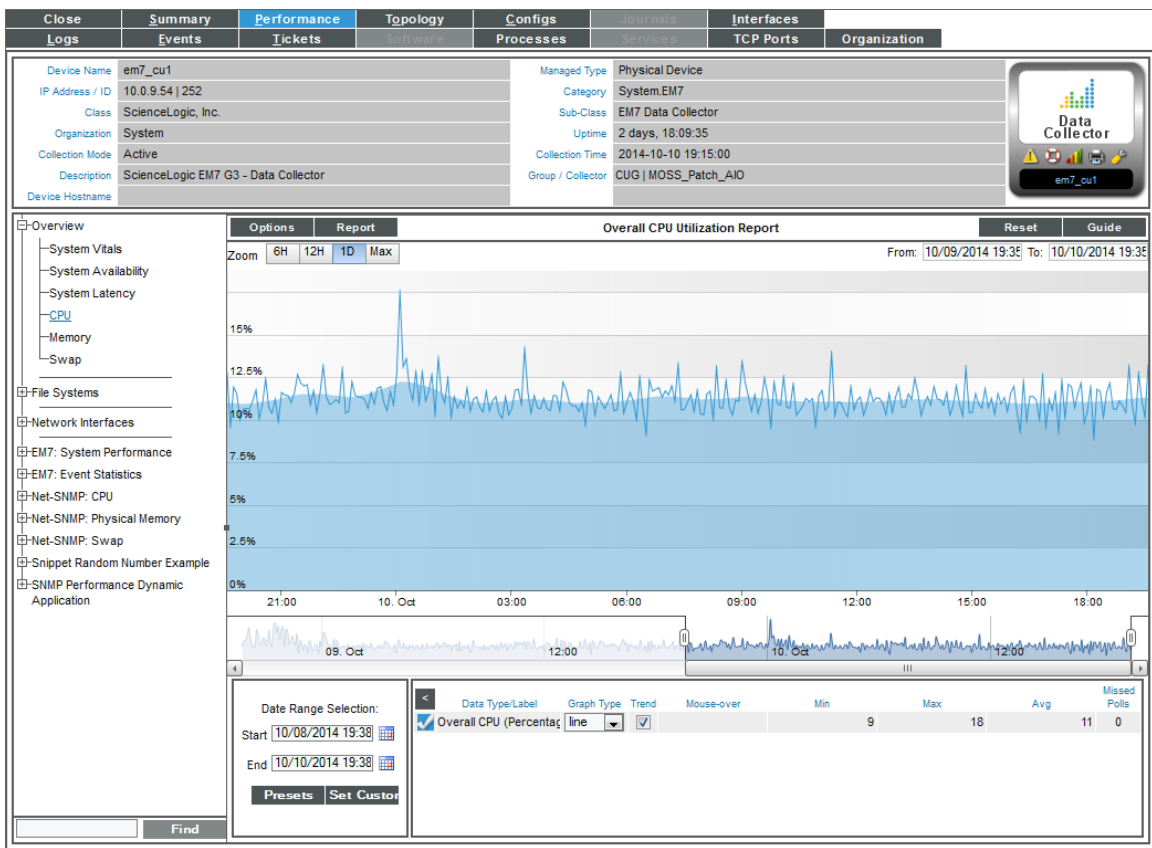
Viewing CPU Utilization

For each device for which SL1 discovered a CPU, you can view a CPU Utilization report.

The CPU Utilization report displays the device's total CPU usage, in percentage. If a device contains multiple CPUs, the report displays the total combined CPU usage, in percent.

To view the CPU Utilization report for a device:

1. You can access the CPU Utilization report from two places:
 - Go to the **Device Manager** page (Devices > Device Manager), find the device where the CPU resides, and select its bar graph icon (📊).
 - Go to the **Device Hardware** page (Devices > Hardware), filter by CPU, find the device where the CPU resides, and select its bar graph icon (📊).
2. When the **Device Reports** panel appears, select the Performance tab.
3. In the **Device Performance** page, go to the NavBar (list of links in the left pane), expand the **Overview** link, and select **CPU Utilization**.




