Monitoring with the ScienceLogic Agent

ScienceLogic version 8.10.0

Table of Contents

Introduction to ScienceLogic Agents 4
What is a ScienceLogic Agent? 5
What Kind of Data Can an Agent Collect? 5
Supported Operating Systems 6
Agent Architecture 6
Installing a ScienceLogic Agent 8
Getting Started 9
Gathering Installation Information from the Device Manager Page 9
Installing an Agent on a Linux System 10
Checking the Version of an Agent on a Linux System 11
Updating an Agent on a Linux System 11
Uninstalling an Agent on a Linux System 11
Installing an Agent on a Windows System 12
Checking the Version of an agent on a Windows System 12
Uninstalling the agent on a Windows System 12
Viewing the Discovered Device 12
Device Classes for Agent-Only Devices 13
Configuring a ScienceLogic Agent 14
Configuring an Agent 15
Adding the "SL Agent" Column to the Device Manager Page 15
Configuring Agent Settings on a Device 16
Changing the Target Message Collector for the Agent 17
Monitoring Ports Using ScienceLogic Agents 19
What is a Port? 20
Port Security 20
Port Availability 20
Viewing a List of All Open Ports on All Devices 21
Viewing a List of All Open Ports on a Single Device 22
System Settings for Monitoring Port Availability 24
Viewing the TCP/IP Port Monitoring Policies 25
Defining a Monitoring Policy for Port Availability 26
Example Policy for TCP/IP Port Availability 28
Editing a Monitoring Policy for a TCP/IP Port 28
Executing a TCP-IP Port Monitoring Policy 29
Deleting a TCP/IP Port Monitoring Policy 30
Monitoring Processes Using ScienceLogic Agents 31
What is a Process? 32
Viewing the List of Device Processes 32
Generating a Report on Multiple System Processes 34
Generating an Exclusion Report for a Single System Process 36
Viewing the System Process Monitoring Policies 38
Defining a Monitoring Policy for a System Process 39
Example System Process Monitoring Policy 42
Editing a System Process Monitoring Policy 43
Executing a System Process Monitoring Policy 44
Deleting a System Process Monitoring Policy 44
Monitoring Logs Using ScienceLogic Agents 46
Using a Log File Monitoring Policy 47
Viewing the List of Log File Monitoring Policies 47
Filtering the List of Log File Monitoring Policies 48

Creating a Log File Monitoring Policy 49
Editing a Log File Monitoring Policy 51
Deleting Log File Monitoring Policies 51
Viewing the List of Log File Monitoring Policies and Aligned Devices 52
Filtering the List of Log File Monitoring Policies and Aligned Devices 53
Aligning a Log File Monitoring Policy to Devices 53
Unaligning Log File Monitoring Policies from Devices 56
Creating an Event Policy for Agent Logs 56
Monitoring Vitals Using ScienceLogic Agents 60
Viewing System Availability Reports for a Device 61
Changing the Method for Measuring Device Availability 64
Viewing CPU and Memory Utilization for a Device 65
Viewing CPU Utilization 65
Viewing Memory Utilization 67
Changing the Dynamic Application Precedence Settings for CPU and Memory Utilization 69
Troubleshooting ScienceLogic Agents 72
Determine if the Agent Process is Running 73
Determine if the Agent Configuration is Valid 74
Determine if the Agent is Able to Upload Data 74
Check the Agent Upload Directory 74
Run the Agent in Debug Mode (Linux) 75
Determine if SL1 is Receiving Agent Data 75
Determine if SL1 Cannot Process Agent Data 76
Determine if the Number of Processes is Inconsistent with Other Applications 77
Troubleshooting Examples 77
Example /var/log/insight/streamer.log for successful discovery 77
Example /var/log/uwsgi/streamer.log for successful discovery in streamer_prime 78
Save incoming data for a specific device ID (streamer_prime) 78
Save incoming data for a specific device ID (Converged Platform or SL1) 78
Additional Troubleshooting Situations and Best Practices 78

Chapter 1

Introduction to ScienceLogic Agents

Overview

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

This chapter includes the following topics:

What is a ScienceLogic Agent? 5

What Kind of Data Can an Agent Collect? 5

Supported Operating Systems 6

Agent Architecture 6

What is a ScienceLogic Agent?

A ***ScienceLogic 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?

A ScienceLogic 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 ScienceLogic 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

  These metrics can be viewed in the **[Performance]** tab in the **Device Reports** panel.
- ***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

  The collected configuration data can be viewed in the **[Configs]** tab in the **Device Reports** panel.
- ***System Processes***. The agent collects a list of all processes running on the device. The list of processes is displayed in the **[Processes]** tab in the **Device Reports** panel. 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. The list of open ports is displayed in the **[TCP/UDP Ports]** tab in the **Device Reports** panel. 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. Logs collected by the agent appear in 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

A ScienceLogic agent collects data from the device on which it is installed and transfers that data to a Message Collector in a 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.

A ScienceLogic agent requires a Message Collector for a distributed architecture.

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.

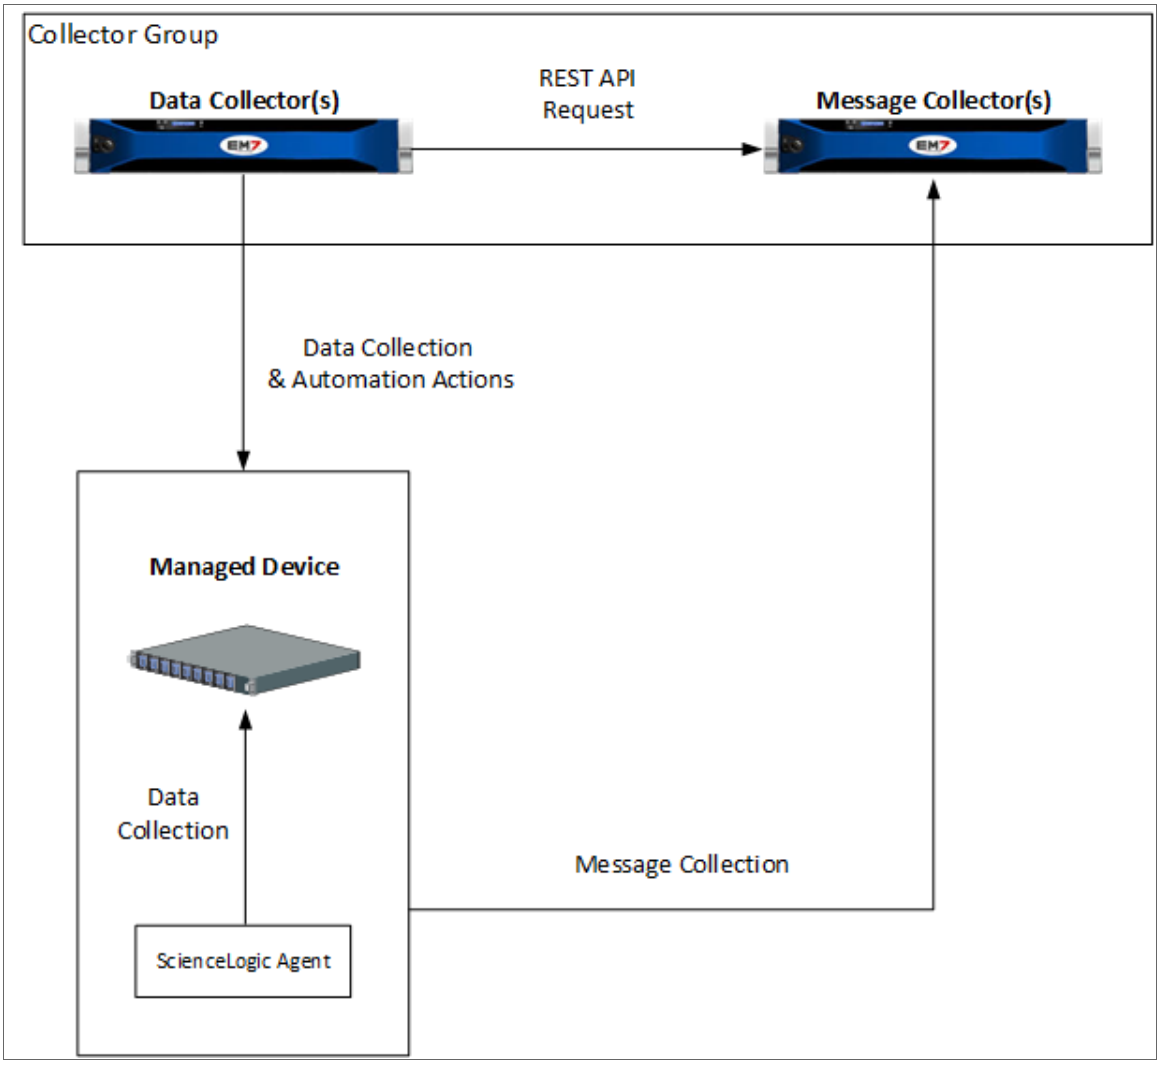

Chapter 2

Installing a ScienceLogic Agent

Overview

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

This chapter includes the following topics:

Getting Started 9

Gathering Installation Information from the Device Manager Page 9

Installing an Agent on a Linux System 10

Checking the Version of an Agent on a Linux System 11

Updating an Agent on a Linux System 11

Uninstalling an Agent on a Linux System 11

Installing an Agent on a Windows System 12

Checking the Version of an agent on a Windows System 12

Uninstalling the agent on a Windows System 12

Viewing the Discovered Device 12

Device Classes for Agent-Only Devices 13

Getting Started

To install a ScienceLogic agent, you must:

- *Gather installation information* from the **Device Manager** page (Registry > Devices > Device Manager). For a Linux system, the **Device Manager** page provides commands that must be executed on the Linux system. For a Windows system, the **Device Manager** page provides an executable file to run on the Windows system.
- Install the agent on the *Linux* or *Windows* device by running the provided commands or executable.

Gathering Installation Information from the Device Manager Page

To gather the necessary commands and/or executable files to install an agent on a device:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. Click **[Actions]** and select *Download/Install Agent*. The **Agent Installation** page appears:

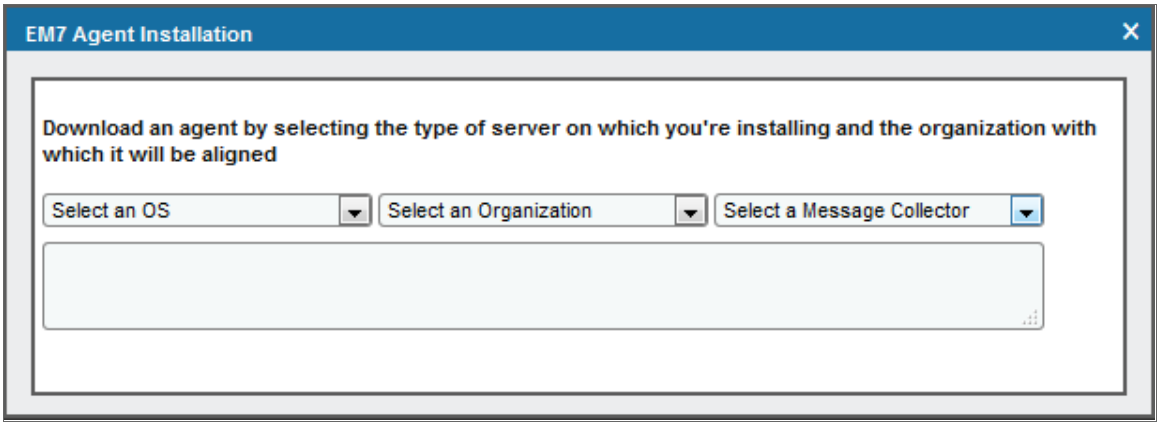

3. Complete the following fields:
   - ***Select an OS***. Select the operating system running on the device on which you want to install the agent.

     NOTE: If you require a FIPS-compliant version of the ScienceLogic agent, select *RedHat/CentOS 64-bit (OS Libs)*.

   - ***Select an Organization***. Select an organization from the list of possible organizations. The list of organizations is dependent on your user account. If the agent discovers a new device, that device will be associated with the organization you select here.

> **NOTE:** If you are installing an agent on a device that has already been discovered, you must select the organization that is already aligned with the existing device.

- ***Select a Message Collector***. Select the Message Collector to which the agent will send its collected data.

4. If you selected a Linux operating system in the ***Select an OS*** field, the **Agent Installation** page displays a list of commands to execute on the Linux system. Copy the commands for use during the *installation on the Linux device*.
5. If you selected a Windows operating system in the ***Select an OS*** field, the **Agent Installation** page displays a ***Download Windows Agent*** link. Click the link and save the executable file for use during the *installation on the Windows device*.

> **TIP:** If you are installing an agent on multiple devices that run the same operating system, are part of the same organization, and connect to the same Message Collector, you can re-use the same commands or executable file on each of those devices.

