Installation and Initial Configuration

ScienceLogic version 10.2.0, revision 1

Table of Contents

Chapter 1

Introduction to Installing SL1

Overview

This manual describes how to install and configure SL1. This manual includes the following topics:

- *Preparing Hardware Appliances for SL1*
- *Preparing Virtual Appliances for SL1*
- *Required Ports for SL1*
- *Installing SL1 Hardware Appliances or Virtual Appliances*
- *Installing SL1 in the Amazon Cloud*
- *Installing SL1 in Microsoft Azure*

This chapter includes the following topics:

What is SL1?

In a ***Distributed*** system, there are four general functions that an SL1 appliance can perform: user interface, Database Server, Data Collector, and Message Collectors. In large SL1 systems, dedicated appliances perform each function. In smaller systems, some appliances perform multiple functions. In the ***All-In-One Appliance*** system, a single SL1 appliance performs all four functions.

User Interface

Administrators and users access the user interface through a web browser. In the user interface, you can view collected data and reports, define organizations and user accounts, define policies, view events, and create and view tickets, among other tasks. The appliance that provides the user interface also generates all scheduled reports and provides access to the ScienceLogic API. The following appliances provide the user interface:

- ***All-In-One Appliance***. An ***All-In-One Appliance*** performs all functions, including providing the user interface.
- ***Database Server***. A ***Database Server*** can provide the user interface in addition to its database function.
- ***Administration Portal***. A dedicated ***Administration Portal*** appliance can provide the user interface.

NOTE: The Administration Portal communicates only with the Database Server and no other SL1 appliance. All connections between the Administration Portal and the Database Server are encrypted in both directions.

Database Server

The appliance that provides the database function is responsible for:

- Storing all configuration data and policy data.
- Storing performance data collected from managed devices.
- In a distributed system, pushing data to and retrieving data from the appliances responsible for collecting data and collecting messages.
- Processing and normalizing collected data.
- Allocating tasks to the other appliances in the SL1 System.
- Executing some automation actions in response to events.
- Sending all Email generated by the system.
- Receiving all inbound Email for events, ticketing, and round-trip Email monitoring.

The following appliances can perform these database functions:

- ***All-In-One Appliance***. An ***All-In-One Appliance*** performs all functions.
- ***Database Server***. A dedicated ***Database Server*** provides all database functions.

Data Collection

The SL1 appliances that retrieves data from monitored devices . In a distributed system, appliances that perform the data collection function also perform some pre-processing of collected data and execute automation actions.

The following appliances can perform the collection function:

- ***All-In-One Appliance***. An ***All-In-One Appliance*** performs all functions.
- ***Data Collector***. One or more Data Collectors care configured in ***collector groups*** for resilience. A collector group can be configured such that if an individual collector fails, other members of the group will pick up and share the load (N+1). A Data Collector can also perform the message collection function.

NOTE: The SL1 Agent can also be used to collect data from devices on which it can be installed. See the System Requirements page of the Support Site for a complete list of operating systems and versions supported by the agent. You can collect data from devices using only Data Collectors, using only the SL1 Agent, or using a combination of both.

Message Collection

The SL1 appliances that receive and process inbound, asynchronous syslog and trap messages from monitored devices.

The following appliances can perform the message collection function:

- ***All-In-One Appliance***. An ***All-In-One Appliance*** performs all functions.
- ***Message Collector***. A dedicated ***Message Collector*** receives and processes inbound, asynchronous syslog and trap messages from monitored devices.
    - In distributed systems that use the SL1 agent, the Message Collector passes agent data to the Database server. On these distributed systems, the ***Message Collector*** must be a stand-alone appliance, not a combination ***Data Collector/Message Collector***.
- ***Data Collector***. A Data Collector can also perform the message collection function in addition to the data collection function.

What is SL1 Extended?

SL1 Extended Architecture includes additional types of SL1 Appliances. The following SL1 features require the SL1 Extended Architecture:

- ***Expanded Agent Capabilities***. You can configure the SL1 Agent to communicate with SL1 via a dedicated Message Collector. However, this configuration limits the capabilities of the SL1 Agent. If you configure the SL1 Agent to communicate with SL1 via a Compute Cluster, you expand the capabilities of the SL1 Agent to include features like extensible collection and application monitoring.
- ***Data Pipelines***. Data pipelines transport and transform data. Data transformations include enrichment with metadata, data rollup, and pattern-matching for alerting and automation. The Data Pipelines provide an alternative to the existing methods of data transport (data pull, config push, streamer, and communication via encrypted SQL) in SL1. Data pipelines introduce message queues and communicate using encrypted web services.
- ***Publisher***. Publisher enables the egress of data from SL1. Publisher can provide data for long-term storage or provide input to other applications the perform analysis or reporting.

