Monitoring with the SL1 Agent

SL1 version 8.14.0

Table of Contents

Chapter

1

Introduction to the SL1 Agent

Overview

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

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

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

This chapter includes the following topics:

What is an SL1 Agent?

The ***SL1 Agent*** is a program that you can optionally install on a device or element monitored by SL1. The SL1 Agent collects data from the device, interface, or other element and pushes that data back to SL1. You can install and use multiple SL1 Agents, as needed.

Because an agent is always running on a device, the SL1 Agent can collect more granular data than can be collected by polling the device periodically with a Data Collector. You can collect data from devices using only the SL1 Agent or using a combination of the SL1 Agent and Data Collectors.

Agent PowerPacks

SL1 includes two PowerPacks that can be used to collect agent-based system configuration and performance data:

- The *ScienceLogic: Agent* PowerPack, which collects agent-based data for devices on SL1 systems running on the SL1 extended architecture
- The *Host Agent* PowerPack, which collects agent-based data for devices on SL1 systems running on a distributed architecture

Both of these PowerPacks are installed by default on your SL1 system. 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

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

What Kind of Data Can an Agent Collect?

In SL1 ***Extended***, the SL1 Agent collects the following data:

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

  - Disk Utilization
  - Memory Utilization
  - Network Bytes Read
  - Network Bytes Written
  - 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 Interfaces***. 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** panel. This list includes attributes such as the interface MAC address, IP address, position, and speed.
- ***System Processes***. The SL1 Agent collects a list of all processes running on the device. You can view the list of processes on the **[Processes]** tab of the **Device Reports** panel and the **[Processes]** tab of the **Device Investigator** page. Monitoring policies can be configured to trend and alert on process availability, process CPU usage, and process memory usage.
- ***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** panel.
- ***Logs***. The SL1 Agent can be configured to push logs that match specific criteria from a log file or the Windows Event Log to SL1. You can view logs collected by the SL1 Agent on the **Device Investigator** page and the **Device Logs** page (Registry > Monitors > Logs) for a device. Logs can be configured to trigger events.

NOTE: For more information about the data that SL1 agent can collect, see the ***Monitoring Device Infrastructure Health*** manual.

Supported Operating Systems

You can install agents on the following operating systems:

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

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

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.

Agent Versions

The following table indicates the Windows and Linux agent versions that are bundled with each version of SL1.

NOTE: The SL1agent is backwards- and forwards-compatible. Therefore, while these agent versions are bundled with the SL1 versions listed below, you can use older versions of the agent with newer versions of SL1 and vice-versa.

SL1 Version	Windows Agent Version	Linux Agent Version
8.14.0	125	167
8.12.0	119	163
8.11.0	115	159
8.10.0	114	157
8.9.0	112	152
8.8.2	110	150
8.8.1	110	150
8.8.0	107	148
8.7.1.2	107	148
8.7.0	107	146

Agent Architecture

The following sections describe how the SL1 agent works in a distributed architecture and in an extended architecture.

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 the Message Collector and the device on which an agent is installed.

In a distributed architecture, the SL1 agent requires a standalone Message Collector. The Message Collector does not need to be dedicated to the agent. 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.