Installing an Agent on a Linux System

To install 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. Execute the commands that you copied from the **Agent Installation** page in SL1. If the installation was successful, the output will look similar to the following:

```
[em7admin@em7ao ~]$ sudo wget --no-check-certificate
https://10.64.68.16/packages/initial/0/silo-agent-x86_64.rpm
[sudo] password for em7admin:
--2016-11-15 21:10:28-- https://10.64.68.16/packages/initial/0/silo-agent-x86_
64.rpm
Connecting to 10.64.68.16:443... connected.
WARNING: cannot verify 10.64.68.16's certificate, issued by
'/C=US/ST=Silo/L=Reston/O=Silo/CN=10.64.68.16':
Self-signed certificate encountered.
HTTP request sent, awaiting response... 200 OK
Length: 2018317 (1.9M) [application/x-rpm]
Saving to: 'silo-agent-x86_64.rpm'
100%[======================================>] 2,018,317 --.-K/s in 0.01s
2016-11-15 21:10:28 (169 MB/s) - 'silo-agent-x86_64.rpm' saved [2018317/2018317]
[em7admin@em7ao ~]$ sudo rpm -ihv silo-agent-x86_64.rpm
Preparing... ################################# [100%]
Updating / installing...
1:scilogd-0.128-0 ################################# [100%]
Created symlink from /etc/systemd/system/multi-user.target.wants/scilogd.service
to /etc/systemd/system/scilogd.service.
```

Checking the Version of an Agent on a Linux System

To check the version number of 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. Execute the following command:

```
grep Version /var/log/scilogd.log
```

Updating an Agent on a Linux System

To update the agent on a Linux system:

1. Follow the steps listed in the *Gathering Installation Information from the Device Manager Page* section.
2. Log in to the Linux system via the console or SSH as a user that has sudo administrator permissions.
3. Execute the **first** command that you copied from the **Agent Installation** page.
4. Do one of the following:

   - For RedHat-based Linux distros, execute the following command:

```
sudo rpm -Uvh silo-agent-x86_64.rpm
```

   - For Ubuntu-based Linux distros, execute the following command:

```
sudo dpkg -i silo-agent-x86_64.deb
```

Uninstalling an Agent on a Linux System

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 RedHat-based Linux distros, execute the following command:

```
rpm -e scilogd-0.128-0.[ARCH].rpm where [ARCH] = i386 or x86_64
```

   - For Ubuntu-based Linux distros, execute the following command:

```
dpkg --purge silo-agent-[ARCH].deb where [ARCH] = i386 or x86_6
```

3. Remove the agent configuration directory from the Linux system. The configuration directory can be found at:

```
/etc/scilog
```

Installing an Agent on a Windows System

To install an agent on a Windows system:

1. Copy the SiloAgent-install.exe file you downloaded from the **Agent Installation** page to the Windows system. You can go to the console of the Windows system or use a utility like WinSCP.
2. Run the following command as an Administrator:

   ```
   SiloAgent-install.exe tenant=0 urlfront=<URL_for_your_SL1_system>
   ```

3. 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.

Checking the Version of an agent on a Windows System

To check the version number of the agent on a Windows System:

1. On the Windows system, navigate to C:\Program Files\ScienceLogic\SiloAgent\bin in the File Explorer.
2. Right click on the "SiloAgent" file and select *Properties*. The version number is displayed in the ***Product Version*** field.

Uninstalling the agent on a Windows System

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
   ```

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.

Chapter 3

Configuring a ScienceLogic Agent

Overview

This chapter describes how to configure agent settings on a device and the settings on the Message Collector with which the agent communicates.

This chapter includes the following topics:

Configuring an Agent *15*

Adding the "SL Agent" Column to the Device Manager Page *15*

Configuring Agent Settings on a Device *16*

Changing the Target Message Collector for the Agent *17*

Configuring an Agent

You can control how an agent runs on a device by configuring the following agent settings:

> **NOTE:** To configure agent settings, you must first add the ***SL Agent*** column to the **Device Manager** page in the classic user interface. For more information about adding the ***SL Agent*** column, see *Adding the SL Agent Column to the Device Manager Page*.

- ***Disk Space***. Controls the amount of disk space that the agent can use to store data. If an agent loses connectivity to SL1, this disk space will be used to store collected data until the connection to SL1 is restored.
- ***Data Directory***. Defines the directory in which the agent will store temporary data.
- ***Excludes***. Defines the list of processes and directories to explicitly exclude from monitoring by the agent.
- ***Includes***. Defines the list of processes and directories that must be explicitly monitored by the agent. Use the ***Includes*** field to ensure 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.

Adding the "SL Agent" Column to the Device Manager Page

The ***SL Agent*** column allows you to access the configuration settings for the agent on a device. For more information about agent configuration settings, see *Configuring Agent Settings on a Device*. By default, the ***SL Agent*** column is not displayed in the **Device Manager** page (Registry > Devices > Device Manager).

To add the ***SL Agent*** column to the **Device Manager** page:

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

2. Click **[Actions]**, and then select *Device Manager Preferences*. The **Edit Device Manager Preferences** modal page appears:

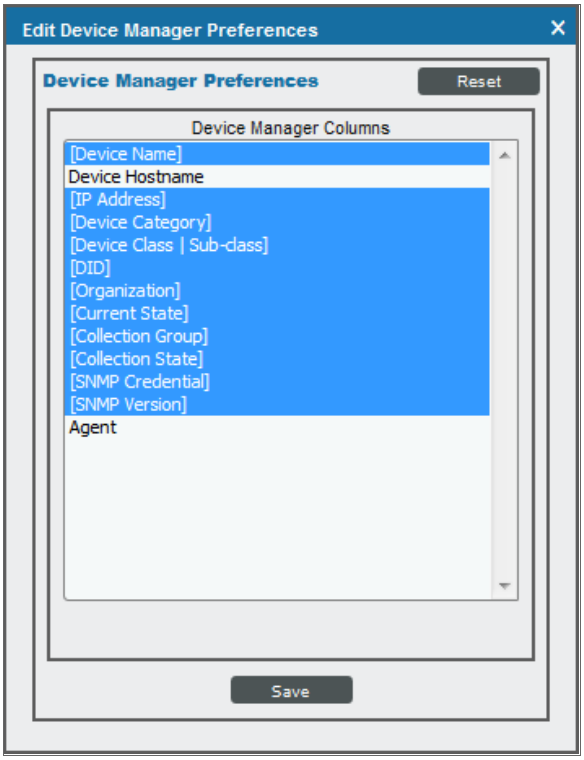

3. In the ***Device Manager Columns*** field, control-click *Agent*.
4. Click **[Save]**.

Configuring Agent Settings on a Device

To configure agent settings, you must first add the ***SL Agent*** column to the **Device Manager** page. For more information about adding the ***SL Agent*** column, see *Adding the SL Agent Column to the Device Manager Page*.

To configure agent settings on a device:

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

2. Find the device for which you want to edit agent settings. In the ***SL Agent*** column, click the gear icon ( ) for the device. The **Agent Configuration** page appears:

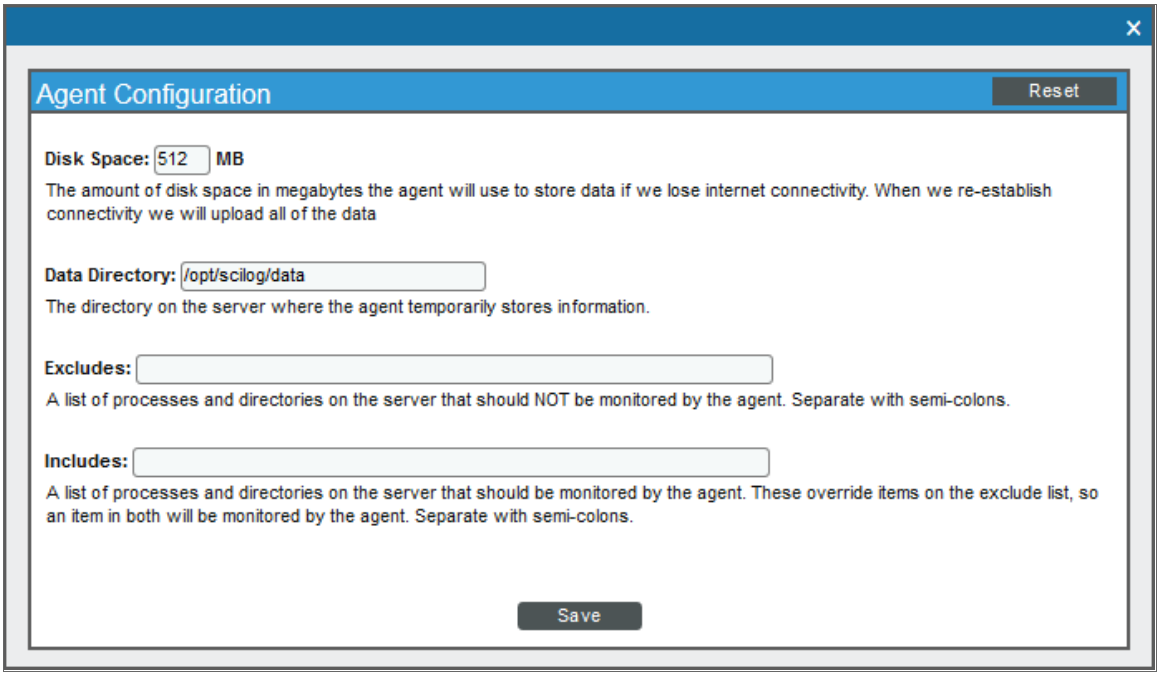

3. Supply values in the following fields:
   - ***Disk Space***. Enter the amount of disk space that the agent can use to store data. If the agent loses connectivity to SL1, this disk space will be used to store collected data until the connection to SL1 is restored.
   - ***Data Directory***. Enter the directory in which the agent will store temporary data.
   - ***Excludes***. Enter a semi-colon delimited list of processes and directories to explicitly exclude from monitoring by the agent.
   - ***Includes***. Enter a semi-colon delimited list of processes and directories that must be monitored by the agent. Use the ***Includes*** field to ensure 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.

4. Click **[Save]**.

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"
   ```

Chapter 4

Monitoring Ports Using ScienceLogic Agents

Overview

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

This chapter includes the following topics:

What is a Port? 20
Port Security 20
Port Availability 20
Viewing a List of All Open Ports on All Devices 21
Viewing a List of All Open Ports on a Single Device 22
System Settings for Monitoring Port Availability 24
Viewing the TCP/IP Port Monitoring Policies 25
Defining a Monitoring Policy for Port Availability 26
Example Policy for TCP/IP Port Availability 28
Editing a Monitoring Policy for a TCP/IP Port 28
Executing a TCP-IP Port Monitoring Policy 29
Deleting a TCP/IP Port Monitoring Policy 30

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 (Registry > Devices > Device Manager > bar-graph icon > Performance) displays a list of all open ports on a device.

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

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

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

For SNMP and pingable devices, the platform scans all the ports of each managed device every day. If any new ports are opened, the platform 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, the platform will not notify you of newly opened ports on the device.

Port Availability

The ScienceLogic platform can monitor ports for availability. When a port monitor is created, the platform monitors the port for availability every five minutes. You can choose whether a policy is executed by the platform 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, the platform 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 the ScienceLogic platform using NMAP and the ScienceLogic 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).

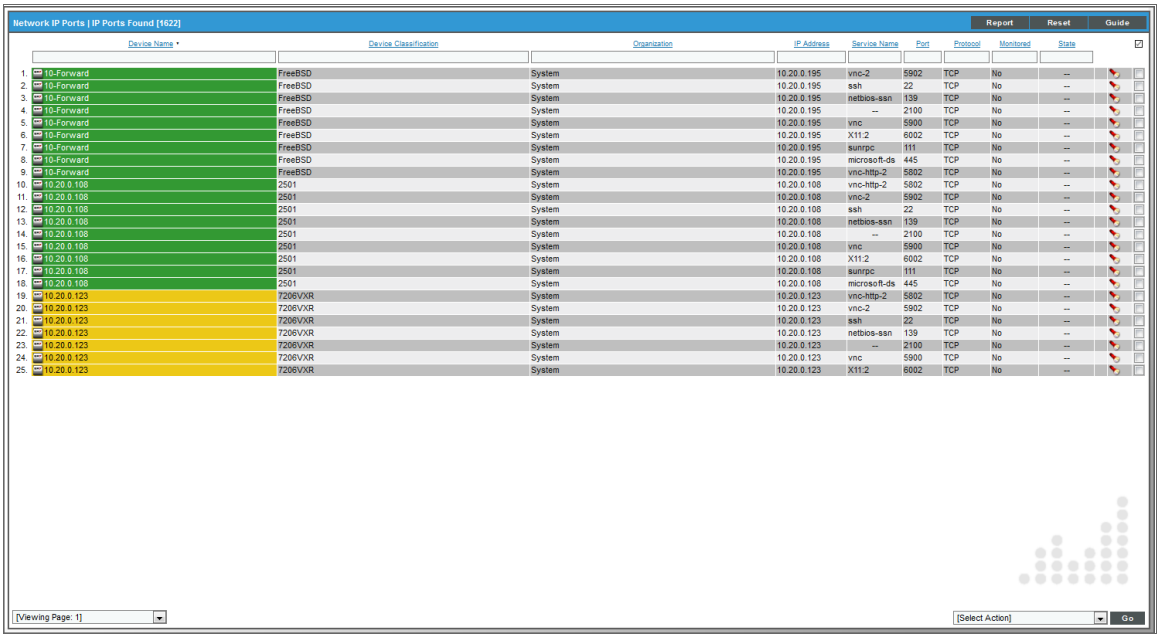

2. The **Network IP Ports** page displays a list of all discovered ports. For each port, the **Network IP Ports** page displays the following:

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 the platform 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. The ScienceLogic platform monitors the port and collects availability data about the port.
    - *Disabled*. The port-monitoring policy has not been activated. The ScienceLogic platform 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 (Registry > 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.