4. The Overall CPU Utilization report displays total CPU usage and average CPU usage over time. If a device contains multiple CPUs, the report displays the total combined CPU usage, in percent, and the combined average CPU usage, in percent. The graph displays CPU usage for the selected date and time range.
 - The y-axis displays usage, in percent to the left.
 - The x-axis displays time. The increments vary, depending upon the selected data type (from the **[Options]** menu) and the date range (from the **Date Range Selection** pane).

- Mousing over any point in any line displays (in the Data Table pane) the high, low, and average value at the select time-point.
 - You can use your mouse to scroll the report to the left and right.
 - In a graph of normalized data, clicking on a data point zooms in on that time period and shows the non-normalized data.
5. The **[Options]** menu in the upper left of the report displays a menu of options you can apply to data in the current report.
 6. The **[Reports]** menu in the upper left of the report allows you to export and save the current data and graph as a report, and displays a list of formats for saving the report.
 7. The Data Table at the bottom of each report allows you to view details about each data point and view information about the entire report. The data table includes the following:
 - **Data Type/Label.** For graphs that include multiple types of data on a single graph (for example, availability and latency), each data type has its own row in this table. This column displays the type of data and how it is color coded in the report. Clicking on the checkmark toggles on and off the data in the report.
 - **Graph Type.** For selected reports, allows you to specify how you want the data type to be represented in the report. Choices include candlestick, line, stepline, column, area, or stacked. For some reports, the graph type is static and you cannot select a graph type.
 - **Trend.** Toggles on and off a trendline. The trendline shows a bi-directional weighted average, which "smooths" the data for easier consumption. This trending appears as a shaded area superimposed over the graph.
 - **Mouseover.** When you mouseover the graph, this column displays the exact value for each data type at that time point on the graph.
 - **Min.** This column displays the minimum value for the data type in the report.
 - **Max.** This column displays the maximum value for the data type in the report.
 - **Avg.** This column displays the average value for the data type in the report.
 - **Missed Polls.** This column displays the number of times SL1 was unable to collect the data within the time span of the report.


Viewing Memory Utilization

You can view an Overall Memory Utilization report for each device for which SL1 has discovered physical memory. The Overall Memory Utilization Report displays total memory usage and average memory usage over time.

To view the Overall Memory Utilization report for a device:

1. You can access the Memory Utilization report from two places:
 - Go to the **Device Manager** page (Devices > Device Manager), find the device where the memory resides, and select its bar graph icon ().

- Go to the **Device Hardware** page (Devices > Hardware), filter by CPU, find the device where the memory resides, and select its bar graph icon (📊).
2. When the **Device Reports** panel appears, select the Performance tab.
 3. In the **Device Performance** page, go to the NavBar (list of links in the left pane), expand the **Overview** link, and select **Memory Utilization**



4. The Overall Memory Utilization report displays total memory usage and average memory usage over time. The graph displays memory usage for the selected date and time range.
 - The y-axis displays memory usage, in percent, to the left.
 - The x-axis displays time. The increments vary, depending upon the selected data type (from the **[Options]** menu) and the date range (from the **Date Range Selection** pane).
 - If the report includes both physical memory and virtual memory, each is represented by a color-coded stack and color-coded line on the graph.
 - The line graph represents actual usage and the stack represents average usage.
 - Mousing over any point in any line (in the Data Table pane) displays the high, low, and average value at the selected time-point.
 - You can use your mouse to scroll the report to the left and right.