- ***Scale-out storage of performance data*** . Extended Architecture includes a non-SQL database (Scylla) for scalable storage of performance data.
- ***Anomaly Detection and future AI/ML developments***. Anomaly detection is a technique that uses machine learning to identify unusual patterns that do not conform to expected behavior. SL1 does this by collecting data for a particular metric over a period of time, learning the patterns of that particular device metric, and then choosing the best possible algorithm to analyze that data. Anomalies are detected when the actual collected data value falls outside the boundaries of the expected value range.

SL1 Extended Architecture includes the following additional SL1 Appliances:

Compute

Compute nodes are the SL1 appliances that transport, process, and consume the data from Data Collectors and the SL1 Agent. SL1 uses Docker and Kubernetes to deploy and manage these services. T

Load Balancer

The SL1 appliance that brokers communication with services running on the Compute Cluster. Services running on the Compute Cluster are managed by Kubernetes. Therefore, a single service could be running on one Compute node in the Compute Cluster; to provide scale, multiple instances of a single service could be running on one, many, or all nodes in the Compute Cluster. To provide scale and resiliency, you can include multiple Load Balancers in your configuration.

Storage

SL1 Extended includes a ***Storage Cluster*** that includes multiple Storage Nodes and a Storage Manager. These SL1 appliances provide a NoSQL alternative to the SL1 relational database. The Storage Cluster can store performance and log data collected by the Data Collectors and the SL1 Agent.

Management

The ***Management Node*** allows administrators to install, configure, and update packages on the Compute Nodes cluster, Storage Nodes , and the Load Balancer. The Management Node also allows administrators to deploy and update services running on the Computer Cluster.

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.

Third-Party Software

ScienceLogic does not support users installing third-party software on SL1 systems or users making unauthorized changes to the configuration of SL1. Doing so voids any warranties, express or implied.

Chapter 2

Preparing Hardware Appliances for SL1

Overview

This chapter describes how to prepare hardware appliances before installing SL1.

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:

Hardware Specifications

For details about supported ScienceLogic Hardware Appliances, see the ScienceLogic Support Site.
https://support.sciencelogic.com/s/system-requirements?tabset-e65a2=2

Prerequisites for SL1 Hardware Appliances

Perform the following steps to prepare an SL1 appliance for configuration:

- Install the SL1 appliance in a server rack and connect the power cables according to the instructions provided with the hardware.
- Connect the SL1 appliance to your network.
- Connect a monitor and keyboard to the SL1 appliance.

Initial Configuration for SL1 Hardware Appliances

You must perform the following tasks during initial configuration of an SL1 hardware appliance shipped by ScienceLogic:

- Change the password for the administrative user ***em7admin***.
- Change the primary IP address of the appliance. You must have already allocated IP addresses for the SL1 appliances.
- Change the netmask for the primary IP address of the appliance.
- Change the IP address for the network gateway.
- Change the IP address for the primary Nameserver.

Changing the Password for em7admin

To change the password for the default administrative user ***em7admin*** for console logins and SSH access:

1. Either go to the console of the SL1 appliance or use SSH to access the server.
2. Log in as user ***em7admin*** with the appropriate password. The default password is ***em7admin***.
3. At the shell prompt, type the following:

   ```
   passwd
   ```

4. When prompted, type and re-type the new password.

Changing Network Settings

To change the IP address, Netmask, Gateway address, and DNS Server for an appliance in the ifconfig file:

1. Either go to the console of the SL1 appliance or use SSH to access the server.
2. Login as user ***em7admin*** with the appropriate password.
3. Enter the following at the command line:

   ```
   sudo ifconfig
   ```

4. Your output will look like this:

```
ens32: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 10.64.68.20 netmask 255.255.255.0 broadcast 10.64.68.255
inet6 fe80::250:56ff:fe84:455f prefixlen 64 scopeid 0x20<link>
ether 00:50:56:84:45:5f txqueuelen 1000 (Ethernet)
RX packets 1774927 bytes 161985469 (154.4 MiB)
RX errors 0 dropped 861 overruns 0 frame 0
TX packets 1586042 bytes 158898786 (151.5 MiB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
inet 127.0.0.1 netmask 255.0.0.0
inet6 ::1 prefixlen 128 scopeid 0x10<host>
loop txqueuelen 0 (Local Loopback)
RX packets 13406577 bytes 4201274223 (3.9 GiB)
RX errors 0 dropped 0 overruns 0 frame 0
TX packets 13406577 bytes 4201274223 (3.9 GiB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

5. Examine the output, find the first interface in the output, and note its name.

6. Use the vi editor to edit the settings for the interface. To do this, enter the following at the command line:

```
sudo vi /etc/sysconfig/network-scripts/ifcfg-interface name you noted in step #4
```

   For example, from our output, we could enter:

```
sudo vi /etc/sysconfig/network-scripts/ifcfg-ens32
```

7. Your output will look like this:

```
TYPE=Ethernet
BOOTPROTO=none
DNS1=10.64.20.33
DEFROUTE=yes
IPV4_FAILURE_FATAL=no
IPV6INIT=yes
IPV6_AUTOCONF=yes
IPV6_DEFROUTE=yes
IPV6_FAILURE_FATAL=no
NAME=ens32
UUID=d471435d-9adf-47c9-b3f3-32f61dccbad8
DEVICE=ens32
ONBOOT=yes
IPADDR=10.64.68.20
PREFIX=24
GATEWAY=10.64.68.1
IPV6_PEERDNS=yes
IPV6_PEERROUTES=yes
```

8. You can edit one or more of the following settings:
    - ***DNS1***=IP address of the DNS server that will be used by the SL1 appliance.
    - ***IPADDR***=IP address of the SL1 appliance.
    - ***PREFIX***=netmask for the SL1 appliance.
    - ***GATEWAY***=IP address of the network gateway that will be used by the SL1 appliance.
9. Save your changes and exit the file (:wq)
10. At the command line, enter the following:

```
sudo service network restart
```

Ports for SL1 Hardware Appliances

See the chapter on ***ports*** to configure firewalls to allow traffic to and from the SL1 appliances.

Chapter 3

Preparing Virtual Machines for SL1

Overview

This chapter describes how to prepare virtual appliances before installing SL1.

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:

Virtual Machine Specifications

For details about supported hypervisors and the requirements and specifications for each SL1 appliance, see the ScienceLogic Support Site: https://support.sciencelogic.com/s/system-requirements

NOTE: You must have already allocated an IP address for each SL1 appliance.

Build Appliances in This Order

For ease of configuration, create appliances in this order:

1. Database Server
2. Administration Portal (if applicable)
3. Data Collectors
4. Message Collectors (if applicable)

Building an Appliance on a VMware System

To deploy a SL1 appliance on a VMware system:

1. Using the vSphere client, connect to your VMware system as a user that has permissions to deploy a new virtual machine and use the Create New Virtual Machine wizard to create a new virtual machine.

3. In the Create New Virtual Machine wizard, select the configuration options that are appropriate for your environment and the current recommended specifications for the appliance type you are installing. For details about requirements and specifications, see the ScienceLogic Support Site: https://support.sciencelogic.com/s/system-requirements
4. On the **Guest Operating System** page, select *Linux* as the ***Guest Operating System***, and then select *Oracle Linux 4/5/6/7 (64-bit)* in the ***Version*** drop-down list.

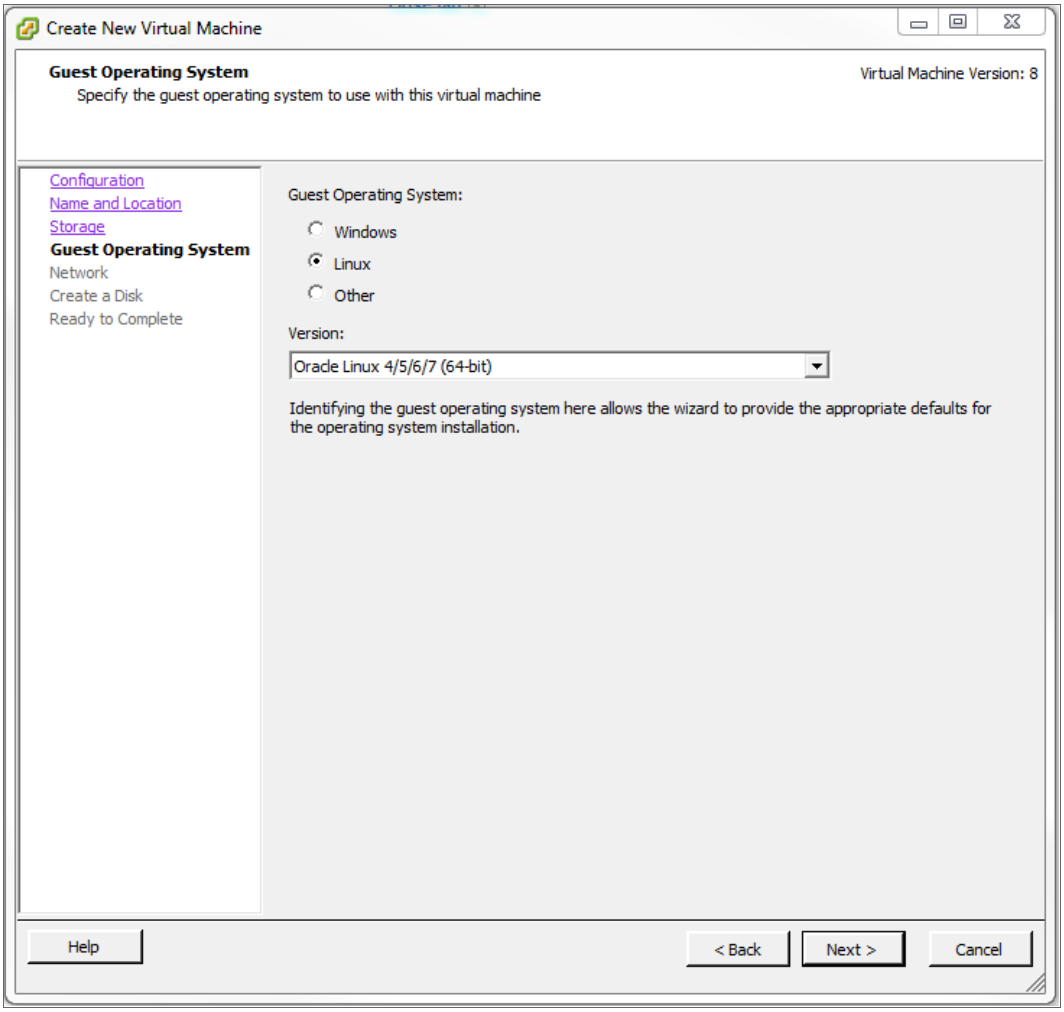

5. On the **Network** page, you must select *VMXNET 3* in the ***Adapter*** field.
6. After completing the Create New Virtual Machine wizard, edit the settings for the virtual machine:
    - Set the CPU and memory allocation to the values recommended in https://support.sciencelogic.com/s/system-requirements
    - Configure the CD/DVD drive to use the SL1 ISO file.
7. Repeat these steps for each appliance in your SL1 system.

Installing VMware Tools

You must install VMware Tools on each Database Server, each Data Collector, and each Message Collector. You can install VMware tools in two ways:

- If your appliance can connect to the Internet, use the yum utility to install the necessary packages.
- If you have an appliance that is not able to reach the Internet, you can retrieve the required packages from a similar appliance that does have Internet access.

Installing VMware Tools Using Yum

To install VMware tools using the yum utility:

1. Log in to the appliance as the em7admin user using the console or SSH.
2. Execute the following command:

   ```
   sudo yum install open-vm-tools
   ```
3. Type the password for the em7admin user when prompted.
4. When prompted to confirm the installation, type "y".
5. Execute the following commands:

   ```
   sudo systemctl start vmtoolsd.service
   sudo systemctl enable vmtoolsd.service
   sudo systemctl status vmtoolsd.service
   ```

   If the installation was successful, the "Active" line in the output indicates VMware tools is "active (running)".

Installing VMware Tools Manually

To install VMware tools manually:

1. Retrieve the required packages from an appliance that has Internet access with the following command:

   ```
   sudo yum install open-vm-tools --downloadonly --downloaddir="/var/tmp/vmtools"
   ```
2. Once the download is complete, gather the downloaded RPM files into an archive file by running the following command, where "vmtools.tgz" can be any filename you choose:

   ```
   cd /var/tmp && tar cvfz vmtools.tgz vmtools
   ```
3. Transfer the archive file to the appliance that does not have Internet access, and extract the RPMs by running the following command:

   ```
   tar zxvf [name of the archive file]
   ```
4. Install the files with the following command:

   ```
   sudo rpm -ivh vmtools/*.rpm
   ```
5. Start the vmtoolsd service with the following command:

   ```
   sudo systemctl start vmtoolsd
   ```

6. To ensure that vmtoolsd starts automatically after a reboot, run the following command:

   ```
   sudo systemctl enable vmtoolsd
   ```

7. Execute the following command:

   ```
   sudo systemctl status vmtoolsd.service
   ```

   If the installation was successful, the "Active" line in the output indicates VMware tools is "active (running)".

Building an Appliance on a Hyper-V System