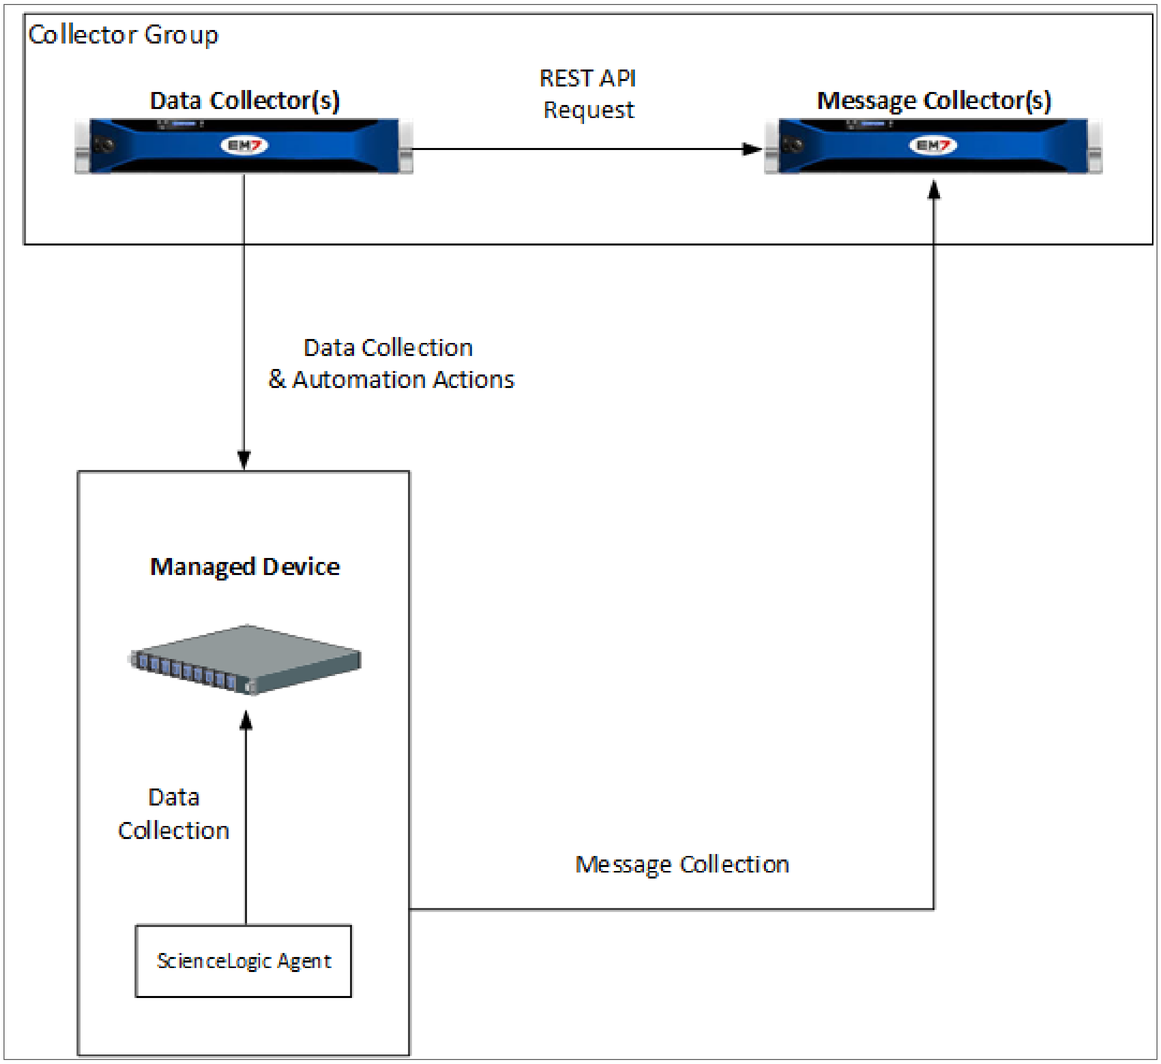

Extended Architecture

In an extended architecture, an SL1 agent collects data about the device on which it is installed and uploads it to the streamer service that is running on the SL1 compute node cluster. The streamer is a web service that SL1 agents and Data Collectors use to deliver data about your monitored devices to the SL1 storage nodes and Database Server using the HTTPS protocol.

Using an SL1 agent in an extended architecture provides more configuration and performance data than using an SL1 agent in a distributed architecture. This additional data includes system vitals.

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 an SL1 Agent

Overview

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

- If you use the ***SL1 extended architecture*** (which includes Compute Nodes, Storage Nodes, and a Management node), you can install the agent from the **Agents** page (Devices > Agent) in the unified user interface. For details, see the section ***Installing an Agent from the Agents Page***.
- If you use the ***SL1 architecture without the extended architecture***, you can install the agent from the **Device Manager** page. This page is available in the unified user interface when you click the menu icon (☰) in the left navigation pane. For details, see the section ***Installing an Agent from the Device Manager Page***.

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

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

This chapter includes the following topics:

Installing an Agent from the Agents Page

When running SL1 on an extended architecture, you can use the **Agents** page (Devices > Agents) to install the SL1 agent on a Linux or Windows device. You do not install the agent from the **Agents** page itself; instead, the **Agents** page (Devices > Agents) enables you to gather the information or files you need to then install the agent on a particular device:

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

NOTE: The **Agents** page appears as an option on the menu ***only*** when you are using the extended SL1 architecture. If you are using a distributed SL1 system or you are using the extended SL1 architecture with the classic SL1 user interface, then you can install the agent from the **Device Manager** page. For instructions, see the section ***Installing an Agent from the Device Manager Page***.

The following sections describe how to install, upgrade, and delete agents from the **Agents** page.

Installing a Linux Agent from the Agents Page

To download and install a Linux agent in SL1:

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

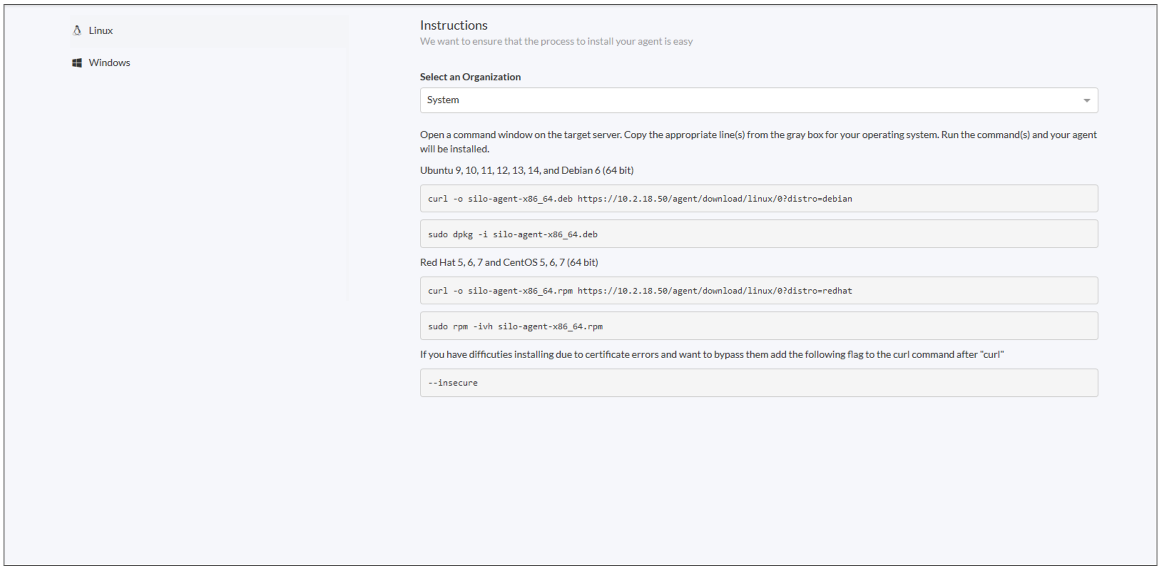

Instructions

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.

Ubuntu 9, 10, 11, 12, 13, 14, and Debian 6 (64 bit)

```
curl -o silo-agent-x86_64.deb https://10.2.18.50/agent/download/linux/0?distro=debian
```

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

Red Hat 5, 6, 7 and CentOS 5, 6, 7 (64 bit)

```
curl -o silo-agent-x86_64.rpm https://10.2.18.50/agent/download/linux/0?distro=redhat
```

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

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

```
--insecure
```

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

Installing a Windows Agent from the Agents Page

To download and install a Windows agent in SL1:

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

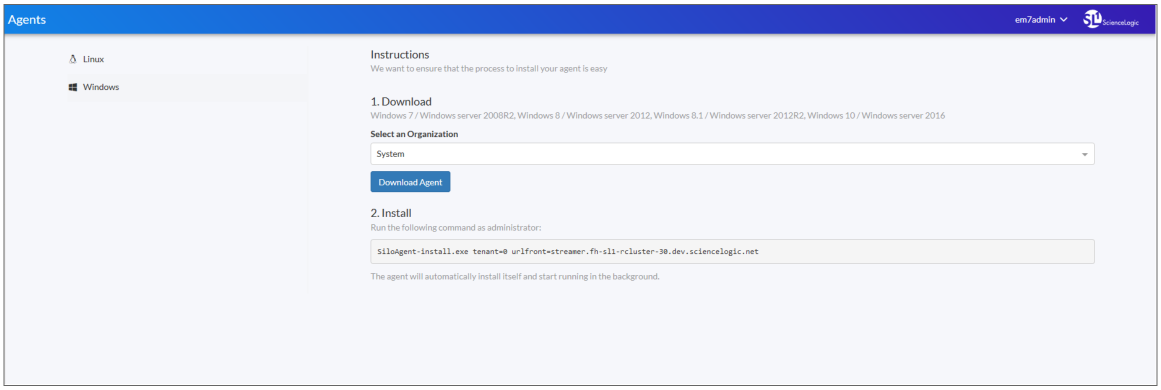

Instructions

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

1. Download

Windows 7 / Windows server 2008R2, Windows 8 / Windows server 2012, Windows 8.1 / Windows server 2012R2, Windows 10 / Windows server 2016

Select an Organization

System

2. Install

Run the following command as administrator:

```
SiloAgent-install.exe tenant=0 urlfront=streamer.fh-sl1-rcluster-30.dev.sciencelogic.net
```

The agent will automatically install itself and start running in the background.

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

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

8. 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: SL1supports 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.

Upgrading an Agent

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

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

Stopping an Agent

You can delete an agent on the **Agents** page. 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*.

> **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.

To delete an agent and its data:

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

Uninstalling an Agent

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

Uninstalling a Linux Agent

To uninstall an agent on a Linux system:

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

     ```
     rpm -e scilogd
     ```

   - For Debian and Ubuntu, execute the following command:

     ```
     dpkg -r scilogd
     ```

3. Optionally, you can remove the agent configuration directory from the Linux system. The configuration directory can be found at:

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

Uninstalling a Windows Agent

To uninstall an agent on a Windows system:

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

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

Using Agents with SELinux

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

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

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

Installing an Agent from the Device Manager Page

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

NOTE: If you are using a distributed SL1 system without the extended architecture (which includes Compute Nodes, Storage Nodes, and a Management Node), you ***must*** use the **Agent Installation** page (which you can access from the **Device Manager** page) to install the agent; you cannot do so from the **Agents** page (Devices > Agents).

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

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

> **TIP:** : In the unified SL1 user interface, you must first click the menu icon (☰) in the left navigation pane. You can then navigate to the **Device Manager** page.

