



---

# Monitoring with the SL1 Agent

SL1 version 10.1.0

---

# Table of Contents

<b>Introduction to the SL1 Agent</b> .....	<b>4</b>
What is the SL1 Agent? .....	5
How the SL1 Agent Works in Different Environments .....	5
What Kind of Data Can an Agent Collect? .....	7
Extensible Collection .....	9
Agent PowerPacks .....	9
Agent-Compatible PowerPacks .....	11
Windows Devices .....	11
Java Management Extensions (JMX) Resources .....	11
Agent Architecture .....	12
SL1 Distributed Architecture .....	12
SL1 Extended Architecture .....	13
<b>Installing the SL1 Agent</b> .....	<b>15</b>
Installing an Agent from the Agents Page .....	16
Installing a Linux Agent .....	17
Using Agents with SELinux .....	18
Installing a Windows Agent .....	19
Installing an AIX Agent .....	21
Installing a Solaris Agent .....	22
Upgrading the Agent .....	23
Stopping the Agent .....	24
Uninstalling the Agent .....	24
Uninstalling a Linux Agent .....	25
Uninstalling a Windows Agent .....	25
Installing an Agent from the Device Manager Page .....	25
Installing the Linux Agent from the Device Manager Page .....	27
Updating an Agent on a Linux System .....	28
Uninstalling an Agent on a Linux System .....	28
Installing the Windows Agent from the Device Manager Page .....	29
Uninstalling the Agent on a Windows System .....	29
<b>Configuring the SL1 Agent</b> .....	<b>30</b>
Configuring Agent Monitoring Based on SL1 Architecture .....	31
Configuring Agent Settings from the Device Investigator Page .....	32
Configuring Extensible Collection .....	35
Configuring an SL1 Agent on the Device Manager Page .....	37
Adding the "SL Agent" Column to the Device Manager Page .....	37
Configuring Agent Settings on a Device .....	38
Configuring the SL1 Agent to Work with PhoneHome Collectors .....	39
<b>Monitoring Logs Using the SL1 Agent</b> .....	<b>43</b>
What is a Log File Monitoring Policy? .....	44
Viewing the List of Log File Monitoring Policies .....	44
Filtering the List of Log File Monitoring Policies .....	45
Creating a Log File Monitoring Policy .....	46
Aligning a Log File Monitoring Policy to Devices .....	48
Unaligning Log File Monitoring Policies from Devices .....	51
Editing a Log File Monitoring Policy .....	51
Deleting Log File Monitoring Policies .....	52
Viewing the List of Log File Monitoring Policies and Aligned Devices .....	52
Filtering the List of Log File Monitoring Policies and Aligned Devices .....	53
Creating an Event Policy for Agent Logs .....	53

Creating an Event Policy for Agent Logs in the Classic User Interface .....	58
<b>Agent Troubleshooting</b> .....	<b>61</b>
Is the Agent Process Running? .....	63
Is the Agent Configuration Valid? .....	64
Has SL1 Discovery Completed? .....	64
Is the Agent Able to Upload Data? .....	65
Check the Agent Upload Directory .....	65
Run the Agent in Debug Mode (Linux) .....	65
Is SL1 Receiving Agent Data? .....	66
Can SL1 Process Agent Data? .....	67
Is the Number of Processes Inconsistent with Other Applications? .....	67
Does MySQL on the Collector Unit Have a Connection Issue? .....	67
Troubleshooting Examples .....	68
Example /var/log/streamer_prime/streamer_prime.log for successful discovery .....	68
Example /var/log/uwsgi/streamer.log for successful discovery in the distributed architecture .....	69
Save incoming data for a specific device ID in the distributed architecture .....	69
Save incoming data for a specific device ID in the SL1 Extended Architecture .....	69
Additional Troubleshooting Situations and Best Practices .....	69
Agent Communication with SL1 .....	70
Return Codes .....	70
Sub-codes .....	71

---

# Chapter

# 1



## Introduction to the SL1 Agent

---

### Overview

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

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

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

This chapter includes the following topics:

<i>What is the SL1 Agent?</i> .....	5
<i>How the SL1 Agent Works in Different Environments</i> .....	5
<i>What Kind of Data Can an Agent Collect?</i> .....	7
<i>Agent PowerPacks</i> .....	9
<i>Agent Architecture</i> .....	12

---

## What is the SL1 Agent?

The **SL1 agent** is a program that you can install on a device monitored by SL1. The SL1 agent collects data from the device and pushes that data back to SL1.

Similar to a Data Collector or Message Collector, the SL1 Agent collects data about infrastructure and applications.

The agent can be configured to communicate with either the Message Collector or Compute Cluster.

---

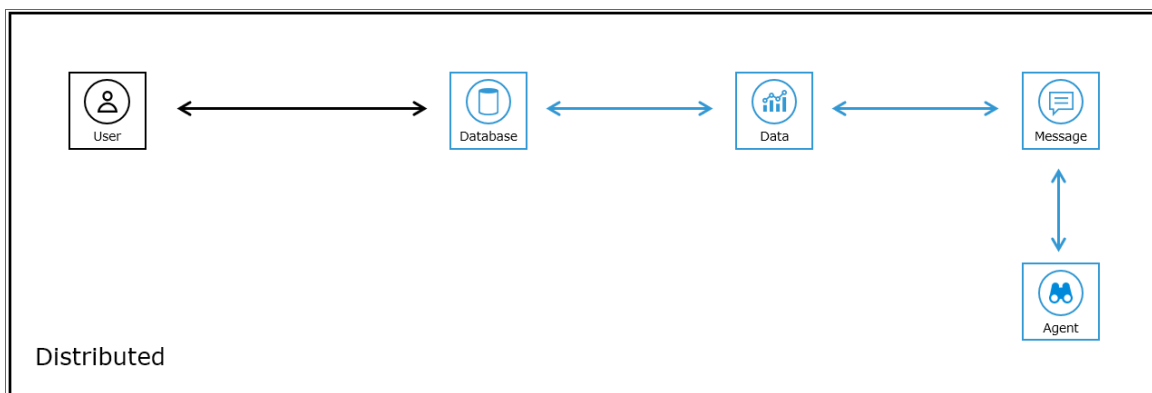
## How the SL1 Agent Works in Different Environments

There are two supported "generations" of the SL1 agent that you can use to gather data: Gen 1 and Gen 3. The agent itself is the same from one generation to the next; the only difference between the generations is the *environment* where the agent is used.

In brief, a **Gen 1 agent** uses a distributed SL1 environment to upload data directly to a Message Collector (MC), while a **Gen 3 agent** uses the SL1 Extended Architecture to upload data to the "Streamer" service running on the SL1 Compute Node cluster.


The following list provides more details about how SL1 uses the different generations of the SL1 agent:

- **Gen 1 agent:** Introduced in SL1 version 8.2.0, the Gen 1 agent uses a distributed environment to upload data directly to a Message Collector (MC). Next, Dynamic Applications running on a Data Collector query data from the Message Collector with an API. The Data Collector (DC) then sends the data to the SL1 Database Server:



The Gen 1 agent provides limited infrastructure health reporting, including Log File Monitoring, Processes, and System Vitals like Availability, CPU Usage, Memory Usage, and Disk Usage. This agent is sometimes called the "PO" agent.

- **Gen 2 agent:** This version of the agent has been deprecated.
- **Gen 3 agent:** Introduced in SL1 version 8.12.0, the Gen 3 agent uses the SL1 Extended Architecture to upload data through a load balancer to the "Streamer" service running on the SL1 Compute Node cluster. The Streamer service then forks and forwards data to other services, and eventually some services will store parsed data in the SL1 Database Server:



The Gen 3 agent provides full infrastructure health reporting, including system vitals (file system, network interface, and Windows service data), log monitoring, and optionally allows you to align PowerShell Dynamic Applications to your device. Also, the Dynamic Applications in the *ScienceLogic: Agent PowerPack* are automatically aligned with the device with the Gen 3 agent installed on it.

If you are already running the Gen 1 agent in a distributed environment and you upgrade your SL1 system to use the SL1 Extended Architecture, any existing devices monitored by SL1 agents will work the same as before (streaming data to a Message Collector). However, if you upgrade to the SL1 Extended Architecture, ScienceLogic recommends that you start streaming your agents to the Compute Node cluster instead of to a Message Collector.

To identify the version of the agent installed on a device:

- **Gen 1.** On the **Devices > Device Investigator > [Settings]** tab for that device, the **Collection Poller** field displays the name of the collector group that includes the Message Collector used by the agent. On the **Device Manager** page (Devices > Device Manager), the **Collection Group** column displays the name of the collector group that includes the Message Collector used by the agent.
- **Gen 3.** On the **Devices > Device Investigator > [Settings]** tab for that device, the **Collection Poller** field displays the label *Agents*. On the **Device Manager** page (Devices > Device Manager), the **Collection Group** column displays the label *Agents*.

The following table provides an overview of the features available when using the Gen 1 or the Gen 3 agent:

Product Capability	Gen 1 Agent	Gen 3 Agent
<b>System Vitals</b>		
Availability and Uptime	Yes	Yes
CPU, Memory, File System, and Network Interface	CPU and Memory only	Yes
Processes and Windows Services	Processes only	Yes
Installed Software	No	Yes
<b>Log Monitoring</b>		
Event Logs	Yes	Yes
Syslog	Yes	Yes
Text Logs	Yes	Yes
<b>Extensible Collection</b>		
PowerShell	No	Yes
JMX	No	Yes

For a current list of supported operating systems and host system requirements for the SL1 agent, see the [System Requirements for the Agent](#) page at the ScienceLogic Support Site.

**NOTE:** Users who are running version 102 or later of the *Microsoft: Windows Server PowerPack* can collect data via the SL1 agent. For more information, see the **Monitoring Windows Systems with PowerShell** manual.

## What Kind of Data Can an Agent Collect?

In the **SL1 Extended Architecture** (which includes Compute Nodes, Storage Nodes, and a Management Node), the Gen 3 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.
- **Logs.** The SL1 agent can be configured to push logs to SL1 that match specific criteria from a log file or the Windows Event Log. You can view logs collected by the SL1 agent on the Logs pane of the **Device Investigator** page. The same logs also appear on the **[Logs]** tab in the **Device Properties** and **Device Summary** pages for that device. You can define event policies that specify how logs collected by an agent will trigger events.

- **Host Performance Metrics.** Using Dynamic Applications, SL1 translates data provided by an SL1 agent to trend the following metrics:
  - Overall CPU Utilization
  - CPU Utilization Breakdown
  - Disk Average Queue Length
  - Disk IO Utilization
  - Memory Utilization
  - Network Bytes Read
  - Network Bytes Written
  - Storage Available
  - Storage Total
  - Storage Utilization
  - Swap Utilization

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

- **Host Configuration.** Using a Dynamic Application, SL1 collects the following configuration data based on data provided by the SL1 Agent:
  - The number and speed of the installed CPUs
  - The amount of installed memory
  - The overall and per-disk storage size
  - The total swap capacity (SL1 Extended Architecture only)

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

- **Network Interface.** The SL1 agent collects a list of the network interfaces running on the device. You can view the list of interfaces on the **[Interfaces]** tab of the **Device Investigator** page and the **Device Summary** page. This list includes attributes such as the interface MAC address, IP address, position, and speed as well as inbound and outbound utilization, number of errors, and discard and usage percentage.
- **File System.** The SL1 agent collects data about the of configuration of the file systems found within a device, such as name, size and, type as well as utilization data such as free space, size, and usage percentage. You can view the file system data on the **[Hardware?]** tab of the **Device Investigator** page and the **Device Summary** page.
- **System Processes.** The SL1 agent collects a list of all processes running on the device, such as name, process ID (PID), and state. You can view the list of processes on the **[Processes]** tab of the **Device Investigator** page and the **[Processes]** tab of the **Device Summary** page. Monitoring policies can be configured to trend and alert on process availability, process CPU usage, and process memory usage.
- **Windows Services.** The SL1 agent collects a list of all Windows services enabled on the device. This list includes attributes such as the service name and run state. You can view the list of Windows Services on the **[Services]** tab of the **Device Investigator** page and the **Device Summary** page.



- **Installed Software.** The SL1 agent collects a list of the software running on the device. This list includes attributes such as software name, version, and installation date. You can view the list of software on the **[Software]** tab of the **Device Investigator** page and the **Device Summary** page.

## Extensible Collection

In addition to the capabilities listed above, you can use the SL1 agent for "extensible collection", where you align the agent with Dynamic Applications to gather metrics and attributes from other infrastructures and applications.

The SL1 Extended Architecture supports aligning PowerShell Dynamic Applications to devices monitored by the SL1 Windows Agent. The SL1 Extended Architecture supports aligning JMX Dynamic Applications to devices monitored by the SL1 Linux Agent.

For more information, see [Configuring Extensible Collection](#).

---

## Agent PowerPacks

SL1 includes two PowerPacks that can be used to collect agent-based system configuration and performance data: the **ScienceLogic: Agent PowerPack** and the **Host Agent PowerPack**.

Both PowerPacks are installed by default on your SL1 system, and they include the following features:

- Dynamic Applications that collect configuration data and performance metrics from devices that are using agent-based collection
- Event Policies and alerts that are triggered when devices that are using agent-based collection meet certain status criteria

The **ScienceLogic: Agent PowerPack** collects agent-based data for devices on SL1 systems running on the SL1 Extended Architecture (Gen 3 agents). This PowerPack contains two Dynamic Applications:

- The "ScienceLogic Agent: System Configuration" Dynamic Applications collects the following data:
  - CPU
  - CPU Information
  - CPUs
  - Hardware Totals
  - Memory
  - Speed (MHz)
  - Swap Capacity
- The "ScienceLogic Agent: System Performance" Dynamic Applications collects the following data:
  - CPU Name
  - CPU Utilization
  - CPU Utilization Breakdown

- Disk Average Queue Length
- Disk IO Utilization
- Disk Name
- Memory Utilization
- Network Read
- Network Write
- Sample Time
- Swap Utilization

The **Host Agent PowerPack**, which collects agent-based data for devices on SL1 systems running on a distributed architecture (Gen 1 agents). This PowerPack contains two Dynamic Applications:

- The "Host Agent: System Config" Dynamic Applications collects the following data:
  - CPU
  - CPU Information
  - CPUs
  - Disk
  - Disk Information
  - Disk Space
  - Disks
  - Hardware Totals
  - Memory
  - Size
  - Speed (MHz)
- The "Host Agent: System Perf" Dynamic Applications collects the following data:
  - CPU Name
  - CPU Utilization
  - CPU Utilization Breakdown
  - Disk Average Queue Length
  - Disk Name
  - Disk Utilization
  - Memory Utilization
  - Network Read
  - Network Write
  - Sample Time
  - Storage Available

- Storage Name
- Storage Total
- Storage Utilization

**NOTE:** Because the *ScienceLogic: Agent PowerPack* is required to collect data from devices that are using agent-based collection, SL1 does not enable you to delete or modify this PowerPack.