- In a graph of normalized data, clicking on a data point zooms in on that time period and shows the non-normalized data.
5. The **[Options]** menu in the upper left of the report displays a menu of options you can apply to data in the current report.
 6. The **[Reports]** menu in the upper left of the report allows you to export and save the current data and graph as a report, and displays a list of formats for saving the report.
 7. The Data Table at the bottom of each report allows you to view details about each data point and view information about the entire report. The data table includes the following:
 - **Data Type/Label.** For graphs that include multiple types of data on a single graph (for example, availability and latency), each data type has its own row in this table. This column displays the type of data and how it is color coded in the report. Clicking on the checkmark toggles on and off the data in the report.
 - **Graph Type.** For selected reports, allows you to specify how you want the data type to be represented in the report. Choices include candlestick, line, stepline, column, area, or stacked. For some reports, the graph type is static and you cannot select a graph type.
 - **Trend.** Toggles on and off a trendline. The trendline shows a bi-directional weighted average, which "smooths" the data for easier consumption. This trending appears as a shaded area superimposed over the graph.
 - **Mouseover.** When you mouseover the graph, this column displays the exact value for each data type at that time point on the graph.
 - **Min.** The column displays the minimum value for the data type in the report.
 - **Max.** This column displays the maximum value for the data type in the report.
 - **Avg.** This column displays the average value for the data type in the report.
 - **Missed Polls.** This column displays the number of times SL1 was unable to collect the data within the time span of the report.

Changing the Dynamic Application Precedence Settings for CPU and Memory Utilization

SL1 collects CPU and memory utilization metrics using Dynamic Applications. If a SNMP device is monitored using the agent, multiple Dynamic Applications can collect CPU and memory utilization metrics. When multiple Dynamic Applications collect CPU and/or memory utilization for a device, SL1 evaluates precedence settings to determine which Dynamic Application will be used to represent CPU and memory utilization for that device.


By default, the precedence settings are configured so the Dynamic Applications that poll the device (using methods other than the agent) represent CPU and memory utilization for that device.

You can change the precedence settings so the Dynamic Applications that use data collected by the agent represent CPU and memory utilization:

- For all applicable devices discovered in the future
- Per-device

To change the precedence settings for all applicable devices discovered in the future:

1. Go to the **Collection Labels** page (System > Manage > Collection Labels).
2. The **Collection Labels** page includes entries for CPU Utilization and Memory Utilization. Select the icon in the **Aligned Presentations** column (🔗) for the utilization metric for which you want to adjust precedence. The **Aligned Presentations** page appears.



3. Locate the entry for the **Host Agent: System Perf** Dynamic Application. Select its checkbox.
4. In the **Select Action** drop-down list, select 0 in the *Change Precedence* section.
5. Click **[Go]**.

To change the precedence settings per-device:

1. Go to the **Collection Labels** page (System > Manage > Collection Labels).

- The **Collection Labels** page includes entries for CPU Utilization and Memory Utilization. Select the icon in the **Duplicates** column (📄) for the utilization metric for which you want to adjust precedence. The **Current Duplicates** page appears.

The screenshot shows a window titled 'Current Duplicates | Label: Memory | Records Found [10]'. The table below lists the duplicate records:

ID	Device Name	Organization	Device IP Address	Device Class Sub-class	Dynamic Application Name	Presentation Name	Align ed
1	em7ao	System	10.64.68.20	ScienceLogic, Inc. EM7	Host Agent: System F...	Memory Utilization	Yes
1	em7ao	System	10.64.68.20	ScienceLogic, Inc. EM7	Net-SNMP: Physical M...	Physical Memory UtilizaNo	No
5	em7gm	System	10.64.68.19	ScienceLogic, Inc. EM7	Host Agent: System F...	Memory Utilization	No
5	em7gm	System	10.64.68.19	ScienceLogic, Inc. EM7	Net-SNMP: Physical M...	Physical Memory UtilizaYes	Yes
6	MS12R2-EXCH13.QA-	System	10.1.0.141	Microsoft Windows Serv...	Fortinet: Memory	Physical Memory UtilizaNo	No
6	MS12R2-EXCH13.QA-	System	10.1.0.141	Microsoft Windows Serv...	Citrix NetScaler: Mem...	Memory Usage	No
6	MS12R2-EXCH13.QA-	System	10.1.0.141	Microsoft Windows Serv...	Cisco: WLC Memory	Memory Utilization	No
6	MS12R2-EXCH13.QA-	System	10.1.0.141	Microsoft Windows Serv...	Netscreen: Memory C...	Overall Physical MemorNo	No
6	MS12R2-EXCH13.QA-	System	10.1.0.141	Microsoft Windows Serv...	Juniper: Memory	Total Physical Memory	Yes
6	MS12R2-EXCH13.QA-	System	10.1.0.141	Microsoft Windows Serv...	Net-SNMP: Physical M...	Physical Memory UtilizaNo	No

At the bottom right, a 'Select Action' dropdown menu is open, showing 'Administration:' and 'Align Presentation for Device' as the selected option. A 'Go' button is also visible.