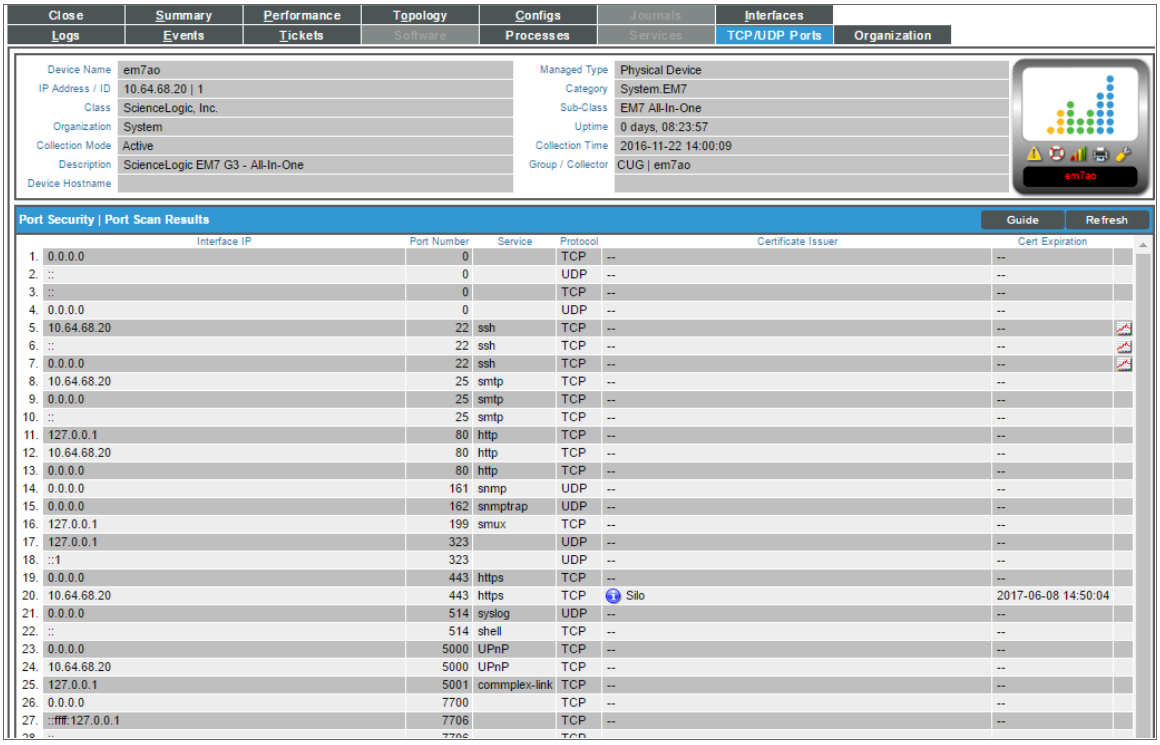

3. For each open port on the device, the **Port Security** page displays the following information:

   - ***Interface IP***. IP address through which the ScienceLogic platform 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:

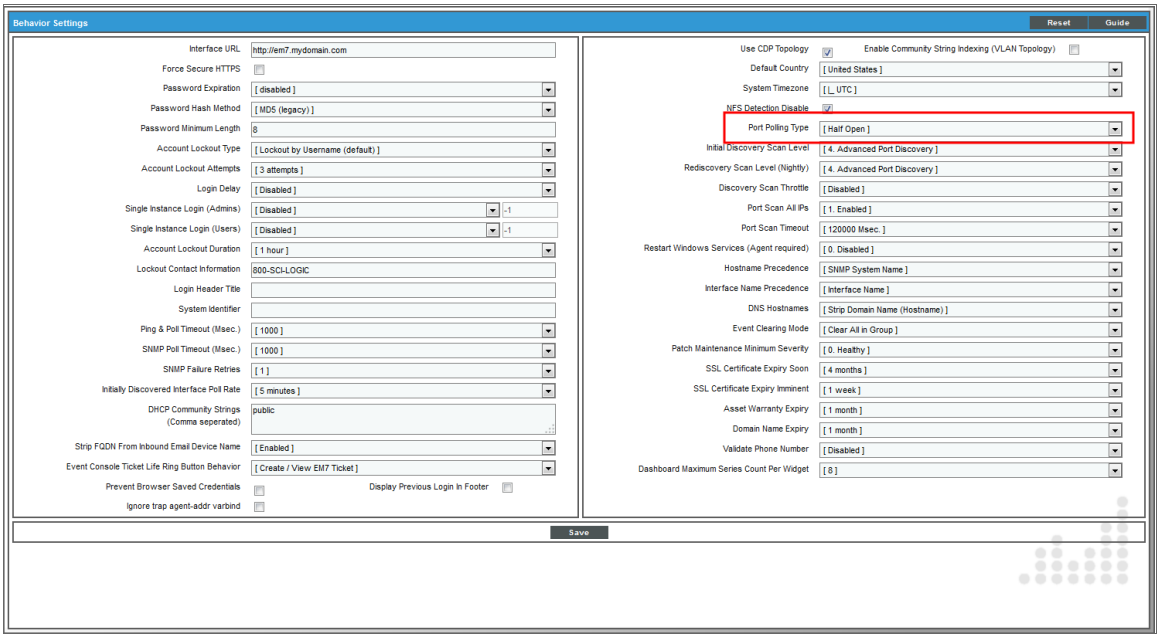

- ***Port Polling Type***. Specifies how the ScienceLogic platform should poll ports for availability using NMAP. The choices are:
    - *Half Open*. Uses a faster TCP/IP connection method (a TCP SYN scan, nmap -sS) and does not appear on device's logs.
    - *Full Connect*. Uses the standard TCP/IP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

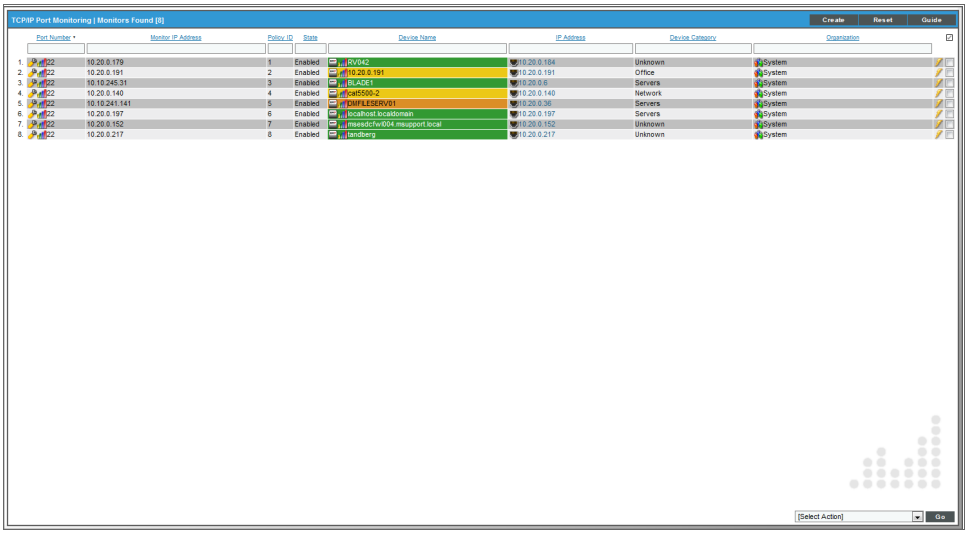

- ***TCP/IP Port Number***. Port number of the port to be monitored.
- ***Monitor IP Address***. IP address associated with the port to be monitored. For devices with multiple IP addresses, the IP address for the port policy might be different than the IP address used by the ScienceLogic platform to communicate with the device.
- ***Policy ID***. Unique, numeric ID, assigned to the policy automatically by the platform.
- ***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 the platform 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 (Registry > 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 (Registry > 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:

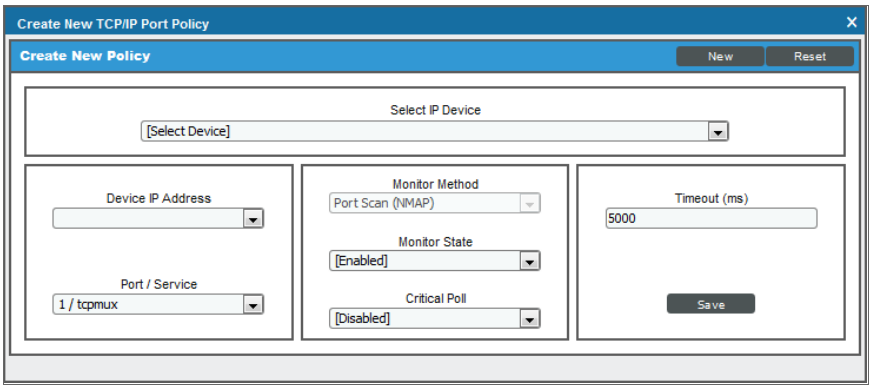

- ***Select Device***. Select a device from this drop-down list to align with this policy. By default, the current device is selected in this field.
- ***Device IP Address***. IP address through which the ScienceLogic platform communicates with the device.
- ***Port/Service***. Port number and the corresponding 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 the platform should start collecting data specified in this policy from the device. Choices are:
  - *Enabled*. The platform 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*. The platform 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 the platform should "ping" the device. If the device does not respond, the platform creates an event. The choices are:
  - *Disabled*. The platform will not ping the device.
  - *Enabled*. The platform will ping the device every 15, 30, 60, or 120 seconds, as specified.

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

3. Click **[Save]**.

Example Policy for TCP/IP Port Availability

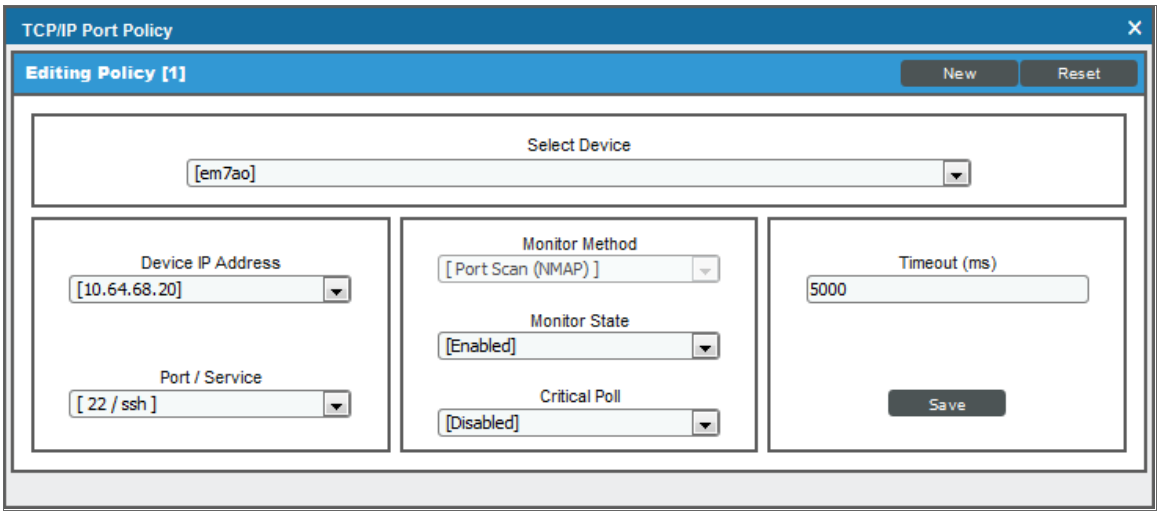

- This policy monitors a TCP/IP port on the device "cisco_10.2.1.29", at IP address 10.1.0.205.
- The policy will monitor port 22 for availability.

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 (Registry > 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 (Registry > 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.

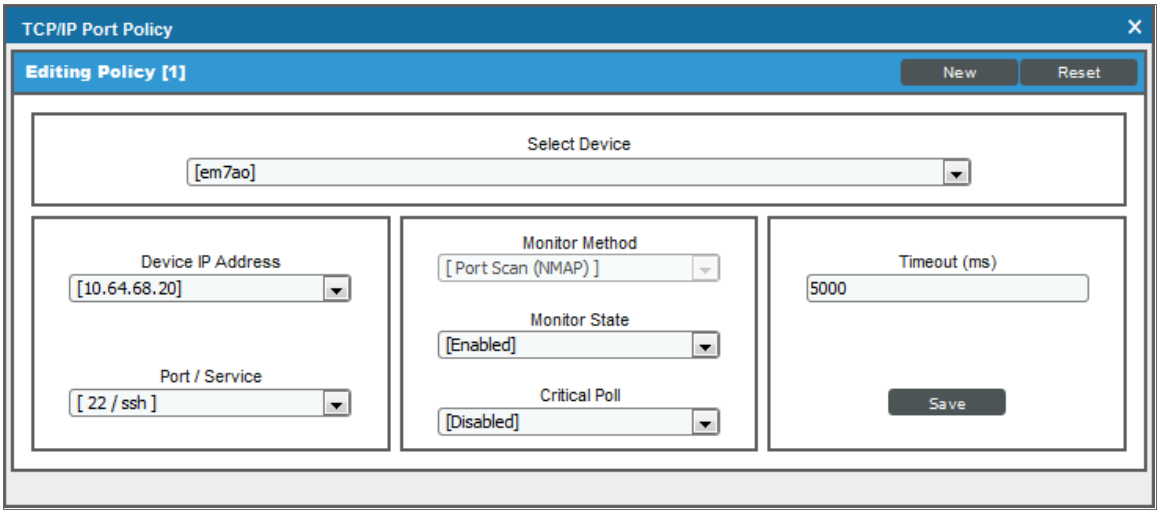

3. Click **[Save]** when done.

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, the ScienceLogic platform or the ScienceLogic 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, the ScienceLogic platform 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, the ScienceLogic platform 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*.

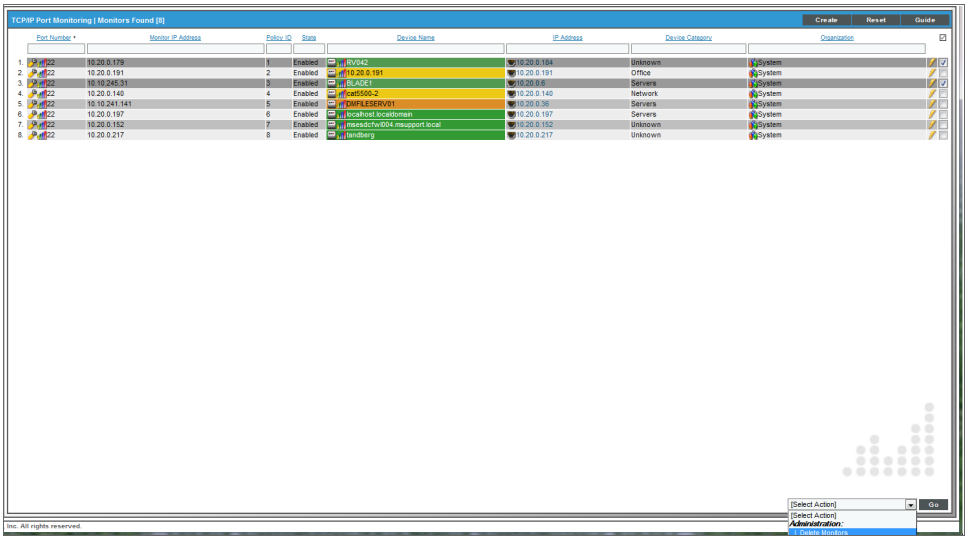

4. Click **[Go]** to delete the port monitoring policy.
5. The policy is deleted from the ScienceLogic platform. The associated reports (from the Device Reports > **[Performance]** tab) are also deleted.