## Agent-Compatible PowerPacks

In addition to the *ScienceLogic: Agent PowerPack* and the *Host Agent PowerPack*, there are several other Agent-compatible PowerPacks you can use to collect data from specific device types.

### Windows Devices

The following PowerPacks include the SL1 Agent PowerShell Default credential and SL1 Agent device template, which you can use to execute the SL1 Agent on Windows devices with PowerShell:

- Microsoft: Windows Server
- SL1 Agent Templates for Microsoft PowerPacks, which includes templates for the following:
  - Microsoft: DHCP Server
  - Microsoft: DNS Server
  - Microsoft: Exchange Server
  - Microsoft: IIS Server
  - Microsoft: Lync Server
  - Microsoft: SharePoint Server
  - Microsoft: SQL Server
  - Microsoft: Windows Server

**NOTE:** For more information, see the *Monitoring Windows with PowerShell* manual.

### Java Management Extensions (JMX) Resources

You can also use the *JMX Base Pack \*BETA\** PowerPack to monitor JMX resources with the SL1 agent.

For more information, see the *Monitoring Java Management Extensions (JMX)* manual.

---

## Agent Architecture

The following sections describe how the SL1 agent works in a distributed architecture and in the SL1 Extended Architecture.

### SL1 Distributed Architecture


In a distributed architecture, the SL1 Agent collects data from the device on which it is installed and transfers that data to a Message Collector in an SL1 system using the HTTPS protocol. 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 Message Collector and the device on which an agent is installed.

In a distributed architecture, the SL1 agent requires a standalone, dedicated Message Collector. The Message Collector does not need to be dedicated to agent usage, but the Message Collector cannot be a Data Collector that also performs message collection

**NOTE:** Message Collectors that process data from the SL1 agent have different system requirements than Message Collectors that do not process data from the SL1 agent. For more information about the system requirements when running SL1 agents in a distributed architecture, see <https://support.sciencelogic.com/s/system-requirements>.

The diagram below shows the collection layer of a distributed system containing both Data Collectors and Message Collectors in which the SL1 Agent is installed on a managed device.



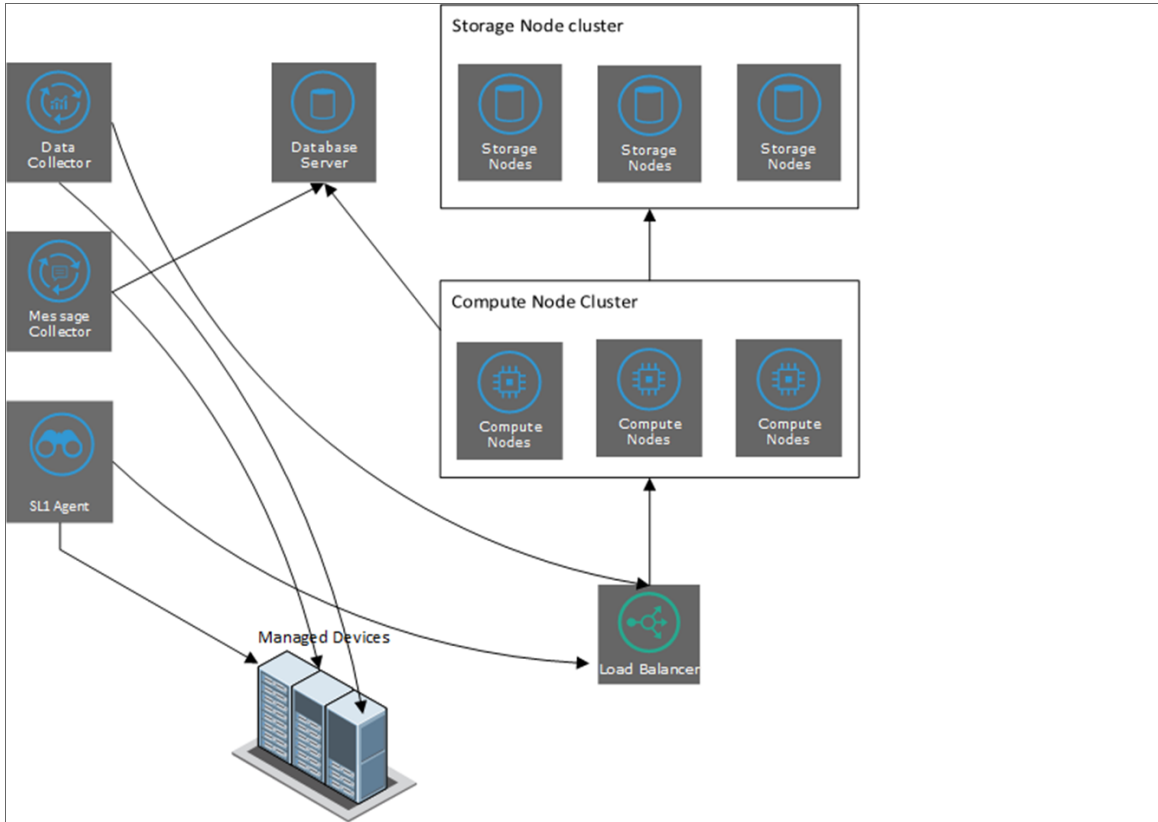
## SL1 Extended Architecture

In the SL1 Extended Architecture, an **SL1 agent** collects data from the device on which it is installed and sends that data to a Load Balancer in front of a Compute Cluster. The Compute Cluster transforms the data and stores high-volume performance data in the Storage Cluster and other performance and configuration data in the Database Server.

If required, agents can use an HTTP proxy server as an intermediate step in sending data to SL1.

In the diagram below:

- The **SL1 agent** collects data from managed devices and sends the data to the Load Balancer and Compute Node cluster for processing.
- The optional **Message Collector** collects asynchronous traps and syslog messages and sends them to the Database Server.
- The **Data Collector** collects data from managed devices and sends the data to the Load Balancer and Compute Node cluster for processing and then storage.



Using an SL1 agent in the SL1 Extended Architecture provides more configuration and performance data than using an SL1 agent in a distributed architecture. This additional data includes system vitals, log data, and extensible collection.

**NOTE:** For more information about the system requirements when running SL1 agents in an extended architecture, see <https://support.sciencelogic.com/s/system-requirements>.

---

# Chapter

# 2

## Installing the SL1 Agent

---

### Overview

This chapter describes how to install, upgrade, and uninstall SL1 agents for Windows, Linux, AIX, and Solaris operating systems. All agent installations, upgrades, and uninstallations require command-line interface access to the targeted device.

- If you use the **SL1 Extended Architecture** (which includes Compute Nodes, Storage Nodes, and a Management node), install the agent from the **Agents** page (Devices > Agent). For details, see [Installing an Agent from the Agents Page](#).
- If you use a **distributed SL1 architecture** (where the SL1 Agent collects and transfers data to a Message Collector), install the agent from the **Device Manager** page (Devices > Device Manager). For details, see [Installing an Agent from the Device Manager Page](#).

**NOTE:** If you are using PhoneHome Message Collectors in a distributed architecture, see [Configuring the SL1 Agent to Work with PhoneHome Collectors](#) before you install the SL1 agent.

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

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

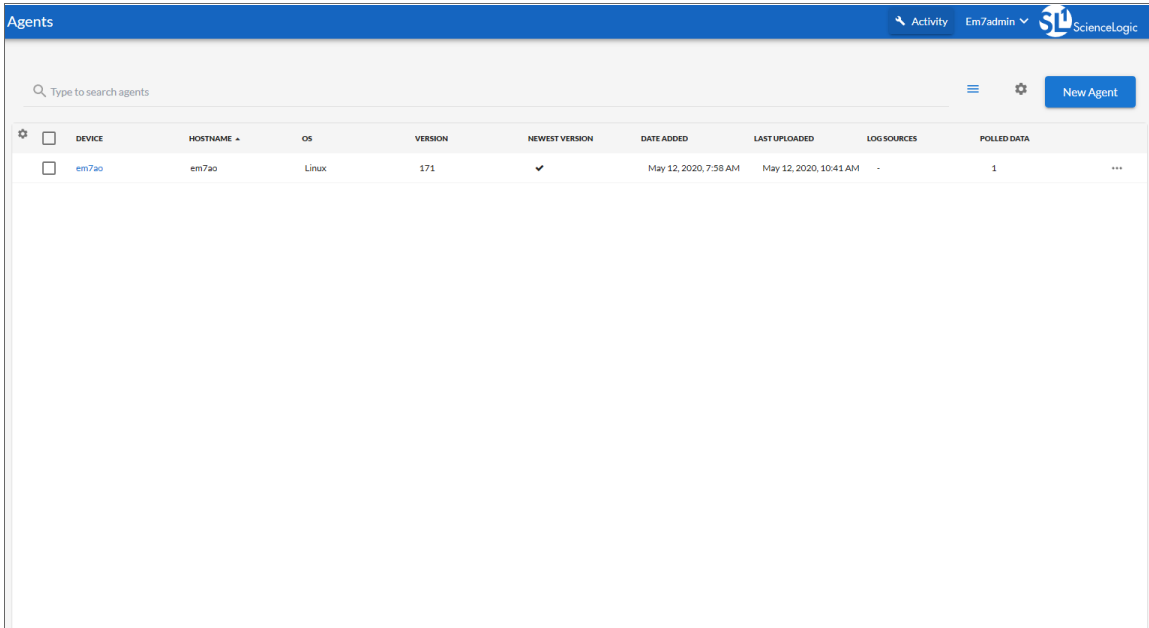
This chapter includes the following topics:

<a href="#">Installing an Agent from the Agents Page</a> .....	16
<a href="#">Upgrading the Agent</a> .....	23
<a href="#">Stopping the Agent</a> .....	24

<a href="#">Uninstalling the Agent</a> .....	24
<a href="#">Installing an Agent from the Device Manager Page</a> .....	25

## Installing an Agent from the Agents Page

When running SL1 on the SL1 Extended Architecture, you can use the **Agents** page (Devices > Agents) to help you install the SL1 agent on a Linux, Windows, AIX, or Solaris device:



The screenshot shows the 'Agents' page in the SL1 interface. At the top, there is a search bar with the placeholder text 'Type to search agents'. To the right of the search bar are a hamburger menu icon, a settings gear icon, and a 'New Agent' button. Below the search bar is a table with the following columns: DEVICE, HOSTNAME, OS, VERSION, NEWEST VERSION, DATE ADDED, LAST UPLOADED, LOG SOURCES, and POLLED DATA. There is one row of data in the table:

DEVICE	HOSTNAME	OS	VERSION	NEWEST VERSION	DATE ADDED	LAST UPLOADED	LOG SOURCES	POLLED DATA
em7ao	em7ao	Linux	171	✓	May 12, 2020, 7:58 AM	May 12, 2020, 10:41 AM	-	1

**TIP:** The **Agents** page appears as an option on the **Devices** menu *only* when you are using the SL1 Extended Architecture. If your version of SL1 does not have an **Agents** page, see [Installing an Agent from the Device Manager Page](#).

**NOTE:** If you are using an SL1 system with the SL1 Extended Architecture, you **must** use the **Agents** page (Devices > Agents) to install agents, including AIX and Solaris agents. AIX and Solaris agents are not available for non-SL1 Extended systems.



You can also upgrade and delete agents from SL1 on the **Agents** page.

**TIP:** If you are looking for a very specific set of agents, click the gear icon (⚙️) to the right of the **Search** field and select *Advanced*. In this mode you can create an advanced search using "AND" or "OR" for multiple search criteria. For more information, see the "Performing an Advanced Search" topic in the *Introduction to SL1* manual.

You do not install the agent from the **Agents** page itself. Instead, the **Agents** page enables you to gather the information or files you need to then install the agent on a particular device.

The instructions on the relevant tab of the **Agents** page also include command line parameters such as **URLFRONT**, which is used to define the Streamer service you are using as part of the SL1 Extended Architecture.

If the agent installation is successful, 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.

During initial discovery, the agent returns operating system type and version information to SL1. Based on this information, SL1 assigns the corresponding device classes to a device monitored only by an agent. If a device is monitored by an agent and via SNMP, the device class assigned by SNMP discovery will take precedence.