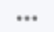
- The **Current Duplicates** page displays multiple rows for each device; each row specifies a device and Dynamic Application metric pair. For each group of rows for a device, use the radio button to the right of the page to select the Dynamic Application metric you want to use for that device.
- In the **Select Action** drop-down list, select *Align Presentation for Device*.
- Click **[Go]**.

Troubleshooting SL1 Agents

Overview

The **Agent Investigator** page on the **Agents** page (Devices > Agents) is created when you add a Linux or Windows agent to SL1. The **Agent Investigator** page give you access to all of the data associated with that agent, and on that page you can configure the agent settings.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon ().
- To view a page containing all of the menu options, click the Advanced menu icon ().

This chapter includes the following topics:

<i>Determine if the Agent Process is Running</i>	79
<i>Determine if the Agent Configuration is Valid</i>	80
<i>Determine if the Agent is Able to Upload Data</i>	80
<i>Check the Agent Upload Directory</i>	80
<i>Run the Agent in Debug Mode (Linux)</i>	81
<i>Determine if SL1 is Receiving Agent Data</i>	82
<i>Determine if SL1 Cannot Process Agent Data</i>	82
<i>Determine if the Number of Processes is Inconsistent with Other Applications</i>	83
<i>Troubleshooting Examples</i>	83
Example /var/log/insight/streamer.log for successful discovery	83
Example /var/log/uwsgi/streamer.log for successful discovery in streamer_prime	84
Save incoming data for a specific device ID (streamer_prime)	84

Save incoming data for a specific device ID (Converged Platform or SL1)	84
<i>Additional Troubleshooting Situations and Best Practices</i>	84

Determine if the Agent Process is Running

To determine if the agent process is running:

1. Check the Windows Task Manager or run the "tasklist" or "top" command, and look for **SiloAgent.exe** (Windows) or **scilogd** (Linux).
2. If **SiloAgent.exe** is not running, check the "Application" event log for events with `source=SiloAgent`.
3. If **scilogd** is not running, check `/var/log/messages` or `/var/log/syslog` for relevant log messages.

If you are using the new user interface for SL1 or the converged platform for the agent, determine if the agent was deleted from the **[Agents]** tab instead of uninstalling the agent.

If the agent was deleted, SL1 shuts down the agent instead of uninstalling the agent. You should re-install the agent that you deleted in the new user interface.

To re-install the agent that was shut down:

1. Uninstall the agent that you shut down.
2. Delete that agent's configuration from one of the following locations:
 - Windows: **C:\Program Files\ScienceLogic\SiloAgent\conf\scilog.conf**
 - Linux: **/etc/scilogd/scilog.conf**
3. Install a new agent.

If the agent was not deleted, then the issue could be with the agent. You should generate diagnostics information to share with your ScienceLogic contact.

To generate diagnostics information for an agent:

1. From an administrator command prompt, run one of the following commands:
 - Windows: `C:\Program Files\ScienceLogic\SiloAgent\bin\SiloAgent.exe -diag`
 - Linux: `/usr/bin/scilogd --diag`
2. Share the contents of the newly created diagnostic file in the current directory with your ScienceLogic contact. Depending on your operating system, the file name is:
 - Windows: **scilog-<current date>.diag.tgz**
 - Linux: **sl-diag.tar.gz**