Chapter 5

Monitoring Processes Using ScienceLogic Agents

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.

This chapter includes the following topics:

What is a Process? 32
Viewing the List of Device Processes 32
Generating a Report on Multiple System Processes 34
Generating an Exclusion Report for a Single System Process 36
Viewing the System Process Monitoring Policies 38
Defining a Monitoring Policy for a System Process 39
Example System Process Monitoring Policy 42
Editing a System Process Monitoring Policy 43
Executing a System Process Monitoring Policy 44
Deleting a System Process Monitoring Policy 44

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

The ScienceLogic platform allows you to create policies that monitor system processes every five minutes:

- If a device is not monitored using the ScienceLogic agent, the policy collection is performed using SNMP.
- If a device is monitored using the ScienceLogic 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 the ScienceLogic platform on all devices.

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

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

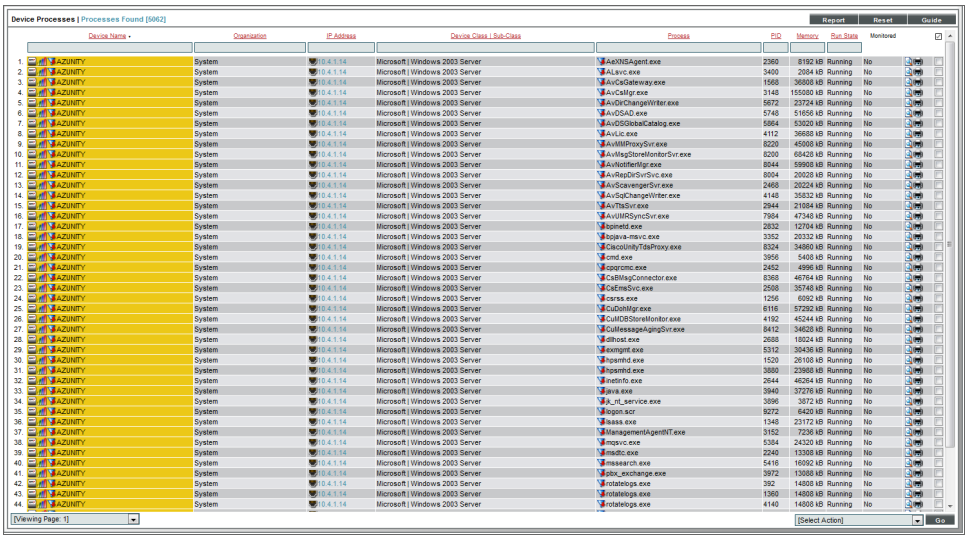

2. The **Device Processes** page displays the following about each process:

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 the ScienceLogic platform monitors the process:
  - *Yes*. The ScienceLogic platform currently monitors this process.
  - *No*. The ScienceLogic platform 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 (Registry > Devices > Processes) you can generate a report on all, multiple, or a single process in the ScienceLogic platform.

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

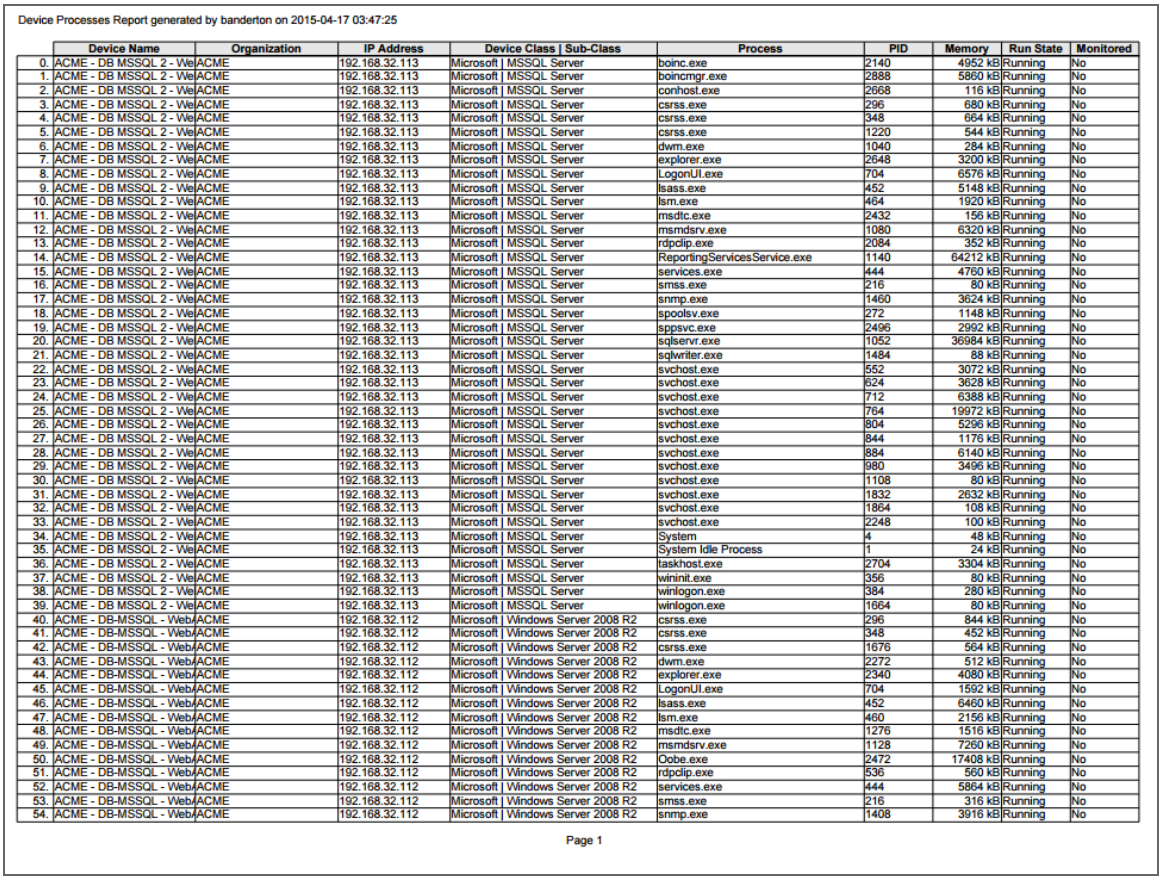

Device Processes Report generated by banderton on 2015-04-17 03:47:25

	Device Name	Organization	IP Address	Device Class | Sub-Class	Process	PID	Memory	Run State	Monitored
0.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	boinc.exe	2140	4952 kB	Running	No
1.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	boincmgr.exe	2888	5860 kB	Running	No
2.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	conhost.exe	2668	116 kB	Running	No
3.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	csrss.exe	296	680 kB	Running	No
4.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	csrss.exe	348	664 kB	Running	No
5.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	csrss.exe	1220	544 kB	Running	No
6.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	dwm.exe	1040	284 kB	Running	No
7.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	explorer.exe	2648	3200 kB	Running	No
8.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	LogonUI.exe	704	6576 kB	Running	No
9.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	lsass.exe	452	5148 kB	Running	No
10.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	lsm.exe	464	1920 kB	Running	No
11.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	msdtc.exe	2432	156 kB	Running	No
12.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	msmdsrv.exe	1080	6320 kB	Running	No
13.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	rdpclip.exe	2084	352 kB	Running	No
14.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	ReportingServicesService.exe	1140	64212 kB	Running	No
15.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	services.exe	444	4760 kB	Running	No
16.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	smss.exe	216	80 kB	Running	No
17.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	snmp.exe	1460	3624 kB	Running	No
18.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	spoolsv.exe	272	1148 kB	Running	No
19.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	sppsvc.exe	2496	2992 kB	Running	No
20.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	sqlservr.exe	1052	36984 kB	Running	No
21.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	sqlwriter.exe	1484	88 kB	Running	No
22.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	svchost.exe	552	3072 kB	Running	No
23.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	svchost.exe	624	3628 kB	Running	No
24.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	svchost.exe	712	6388 kB	Running	No
25.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	svchost.exe	764	19972 kB	Running	No
26.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	svchost.exe	804	5296 kB	Running	No
27.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	svchost.exe	844	1176 kB	Running	No
28.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	svchost.exe	884	6140 kB	Running	No
29.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	svchost.exe	980	3496 kB	Running	No
30.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	svchost.exe	1108	80 kB	Running	No
31.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	svchost.exe	1832	2632 kB	Running	No
32.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	svchost.exe	1864	108 kB	Running	No
33.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	svchost.exe	2248	100 kB	Running	No
34.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	System	4	48 kB	Running	No
35.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	System Idle Process	1	24 kB	Running	No
36.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	taskhost.exe	2704	3304 kB	Running	No
37.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	wininit.exe	356	80 kB	Running	No
38.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	winlogon.exe	384	280 kB	Running	No
39.	ACME - DB MSSQL 2 - We	ACME	192.168.32.113	Microsoft | MSSQL Server	winlogon.exe	1664	80 kB	Running	No
40.	ACME - DB-MSSQL - Web	ACME	192.168.32.112	Microsoft | Windows Server 2008 R2	csrss.exe	296	844 kB	Running	No
41.	ACME - DB-MSSQL - Web	ACME	192.168.32.112	Microsoft | Windows Server 2008 R2	csrss.exe	348	452 kB	Running	No
42.	ACME - DB-MSSQL - Web	ACME	192.168.32.112	Microsoft | Windows Server 2008 R2	csrss.exe	1676	564 kB	Running	No
43.	ACME - DB-MSSQL - Web	ACME	192.168.32.112	Microsoft | Windows Server 2008 R2	dwm.exe	2272	512 kB	Running	No
44.	ACME - DB-MSSQL - Web	ACME	192.168.32.112	Microsoft | Windows Server 2008 R2	explorer.exe	2340	4080 kB	Running	No
45.	ACME - DB-MSSQL - Web	ACME	192.168.32.112	Microsoft | Windows Server 2008 R2	LogonUI.exe	704	1592 kB	Running	No
46.	ACME - DB-MSSQL - Web	ACME	192.168.32.112	Microsoft | Windows Server 2008 R2	lsass.exe	452	6460 kB	Running	No
47.	ACME - DB-MSSQL - Web	ACME	192.168.32.112	Microsoft | Windows Server 2008 R2	lsm.exe	460	2156 kB	Running	No
48.	ACME - DB-MSSQL - Web	ACME	192.168.32.112	Microsoft | Windows Server 2008 R2	msdtc.exe	1276	1516 kB	Running	No
49.	ACME - DB-MSSQL - Web	ACME	192.168.32.112	Microsoft | Windows Server 2008 R2	msmdsrv.exe	1128	7260 kB	Running	No
50.	ACME - DB-MSSQL - Web	ACME	192.168.32.112	Microsoft | Windows Server 2008 R2	Oobe.exe	2472	17408 kB	Running	No
51.	ACME - DB-MSSQL - Web	ACME	192.168.32.112	Microsoft | Windows Server 2008 R2	rdpclip.exe	536	560 kB	Running	No
52.	ACME - DB-MSSQL - Web	ACME	192.168.32.112	Microsoft | Windows Server 2008 R2	services.exe	444	5864 kB	Running	No
53.	ACME - DB-MSSQL - Web	ACME	192.168.32.112	Microsoft | Windows Server 2008 R2	smss.exe	216	316 kB	Running	No
54.	ACME - DB-MSSQL - Web	ACME	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 the ScienceLogic platform:

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

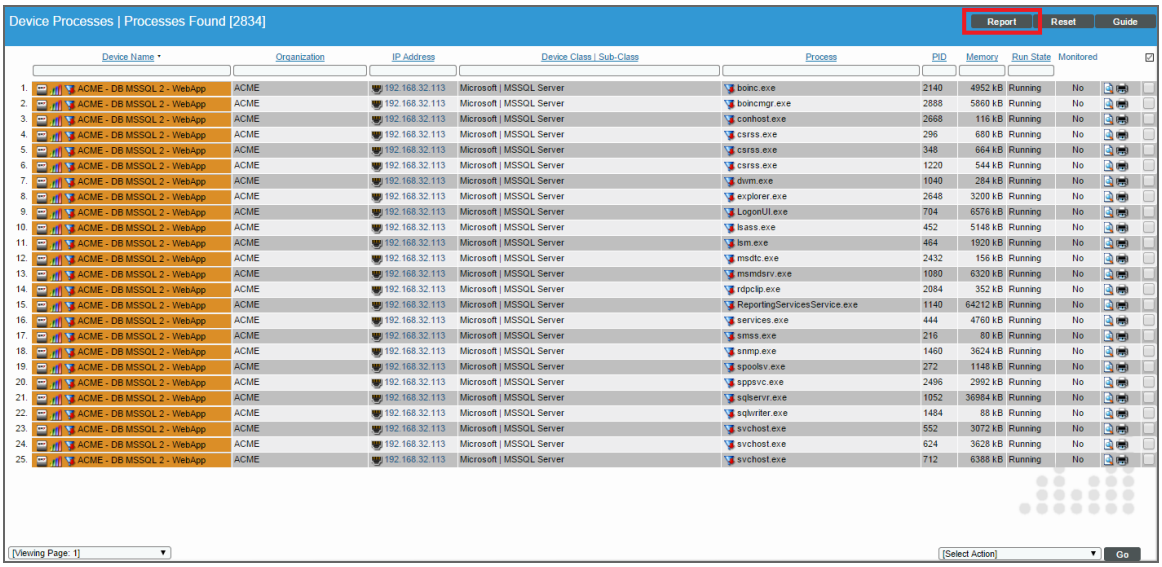

NOTE: If you want to include only certain processes in the report, use the "search as you type" fields at the top of each column. You can filter the list by one or more column headings. You can then select the **[Report]** button, and only the processes displayed in the **Device Processes** page will 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 the ScienceLogic platform 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 (Registry > Devices > Processes), you can generate an exclusion report for a process. The ScienceLogic platform 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. The ScienceLogic platform lists only appropriate servers in this report. For example, Linux servers would not appear in a report for Windows-based processes.