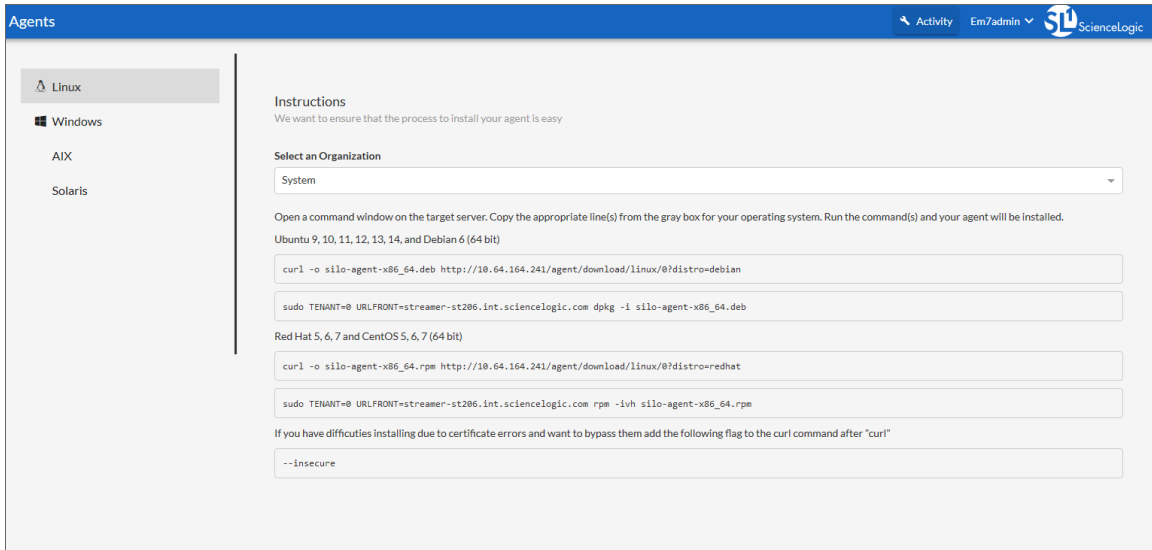
## Installing a Linux Agent

For a Linux system, the **Agents** page provides version-specific commands that you will need to run on the Linux system you want to monitor.

**NOTE:** If you are installing on SELinux, see [Using Agents with SELinux](#).

To install a Linux agent:

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



2. Click the **[Linux]** tab at top left.
3. From the **Select an Organization** drop-down list, select an organization for the device on which you are installing the agent. The **TENANT** value in the commands on this page is updated with the SL1 Organization ID for the organization you selected.

**TIP:** The **URLFRONT** value is the URL for the Streamer service you are using as part of the SL1 Extended Architecture.

4. Based on the version of Linux running on the device you want to monitor, run the relevant commands from the **[Linux]** tab on the Linux device (the "target server"). After the commands complete, the Linux agent starts running in the background. The device on which you installed the agent appears on the **Agents** page (Devices > Agents), and the device is also added to the list of devices on the **Devices** page.

**NOTE:** SL1 supports the use of proxy server connections when using the SL1 agent on Linux systems. To do so, open a command window on the target server and first configure curl to use a proxy server using the **CURLOPT\_PROXY** option, and then to use a username and password combination for that proxy server using the **CURLOPT\_PROXYUSERPWD** option.

## Using Agents with SELinux

The Linux agent supports running in SELinux for CentOS or Red Hat in both permissive and enforcing modes. The RPM installer for Linux will detect if SELinux is active (permissive/enforcing) upon installation. The agent automatically creates and installs an appropriate SELinux policy for use with the agent daemon.

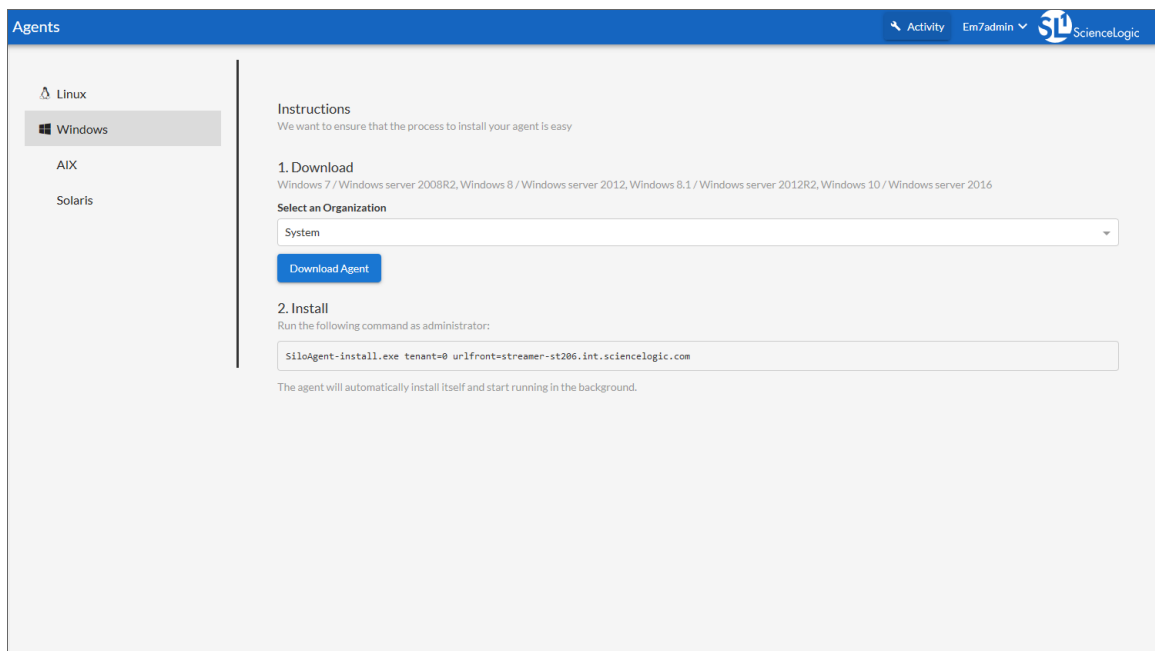
With the Linux Agent, a shared library gets injected into other processes. For non-root processes, this shared library should not generate any errors or audit logs. However, for root processes, this shared library might generate some additional audit logs where some of the shared library functionality is restricted by SELinux. If this is the case, you can update the agent configuration to exclude pre-loading into these root processes.

## Installing a Windows Agent

For a Windows system, the **Agents** page provides an executable file that you will need to run on the Windows system you want to monitor.

To download and install a Windows agent in SL1 :

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



3. From the **Select an Organization** drop-down list, select an organization for the new agent. The **tenant** value in the command on this page is updated with the SL1 Organization ID for the organization you selected.
4. Click **[Download Agent]**. An **Opening SiloAgent-install.exe** dialog appears.
5. Save the **SiloAgent-install.exe** file for installing the agent.
6. Copy the **SiloAgent-install.exe** file you downloaded to the Windows system on which you want to run the agent. To do so, you can either go to the console of the Windows system or use a utility like WinSCP.

7. To install the agent, use Windows PowerShell as an Administrator to run the command on the Windows system (the command is listed in step 2 on the **Windows** tab on the **Agents** page):

```
.\SiloAgent-install.exe tenant=<organization_id> urlfront=<streamer_URL>
```

where <organization\_id> is Organization ID for the organization you selected in step 3, and <streamer\_URL> is the URL for the Streamer service you are using as part of the SL1 Extended Architecture. For example:

```
.\SiloAgent-install.exe tenant=1 urlfront=streamer-2020.int.sciencelogic.com
```

8. When the agent installation completes, the PowerShell window displays "info: install complete", and the Windows agent starts running in the background. The device on which you installed the agent appears on the **Agents** page (Devices > Agents), and the device is also added to the list of devices on the **Devices** page.
9. 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.

**NOTE:** SL1 supports the use of proxy server connections when using the SL1 agent on Windows systems. In Windows 10, go to Settings > Network & Internet > Proxy to set up a proxy; in older versions of Windows, you can do this by clicking on the Control Panel and going to Internet Options > Connections > LAN Settings. If a proxy server requires authentication, you will need to provide a username and password before traffic is allowed through the proxy server and to the Internet.

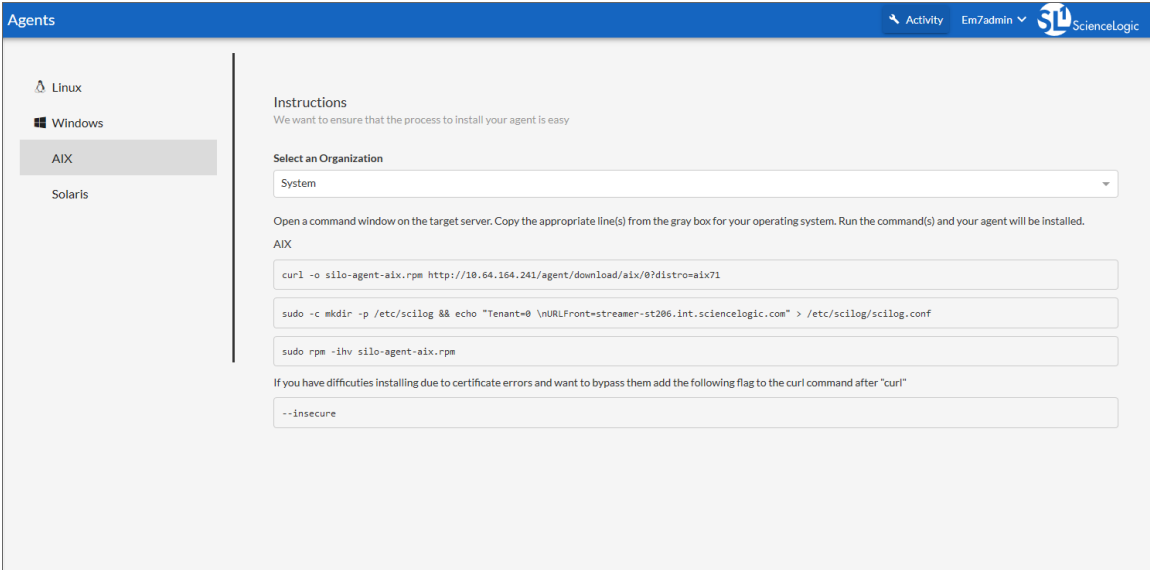
## Installing an AIX Agent

For an AIX system, the **Agents** page provides a set of commands that you will need to run on the AIX system you want to monitor.

**NOTE:** The AIX agent is a Beta feature in SL1 version 10.1.0.

To download and install an AIX agent:

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



The screenshot shows the 'Agents' page in the ScienceLogic interface. The left sidebar has tabs for Linux, Windows, AIX (selected), and Solaris. The main content area is titled 'Instructions' and contains the following text and form elements:

**Instructions**  
We want to ensure that the process to install your agent is easy

**Select an Organization**  
System

Open a command window on the target server. Copy the appropriate line(s) from the gray box for your operating system. Run the command(s) and your agent will be installed.

**AIX**

```
curl -o silo-agent-aix.rpm http://10.64.164.241/agent/download/aix/?distro=aix71
```

```
sudo -c mkdir -p /etc/scilog && echo "Tenant=0 \nURLFront=streamer-st206.int.sciencelogic.com" > /etc/scilog/scilog.conf
```

```
sudo rpm -ihv silo-agent-aix.rpm
```

If you have difficulties installing due to certificate errors and want to bypass them add the following flag to the curl command after "curl"

```
--insecure
```

2. Click the **[AIX]** tab.
3. From the **Select an Organization** drop-down list, select an organization for the device on which you are installing the agent. The **Tenant** value in the commands on this page is updated with the SL1 Organization ID for the organization you selected.

**TIP:** The **URLFront** value is the URL for the Streamer service you are using as part of the SL1 Extended Architecture.

4. Run the relevant commands from the **[AIX]** tab on the AIX device (the "target server"). After the commands complete, the AIX agent starts running in the background. The device on which you installed the agent appears on the **Agents** page (Devices > Agents), and the device is also added to the list of devices on the **Devices** page.

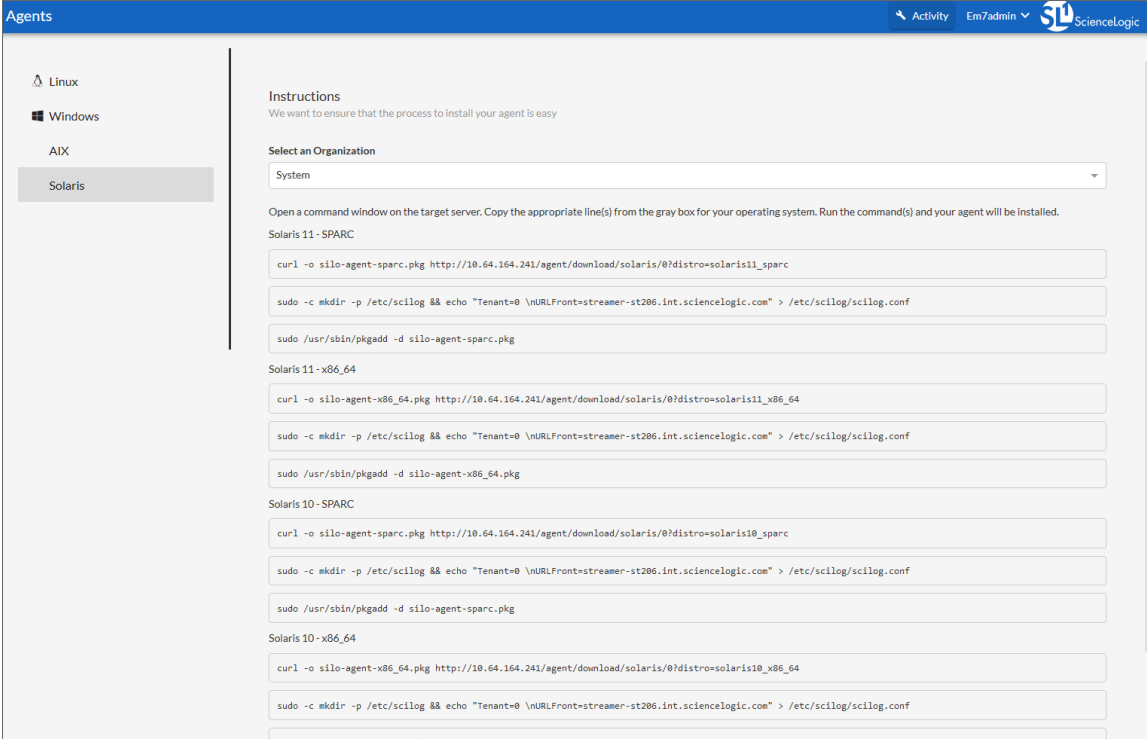
## Installing a Solaris Agent

For a Solaris system, the **Agents** page provides version-specific commands that will need to run on the Solaris system you want to monitor.