2. Click **[Actions]** and select *Download/Install Agent*. The **Agent Installation** page appears:

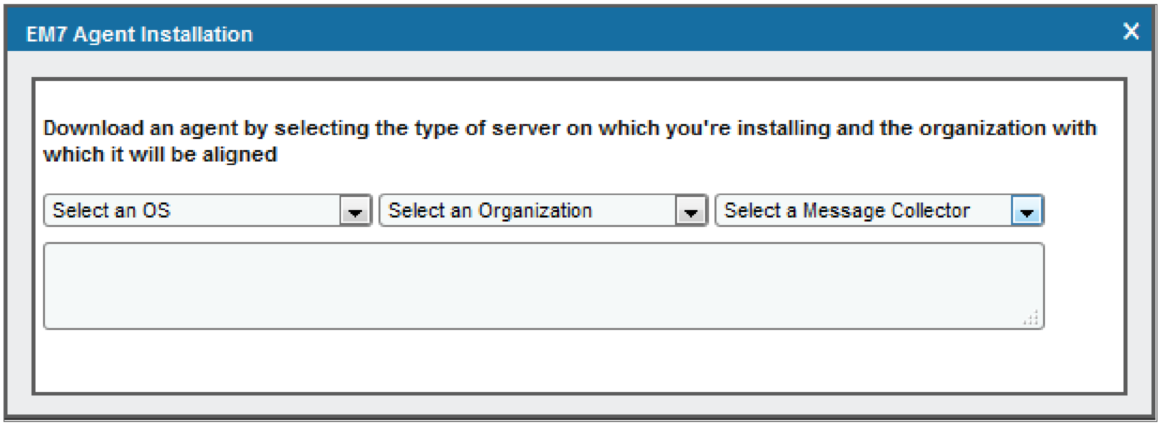

EM7 Agent Installation

Download an agent by selecting the type of server on which you're installing and the organization with which it will be aligned

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 Classic Linux Agent

To install an agent on a Linux system:

1. Log in to the Linux system via the console or SSH as a user that has sudo administrator permissions.
2. Execute the commands that you copied from the **Agent Installation** page in SL1. If the installation was successful, the output will look similar to the following:

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

NOTE: SL1supports 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.

Checking the Version of an Agent on a Linux System

To check the version number of an agent on a Linux system:

1. Log in to the Linux system via the console or SSH as a user that has sudo administrator permissions.
2. Execute the following command:

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

Updating an Agent on a Linux System

To update the agent on a Linux system:

1. 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_6
   ```

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

   ```
   /etc/scilog
   ```

Installing the Classic Windows Agent

To install an agent on a Windows system:

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

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

3. To verify that the installation was successful, open the Windows Task Manager or enter the TASKLIST command to view running processes. The SiloAgent process will be running on the Windows machine.

NOTE: SL1supports 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.

Checking the Version of an Agent on a Windows System

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

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

Uninstalling the Agent on a Windows System

To uninstall an agent on a Windows system:

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

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

Viewing the Discovered Device

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

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

Device Classes for Agent-Only Devices

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

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

- Microsoft Windows Workstation

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

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

Chapter

3

Configuring an SL1 Agent

Overview

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

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

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

This chapter includes the following topics:

Using the Agent Investigator

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

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

The Config Tab

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

To configure the agent:

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

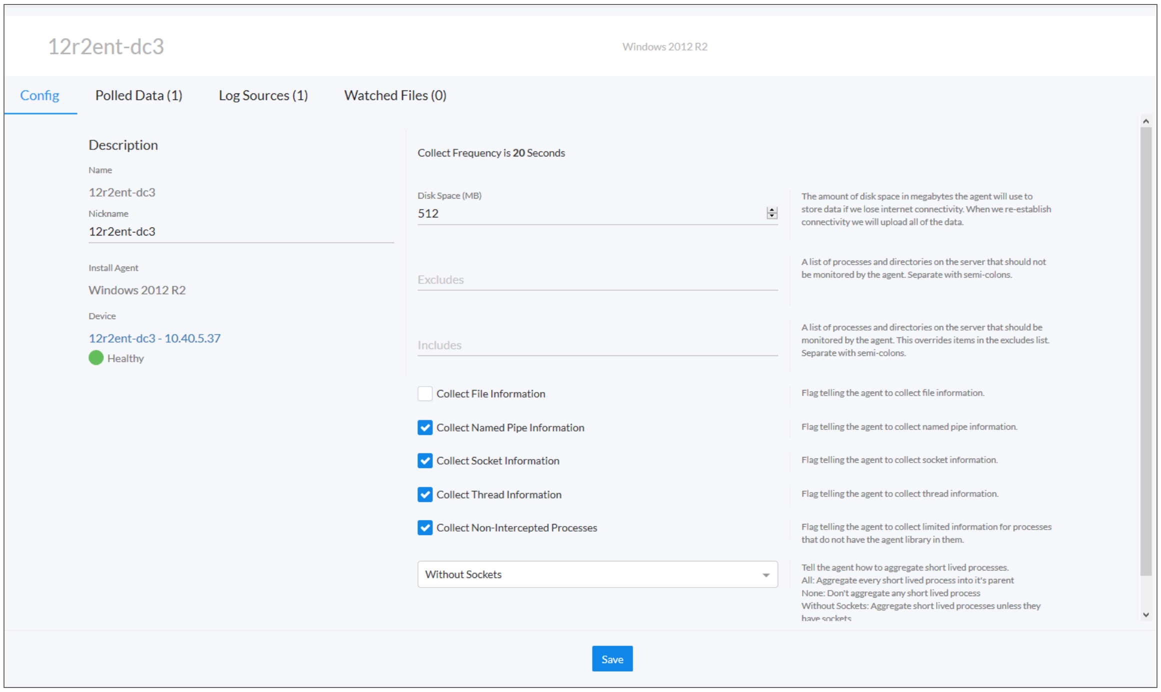

12r2ent-dc3

Windows 2012 R2

Description

Name
12r2ent-dc3

Nickname
12r2ent-dc3

Install Agent
Windows 2012 R2

Device
12r2ent-dc3 - 10.40.5.37
● Healthy

Collect Frequency is **20** Seconds

Disk Space (MB)
512

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.

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. This overrides items in the excludes list. Separate with semi-colons.

□ Collect File Information — Flag telling the agent to collect file information.

☑ Collect Named Pipe Information — Flag telling the agent to collect named pipe information.

☑ Collect Socket Information — Flag telling the agent to collect socket information.

☑ Collect Thread Information — Flag telling the agent to collect thread information.

☑ Collect Non-Intercepted Processes — Flag telling the agent to collect limited information for processes that do not have the agent library in them.

Without Sockets

Tell the agent how to aggregate short lived processes.
All: Aggregate every short lived process into it's parent
None: Don't aggregate any short lived process
Without Sockets: Aggregate short lived processes unless they have sockets

2. Complete in the following fields, as needed:
   - ***Disk Space***. Specify the amount of disk space that the agent can use to store data. If the agent loses connectivity to 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 semicolons, that you do not want the agent to monitor.
   - ***Includes***. Type a list of processes and directories, separated by semicolons, that you want the agent to monitor. This field ensures that specific processes are monitored.

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

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

The Polled Data Tab

The **[Polled Data]** tab of the **Agent Investigator** page lets you configure a new source for ***polled data***, the data you gather by running scripts on this agent over time, and manage existing sources for polled data.

Creating a New Source for Polled Data

To create a new source for polled data:

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