EM7™ Management Systems

Windows Service Exclusion Report

April 17, 2015, 3:49 am

Devices That Have [Reporting Services Service.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
Services Found	0 [on + off]
Services Not Found	0

Report Created By ScienceLogic EM7™

A Process Exclusion Report displays the following:

- Name of the process.
- List of all devices in the ScienceLogic platform where the process is running.
- List of all devices in the ScienceLogic platform where the process is not running. The ScienceLogic platform 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 (Registry > Devices > Processes).

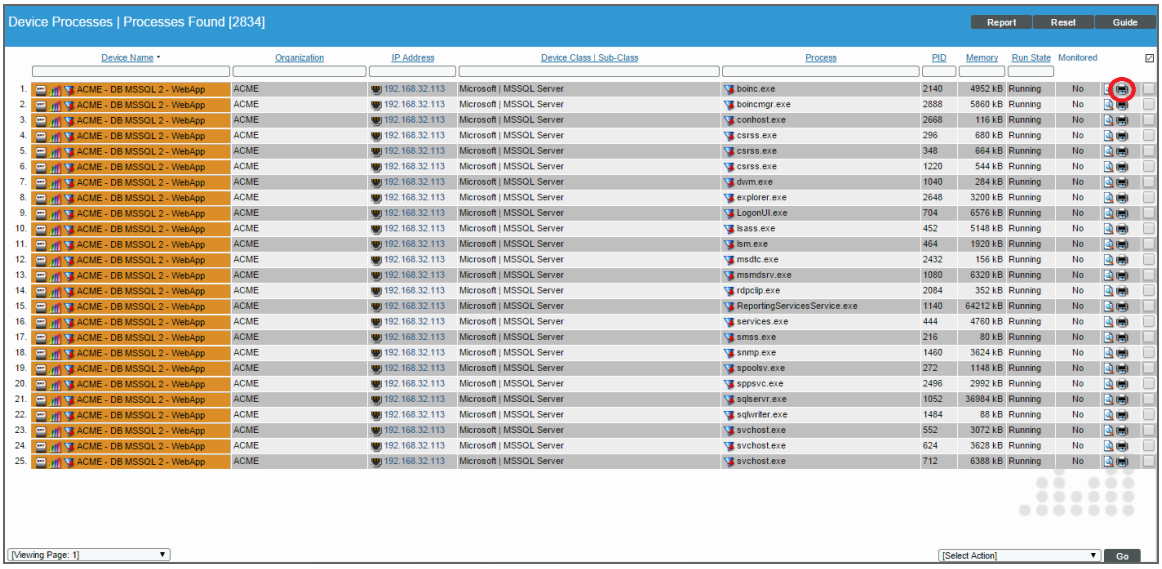

2. In the **Device Processes** page, find an instance of the process you want to generate an exclusion report for. Select its printer icon ( ).
3. You will be prompted to save or view the generated report.

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:

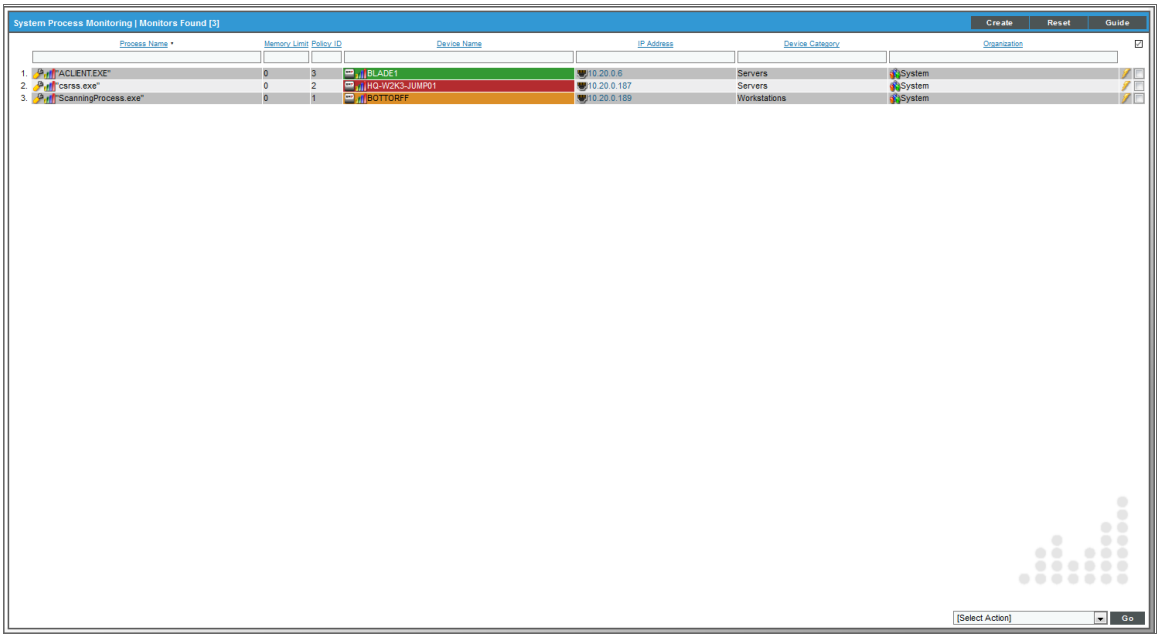

- ***Process Name***. Name of the policy.
- ***Memory Limit***. The maximum amount of memory that can be used or reserved by a single instance of the process, as specified in the process policy.
- ***Policy ID***. Unique, numeric ID, assigned to the policy automatically by the ScienceLogic platform.
- ***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 the ScienceLogic platform 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 (Registry > 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 (Registry > 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 he **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:

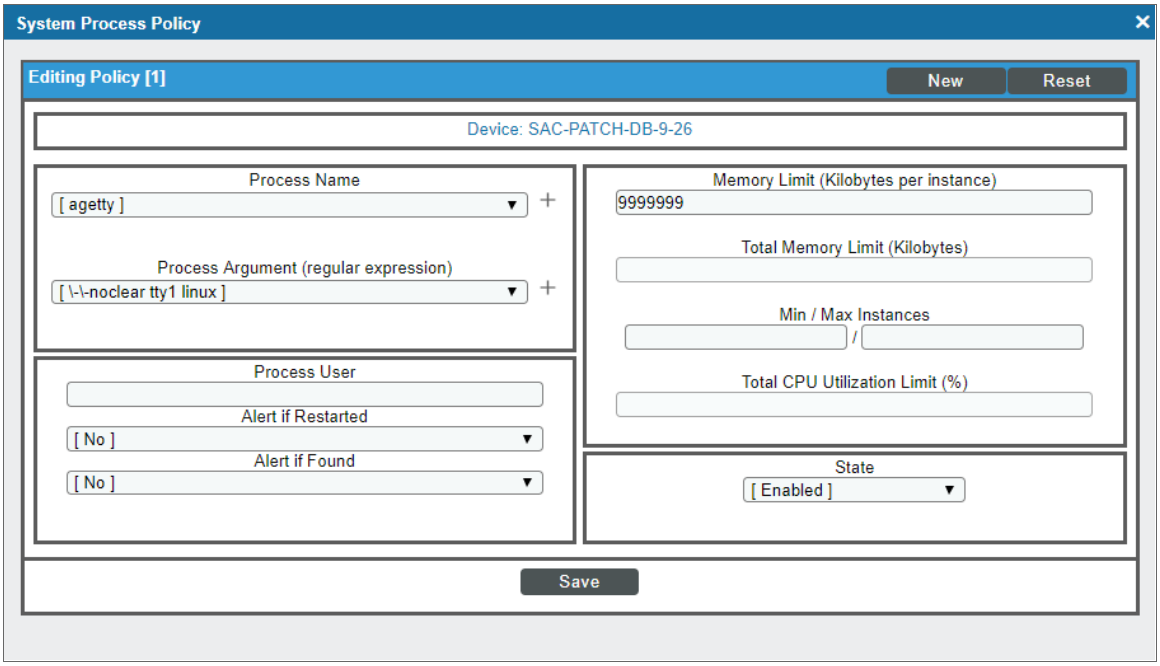

- ***Process Name***. The name of the process. You can either:
    - Select from a list of all processes running on this device.
    - Click on the "+" icon and manually enter the name of a process.
- ***Process Argument (regular expression)***. The arguments with which the process is invoked. This field 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. The ScienceLogic platform 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, the ScienceLogic platform 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. The ScienceLogic platform checks every 5 minutes to determine if a system process has restarted. If the platform finds a restarted system process, it will generate an alert in the Device Log.
    - *No*. Use this setting if you do not want the platform 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 the ScienceLogic platform finds the illicit process (specified in the ***Process Name*** field), the platform will generate an event.
        - If the platform does not find the illicit process running, the platform will not generate an event.
    - *No*. Use this setting to ensure that a required process is running.
        - If the platform finds the required (specified in the ***Process Name*** field) running, the platform does not generate an event.
        - If the platform does not find the required process running, the platform 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 ScienceLogic agent is installed on the selected device. The amount of memory, in kilobytes, you will 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, the ScienceLogic platform 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, the platform 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 ScienceLogic 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 the ScienceLogic platform should start collecting data specified in this policy from the device. Choices are:
    - *Enabled*. The ScienceLogic platform 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: OS Process Check*** process.
    - *Disabled*. The ScienceLogic platform 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

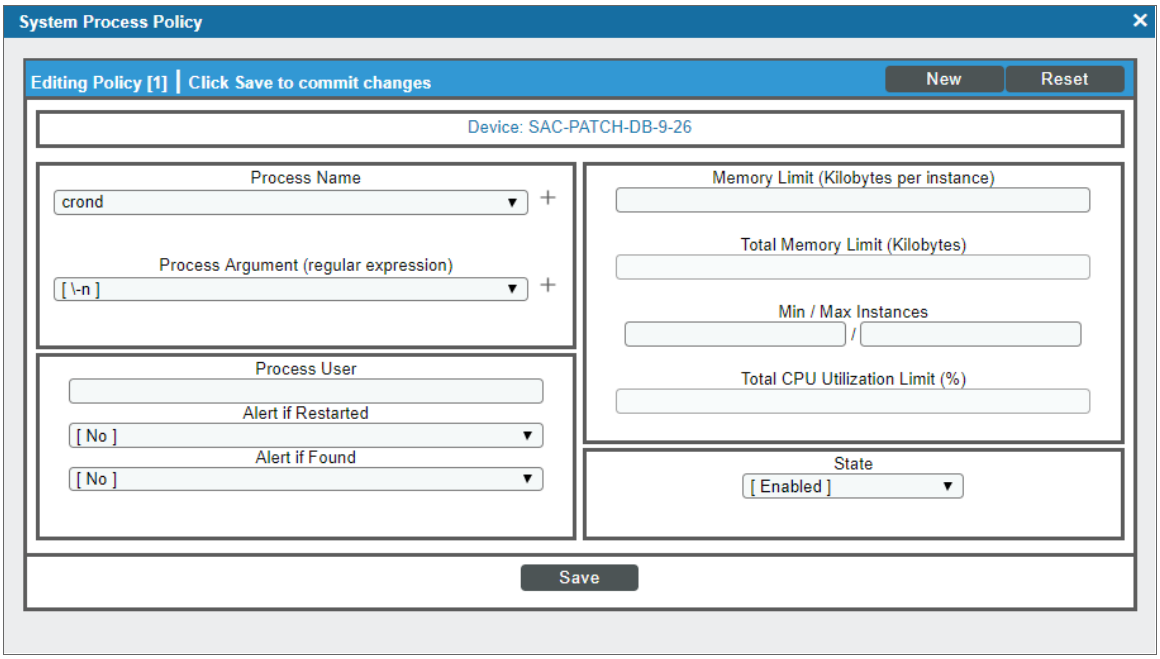

- This policy monitors a system process on the device "em7ao".
- The policy looks for the process "crond".
- If the process is not found running on the device, the ScienceLogic platform generates an event.

Editing a System Process Monitoring Policy

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

1. From the **Device Manager** page (Registry > 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.

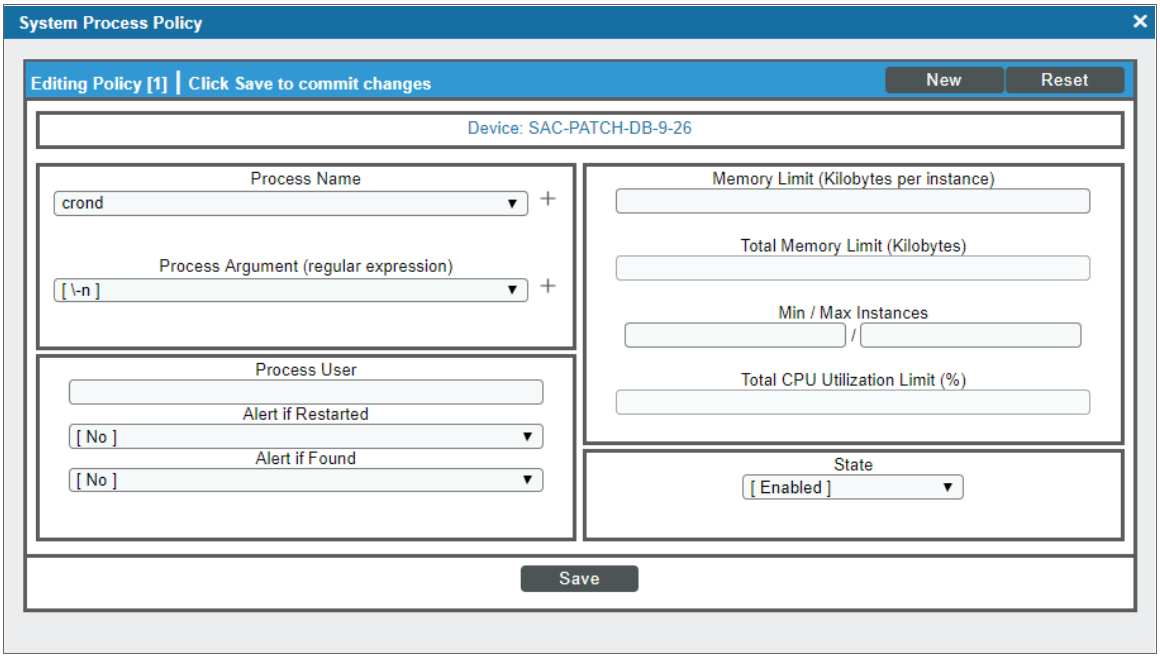

4. In the **System Process Policy** modal page, you can change the values in one or more of the fields described in the section on *Defining a Monitoring Policy for System Processes*.
5. To save your changes to the policy, select the **[Save]** button.