**NOTE:** The Solaris agent is a Beta feature in SL1 version 10.1.0.

To download and install a Solaris agent:

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



The screenshot shows the 'Agents' page in the ScienceLogic console. The sidebar on the left has tabs for Linux, Windows, AIX, and Solaris. The main content area is titled 'Instructions' and contains the following text: 'We want to ensure that the process to install your agent is easy'. Below this is a 'Select an Organization' dropdown menu with 'System' selected. The instructions then state: 'Open a command window on the target server. Copy the appropriate line(s) from the gray box for your operating system. Run the command(s) and your agent will be installed.' There are four sections of commands, each in a gray box, corresponding to different Solaris configurations: Solaris 11-SPARC, Solaris 11-x86\_64, Solaris 10-SPARC, and Solaris 10-x86\_64. Each section contains three lines of commands: a curl command to download the agent package, a mkdir command to create a directory and echo the tenant information into a config file, and a pkgadd command to install the package. The tenant information is represented by a placeholder URL.

3. From the **Select an Organization** drop-down list, select an organization for the device on which you are installing the agent. The **Tenant** value in the commands on this page is updated with the SL1 Organization ID for the organization you selected.

**TIP:** The **URLFront** value is the URL for the Streamer service you are using as part of the SL1 Extended Architecture.

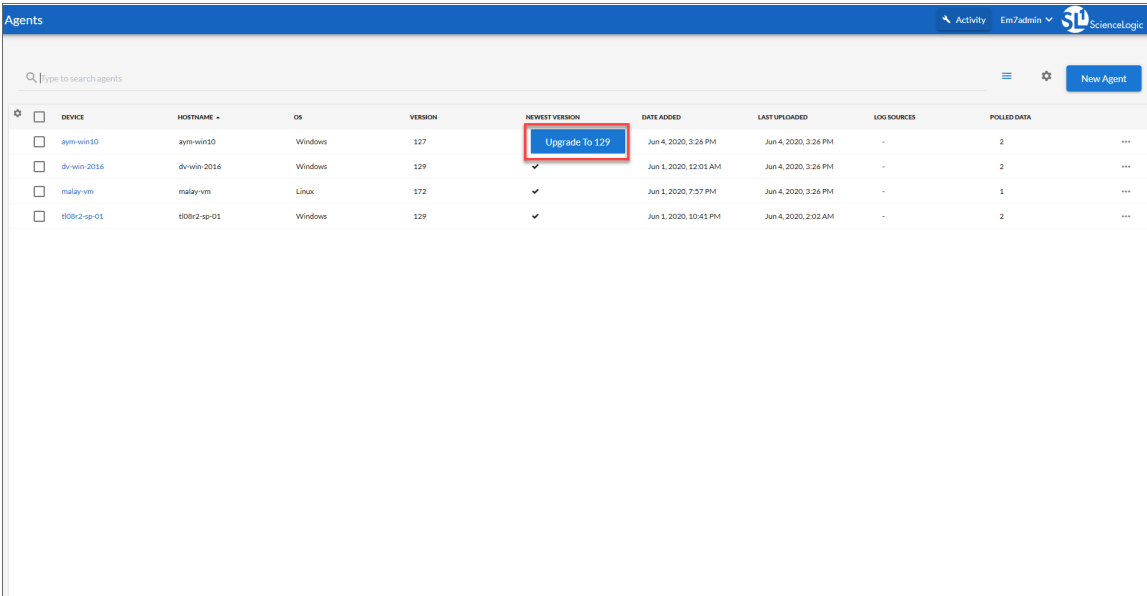
- Based on the version of Solaris running on the device you want to monitor, run the relevant commands from the **[Solaris ]** tab on the Solaris device (the "target server"). After the commands complete, the Solaris agent starts running in the background. The device on which you installed the agent appears on the **Agents** page (Devices > Agents), and the device is also added to the list of devices on the **Devices** page.

## Upgrading the Agent

When you have the latest version of the SL1 agent, a check mark icon ( ✓ ) appears in the **Newest Version** column for that agent. If you do not have the latest version of the agent, an **[Upgrade]** button appears in the **Newest Version** column on the **Agents** page.

To upgrade to the latest version of an agent:

- On the **Agents** page (Devices > Agents), locate the agent you want to upgrade:



DEVICE	HOSTNAME	OS	VERSION	NEWEST VERSION	DATE ADDED	LAST UPLOADED	LOG SOURCES	POLLED DATA
<input type="checkbox"/> aym-win10	aym-win10	Windows	127	<a href="#">Upgrade to 129</a>	Jun 4, 2020, 3:26 PM	Jun 4, 2020, 3:26 PM	-	2
<input type="checkbox"/> dv-win-2016	dv-win-2016	Windows	129	✓	Jun 1, 2020, 12:01 AM	Jun 4, 2020, 3:26 PM	-	2
<input type="checkbox"/> malay-vm	malay-vm	Linux	172	✓	Jun 1, 2020, 7:57 PM	Jun 4, 2020, 3:26 PM	-	1
<input type="checkbox"/> t00r2-sp-01	t00r2-sp-01	Windows	129	✓	Jun 1, 2020, 10:41 PM	Jun 4, 2020, 2:02 AM	-	2

- Click the **[Upgrade]** button for that agent, or click the **[Actions]** button ( ⋮ ) and select *Upgrade*. The agent starts the upgrade process, which might take up to a minute to complete.
- Click the **[Refresh]** button in your browser to ensure that a check mark icon ( ✓ ) appears in the **Newest Version** column for that agent.

---

## Stopping the Agent

You can use the "delete" option on the **Agents** page to stop an agent from running. When you delete an agent, SL1 stops the agent from running on the device, deletes the data gathered by that agent, and removes that agent from the **Agents** page.

**NOTE:** Using the "delete" option for an agent **does not** actually remove the agent from the device. As a best practice, use the delete process to delete the data gathered by the agent (the uninstallation process does not delete this data), and then uninstall that agent, if needed. For uninstallation details, see [Uninstalling an Agent](#). Also, ScienceLogic recommends that you delete the *device* that had the agent installed on it, or at least remove the agent-related Dynamic Applications, to prevent SL1 from attempting to use the data the agent gathered from that device.

To delete an agent and its data:

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

**NOTE:** If you used the **Agents** page to install an agent, then you should also use the **Agents** page to delete that agent. If you attempt to delete the agent using the classic SL1 user interface rather than the **Agents** page, SL1 deletes only **some** of the data gathered by the agent, rather than all of the data.

---

## Uninstalling the Agent

When you uninstall an agent, you remove that agent completely from SL1, but you do not lose the data collected by that agent.

**NOTE:** After you uninstall the agent, the **scilog.conf** file remains on the device. If you think you might want to monitor with that agent again in the future, you should delete the **scilog.conf** file after you run the uninstallation.



## Uninstalling a Linux Agent

To uninstall an agent on a Linux system:

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

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

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

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

## Uninstalling a Windows Agent

To uninstall an agent on a Windows system:

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

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

---

## Installing an Agent from the Device Manager Page

If you are using a distributed SL1 system *without* the SL1 Extended Architecture (which includes Compute Nodes, Storage Nodes, and a Management Node), you **must** use the **Agent Installation** page to install the agent. You can access the **Agent Installation** page from the **Device Manager** page.

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

**NOTE:** If you are using an SL1 system with the SL1 Extended Architecture, you **must** use the **Agents** page (Devices > Agents) to install agents, including AIX and Solaris agents. AIX and Solaris agents are not available for non-SL1 Extended systems.


**NOTE:** If you are *not* using the SL1 Extended Architecture, the **Agents** page (Devices > Agents) will not be available in the SL1 user interface.

To install an agent from the **Device Manager** page, you first need to gather installation information from the **Agent Installation** page:

- For a Linux system, the **Agent Installation** page provides commands that must be executed on that system.
- For a Windows system, the **Agent Installation** page provides an executable file to run on the Windows system.

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

1. Go to the **Device Manager** page (Device > Device Manager in the SL1 user interface, or Registry > Devices > Device Manager in the classic user interface).
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 SL1 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 the Linux Agent from the Device Manager Page

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=Silco/L=Reston/O=Silco/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.
```

**NOTE:** SL1 supports the use of proxy server connections when using the SL1 agent on Linux systems. To do so, open a command window on the target server and first configure curl to use a proxy server using the `CURLOPT_PROXY` option, and then to use a username and password combination for that proxy server using the `CURLOPT_PROXYUSERPWD` option.

## Updating an Agent on a Linux System

To update the 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 **first** command that you copied from the **Agent Installation** page.
3. 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_64
```

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

```
/etc/scilog
```

## Installing the Windows Agent from the Device Manager Page

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=<organization_id> 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.

**NOTE:** SL1 supports the use of proxy server connections when using the SL1 agent on Windows systems. In Windows 10, go to Settings > Network & Internet > Proxy to set up a proxy; in older versions of Windows, you can do this by clicking on the Control Panel and going to Internet Options > Connections > LAN Settings. If a proxy server requires authentication, you will need to provide a username and password before traffic is allowed through the proxy server and to the Internet.