Determine if the Agent Configuration is Valid

1. Check the agent configuration in one of the following locations:
 - Windows: `C:\Program Files\ScienceLogic\SiloAgent\conf\scilog.conf`
 - Linux: `/etc/scilogd/scilog.conf`
2. Check the configuration item **CollectorID**:
 - If there is no **CollectorID** tag, then the agent has not been able to reach the stream or message collector.
 - If the value is 0 or -1, then SL1 discovery has not completed.
 - If **CollectorID** is a GUID similar to `4179b06ef502129c3023a0f8d58f3c37`, then the agent contacted the backend/streamer and "discovery" has completed, and the CollectorID is valid.
3. Check the configuration item **URLfront**, which is where the agent attempts to get the configuration file.
 - Determine if you can ping the **URLfront**.
 - If you are using `streamer_prime`, **URLfront** should be the URL of the message collector. If you are using the new user interface or the converged platform, **URLfront** should be the URL of the streamer container, such as `pod9-streamer0`.
 - If the URL for **URLfront** is not correct, then re-install the agent. See the re-install steps in the previous topic.
 - If the URL for **URLfront** is correct, then determine if you can ping the host portion of **URLfront**.

Determine if the Agent is Able to Upload Data

Check the Agent Upload Directory

Check the upload directory for the agent for directories and files in one of the following locations:

- Windows: `C:\Program Files\ScienceLogic\SiloAgent\data`
- Linux: `opt/scilog/data`

These locations should only contain the cached system file named `_active-scilog.sys.json` (Windows) or `.active-scilog.sys.json` (Linux). You might see other folders or files in this upload directory that are typically transient, and those folders or files should go away within a few seconds.

The agent typically creates a new data folder every 20 seconds, and optionally (depending on configuration) the agent creates log upload files every minute. If there are many items, then the agent is unable to upload.

- If the number of items is decreasing, the agent might have an issue. The agent is slowly catching up, but this situation indicates that a previous issue existed.

- If the number of items continues to increase overall, check the configuration item URL:
 - The URL is the location where the agent attempts to upload files.
 - Determine if the host portion of the URL is reachable. If the host portion is reachable, the name of the oldest item indicates the approximate time of the issue.

NOTE: To prevent consuming the disk with backed-up data, the agent limits the size and count of items in the upload directory.

A procedural note regarding backed-up data:

For a new installation, the agent reaches out to the streamer for a configuration file. If the configuration file can't reach the streamer, the streamer goes into a slow poll mode, waiting for a good configuration file. In the meantime, the streamer does nothing else (it does not generate data or log files). As a result, even though it looks like there is no backup of data files, in reality, there are no data files.

After the streamer receives a valid configuration file:

- After a restart, the agent reaches out to the streamer for a new configuration file.
- If the agent can't reach the streamer, the agent will still generate data files, because it has a valid configuration file from a previous run. In this situation, you will see data files backing up if the streamer is unreachable.

In summary, if you have a valid configuration, you will get data files. If you do not have a backup, streamer can be reached.

Run the Agent in Debug Mode (Linux)

NOTE: You might need to preface the following commands with `sudo` depending on if you are in root-privileged mode.

1. Stop the agent daemon by running the following command:

```
service scilogd stop
```

2. Start the agent from the command line:

```
scilogd -d 2>&1 | tee /tmp/scilogd.log
```

3. Let the agent run for about five minutes.
4. Press **Ctrl+C** and examine the output file.
5. Restart the agent by running the following command:

```
service scilogd start
```

Determine if SL1 is Receiving Agent Data

If you are using streamer_prime:

1. SSH into the message collector and run the following command:

```
sudo tail -n 100 /var/log/uwsgi/streamer_prime_uwsgi
```

2. Look for lines starting with the IP of the server with the agent on it, such as the following:

```
10.2.16.40 - - [19/Apr/2018:17:04:55 +0000] "POST /SaveData.py/save_data HTTP/1.1"
200 59 "-" "Windows SiloAgent : aym-win2012r2-0"
```

3. If there are no matching lines, then the streamer is not getting data from that agent.

If you are using the new user interface or the converged platform:

1. Either SSH into the Compute Node or point your instance to the Rancher cluster.

2. Run the following command to view the logs:.

```
kubectl logs -l app=streamer
```

3. Look for lines starting with the IP of the customer's server.

4. If there are no matching lines, then the streamer is not getting data from that agent.

Determine if SL1 Cannot Process Agent Data

Check the SL1 log files:

1. If you are using streamer_prime, locate the following files from the SL1 message collector and provide the files to your ScienceLogic contact:

- **/var/log/uwsgi/streamer_prime_uwsgi.log**
- **/var/log/streamer_prime/streamer_prime.log**

2. If you are using the new user interface for SL1 or the converged platform, run the following command:

```
kubectl logs -l app=streamer | grep error
```

- **/var/log/uwsgi/streamer.log**
- **/var/log/insight/streamer.log**

3. Contact your ScienceLogic contact with any error messages you find in the log files. If you do not find any error messages, then the issue is most likely with the Dynamic Application that runs on the collector unit.