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, the ScienceLogic platform 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, the ScienceLogic platform 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, the ScienceLogic platform 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*.

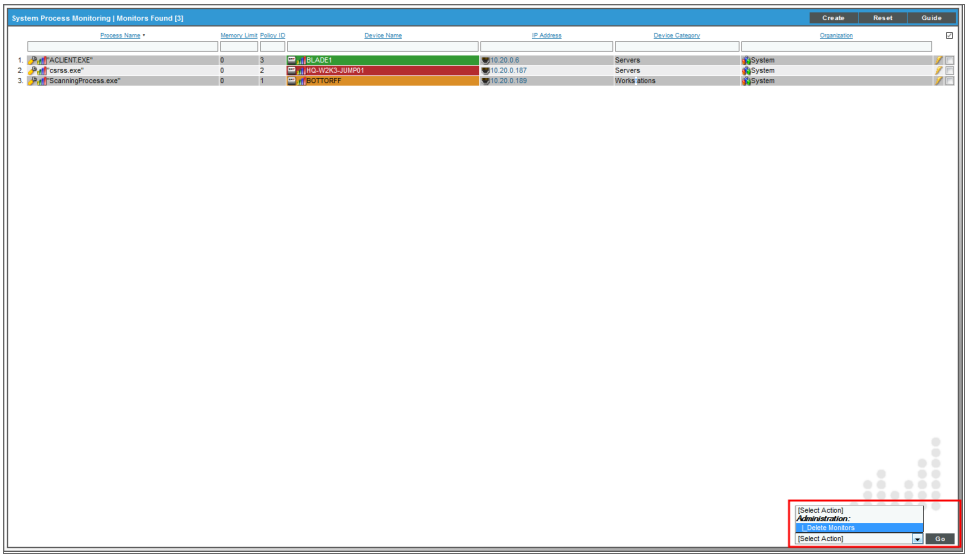

4. Click **[Go]**.
5. The policy is deleted from the ScienceLogic platform. The associated reports (from the Device Reports > **[Performance]** tab) are also deleted.

Chapter 6

Monitoring Logs Using ScienceLogic Agents

Overview

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

This chapter includes the following topics:

Using a Log File Monitoring Policy *47*

Viewing the List of Log File Monitoring Policies *47*

Filtering the List of Log File Monitoring Policies *48*

Creating a Log File Monitoring Policy *49*

Editing a Log File Monitoring Policy *51*

Deleting Log File Monitoring Policies *51*

Viewing the List of Log File Monitoring Policies and Aligned Devices *52*

Filtering the List of Log File Monitoring Policies and Aligned Devices *53*

Aligning a Log File Monitoring Policy to Devices *53*

Unaligning Log File Monitoring Policies from Devices *56*

Creating an Event Policy for Agent Logs *56*

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.

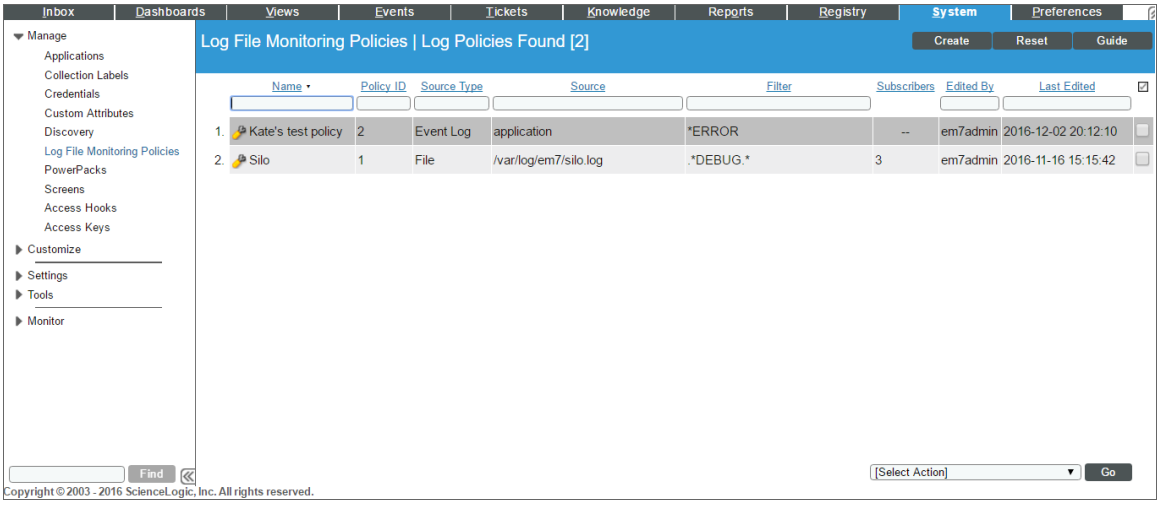

TIP: To sort the list of Log File Monitoring policies, click on a column heading. The list will be sorted by the column value, in ascending order. To sort by descending order, click the column heading again. The ***Last Edited*** column sorts by descending order on the first click; to sort by ascending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

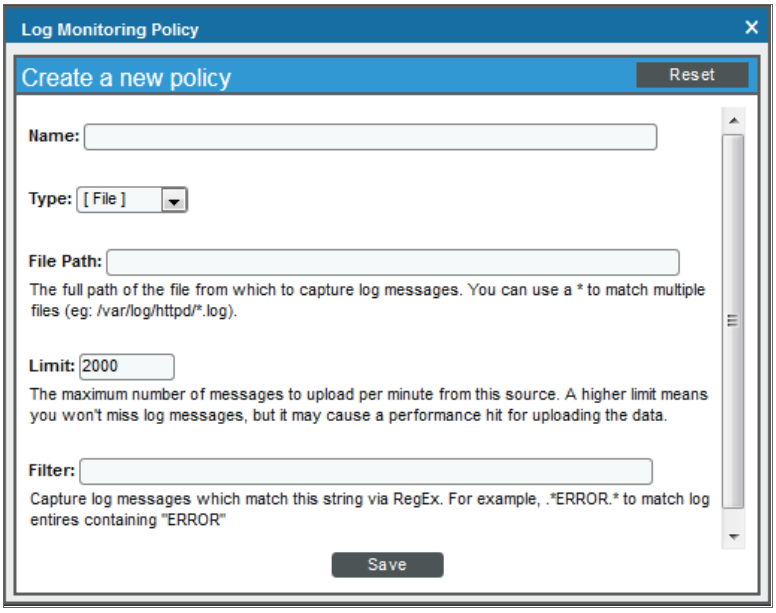

3. Supply values in the following fields:
   - ***Name***. Enter a name for the policy.
   - ***Type***. Select the source of the logs on the monitored device. Choices are:
     - *File*. The agent will monitor a file on the file system of the device(s).
     - *Event Log*. The agent will monitor the 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 ScienceLogic 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.

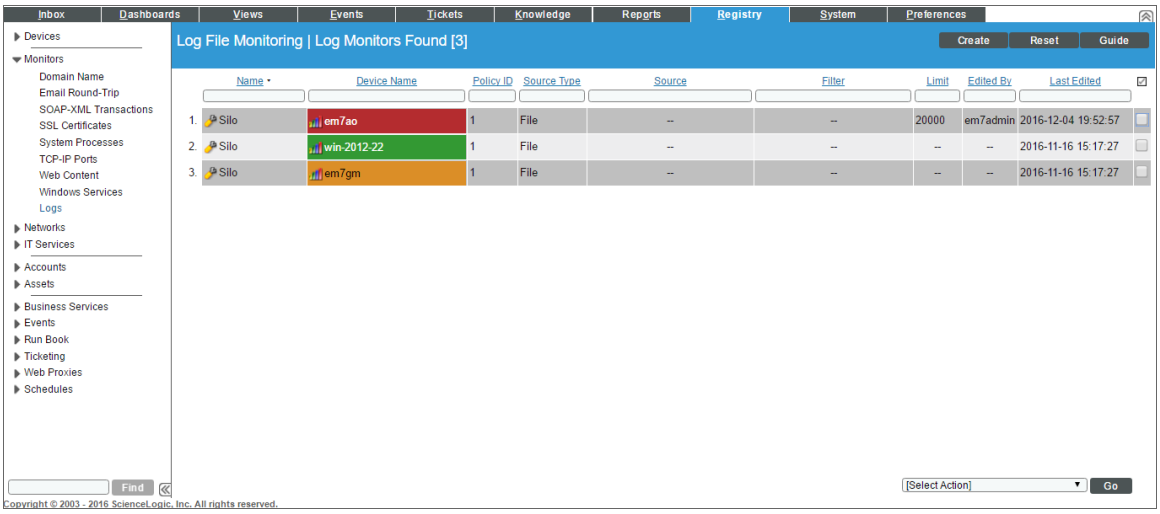

For each aligned Log File Monitoring policy and device, the page displays:

- ***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 the ScienceLogic platform.
- ***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 the platform.

- ***Limit***. The maximum log messages the agent sends to the platform 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. Fore 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.

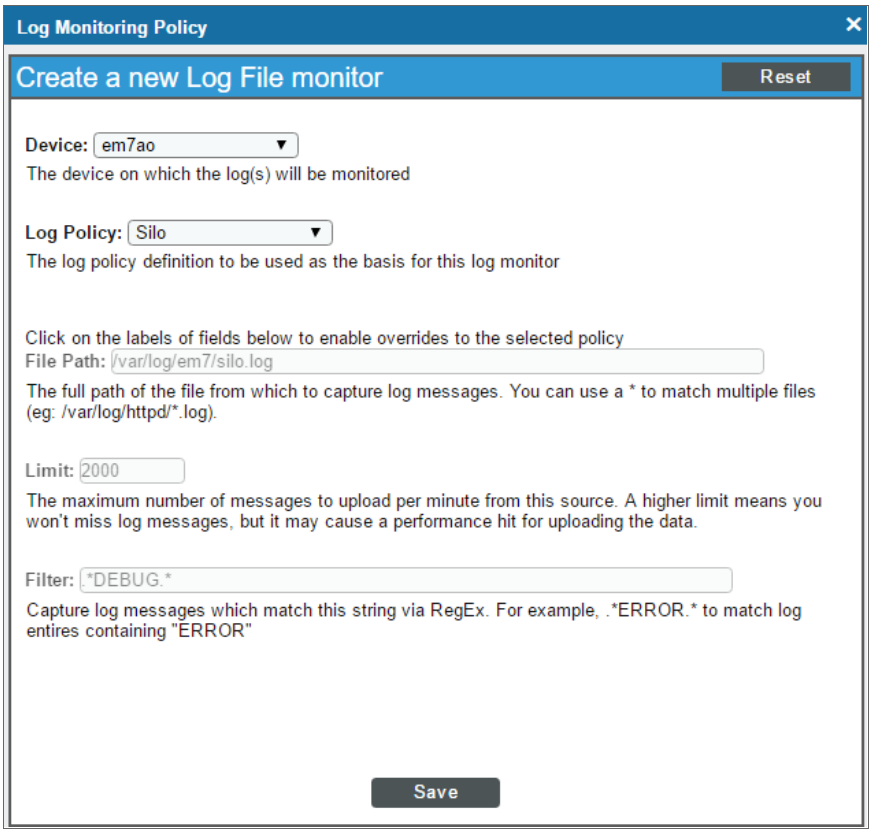

3. In the Log File Monitor modal page, supply values in the following fields:

   - ***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 the platform 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 the ScienceLogic platform.

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.

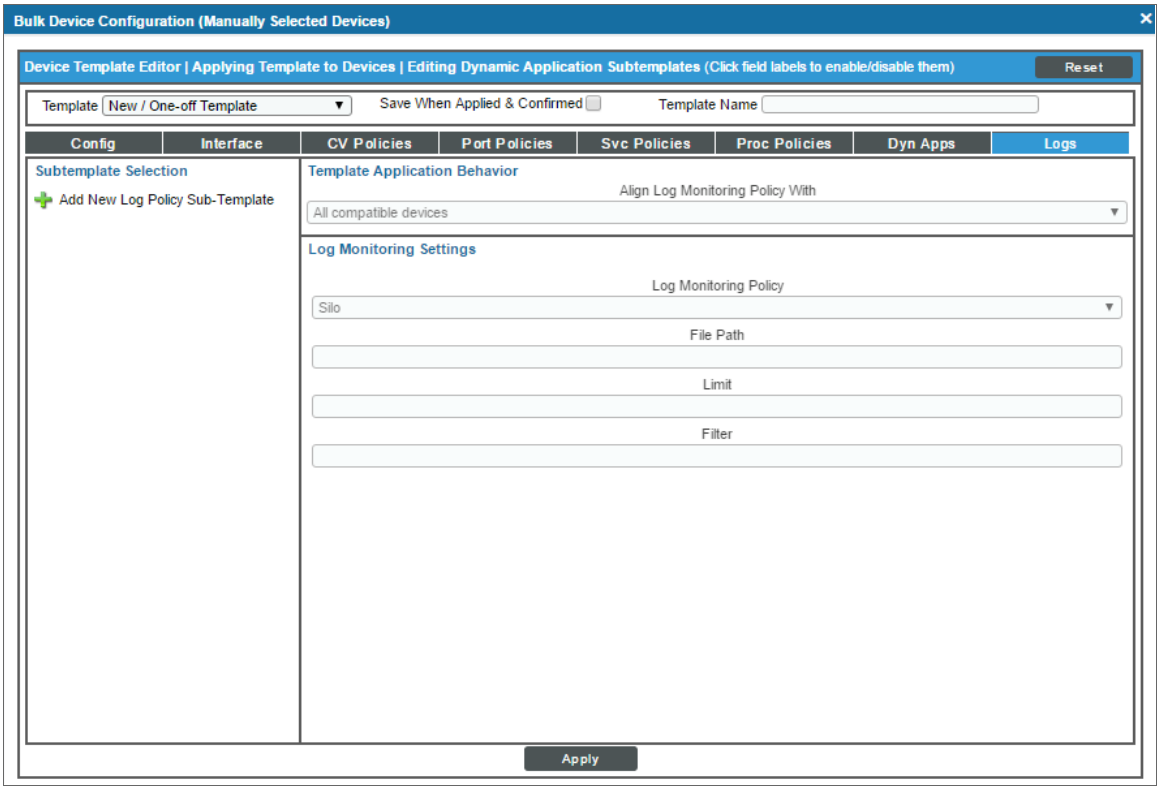

5. Click the **[Logs]** tab.
6. Click the Add New Log Policy Sub-Template icon ( ).
7. Supply values in the following fields:

   - ***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).