## 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
```

---

# Chapter

# 3



## Configuring the SL1 Agent

---

### Overview

This chapter covers how to configure the agent monitoring based on the SL1 architecture you are using. This chapter also describes additional agent configurations, such as enabling extensible collection with an SL1 agent.

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

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

This chapter includes the following topics:

<a href="#">Configuring Agent Monitoring Based on SL1 Architecture</a> .....	31
<a href="#">Configuring Agent Settings from the Device Investigator Page</a> .....	32
<a href="#">Configuring Extensible Collection</a> .....	35
<a href="#">Configuring an SL1 Agent on the Device Manager Page</a> .....	37

---

## Configuring Agent Monitoring Based on SL1 Architecture

Based on your SL1 architecture, you will have access to different workflows for configuring the monitoring settings on an SL1 agent:

- If the SL1 agent is used in an **SL1 Extended Architecture** (which includes Compute Nodes, Storage Nodes, and a Management node), use the following configuration workflows:
  - To configure basic settings related to how the agent collects data, see [Configuring Agent Settings from the Device Investigator Page](#).
  - To configure log file monitoring and Windows event log monitoring with the agent, see [Monitoring Device Logs Using the Agent](#).
  - To configure and view data collected by the agent, see the relevant tabs on the **Device Investigator** page, including the following tabs:
    - [Settings]
    - [Interfaces]
    - [Processes]
    - [Services]
    - [Collections]

For more information about these tabs, see the "Using the Device Investigator" chapter in the **Device Management** manual.

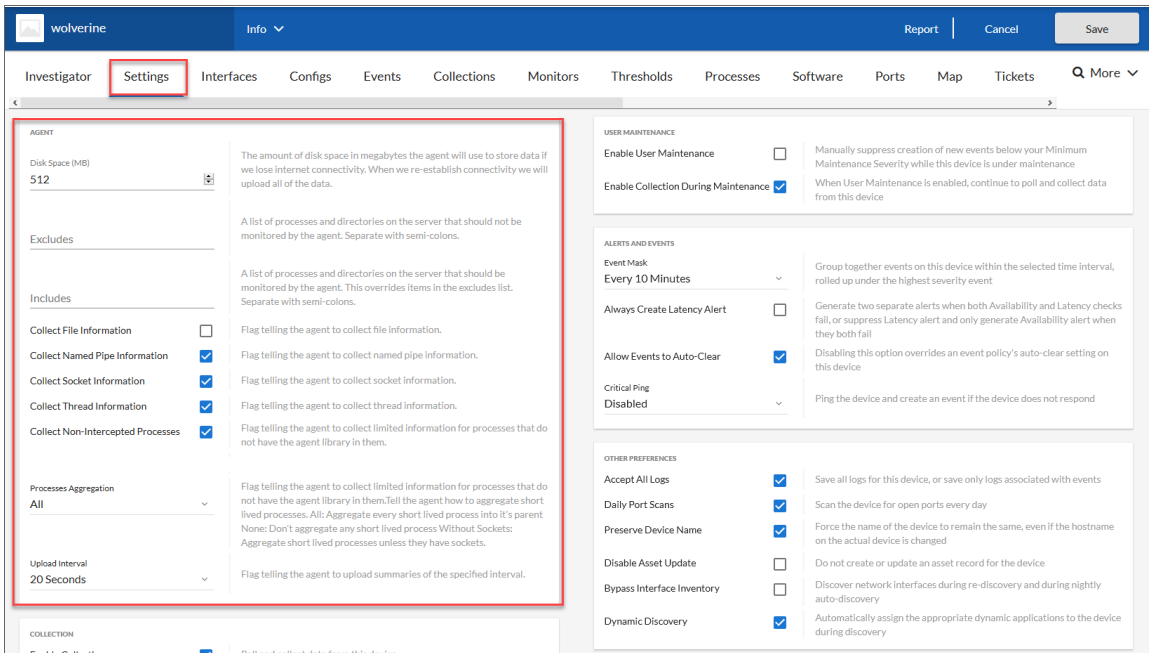
- If the SL1 agent is used in a **distributed SL1 architecture** (where the SL1 Agent collects and transfers data to a Message Collector), use the following configuration workflows:
  - To configure basic settings related to how the agent collects data, see [Configuring an SL1 Agent on the Device Manager Page](#).
  - To configure log file monitoring and Windows event log monitoring with the agent, see [Monitoring Device Logs Using the Agent](#).
  - To configure and view data collected by the agent, see the relevant tabs on the **Device Investigator** page, including the following tabs:
    - [Settings]
    - [Interfaces]
    - [Processes]
    - [Services]
    - [Collections]

For more information about these tabs, see the "Using the Device Investigator" chapter in the **Device Management** manual.

- To configure system vitals monitoring with the agent, see the "Monitoring Vitals Using an Agent" chapter of the *Device Infrastructure Health* manual.
- To configure port monitoring with the agent, see the "Monitoring Ports" chapter of the *Device Infrastructure Health* manual.
- To configure process monitoring with the agent, see the "Monitoring Processes Using an Agent" chapter of the *Device Infrastructure Health* manual.

## Configuring Agent Settings from the Device Investigator Page

If a device does *not* have an agent running on it, the **[Settings]** tab of the **Device Investigator** page for the device will display the **Collection** section at top left. However, if the device has an agent running on it, the **Agents** section appears above the **Collection** section:



In the **Agents** section, you can configure the disk space, excludes, includes, and other settings related to how you want the agent to collect data.

**NOTE:** The **Agents** section appears on the **[Settings]** tab *only* when you are using the SL 1 Extended Architecture. If your version of SL 1 does *not* have an **Agents** section, see [Configuring an SL1 Agent on the Device Manager Page](#).



To configure agent settings:

1. From the **Devices** page, select the device with the agent that you want to configure. The **Device Investigator** page appears.

**TIP:** To quickly find all devices on the **Devices** page that have agents installed on them, type "agent" in the **Search** field and select the "ANY: agent" search.

2. Click the **[Settings]** tab and click **[Edit.]**
3. In the **Agent** section, complete the following fields, as needed:

- **Disk Space.** Specify the amount of disk space in MB 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. When connectivity is re-established, the agent uploads all of its stored data.
- **Excludes.** Type a list of processes and directories, separated by semi-colons, that you do not want the agent to monitor.
- **Includes.** Type a list of processes and directories, separated by semi-colons, that you want the agent to monitor. This field ensures that specific processes are monitored.

**NOTE:** If a process or directory is included in both the **Excludes** field and the **Includes** field, the item in the **Includes** field will override the item in the **Excludes** field.

- **Collect File Information.** Select this option if you want the agent to report the names of files accessed by each monitored process.
- **Collect Named Pipe Information.** Select this option if you want the agent to collect named pipe information.
- **Collect Socket Information.** Select this option if you want the agent to collect socket information.
- **Collect Thread Information.** Select this option if you want the agent to collect thread information.
- **Collect Non-Intercepted Processes.** Select this option if you want the agent to collect limited information for processes that do not contain the agent library.
- **Processes Aggregation.** Specify how you want the agent to collect limited information for processes that do not have the agent library in them, and how to aggregate short-lived processes. Your options include the following:
  - *All:* Aggregate every short-lived process into its parent.
  - *None:* Do not aggregate any short-lived process.
  - *Without Sockets:* Aggregate short-lived processes unless those processes have sockets.
- **Upload Interval.** Specify how often the agent should upload data. Your options include the following:
  - **20 Seconds.** Upload a data snapshot every 20 seconds. This is the default setting, and it is the same interval used with the distributed SL1 architecture.
  - **60 Seconds.** Upload a data summary every 60 seconds. This option uses an improved data format that requires fewer SL1 resources. The SL1 agent continues to internally collect and poll data every 20 seconds, but the agent summarizes and uploads that data every 60 seconds. There is no data loss even though the data is uploaded less frequently.

**NOTE:** The **Upload Interval** option is currently available only with version 132 or later of the Windows SL1 agent. The Linux agent only uploads snapshots, not summaries.

4. Click the **[Save]** button to save your configuration settings for the agent.

**TIP:** You can view and configure additional agent settings from the **Device Investigator** page for the device on which the agent is installed, including the following tabs: **[Interfaces]**, **[Processes]**, **[Services]**, and **[Collections]**. For more information about these tabs, see the *Using the Device Investigator* chapter in the **Device Management** manual.

---

## Configuring Extensible Collection

You can use the SL1 agent for "extensible collection", where you align the SL1 agent with Dynamic Applications to gather metrics and attributes from other infrastructures and applications.

**Dynamic Applications** are customizable policies, created for a specific vendor and a specific type of device or system, that tell SL1:

- What data to collect from devices
- How to present the data that has been collected
- When to generate alerts and events based on the data that has been collected

SL1 includes Dynamic Applications for the most common hardware and software. You can customize these default Dynamic Applications to best work in your environment. You can also create custom Dynamic Applications.

You can align the agent with PowerShell Dynamic Applications (for Windows agents) and JMX Dynamic Applications (for Linux agents).

**TIP:** Before starting the configuration process, make sure that the **Account Type** is set to *Local* for the credential you are using to align the Dynamic Application to the agent. Do not use the Device SNMP credential.

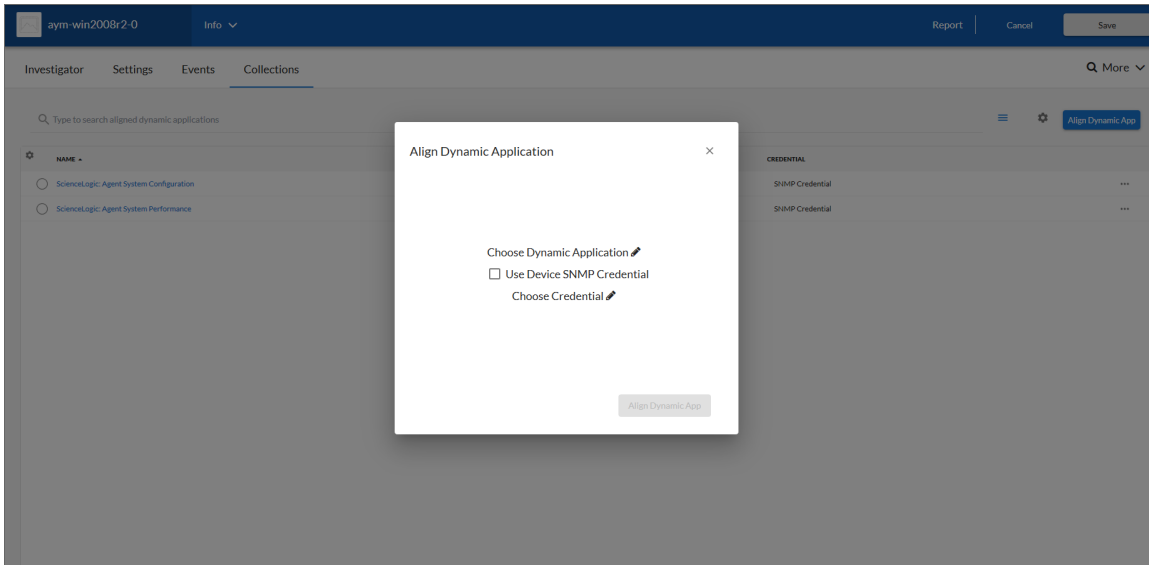
To enable extensible collection:

1. On the **Devices** page, search for the device with the agent installed on it.

**TIP:** To quickly find all devices that have agents installed on them, type "agent" in the **Search** field on the **Devices** page, and select the "ANY: agent" search.

2. Select the device. The **Device Investigator** page appears.

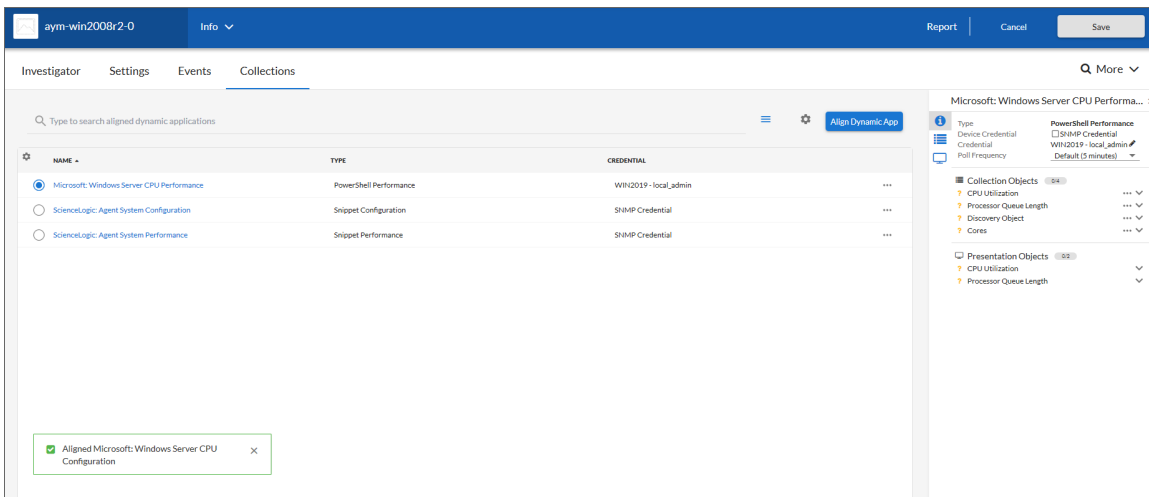
- Click the **[Collections]** tab, click **[Edit]**, and click **[Align Dynamic App]**. The **Align Dynamic Application** window appears:



- Click **Choose Dynamic Application** and search for the Dynamic Application you want to align.
- Select the Dynamic Application and click **[Select]**. The **Align Dynamic Application** window appears again.
- Click **Choose Credential** to search for a credential for the agent. The **Choose Credential** page appears.

**NOTE:** The **Account Type** for the credential must be *Local* for the agent. Do not use the SNMP credential.

- Select the credential and click **[Select]**. The **Align Dynamic Application** window appears again.
- Click **[Align Dynamic App]**. The **[Collections]** tab appears, and a pop-up message displays when the Dynamic Application is aligned.



9. Click **[Save]**.
10. To view the metrics from that newly aligned Dynamic Application, click the **[Investigator]** tab.
11. In the **Add a metric** field, select a metric from the new Dynamic Application to view its data in the right-hand pane.

**NOTE:** The metrics from the new Dynamic Application might take a minute or two to display in the list of metrics. Click **[Refresh]** in your browser if needed.

---

## Configuring an SL1 Agent on the Device Manager Page

**NOTE:** To configure agent settings for an agent, 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](#).

**NOTE:** If you are using the SL1 Extended Architecture, you can access these settings from the **[Settings]** tab of the **Device Investigator** page for the device on which the agent is installed. For more information, see [Configuring Agent Settings from the Device Investigator Page](#).

You can control how an agent runs on a device by configuring the following agent settings on the **Device Manager** page (Devices > Device Manager or Registry > Devices > Device Manager):

- **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.

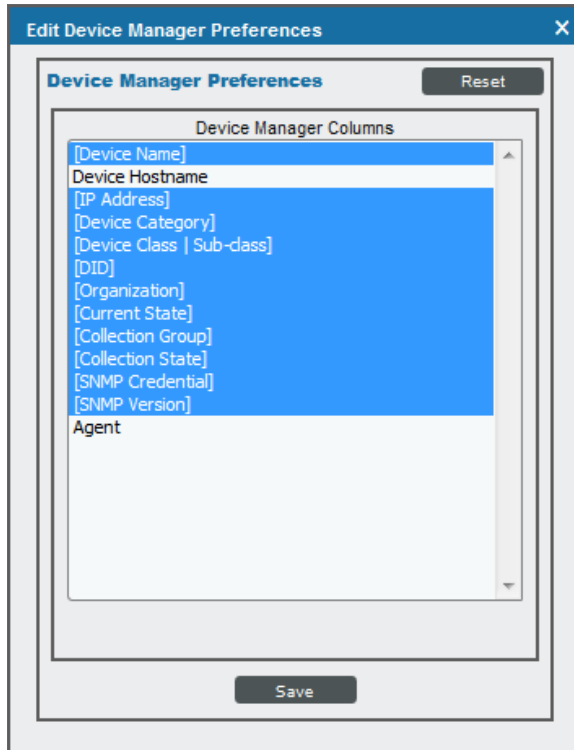
**NOTE:** If a process or directory is included in both the **Excludes** field and the **Includes** field, that process or directory will be monitored by the agent.

## 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.

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

1. Go to the **Device Manager** page (Devices > Device Manager or Registry > Devices > Device Manager in the classic user interface).
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 (Devices > Device Manager or 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:

**Agent Configuration** Reset

**Disk Space:** 512 MB  
The amount of disk space in megabytes the agent will use to store data if we lose internet connectivity. When we re-establish connectivity we will upload all of the data

**Data Directory:** /opt/scilog/data  
The directory on the server where the agent temporarily stores information.

**Excludes:**   
A list of processes and directories on the server that should NOT be monitored by the agent. Separate with semi-colons.

**Includes:**   
A list of processes and directories on the server that should be monitored by the agent. These override items on the exclude list, so an item in both will be monitored by the agent. Separate with semi-colons.

Save

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]**.

## Configuring the SL1 Agent to Work with PhoneHome Collectors

If you are using PhoneHome Message Collectors in a distributed architecture, you will need to configure the settings for the Message Collector to prevent SL1 from incorrectly configuring the SL1 agent to use the loopback IP address when it should be using the local IP address for the Message Collector.

If you do not perform the configuration steps below, when you try to install the agent from a Message Collector when it is configured as a PhoneHome Collector, the URLs for both the download and the internal configuration of the agent reference the loopback address instead of the accessible public/private IP address of the collector.

**NOTE:** For more information about PhoneHome Collectors, see the "Configuring SL1 for PhoneHome Communication" chapter in the *Installation* manual.

To configure the SL1 agent to work with a PhoneHome Collector:

1. Go to the **Appliance Manager** page (System > Settings > Appliances):

The screenshot shows the Appliance Manager interface. At the top, there are tabs for 'System Status', 'Output', 'Reset', and 'Guide'. Below these is a configuration form with fields for 'Host Name', 'IP Address', 'Model type', 'Database [2]', and 'Description'. There are 'Save' and 'Save As' buttons. Below the form is a table listing various appliances.