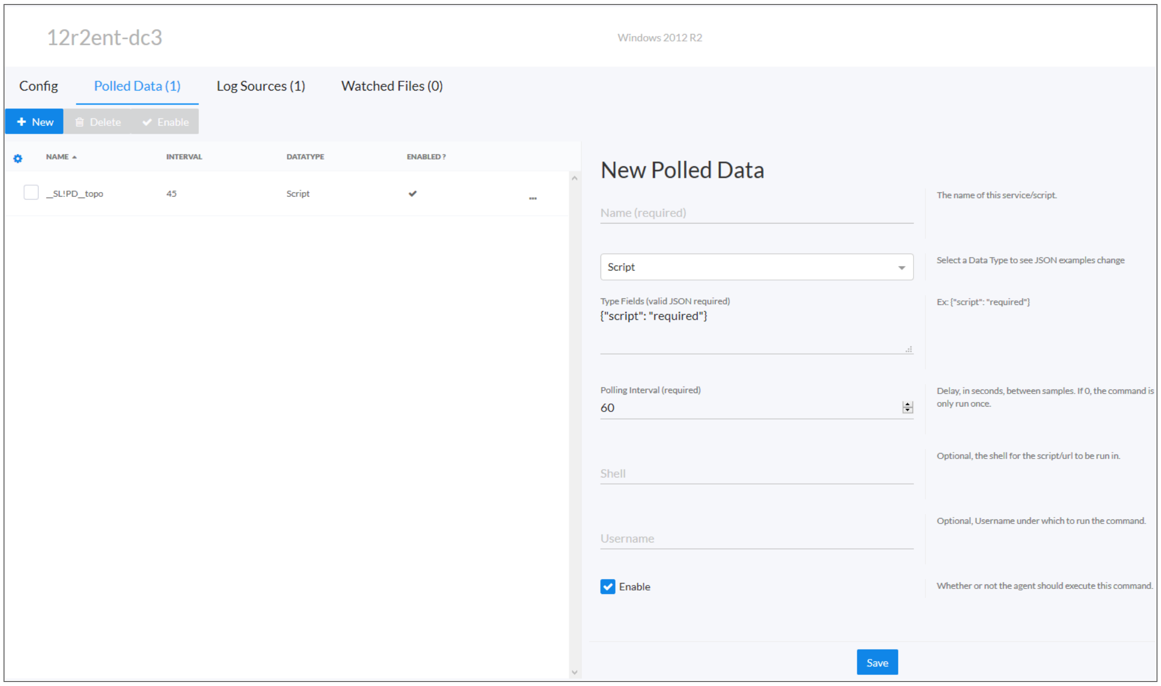

	NAME ▲	INTERVAL	DATATYPE	ENABLED ?	
□	_SL!PD_topo	45	Script	✓	...

New Polled Data

Name (required) — The name of this service/script.

Script — Select a Data Type to see JSON examples change

Type Fields (valid JSON required)
{"script": "required"} — Ex: {"script": "required"}

Polling Interval (required)
60 — Delay, in seconds, between samples. If 0, the command is only run once.

Shell — Optional, the shell for the script/url to be run in.

Username — Optional, Username under which to run the command.

☑ Enable — Whether or not the agent should execute this command.

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

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

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

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

  You can use the wildcard "*" for the process value when querying JMX beans in the Linux agent. If the JMX bean exists in multiple Java Virtual Machines with names matching the specified regular expression (regex), the results appear in a JVL format and are uploaded as a separate file.

  In the Linux agent, you can also specify that multiple polled data commands be grouped together by assigning them a common groupid.

- *Windows Performance Counter*. Use the specified Windows performance counter for polled data.
- *Event*. Use the specified event for polled data.

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

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

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

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

Configuring Nagios Plug-ins as Polled Data Sources

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

1. Download the Nagios plug-ins you want to use.
2. In SL1, navigate to the **Agent Investigator** page (Inventory > Agents > agent record) for the agent you want to use with the Nagios plug-ins.

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

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

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

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

Editing or Deleting Polled Data Sources

To edit or delete existing polled data sources:

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

The Log Sources Tab

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

Creating a New Log Source

To create a new log source:

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

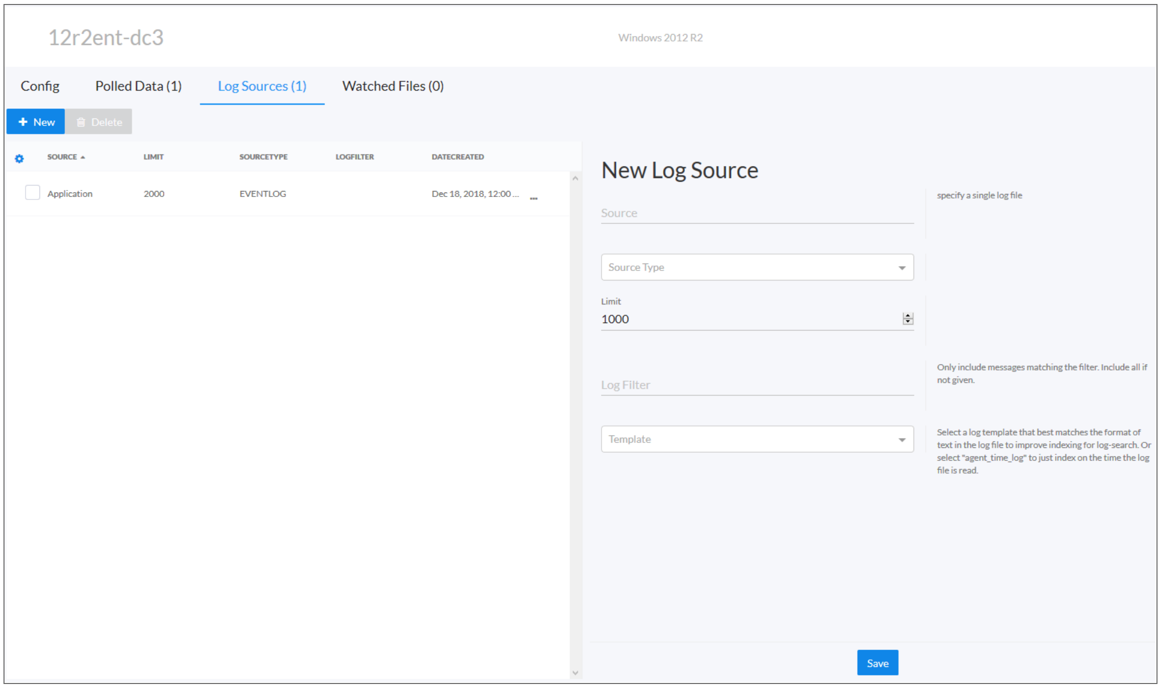

SOURCE	LIMIT	SOURCETYPE	LOGFILTER	DATECREATED
Application	2000	EVENTLOG		Dec 18, 2018, 12:00 ...

New Log Source

Source — specify a single log file

Source Type

Limit: 1000

Log Filter — Only include messages matching the filter. Include all if not given.

Template — Select a log template that best matches the format of text in the log file to improve indexing for log-search. Or select "agent_time_log" to just index on the time the log file is read.

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

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

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

Editing or Deleting Log Sources

To edit or delete existing log sources:

1. From the list of log sources on the **[Log Sources]** tab, select the source you want to edit or delete.
2. Click the **[Actions]** button (⋯) and select an option:

   - *Edit*. When you select this option, the fields for this log source appear in the right-hand panel, and you can update the fields as needed.
   - *Delete*. When you select this option, the log source is immediately deleted.

The Watched Files Tab

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

Adding Watched Files

To add watched files:

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

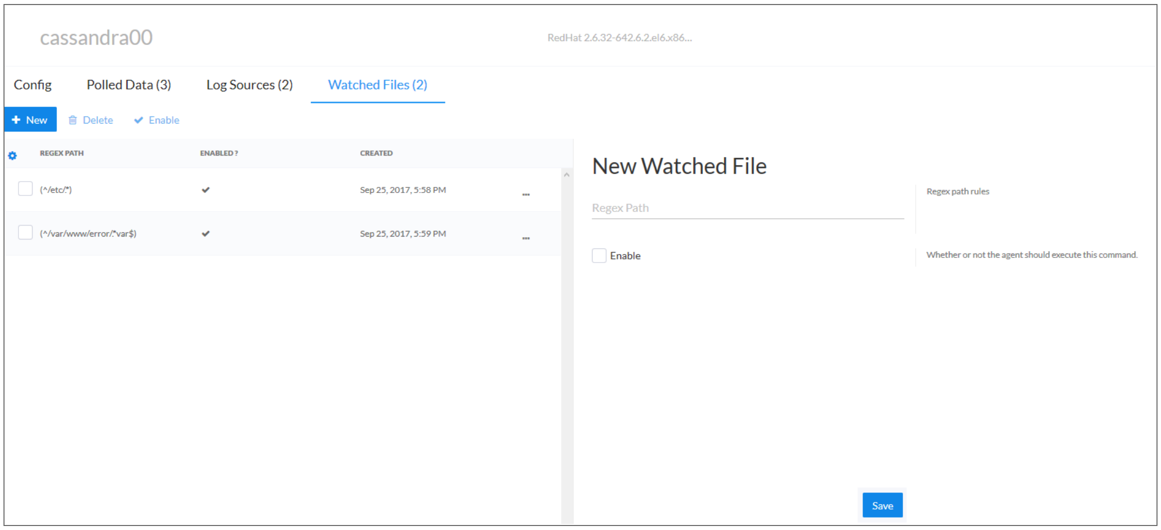

	REGEX PATH	ENABLED ?	CREATED	
□	(^/etc/.*)	✓	Sep 25, 2017, 5:58 PM	...
□	(^/var/www/error/.*var$)	✓	Sep 25, 2017, 5:59 PM	...

New Watched File

Regex Path — Regex path rules

□ Enable — Whether or not the agent should execute this command.

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