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

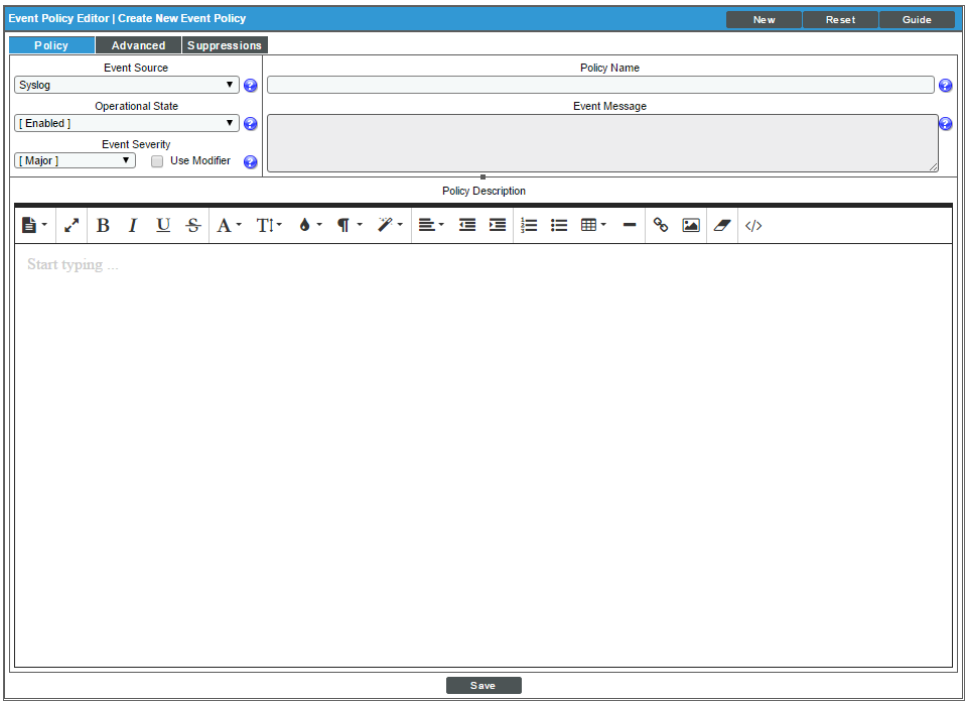

3. In the **Event Policy Editor** page and set of tabs, you can define a new event. The **Event Policy Editor** page contains three tabs:

   - ***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.

4. 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]**.

Chapter 7

Monitoring Vitals Using ScienceLogic Agents

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 the ScienceLogic platform, see the ***Device Management*** manual.

This chapter includes the following topics:

Viewing System Availability Reports for a Device 61
Changing the Method for Measuring Device Availability 64
Viewing CPU and Memory Utilization for a Device 65
Viewing CPU Utilization 65
Viewing Memory Utilization 67
Changing the Dynamic Application Precedence Settings for CPU and Memory Utilization 69

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 the platform 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 the ScienceLogic platform discovers with the discovery tool (System > Manage > Discovery), the ScienceLogic platform determines availability by checking the status of the port specified in the ***Availability Port*** field in the **Device Properties** page. The ScienceLogic platform collects device-availability data every five minutes, as specified in the process "Data Collection: Availability" (in the **Process Manager** page).

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

For devices that the ScienceLogic platform discovers with the agent, the ScienceLogic platform 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 (Registry > 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 ( ).

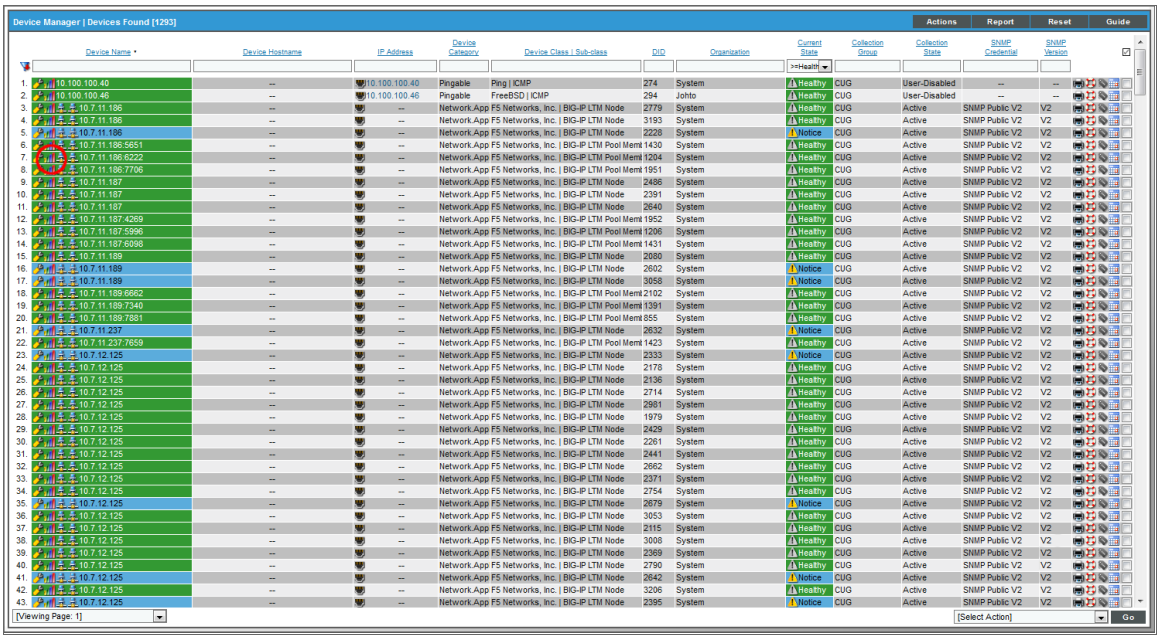

3. In the **Device Reports** panel, select the Performance tab.

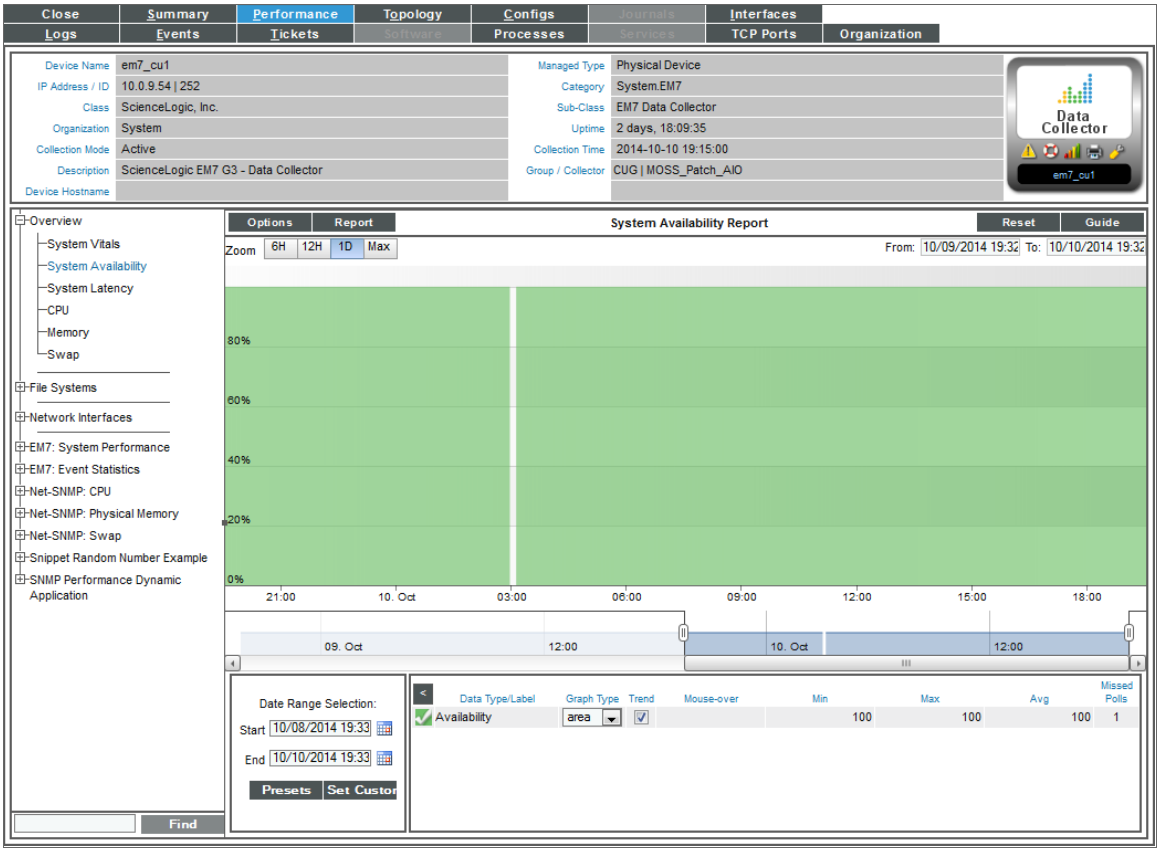

4. In the Performance tab, go to the NavBar (list of links in the left pane), expand the ***Overview*** link, and select ***System Availability***.
5. The System Availability report displays system availability for the selected date and time range.

   - 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.

6. The **[Options]** menu in the upper left of the report displays a menu of options you can apply to data in the current report.
7. 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 the ScienceLogic platform 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 the platform 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 the platform.
- 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 the platform uses to monitor availability:

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

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

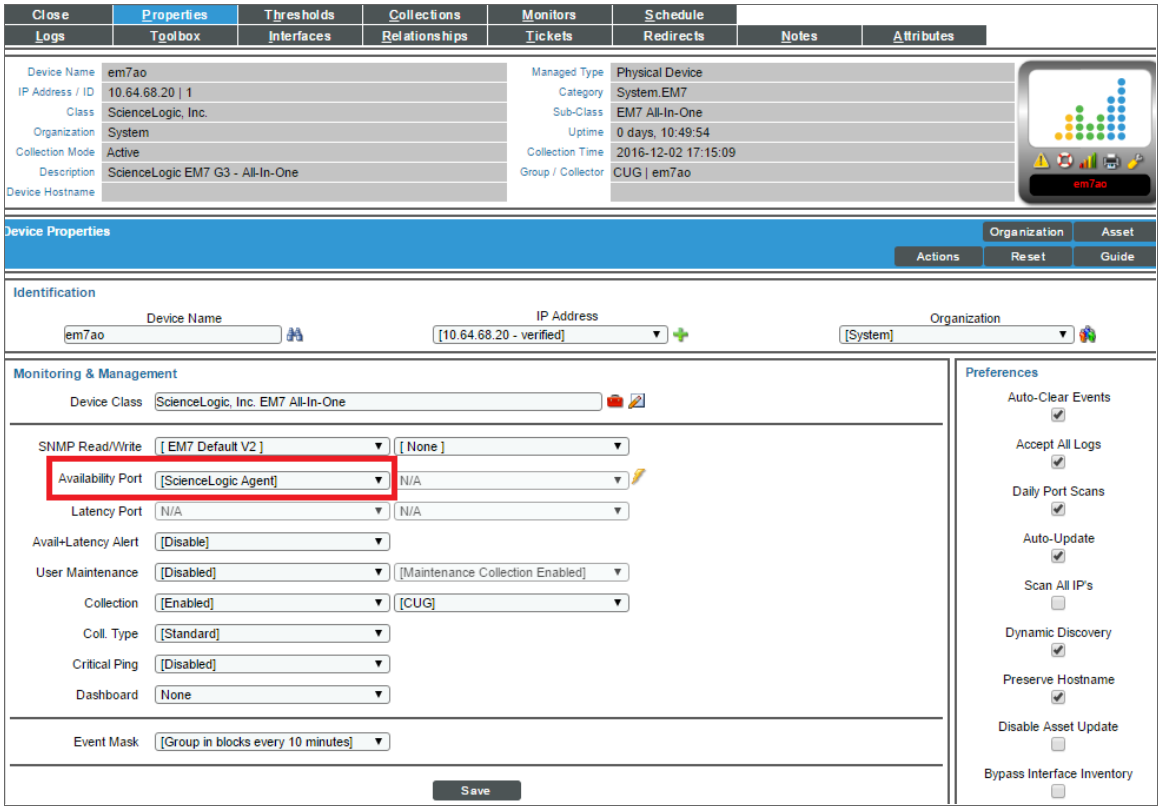

3. In the ***Availability Port*** field, select the method you want to use to monitor availability:
   - *TCP*. Availability is based on whether the the platform can connect to the device using the specified TCP port.
   - *ICMP*. Availability is based on whether the device responds to an ICMP ping request from the platform.
   - *SNMP*. Availability is based on whether the device responds to an SNMP get request from the platform.
   - *ScienceLogic Agent*. Availability is based on whether the agent is reporting data to the platform.
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 the ScienceLogic platform 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 (Registry > Devices > Device Manager), find the device where the CPU resides, and select its bar graph icon ( ).
   - Go to the **Device Hardware** page (Registry > 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***.