Determine if the Number of Processes is Inconsistent with Other Applications

- On Linux, many outputs from the `ps` command list the kernel threads (the processes listed in square brackets). Because the agent is not in the kernel, it will not list kernel threads.
- Be aware that the agent reports processes that are running as well as processes that started and may have stopped, while `top` or `ps` commands show processes that exist when they are executed.
- Check the agent configuration. Due to back-end space limitations, many configuration combinations can limit what data the agent sends. A combination of parameters to get all processes include the following:
 - **NIPD True.** The agent library can not get into all processes at times, often on install. Non-intercepted process discovery reports processes that are not intercepted via the library.
 - **SLPAggregation.** This parameter takes short-lived processes that exist for less than 80 seconds and rolls information about the processes into the information for their parents. As a result, the short-lived processes will not be seen.

Troubleshooting Examples

Example `/var/log/insight/streamer.log` for successful discovery

```
Apr 19 17:50:13 sb-pod IN_STR:146|logger:log_info:132|INFO|Agent config request
received with init flag set to True. Generated new CID:
0a597bc38ae3a15ed96d9310163cba9e. Request: <WSGIRequest: GET
'/api/collector/config/?collector_key=aEf34$aq3TGSDdf&tenant_id=0&host_name=adam-vm-
win7&init=&os=windows&collector_id=0'>

Apr 19 17:50:13 sb-pod IN_STR:115635|logger:log_warning:127|WARNING|Can't check
update version, agent not found in DB: 0a597bc38ae3a15ed96d9310163cba9e

Apr 19 17:50:14 sb-pod IN_STR:115635|logger:log_warning:127|WARNING|System file
received from adam-vm-win7

Apr 19 17:50:14 sb-pod IN_STR:115635|logger:log_info:132|INFO|Agent aid:
0a597bc38ae3a15ed96d9310163cba9e's CID not found; assuming new. CID: adam-vm-win7.

Apr 19 17:50:14 sb-pod IN_STR:115635|logger:log_warning:127|WARNING|New agent is
created as 3: adam-vm-win7 - Windows 7 SP 1 - v109

Apr 19 17:50:14 sb-pod IN_STR:115635|logger:log_info:132|INFO|Agent 3 current pod ID
set to 1

Apr 19 17:50:14 sb-pod IN_STR:115635|logger:log_info:132|INFO|Agent 3: adam-vm-win7
- Windows 7 SP 1 - v109 current (time stamp: 1524160214.0391028) pod ID set to 1

Apr 19 17:50:14 sb-pod IN_STR:115635|logger:log_info:132|INFO|Agent Agent: 3 Pod: 1
current pod ID set to 1

Apr 19 17:50:14 sb-pod IN_STR:115635|logger:log_info:132|INFO|New agent created in
db: 3: adam-vm-win7 - Windows 7 SP 1 - v109
```

```
Apr 19 17:50:14 sb-pod IN_STR:115635|logger:log_warning:127|WARNING|Agent device id does not exists, creating EM7 record for agent: 3
```

```
Apr 19 17:50:14 sb-pod IN_STR:149|logger:log_warning:127|WARNING|System file received from adam-vm-win7
```

Example /var/log/uwsgi/streamer.log for successful discovery in streamer_prime

```
10.234.196.19 - - [29/Sep/2017:14:04:52 +0000] "POST /api/update_agent/agent/HTTP/1.1" 200 59 "-" "python-requests/2.7.0 CPython/2.7.5 Linux/3.10.0-514.10.2.el7.x86_64"
```