To deploy an SL1 appliance on a Hyper-V system:

1. Follow the instructions from Microsoft:

   https://docs.microsoft.com/en-us/virtualization/hyper-v-on-windows/quick-start/create-virtual-machine#create-a-virtual-machine-with-hyper-v-manager

2. When prompted to **select a Generation** for the VM:
   - ***Generation 1***. Fully supports Oracle Linux and SL1.
   - ***Generation 2***. To support Oracle Linux and SL1, you must disable the "secure boot" feature.
3. When prompted to **Assign Memory** and **Connect Virtual Hard Disk**, enter the hardware requirements as specified here:

   https://support.sciencelogic.com/s/system-requirements

4. In the **Installation Options** wizard, select ***Install an operating system later***
5. Click **[Finish]**.
6. If you selected a ***Generation 2*** virtual machine, open a PowerShell session on the Hyper-V Manager host and execute the following PowerShell cmdlet to disable secure boot on the VM:

   ```
   Set-VMFirmware "Test VM" -EnableSecureBoot Off
   ```

7. Follow the steps specified here to install the Operating System (Oracle Linux 64 bit)

   https://docs.microsoft.com/en-us/virtualization/hyper-v-on-windows/quick-start/create-virtual-machine#complete-the-operating-system-deployment

8. Repeat these steps for each appliance in your SL1 system.
9. To install SL1 on the Hyper-V virtual machines, see the chapter ***Installing SL1 on Hardware Appliances and Virtual Appliances***.

Ports for Virtual Appliances

See the chapter on ***ports*** to configure firewalls to allow traffic to and from the SL1 appliances.

Chapter 4

Required Ports

Overview

This chapter describes the required open ports on each SL1 appliance. These open ports allow communication between appliances in an SL1 system.

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:

Open Ports on the ScienceLogic All-In-One Appliance

Name	Description	Protocol	Port
HTTP Interface	HTTP from browser session on user workstation. ScienceLogic recommends disabling HTTP during deployment.	TCP	80

Name	Description	Protocol	Port
HTTPS Secure Interface	Used for browser sessions on a user workstation, API requests from external systems, and requests from the ScienceLogic Agent running on a monitored device.	TCP	443
Database Web Admin	Optional. Administrative Web Interface (phpMyAdmin) from browser session on user workstation to Database.	TCP	8008
SSH	Optional. For ssh sessions from user workstation.	TCP	22
Web Configurator	Configuration Utility from browser session on user workstation.	TCP	7700
SNMP	Optional. SNMP information about the All-In-One Appliance can be collected by SL1.	UDP	161
SNMP Traps	Optional. Can receive SNMP traps from managed devices.	UDP	162
Syslog messages	Optional. Can receive syslog messages from managed devices.	UDP	514
SMTP	Optional. To receive inbound Email for tickets, events, and email round-trip monitoring.	TCP	25
DRBD Replication	This port is open only if your All-In-One Appliance is configured for Disaster Recovery.	TCP	7788

Open Ports on the ScienceLogic Database Server Appliance

Name	Description	Protocol	Port
HTTP Interface	Optional. Can be used if the Database Server also serves as an Administration Portal.	TCP	80
HTTPS Secure Interface	Optional. Can be used if the Database Server also serves as an Administration Portal.	TCP	443
Database Web Admin	Optional. Administrative Web Interface (PHPMyAdmin) from browser session on user workstation.	TCP	8008
MariaDB	Communication from Administration Portal. Communication from HA-secondary and DR to HA primary. **NOTE:** If you are using HA/DR, you must keep this port open. This port is required for communication between the HA-secondary and DR to the HA-primary appliance. If this port is blocked, the em7service on these databases will fail and could lead to issues such as DR backup not working or inability to license the appliances.	TCP	7706
SSH	Optional. Can be manually closed. For ssh sessions from user workstation.	TCP	22
Web Configurator	Configuration Utility from browser session on user workstation.	TCP	7700

Name	Description	Protocol	Port
SNMP	Optional. SNMP information about the Database Server can be collected by SL1.	UDP	161
ScienceLogic HA	Optional. Communication between Database Server and other Database Server(s) in a high-availability cluster.	TCP	694
SMTP	Optional. Can be manually closed. To receive inbound email for tickets, events, and email round-trip monitoring.	TCP	25
High Availability	One of two ports used by the cluster management process to test cluster availability. This port is open only if your Database Server appliance is configured for High Availability.	UDP	5555