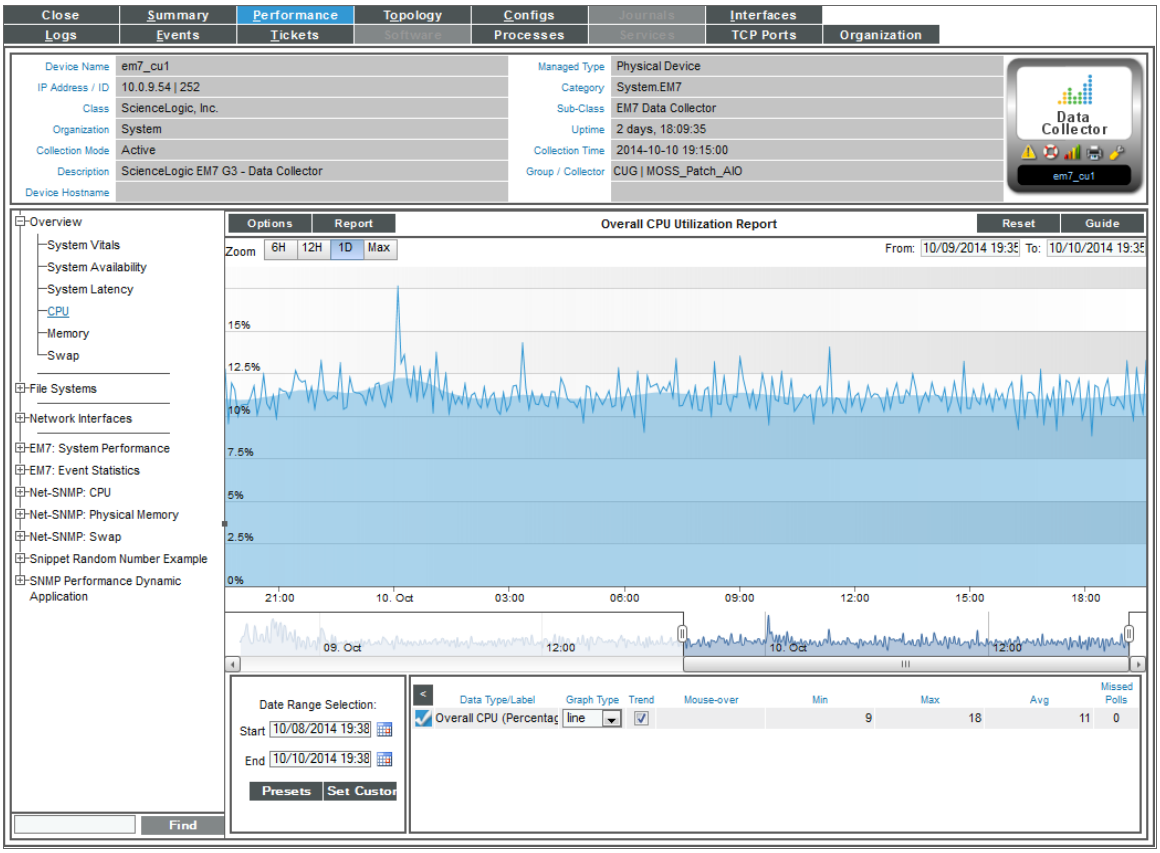

4. The Overall CPU Utilization report displays total CPU usage and average CPU usage over time. If a device contains multiple CPUs, the report displays the total combined CPU usage, in percent, and the combined average CPU usage, in percent. The graph displays CPU usage for the selected date and time 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 the ScienceLogic platform 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 the ScienceLogic platform 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 (Registry > Devices > Device Manager), find the device where the memory resides, and select its bar graph icon ( ).
    - Go to the **Device Hardware** page (Registry > 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***

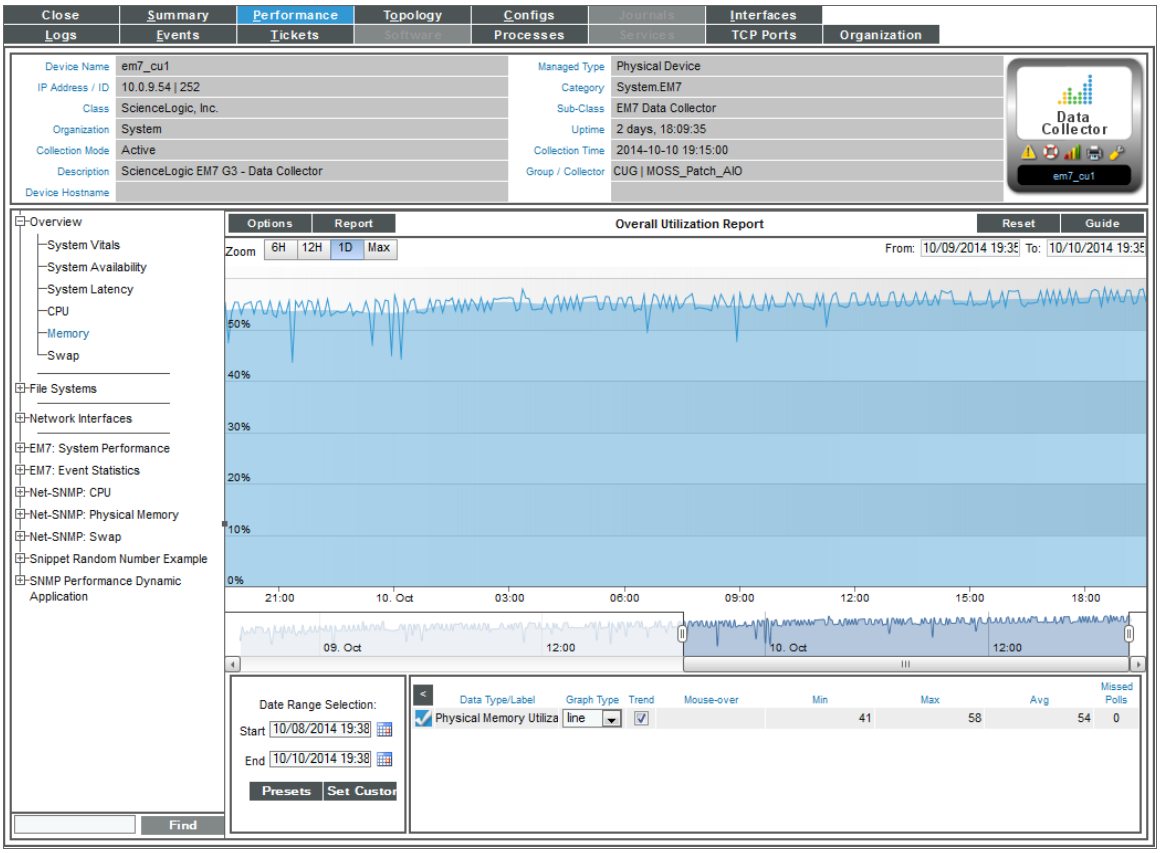

4. The Overall Memory Utilization report displays total memory usage and average memory usage over time. The graph displays memory usage for the selected date and time range.
    - The y-axis displays memory usage, in percent, to the left.
    - The x-axis displays time. The increments vary, depending 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 the ScienceLogic platform was unable to collect the data within the time span of the report.

Changing the Dynamic Application Precedence Settings for CPU and Memory Utilization

The ScienceLogic platform 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, the platform 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.

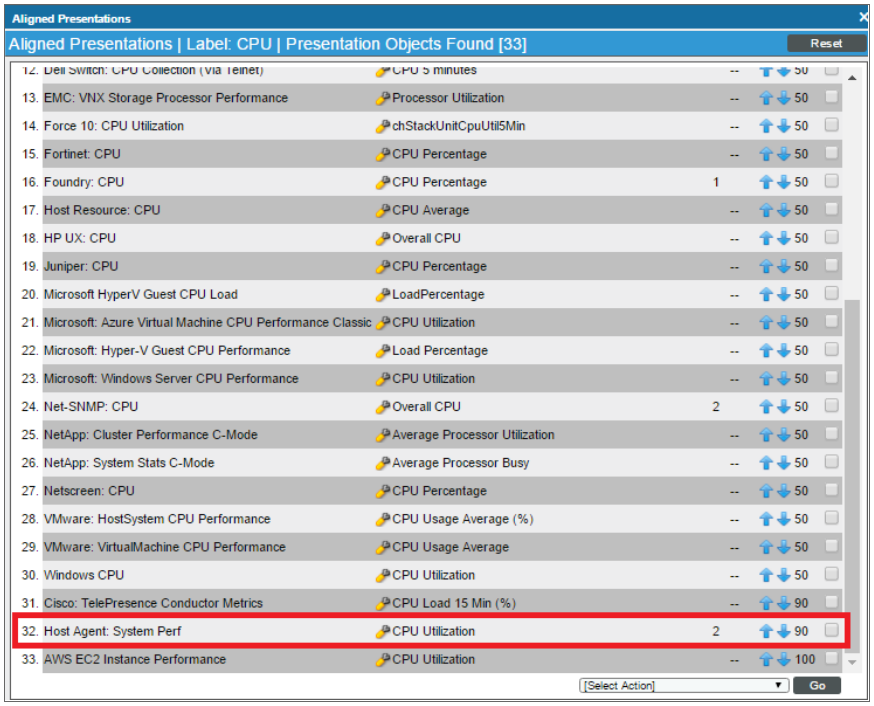

3. Locate the entry for the ***Host Agent: System Perf*** Dynamic Application. Select its checkbox.
4. In the ***Select Action*** drop-down list, select *0* in the *Change Precedence* section.
5. Click **[Go]**.

To change the precedence settings *per-device*:

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

2. 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.

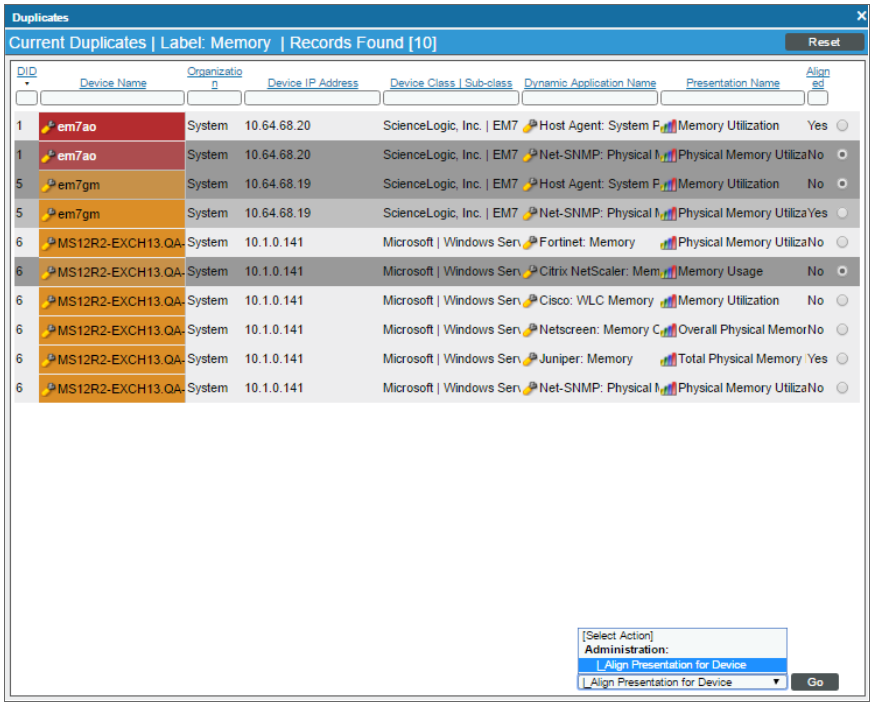

3. The **Current Duplicates** page displays multiple rows for each device; each row specifies a device and Dynamic Application metric pair. For each group of rows for a device, use the radio button to the right of the page to select the Dynamic Application metric you want to use for that device.
4. In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

Chapter 8

Troubleshooting ScienceLogic Agents

Overview

To troubleshoot potential issues with SL1 agents, perform the following procedures, in the following order.

Determine if the Agent Process is Running 73
Determine if the Agent Configuration is Valid 74
Determine if the Agent is Able to Upload Data 74
Check the Agent Upload Directory 74
Run the Agent in Debug Mode (Linux) 75
Determine if SL1 is Receiving Agent Data 75
Determine if SL1 Cannot Process Agent Data 76
Determine if the Number of Processes is Inconsistent with Other Applications 77
Troubleshooting Examples 77
Example /var/log/insight/streamer.log for successful discovery 77
Example /var/log/uwsgi/streamer.log for successful discovery in streamer_prime 78
Save incoming data for a specific device ID (streamer_prime) 78
Save incoming data for a specific device ID (Converged Platform or SL1) 78
Additional Troubleshooting Situations and Best Practices 78

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.

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 commands:

   ```
   "cd /var/log/uwsgi".
   "sudo tail -n 100 streamer_prime_uwsgi.log"
   ```

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. SSH into the general-compute VM.
2. Run the following command to see the list of containers, and look for a container with "streamer" in its name.

```
sudo docker ps
```

3. Run the following commands:

```
sudo docker exec -it pod1-streamer0 bash"
"cd /var/log/uwsgi"
"sudo tail -n 100 streamer.log"
```

4. Look for lines starting with the IP of the customer's server.
5. If there are no matching lines, then the streamer is not getting data from that agent.

Determine if SL1 Cannot Process Agent Data

First, review the uploaded information for the agent on the server:

1. To make the agent copy the uploaded information to a known location after uploading, add the following code to the agent configuration:

```
DataKeep 30
```

NOTE: The "30" indicates that the agent will keep the last 30 uploads.

2. After you make the above change, restart the agent. The agent copies its uploads to one of the following locations:
   - Windows: C:\Program Files\ScienceLogic\SiloAgent\logs
   - Linux: /tmp/data

NOTE: You can also check uploads for the agent from the streamer by running the following command:

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

Next, check the SL1 log files:

1. If you are using streamer_prime, locate the following files from the SL1 message-collector):
   - /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 commands and locate the following files:

```
sudo docker ps
```

```
sudo docker exec -it pod1-streamer0 bash
```

- /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.

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 the 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)

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

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.	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. To address this issue, you can ***merge*** the device records using the existing ("classic") user interface. For more information, see the ***Device Management*** manual.
The SNMP device record has IPv4, but the agent device record has IPv6.	The agent reports all network interfaces to the message collector. The message collector uses the first "bound" IP address reported by the agent. To address this issue, you can manually edit the agent device record in the "classic" user interface and update the IP address.

Situation	Cause / Resolution
If you uninstall an agent and then run a different installation executable file, you still see the same organization ID for the agent record.	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. 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.

© 2003 - 2018, 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

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

International: +1-703-354-1010