Name	IP Address	Module Type	Collector Group	Description	Build	MariaDB	Capacity	Allocation	ID	Validated Endpoint	Edit Date	Edit User	Create Date	
1. st206-cdb	10.64.164.241	Database	n/a	Database: 10.64.164.241	10.1.0.BETA3 r2097	10.4.12	7,000	n/a	1	Yes	--	2020-05-19 18:50:56	em7admin	2020-05-19
2. st206-ap	10.64.164.242	Administration Portal	n/a	Application Server: 10.64.164.242	10.1.0.BETA3 r2097	10.4.12	n/a	n/a	2	Yes	--	2020-05-19 18:49:02	em7admin	2020-05-19
3. st206-cu-benedict	10.64.164.253	Data Collection Unit	CUG-Benedict	collector unit: 10.64.164.253	10.1.0.BETA3 r2097	10.4.12	n/a	9	19	Yes	--	2020-05-19 18:52:07	em7admin	2020-05-19
4. st206-cu-moss1	10.64.164.251	Data Collection Unit	CUG_MOSS	collector unit: 10.64.164.251	10.1.0.BETA3 r2097	10.4.12	n/a	0	14	Yes	--	2020-05-19 18:51:58	em7admin	2020-05-19
5. st206-cu-moss2	10.64.164.252	Data Collection Unit	CUG_MOSS	collector unit: 10.64.164.252	10.1.0.BETA3 r2097	10.4.12	n/a	1	16	Yes	--	2020-05-19 18:52:01	em7admin	2020-05-19
6. st206-cu-racecondition1	10.64.164.246	Data Collection Unit	CUG-RaceCondition	collector unit: 10.64.164.246	10.1.0.BETA3 r2097	10.4.12	n/a	2	8	Yes	--	2020-05-19 18:51:36	em7admin	2020-05-19
7. st206-cu-racecondition2	10.64.164.247	Data Collection Unit	CUG-RaceCondition	collector unit: 10.64.164.247	10.1.0.BETA3 r2097	10.4.12	n/a	3	9	Yes	--	2020-05-19 18:51:37	em7admin	2020-05-19
8. st206-cu-racecondition3	10.64.164.248	Data Collection Unit	CUG-RaceCondition	collector unit: 10.64.164.248	10.1.0.BETA3 r2097	10.4.12	n/a	1	10	Yes	--	2020-05-19 18:51:35	em7admin	2020-05-19
9. st206-cu-rngs	10.64.164.245	Data Collection Unit	CUG-Solutions	collector unit: 10.64.164.245	10.1.0.BETA3 r2097	10.4.12	n/a	2	7	Yes	--	2020-05-19 18:51:39	em7admin	2020-05-19
10. st206-cu-scrumandcoke	10.64.164.254	Data Collection Unit	CUG-SAC	collector unit: 10.64.164.254	10.1.0.BETA3 r2097	10.4.12	n/a	0	21	Yes	--	2020-05-19 18:52:04	em7admin	2020-05-19
11. st206-cu-shared	10.64.164.244	Data Collection Unit	CUG-Solutions	collector unit: 10.64.164.244	10.1.0.BETA3 r2097	10.4.12	n/a	7	6	Yes	--	2020-05-19 18:51:40	em7admin	2020-05-19
12. st206-cu-usualsuspects1	10.64.164.249	Data Collection Unit	CUG-UsualSuspects	collector unit: 10.64.164.249	10.1.0.BETA3 r2097	10.4.12	n/a	2	11	Yes	--	2020-05-19 18:51:39	em7admin	2020-05-19
13. st206-cu-usualsuspects2	10.64.164.250	Data Collection Unit	CUG-UsualSuspects	collector unit: 10.64.164.250	10.1.0.BETA3 r2097	10.4.12	n/a	2	12	Yes	--	2020-05-19 18:51:51	em7admin	2020-05-19
14. st206-mc	10.64.164.243	Message Collection Unit	--	collector unit: 10.64.164.243	10.1.0.BETA3 r2097	10.4.12	n/a	n/a	5	Yes	10.64.164.243	2020-05-19 18:52:06	em7admin	2020-05-19
15. st206-cn01	10.64.164.233	Compute Node	n/a	Compute Node: 10.64.164.233	Unknown		n/a	n/a	54	Yes	--		em7admin	
16. st206-cn02	10.64.164.234	Compute Node	n/a	Compute Node: 10.64.164.234	Unknown		n/a	n/a	55	Yes	--		em7admin	
17. st206-cn03	10.64.164.235	Compute Node	n/a	Compute Node: 10.64.164.235	Unknown		n/a	n/a	56	Yes	--		em7admin	



2. In the **Message Collection Unit** row, click the Agent Endpoint Configuration icon (🌐) for the agent. The Agent Endpoint Configuration modal appears:

Agent Endpoint Configuration

Endpoint [5] Reset

**Appliance Details**  
collector unit: 10.64.164.243: 10.64.164.243

**Download Link Visibility**  
[ Enabled ]

**Endpoint Name for Download Link**  
st206-mc

**Hostname/IP for Agent Download**  
10.64.164.243

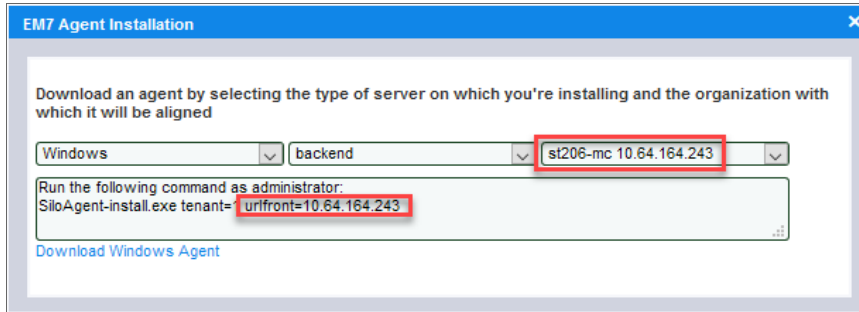
**Hostname/IP for Data Collection**  
(Leave blank to use value specified in streamer settings on appliance)

**Default CUG Alignment**  
[ CUG\_MOSS ]

Save

3. Make sure that the configuration information under the **Appliance Details** field is accurate.
4. Complete the following fields:

- **Download Link Visibility.** Select whether you want to display the download link in the **Agent Installation** modal page.
  - If you select *Enabled*, the link to the Message Collector appears on the **Agent Installation** modal, which is available on the **Device Manager** page (Devices > Device Manager or Registry > Device > Device Manager in the classic user interface) when you click the **[Actions]** button and select *Download/Install Agent*:



- If you select *Disabled*, the link to the Message Collector does not appear in the **Agent Installation** modal page.
- **Endpoint Name for Download Link.** This field contains the Message Collector hostname. If the hostname is not populated by default, you can enter the hostname.
  - **Hostname/IP for Agent Download.** This field contains the hostname or IP address for the agent download. If the hostname or IP address is not populated by default, you can enter the hostname or IP address.
  - **Hostname/IP for Data Collection.** You can add a hostname or IP address for the PhoneHome Collector, or you can leave this field blank to use a value specified in the Steamer container settings on the appliance.
  - **Default CUG Alignment.** Select the Collector Group aligned with the agent.
5. Click **Save**. The agent is now configured to work with a Message Collector when it is configured as a Phone Home Collector.



## Monitoring Logs Using the SL1 Agent

---

### Overview

This chapter describes how to use the agent to monitor logs with Log File Monitoring policies. SL1 supports multiple methods for ingesting log data, which you can use to generate events.

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

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

This chapter includes the following topics:

<i>What is a Log File Monitoring Policy?</i> .....	44
<i>Viewing the List of Log File Monitoring Policies</i> .....	44
<i>Creating a Log File Monitoring Policy</i> .....	46
<i>Aligning a Log File Monitoring Policy to Devices</i> .....	48
<i>Editing a Log File Monitoring Policy</i> .....	51
<i>Deleting Log File Monitoring Policies</i> .....	52
<i>Viewing the List of Log File Monitoring Policies and Aligned Devices</i> .....	52
<i>Creating an Event Policy for Agent Logs</i> .....	53
<i>Creating an Event Policy for Agent Logs in the Classic User Interface</i> .....	58

---

## What is 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 (System > Manage > Log File Monitoring Policies). 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.

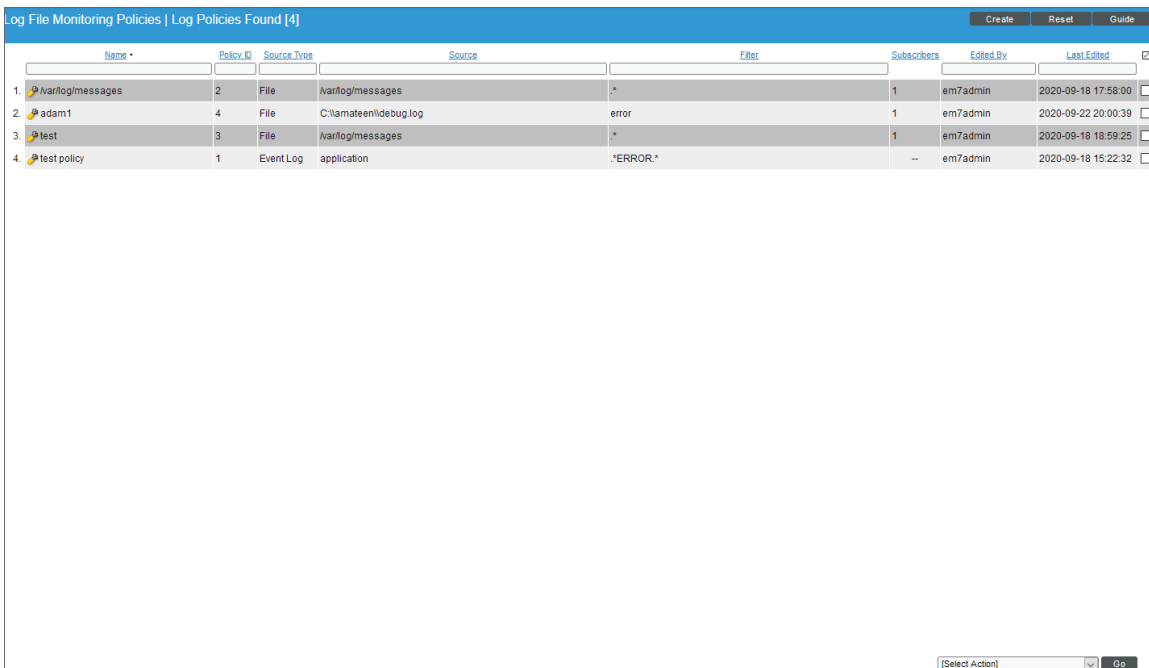
You can view logs collected by the SL1 agent on the **Logs** pane on the **Device Investigator** page for the device on which the agent is installed. The same logs also appear on the **[Logs]** tab in the **Device Properties** and **Device Summary** pages for that device. You can define event policies 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.



The screenshot shows the 'Log File Monitoring Policies' page with a table of 4 policies. The table has columns for Name, Policy ID, Source Type, Source, Filter, Subscribers, Edited By, and Last Edited. There are also buttons for 'Create', 'Reset', and 'Guide' at the top right.

Name	Policy ID	Source Type	Source	Filter	Subscribers	Edited By	Last Edited	
1. /var/log/messages	2	File	/var/log/messages	*	1	em7admin	2020-09-18 17:58:00	<input type="checkbox"/>
2. adam1	4	File	C:\amateen\debug.log	error	1	em7admin	2020-09-22 20:00:39	<input type="checkbox"/>
3. test	3	File	/var/log/messages	*	1	em7admin	2020-09-18 18:59:25	<input type="checkbox"/>
4. test policy	1	Event Log	application	.*ERROR.*	--	em7admin	2020-09-18 15:22:32	<input type="checkbox"/>

**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:

The screenshot shows a modal window titled "Log Monitoring Policy" with a close button (X) in the top right corner. The window has a blue header bar with the text "Create a new policy" and a "Reset" button on the right. Below the header, there are several form fields: "Name:" with a text input; "Type:" with a dropdown menu showing "[ File ]"; "File Path:" with a text input and a descriptive note: "The full path of the file from which to capture log messages. You can use a \* to match multiple files (eg: /var/log/httpd/\*.log)."; "Limit:" with a text input containing "2000" and a descriptive note: "The maximum number of messages to upload per minute from this source. A higher limit means you won't miss log messages, but it may cause a performance hit for uploading the data."; and "Filter:" with a text input and a descriptive note: "Capture log messages which match this string via RegEx. For example, .\*ERROR.\* to match log entries containing \"ERROR\"". A "Save" button is located at the bottom center of the form.

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. You can use a \* to match multiple files, such as `/var/log/em7/*.log`.
- **Source**. If you selected *Event Log* in the **Type** field, this field is displayed. Select the Windows log to monitor. Choices are:
  - *application*
  - *system*
  - *security*
- **Limit**. The maximum log messages the agent sends to SL1 per minute. If the number of matching logs exceeds this value, the agent will stop sending logs to the platform for the remainder of the minute. The limit resets at the beginning of the next minute. For example, suppose you set this field to *10,000*. Suppose the agent monitors a device that has 30,000 log messages. The agent will retrieve 10,000 logs and then wait until the beginning of the next minute. The agent will then retrieve the next 10,000 logs and then wait until the beginning of the next minute. The agent will continue to retrieve 10,000 logs per minute until it has retrieved all the logs from the device.

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

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

4. Click **[Save]**.
5. Before you can use this Log File Monitoring policy, you will need to align the policy with one or more devices. For more information, see [Aligning a Log File Monitoring Policy to Devices](#).

---

## Aligning a Log File Monitoring Policy to Devices

You can align Log File Monitoring policies 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:

The screenshot shows a modal window titled "Log Monitoring Policy" with a sub-header "Create a new Log File monitor" and a "Reset" button. The form contains the following fields and labels:

- Device:** A dropdown menu with "em7ao" selected. Below it, the text reads: "The device on which the log(s) will be monitored".
- Log Policy:** A dropdown menu with "Silo" selected. Below it, the text reads: "The log policy definition to be used as the basis for this log monitor".
- File Path:** A text input field containing "/var/log/em7/silo.log". Above it, the text reads: "Click on the labels of fields below to enable overrides to the selected policy". Below it, the text reads: "The full path of the file from which to capture log messages. You can use a \* to match multiple files (eg. /var/log/httpd/\*.log)".
- Limit:** A text input field containing "2000". Below it, the text reads: "The maximum number of messages to upload per minute from this source. A higher limit means you won't miss log messages, but it may cause a performance hit for uploading the data".
- Filter:** A text input field containing "\*DEBUG.\*". Below it, the text reads: "Capture log messages which match this string via RegEx. For example, \*.ERROR.\* to match log entieres containing 'ERROR'".

A "Save" button is located at the bottom center of the modal.

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 the names of the following fields to enable and edit them. These fields allow you to override settings of the policy you selected in the **Log Policy** field for the device selected in the **Device** field:
  - **File Path.** Enter the full file path or the file name to monitor. This field appears only if the type of the policy is *File*.
  - **Limit.** The maximum log messages the agent sends to SL1 per minute. If the number of matching logs exceeds this value, the agent will stop sending logs to the platform for the remainder of the minute. The limit resets at the beginning of the next minute. For example, suppose you set this field to *10,000*. Suppose the agent monitors a device that has 30,000 log messages. The agent will retrieve 10,000 logs and then wait until the beginning of the next minute. The agent will then retrieve the next 10,000 logs and then wait until the beginning of the next minute. The agent will continue to retrieve 10,000 logs per minute until it has retrieved all the logs from the device.

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

5. Click **[Save]**.

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

1. Go to the **Device Manager** page (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 checkboxes for the devices from which the policy must be unaligned.
3. In the **Select Action** drop-down menu, choose *Delete Log File Monitors*.


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

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

---

## 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

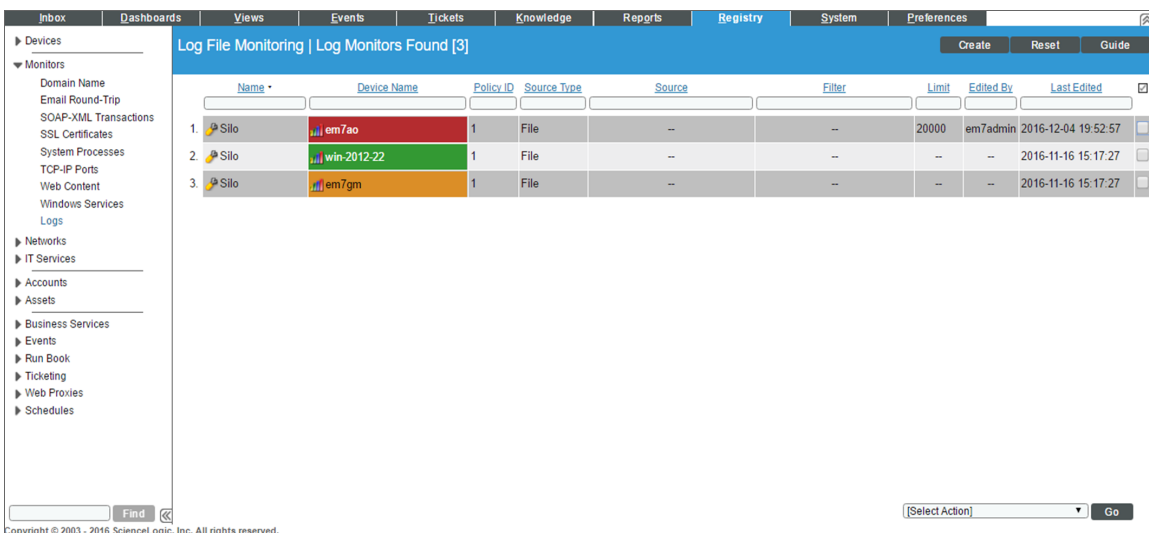
**NOTE:** Before you delete a Log File Monitoring Policy, you must unalign that policy from all devices. [See \*Unaligning Log File Monitoring Policies\*](#) for more information.

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.



The screenshot shows the 'Log File Monitoring | Log Monitors Found [3]' page. The table lists three aligned policies:

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

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

- **Name.** The name of the Log File Monitoring policy.
- **Device Name.** The name of the device aligned to the Log File Monitoring policy.