Save incoming data for a specific device ID (streamer_prime)

```
PYTHONPATH=/opt/em7/lib/python3:/opt/streamer_prime python3 /opt/streamer_prime/streamer_prime/manage.py agent_save_xml -d <agent guid> -e true
```

Save incoming data for a specific device ID (Converged Platform or SL1)

```
kubectl exec -it $(kubectl get pods -l app=streamer -o jsonpath="{.items[0].metadata.name}") -- python -m streamer agent_save_data --host_id <host id> --enable true
```

You can find the host id from the ADS url, such as https://<sl1_address>/ads/servers/13/system). You can locate the files in the `/tmp/save_agent_data` directory.

Additional Troubleshooting Situations and Best Practices

The following situations might occur while configuring or working with agents:

Situation	Cause / Resolution
Two device records exist in the new user interface for SL1 for the same device.	<p>This situation occurs when the new user interface first discovered this device with SNMP, and then the agent was installed and started polling that device. This duplication of records also occurs if the agent was installed first, and then you ran an SNMP discovery.</p> <p>To address this issue, you can merge the device records using the existing ("classic") user interface. For more information, see the Device Management manual.</p>

Situation	Cause / Resolution
<p>The SNMP device record has IPv4, but the agent device record has IPv6.</p>	<p>The agent reports all network interfaces to the message collector. The message collector uses the first "bound" IP address reported by the agent.</p> <p>To address this issue, you can manually edit the agent device record in the "classic" user interface and update the IP address.</p>
<p>If you uninstall an agent and then run a different installation executable file, you still see the same organization ID for the agent record.</p>	<p>After you uninstall the agent, the scilog.conf file is left on the server in case the agent is reinstalled. The new user interface can reuse the same device record and maintain historical performance data for that agent.</p> <p>To address this issue, delete the file after you run the uninstallation. If you install this agent again, the new user interface assigns a new organization ID to the agent and creates a new device record.</p>

© 2003 - 2019, ScienceLogic, Inc.

All rights reserved.

LIMITATION OF LIABILITY AND GENERAL DISCLAIMER

ALL INFORMATION AVAILABLE IN THIS GUIDE IS PROVIDED "AS IS," WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED. SCIENCELOGIC™ AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT.

Although ScienceLogic™ has attempted to provide accurate information on this Site, information on this Site may contain inadvertent technical inaccuracies or typographical errors, and ScienceLogic™ assumes no responsibility for the accuracy of the information. Information may be changed or updated without notice. ScienceLogic™ may also make improvements and / or changes in the products or services described in this Site at any time without notice.

Copyrights and Trademarks

ScienceLogic, the ScienceLogic logo, and EM7 are trademarks of ScienceLogic, Inc. in the United States, other countries, or both.

Below is a list of trademarks and service marks that should be credited to ScienceLogic, Inc. The ® and ™ symbols reflect the trademark registration status in the U.S. Patent and Trademark Office and may not be appropriate for materials to be distributed outside the United States.

- ScienceLogic™
- EM7™ and em7™
- Simplify IT™
- Dynamic Application™
- Relational Infrastructure Management™

The absence of a product or service name, slogan or logo from this list does not constitute a waiver of ScienceLogic's trademark or other intellectual property rights concerning that name, slogan, or logo.

Please note that laws concerning use of trademarks or product names vary by country. Always consult a local attorney for additional guidance.

Other

If any provision of this agreement shall be unlawful, void, or for any reason unenforceable, then that provision shall be deemed severable from this agreement and shall not affect the validity and enforceability of any remaining provisions. This is the entire agreement between the parties relating to the matters contained herein.

In the U.S. and other jurisdictions, trademark owners have a duty to police the use of their marks. Therefore, if you become aware of any improper use of ScienceLogic Trademarks, including infringement or counterfeiting by third parties, report them to Science Logic's legal department immediately. Report as much detail as possible about the misuse, including the name of the party, contact information, and copies or photographs of the potential misuse to: legal@sciencelogic.com



ScienceLogic

800-SCI-LOGIC (1-800-724-5644)

International: +1-703-354-1010