```
(^/etc/.*)

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

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

Editing or Deleting Watched Files

To edit or delete existing watched files:

1. From the list of watched files on the **[Watched Files]** tab, select the file you want to edit or delete.
2. Click the **[Actions]** button (⋯) and select an option:

    - *Edit*. When you select this option, the fields for this watched file appear in the right-hand panel, and you can update the fields as needed.
    - *Delete*. When you select this option, the watched file is immediately deleted.

Configuring a Classic Agent

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

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

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

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

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

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

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

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. Click **[Actions]**, and then select *Device Manager Preferences*. The **Edit Device Manager Preferences** modal page appears:

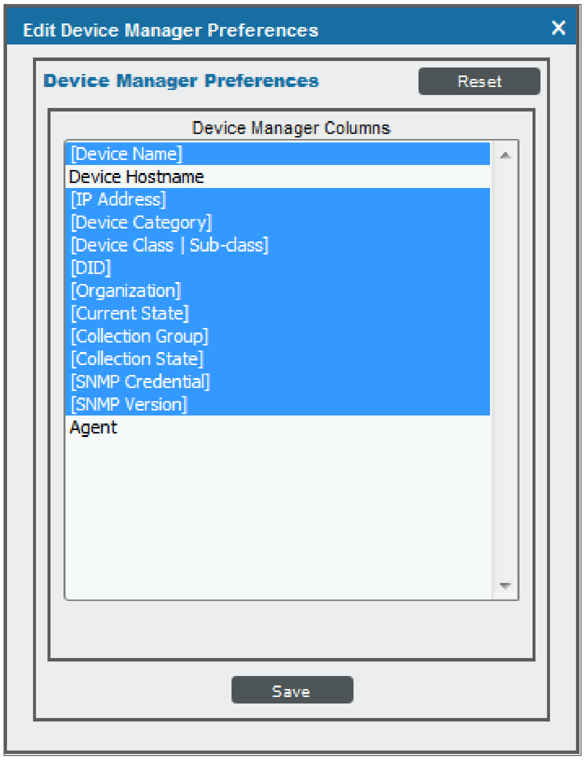

3. In the ***Device Manager Columns*** field, control-click *Agent*.
4. Click **[Save]**.

Configuring Agent Settings on a Device

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

To configure agent settings on a device:

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

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

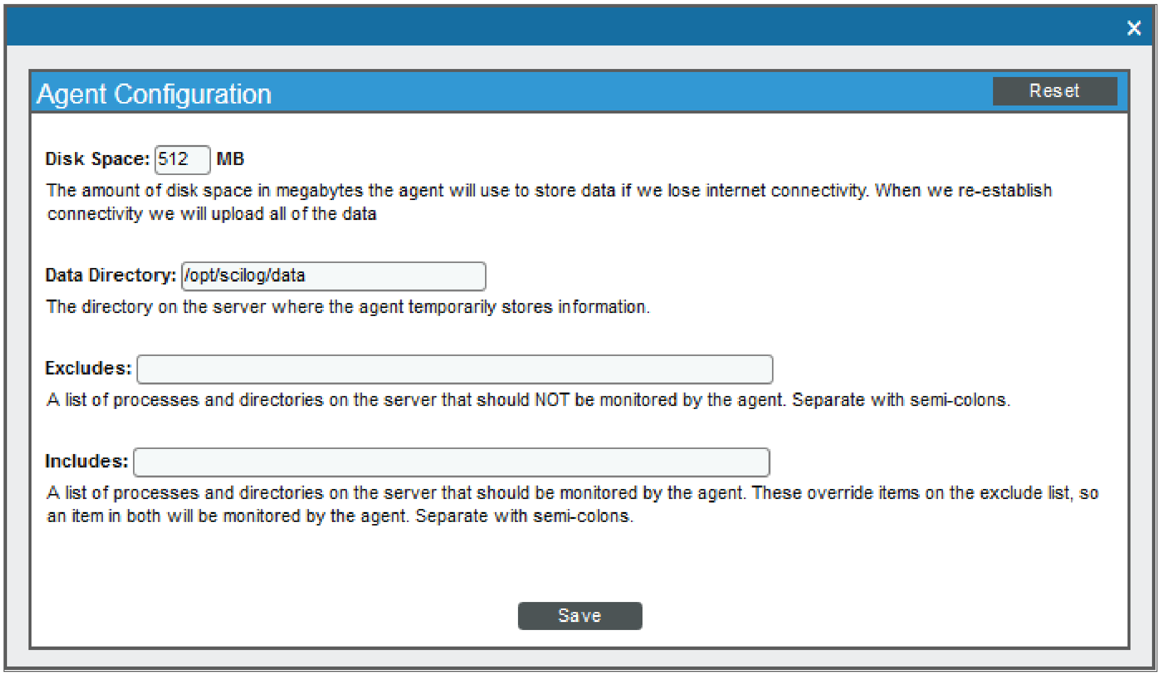

Disk Space: 512 MB

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.

3. Supply values in the following fields:

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

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

4. Click **[Save]**.

Changing the Target Message Collector for the Agent

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

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

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

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

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

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

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

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

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

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

   ```
   URLfront <IP address>
   ```

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

   ```
   sudo /etc/init.d/scilogd restart
   ```

7. On a Windows system, restart the SiloAgent Service service.

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

Chapter 4

Troubleshooting SL1 Agents

Overview

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

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

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

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

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

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

Determine if the Agent Process is Running

To determine if the agent process is running:

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

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

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

To re-install the agent that was shut down:

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

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

To generate diagnostics information for an agent:

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

Determine if the Agent Configuration is Valid

1. Check the agent configuration in one of the following locations:

   - Windows: **C:\Program Files\ScienceLogic\SiloAgent\conf\scilog.conf**
   - Linux: **/etc/scilogd/scilog.conf**

2. Check the configuration item **CollectorID**:

   - If there is *no* **CollectorID** tag, then the agent has not been able to reach the stream or message collector.

3. Check the configuration item **URLfront**, which is where the agent attempts to get the configuration file.

   - Determine if you can ping the **URLfront**.
   - If you are using streamer_prime, **URLfront** should be the URL of the message collector. If you are using the new user interface or the converged platform, **URLfront** should be the URL of the streamer container, such as *pod9-streamer0*.
     - If the URL for **URLfront** is not correct, then re-install the agent. See the re-install steps in the previous topic.
     - If the URL for **URLfront** is correct, then determine if you can ping the host portion of **URLfront**.

Determine if the Agent is Able to Upload Data

Check the Agent Upload Directory

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

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

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

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

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

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

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

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

A procedural note regarding backed-up data:

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

After the streamer receives a valid configuration file:

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

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

Run the Agent in Debug Mode (Linux)

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

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

   ```
   service scilogd stop
   ```

2. Start the agent from the command line:

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

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

   ```
   service scilogd start
   ```

Determine if SL1 is Receiving Agent Data

If you are using streamer_prime:

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

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

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

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

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

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

1. Either SSH into the Compute Node or point your instance to the Rancher cluster.
2. Run the following command to view the logs:.

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

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

Determine if SL1 Cannot Process Agent Data

Check the SL1 log files:

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

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

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

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

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

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

Determine if the Number of Processes is Inconsistent with Other Applications

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

Troubleshooting Examples

Example /var/log/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 streamer_prime

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

Save incoming data for a specific device ID (streamer_prime)

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

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

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

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

Additional Troubleshooting Situations and Best Practices

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

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

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

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

SL1 ScienceLogic

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

International: +1-703-354-1010