- **ID.** The unique numeric ID of the Log File Monitoring policy. The ID is automatically assigned by SL1.

- **Source Type.** The source of the logs in the monitored device. The possible values are:
  - File. The agent monitors a file on the file system of the device. Usually, this is used to monitor Linux log files.
  - Event Log. The agent monitors to Windows log on the device.
- **Source.** The full path of the log file or the name of the Windows log that the agent monitors.
- **Filter.** The regular expression the agent uses to determine if a log should be sent to SL1.
- **Limit.** The maximum log messages the agent sends to SL1 per minute. If the number of matching logs exceeds this value, the agent will stop sending logs to the platform for the remainder of the minute. The limit resets at the beginning of the next minute. For example, suppose you set this field to *10,000*. Suppose the agent monitors a device that has 30,000 log messages. The agent will retrieve 10,000 logs and then wait until the beginning of the next minute. The agent will then retrieve the next 10,000 logs and then wait until the beginning of the next minute. The agent will continue to retrieve 10,000 logs per minute until it has retrieved all the logs from the device.
- **Edited By.** The user who created or last edited the alignment between the device and Log File Monitoring policy.
- **Last Edited.** The date and time the alignment between the device and Log File Monitoring policy was created or last edited.

## Filtering the List of Log File Monitoring Policies and Aligned Devices

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

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

---

## Creating an Event Policy for Agent Logs

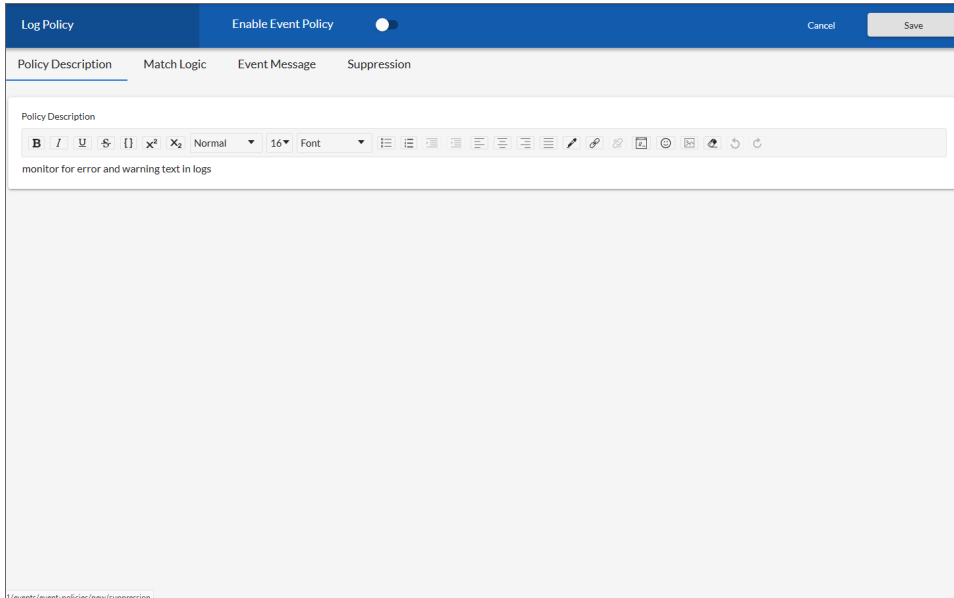
To trigger events based on log messages collected by the agent, you must create an event policy on the **Event Policies** page, and then associate that event policy with a Log File Monitoring policy.

**NOTE:** If you are using the classic user interface for SL1, see [Creating an Event Policy for Agent Logs in the Classic User Interface](#).

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

1. Go to **Event Policies** page (Events > Event Policies).

2. On the **Event Policies** page, click the **[Create Event Policy]** button. The **Event Policy Editor** page appears:



3. On the **Event Policy Editor** page and set of tabs, you can define a new event. The **Event Policy Editor** page contains the following fields and tabs:
  - **Policy Name**. At the top left of the **[Policy Description]** tab, type a name for the event policy in this field.
  - **Enable Event Policy**. This toggle allows you to enable and disable the event policy.
  - **Policy Description**. In this field, type a description of the event policy.

4. Click the **[Match Logic]** tab to define pattern-matching for the event:

The screenshot shows the 'Log Policy' configuration window with the 'Match Logic' tab selected. The window has a blue header with 'Log Policy' on the left, 'Enable Event Policy' with a radio button in the center, and 'Cancel' and 'Save' buttons on the right. Below the header are four tabs: 'Policy Description', 'Match Logic', 'Event Message', and 'Suppression'. The 'Match Logic' tab is active. On the left side, under 'MATCH CRITERIA', there are two dropdown menus: 'Event Source' set to 'ScienceLogic Agent' and 'Log Policy' set to 'event log errors'. On the right side, there are several fields and options: a 'String' dropdown menu, a 'Match String' text input field, a 'Second Match String (Optional)' text input field, two toggle switches for 'Allow event to expire if it doesn't reoccur within a time frame' and 'Require multiple triggers within a time frame', a 'MATCHING' section with a 'Detection Weight' input field set to '0', and two checkboxes: 'Multi Match' (checked) and 'Message Match' (unchecked). Below the 'Message Match' checkbox is the text 'Create new event if entire message does not match'.

5. Complete the following fields:
  - **Event Source.** Select *ScienceLogic Agent* as the source for the event. The fields below this field will change based on your selection.
  - **Log Policy.** Select the Log File Monitoring Policy the agent will use to collect the log message.
6. After selecting and defining your **Event Source**, enter values in the fields on the right side of the **Match Logic** tab specify the text that must appear in the log message for the event policy to trigger:
  - **String/Regular Expression.** Use this dropdown to select *String* or *Regular Expression*.
  - **Match String.** A string used to correlate the event with a log message. Can be up to 512 characters in length. To match this event policy, the text of a log message or alert must match the value you enter in this field. Can be any combination of alpha-numeric and multi-byte characters. SL1's expression matching is case-sensitive. This field is required for events generated with a source of Syslog, API, and Email.
  - **Second Match String (Optional).** A secondary string used to match against the originating log message. Can be up to 512 characters in length. Can be any combination of alpha-numeric and multi-byte characters. To match this event policy, the text of a log message or alert must match the value you enter in this field and the value you entered in the **Match String** field. This field is optional.
  - **Allow event to expire if it doesn't reoccur within a time frame.** If toggled on, enter the time in which an active event will be cleared automatically if there is no reoccurrence of the event in the fields that appear. You can enter time in minutes or hours.
  - **Require multiple triggers within a time frame.** If toggled on, enter the number of events and the time in which an event requires multiple triggers to occur in the fields that appear. You can enter time in minutes or hours.

- **Detection Weight.** If two event definitions are very similar, the weight field specifies the order in which SL1 should match messages against the similar event definitions. The event definition with the lowest weight will be matched first. This field is most useful for events that use expression matching. Options range from 0 (first) - 20 (last).
- **Multi Match.** By default, SL1 will match a log message or alert to only one event policy. If a log message or alert matches multiple event policies, SL1 will use the **Detection Weight** setting to determine which event policy the log message or alert will match. If you select the **Multi Match** checkbox in all event policies that can match the same log message or alert, SL1 will generate an event for every event policy that matches that single log message or alert.
- **Message Match.** If SL1 has generated an event and then a second log message or alert matches the same event policy for the same entity, SL1 will not generate a second event, but will increase the **count** value for the original event on the **Events** page and in the **Viewing Events** page. By default, this behavior occurs regardless of whether the two log messages or alerts contain the same message. If you select the **Message Match** checkbox, this behavior will occur only if the log messages or alerts contain the same message.

7. Click the **[Event Message]** tab to configure the event:

The screenshot displays the configuration page for an event policy in SL1. The top navigation bar includes 'Policy Name', 'Enable Event Policy' (with a toggle switch), and 'Cancel' and 'Save' buttons. Below this, there are four tabs: 'Policy Description', 'Match Logic', 'Event Message' (which is active), and 'Suppression'. The 'Event Message' tab contains the following settings:

- Event Message:** A text input field containing the placeholder '%M'.
- Event Severity:** A dropdown menu currently set to 'Major'.
- Use Modifier:** An unchecked checkbox with the subtext 'Respect severity modifiers if set on devices'.
- Correlate this event with external system:** An unchecked toggle switch.
- Categorize event for external system:** An unchecked toggle switch.
- Extract sub-entity using a regular expression:** An unchecked toggle switch.

On the right side of the configuration area, there are additional settings:

- Autoclear:** A checked toggle switch with the subtext 'Clear other events when this event is triggered' and an 'Add Event Policy' button.
- TOPOLOGY MASKING:** A section header followed by two unchecked checkboxes: 'Mask events on child devices' and 'Maskable under a parent device's event', with an 'Add Category' button below them.

8. Complete the following fields:



- **Event Message.** The message that appears in the **Event Console** page or the **Viewing Events** page when this event occurs. This field defaults to "%M" for new event policies upon creation. The message can be any combination of alphanumeric and multi-byte characters. Variables include the characters "%" (percent) and "|" (bar). You can also use regular expressions and variables that represent text from the original log message to create the **Event Message**.

To include regular expressions in the Event Message, surround the regular expression with %R and %/R. For example, %Rfilename: .\*? %/R would search for the first instance of the string "filename: " (filename-colon-space) followed by any number of any characters up to the line break. The %R indicates the beginning of a regular expression. The %/R indicates the end of a regular expression.

SL1 uses the regular expression to search the log message and use the matching text in the event message. 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 impacts service and requires immediate investigation.
    - *Critical.* Critical events indicate a condition that can seriously impair or curtail service and requires 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*.

9. 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 *Defining and Editing Event Policies* chapter in the **Events** manual.
10. Click **[Save]**.

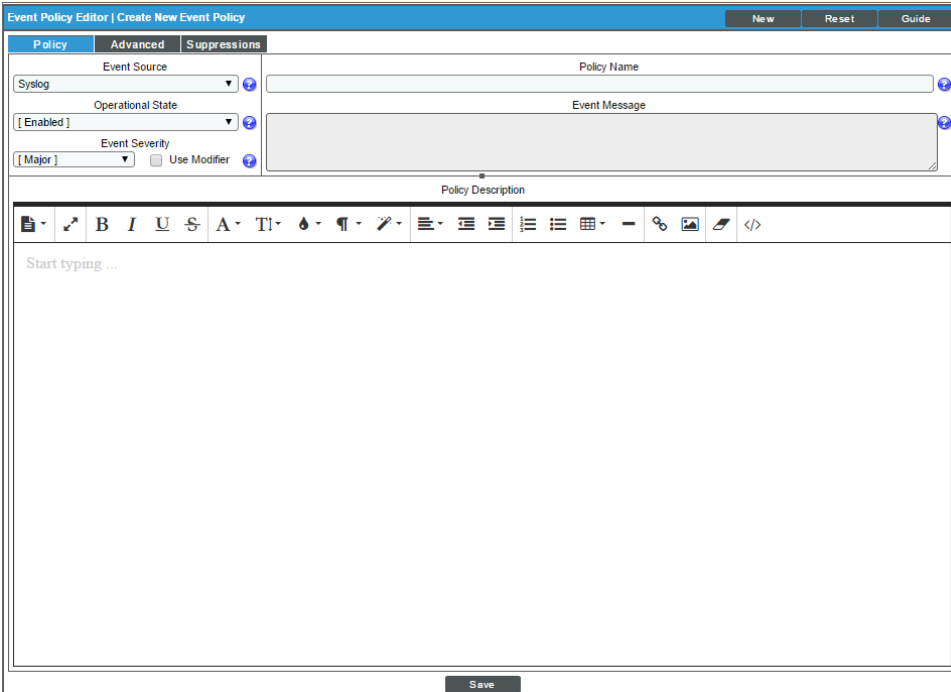
---

## Creating an Event Policy for Agent Logs in the Classic User Interface

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

To create an event policy in the classic user interface based on log data collected by the agent:

1. Go to **Event Policy Manager** page (Registry > Events > Event Manager in the classic user interface).
2. In the **Event Policy Manager** page, click **[Create]**. The **Event Policy Editor** page appears:



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

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

Buttons for 'New', 'Reset', and 'Guide' are in the top right corner, and a 'Save' button is at the bottom center.

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/>.
      - **%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

# 5

## Agent Troubleshooting

---



### Overview

This chapter contains troubleshooting processes that you can use to address issues with the SL1 agent.

**TIP:** On a Windows agent, you can run the following diagnostic command to generate a **\*.tar.gz** file that contains useful information for troubleshooting issues: `...\SiloAgent.exe --diag`

**TIP:** On a Linux agent, use `--diag` for the diagnostic option.

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

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

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

<i>Is the Agent Process Running?</i> .....	63
<i>Is the Agent Configuration Valid?</i> .....	64
<i>Has SL1 Discovery Completed?</i> .....	64
<i>Is the Agent Able to Upload Data?</i> .....	65
<i>Is SL1 Receiving Agent Data?</i> .....	66
<i>Can SL1 Process Agent Data?</i> .....	67
<i>Is the Number of Processes Inconsistent with Other Applications?</i> .....	67

<i>Does MySQL on the Collector Unit Have a Connection Issue?</i> .....	67
<i>Troubleshooting Examples</i> .....	68
<i>Additional Troubleshooting Situations and Best Practices</i> .....	69
<i>Agent Communication with SL1</i> .....	70

---

## Is the Agent Process Running?

As a first step, always locate the following logs when troubleshooting:

- `/var/log/streamer_prime/streamer_prime.log`
- `/var/log/uwsgi/streamer_prime.log`

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).
  - Windows: If **SiloAgent.exe** is not running, check the "Application" event log for events with `source=SiloAgent`.
  - Linux: If **scilogd** is not running, check `/var/log/messages` or `/var/log/syslog` for relevant log messages.

If you are using the SL1 Extended Architecture, determine if the agent was deleted from the **[Agents]** tab instead of uninstalling the agent.

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

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**

---

## Is the Agent Configuration 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.
  - The **CollectorID** should be a GUID similar to "4179b06ef502129c3023a0f8d58f3c37".
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 a distributed architecture, **URLfront** should be the URL of the Message Collector.
  - If you are using the SL1 Extended Architecture, **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.
  - If the URL for **URLfront** is correct, then determine if you can ping the host portion of **URLfront**. If you cannot ping, then there are customer firewall or NAT issues.

---

## Has SL1 Discovery Completed?

To check if SL1 discovery finished:

1. Check the agent record in the distributed architecture by using SSH to communicate with the Message Collector and running the following commands:

```
redis-cli -p 6380
keys *
hgetall agent_<GUID>
```

where <GUID> is the specific value from **scilog.conf**.

2. Look for a field "did". The following line is the EM7-device-id.
3. If the EM7-device-id is not present, then discovery has not completed.
4. Alternately, you can look on the **Devices** page for the device to compare the shown device-id to the value in the agent record.



---

## Is the Agent 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**

If there are many items, then the agent is unable to upload.

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

If the number of items continues to *increase* overall, check the configuration item URL:

- The URL is the location where the agent attempts to upload files.
- Determine if the host portion of the URL is reachable. If the host portion is reachable, the name of the oldest item indicates the approximate time of the issue.

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

A procedural note regarding backed-up data:

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

After the Streamer container receives a valid configuration file:

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

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

### Run the Agent in Debug Mode (Linux)

**NOTE:** You might need to preface the following commands with `sudo` 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
```

---

## Is SL1 Receiving Agent Data?

If you are using a distributed architecture:

1. SSH into the Message Collector and run the following command:

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

2. Look for lines containing the hostname of the monitored device, 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 container is not getting data from that agent.

If you are using the SL1 Extended Architecture:

1. SSH into the Management Node.
2. Run the following command to view the logs:

```
kubect1 logs -l app=streamer
```

3. Look for lines containing the hostname of the monitored device.
4. If there are no matching lines, then the Streamer container is not getting data from that agent.

---

## Can SL1 Process Agent Data?

Check the Message Collector log files or SL1 Streamer log files:

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

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

2. If you are using the SL1 Extended Architecture, run the following command:

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

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

---

## Is the Number of Processes Inconsistent with Other Applications?

- On Linux, many outputs from the `ps` command list the kernel threads (the processes listed in square brackets). Because the agent is not in the kernel, it will not list kernel threads.
- Be aware that the agent reports processes that are running as well as processes that started and may have stopped, while `top` or `ps` commands show processes that exist when they are executed.
- Check the agent configuration. Due to back-end space limitations, many configuration combinations can limit what data the agent sends. A combination of parameters to get all processes include the following:
  - **NIPD True**. The agent library can not get into all processes at times, often on install. Non-intercepted process discovery reports processes that are not intercepted via the library.
  - **SLP Aggregation**. 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.

---

## Does MySQL on the Collector Unit Have a Connection Issue?

If MySQL on the Collector Unit has more connections than the configured limit, the Collector Unit will raise a connection issue and cause the agent connection failure. The following error message is an example of this situation:

```
Unhandled exception: Traceback (most recent call last):
File "/opt/em7/backend/process/proc_pda.py", line 223, in __init__
File "/opt/em7/backend/silo_collect/streamer/main.py", line 353, in main
File "/opt/em7/backend/silo_collect/shared.py", line 108, in run_collection
```

```
File "/opt/em7/backend/silo_collect/shared.py", line 156, in _run_collect_async
File "/usr/lib64/python2.7/multiprocessing/pool.py", line 554, in get raise self._
value ChildProcessException: Traceback (most recent call last):
File "/opt/em7/backend/silo_common/misc.py", line 520, in new_f File
"/opt/em7/backend/silo_collect/streamer/main.py", line 225, in streamer_collect
File "/opt/em7/backend/silo_common/database.py", line 701, in local_db
File "/opt/em7/backend/silo_common/database.py", line 92, in callFunc
File "/opt/em7/backend/silo_common/database.py", line 693, in get_dbc MySQLError:
Error attempting to connect to database with SSL enabled False: (2013, 'Lost
connection to MySQL server at \'reading initial communication packet\', system error:
0 "Internal error/check (Not system error)")
```

To address this connection issue:

1. To determine the current connection limit, run the following command on the Collector Unit:

```
silo_mysql -e 'show variables like "max_connections";'
```

2. To view the high watermark for concurrently active connections, run the following command on the Collector Unit:

```
show status where 'Variable_name' = 'max_used_connections';
```

3. If the value is more than the limit value, try to dynamically but temporarily increase the limit. For example:

```
silo_mysql -e 'set global max_connections=350'
```

---

## Troubleshooting Examples

### Example /var/log/streamer\_prime/streamer\_prime.log for successful discovery

```
2019-01-04T17:07:42.355291+00:00 amateen-em7 journal: SCILO_SP:6954|logger:log_
info:132|INFO|Agent config request received with init flag set to True. Generated
Temp AID: 2ae22a6b4489457abb14373cd3816076. Request: <WSGIRequest: GET
'/api/collector/config/?collector_key=aEf34$aq3TGSDdf&tenant_id=0&host_name=aym-
win2012r2-1&init=&os=windows&collector_id=0'>
```

```
2019-01-04T17:07:42.619082+00:00 amateen-em7 journal: SCILO_SP:6954|logger:log_
info:132|INFO|Calling Agent version with: <QueryDict: {'collector_id':
['2ae22a6b4489457abb14373cd3816076'], 'type': ['windows_64'], 'tenant_id': ['0'],
'host_name': ['aym-win2012r2-1'], 'collector_key': ['aEf34$aq3TGSDdf'], 'version':
['115']}>
```

```
2019-01-04T17:07:43.028457+00:00 amateen-em7 journal: SCILO_SP:16717|logger:log_
warning:127|WARNING|System file received from aym-win2012r2-1
```

```
2019-01-04T17:07:43.032897+00:00 amateen-em7 journal: SCILO_SP:16717|logger:log_
info:132|INFO|Making discovery call for agent 2ae22a6b4489457abb14373cd3816076
```

```
2019-01-04T17:07:43.746284+00:00 amateen-em7 journal: SCILO_SP:30843|logger:log_
warning:127|WARNING|System file received from aym-win2012r2-1
```

```
2019-01-04T17:07:43.750553+00:00 amateen-em7 journal: SCILO_SP:30843|logger:log_
warning:127|WARNING|Discovery call within time threshold, sleeping.
```

```
2019-01-04T17:07:46.676827+00:00 amateen-em7 journal: SCILO_SP:16717|logger:log_
info:132|INFO|Update agent request did: 4, oid: 0, ip: 10.7.6.119, agent id:
2ae22a6b4489457abb14373cd3816076
```

```
2019-01-04T17:07:46.677114+00:00 amateen-em7 journal: SCILO_SP:16717|logger:log_
info:132|INFO|Discovery complete, getting new agent device id. Downloading new
config for device: 2ae22a6b4489457abb14373cd3816076.
```

```
2019-01-04T17:07:47.420509+00:00 amateen-em7 journal: SCILO_SP:6954|logger:log_
warning:127|WARNING|Agent id: 2ae22a6b4489457abb14373cd3816076 being given a return
code: 2
```

## Example /var/log/uwsgi/streamer.log for successful discovery in the distributed architecture

```
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 in the distributed architecture

```
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 in the SL1 Extended Architecture

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

You can find the host id from the ADS url, such as [https://<sl1\\_address>/ads/servers/13/system](https://<sl1_address>/ads/servers/13/system)). You can locate the files in the `/tmp/save_agent_data` directory.

---

## Additional Troubleshooting Situations and Best Practices

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

Situation	Cause / Resolution
Two device records exist for the same device on the <b>Devices</b> page in SL1 .	<p>This situation occurs when SL1 first discovered this device with SNMP, and then the agent was installed and started polling that device. This duplication of records also occurs if the agent was installed first, and then you ran an SNMP discovery.</p> <p>To address this issue, you can <b>merge</b> the device records using the classic user interface. For more information, see the "Managing a Single Device with the Device Administration Panel" chapter in the <b>Device Management</b> manual.</p>
The SNMP device record has IPv4, but the agent device record has IPv6.	<p>The agent reports all network interfaces to the message collector. The Message Collector uses the first "bound" IP address reported by the agent.</p> <p>To address this issue, you can manually edit the agent device record in the "classic" user interface and update the IP address.</p>
If you uninstall an agent and then run a different installation executable file, you still see the same organization ID for the agent record.	<p>After you uninstall the agent, the <b>scilog.conf</b> file is left on the server in case the agent is reinstalled. SL1 can reuse the same device record and maintain historical performance data for that agent.</p> <p>To address this issue, delete the <b>scilog.conf</b> file after you run the uninstallation. If you install this agent again, SL1 assigns a new organization ID to the agent and creates a new device record.</p>

## Agent Communication with SL1

This section covers the various codes that are sent to the SL1 agent when additional actions are needed.

### Return Codes

The core method of communication between the SL1 agent and SL1 involves the data files the agent uploads to the Streamer. When the agent uploads a file to the Streamer, a number of conditions are checked and the return code given to the agent determines what action should be taken, if any.

Return Code	Meaning
200	No action needed.
409	The data file uploaded by the agent has a problem. The Streamer did not accept it and it should be deleted from the host

Return Code	Meaning
503	Data uploads are occurring more quickly than normal. The agent should refrain from sending its next file until a specified amount of time has passed.
505	The agent should stop streaming.

## Sub-codes

The Streamer communicates another set of actions in a JSON response included with the return code to the SL1 agent.

JSON Response Field	Type	Meaning
send_system_file	boolean	Send a system config file as the next upload from this agent.
get_agent_config	boolean	Retrieve an updated agent config from the Streamer.
get_agent_update	boolean	Download a new version of the agent to update.
get_pd_config	boolean	Retrieve an updated polled data config.
get_log_config	boolean	Retrieve an updated log config.
purge_uploads	boolean	Delete any backed-up data and only upload new files.
set_data_backoff	integer	Do not upload a data file until the set number of seconds have passed.
set_log_backoff	integer	Do not upload a log file until the set number of seconds have passed.

© 2003 - 2021, 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](mailto:legal@sciencelogic.com)





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

International: +1-703-354-1010