High Availability	One of two ports used by the cluster management process to test cluster availability. This port is open only if your Database Server appliance is configured for High Availability.	UDP	5556
DRBD Replication	This port is open only if your Database Server appliance is configured for High Availability, Disaster Recovery, or both.	TCP	7788
PhoneHome Configuration	This port is open only if your Database Server appliance is configured for PhoneHome communication from Data Collectors and Message Collectors. The port number is configurable.	TCP	7705

Open Ports on the ScienceLogic Administration Portal Appliance

Name	Description	Protocol	Port
HTTP Interface	HTTP from browser session on user workstation.	TCP	80
HTTPS Secure Interface	Used for browser sessions on a user workstation and API requests from external systems.	TCP	443
SSH	Optional. For ssh sessions from user workstation.	TCP	22
Web Configurator	Configuration Utility from browser session on user workstation.	TCP	7700
SNMP	Optional. SNMP information about the Administration Portal can be collected by SL1.	UDP	161

Open Ports on the ScienceLogic Data Collector Appliance

Name	Description	Protocol	Port
Data Pull	Requests from Database Servers to retrieve collected data. In a Phone Home configuration, this port is accessed via an SSH tunnel created by the Data Collector.	TCP	7707
SSH	Optional. For ssh sessions from user workstation.	TCP	22
Web Configurator	Configuration Utility from browser session on user workstation.	TCP	7700
SNMP	Optional. SNMP information about the Data Collector can be collected by SL1.	UDP	161
SNMP Traps	Optional. Can receive SNMP traps from managed devices.	UDP	162
Syslog messages	Optional. Can receive syslog messages from managed devices.	UDP	514
HTTPS Secure Interface	Optional. Data from the ScienceLogic Agent running on a monitored device.	TCP	443

Open Ports on the ScienceLogic Message Collector Appliance

Name	Description	Protocol	Port
Data Pull	Requests from Database Servers to retrieve collected data. In a Phone Home configuration, this port is accessed via an SSH tunnel created by the Message Collector.	TCP	7707
SSH	Optional. For ssh sessions from user workstation.	TCP	22
Web Configurator	Configuration Utility from browser session on user workstation.	TCP	7700
SNMP	Optional. SNMP information about the Message Collector can be collected by SL1.	UDP	161
SNMP Traps	Optional. Can receive SNMP traps from managed devices.	UDP	162
Syslog messages	Optional. Can receive syslog messages from managed devices.	UDP	514
HTTPS Secure Interface	Optional. Data from the ScienceLogic Agent running on a monitored device.	TCP	443

Chapter 5

Installing SL1 on Hardware Appliances and Virtual Appliances

Overview

This chapter describes how to install SL1 on hardware Appliances or on virtual machines.

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:

Prerequisites

To perform the steps in this section:

- You must have already performed the prerequisites for all the *ScienceLogic Hardware Appliances* in your SL1 stack or for all the *Virtual Appliances* in your SL1 stack.
- You must have a valid customer account that allows you to download the SL1 ISO. For details, contact your Customer Success Manager.
- You must download the SL1 ISO.
- Mount the SL1 ISO on each virtual appliance.
- You must have access to the files for your SSL certificate.
- You must have a valid customer account that allows you to access the Artifactory page on the ScienceLogic Support Site. For details, contact your Customer Success Manager.

CAUTION: ScienceLogic does not support vmotion or VMware Snapshots for backups of data. For backup purposes, ScienceLogic supports only SL1 backups to remote storage. vmotion andVMware Snapshots can cause SL1 outages. For details on SL1 backups, see the chapter on *Backup Management* in the ***System Administration*** manual.

Upgrading

For detailed upgrade instructions, see the chapter on *Upgrading SL1*, in the ***System Administration*** manual.

Installing the Database Server

The Database Server should be the first appliance you install. To do so:

1. Boot the appliance from the SL1 ISO.

NOTE: If you are using Hyper-V, check that the ScienceLogic installation ISO mounted correctly and that the Virtual Machine displays the install screen. To do this, right-click the Virtual Machine in inventory and select *Connect or View* and then *Connect via Console*.

2. The following window appears:

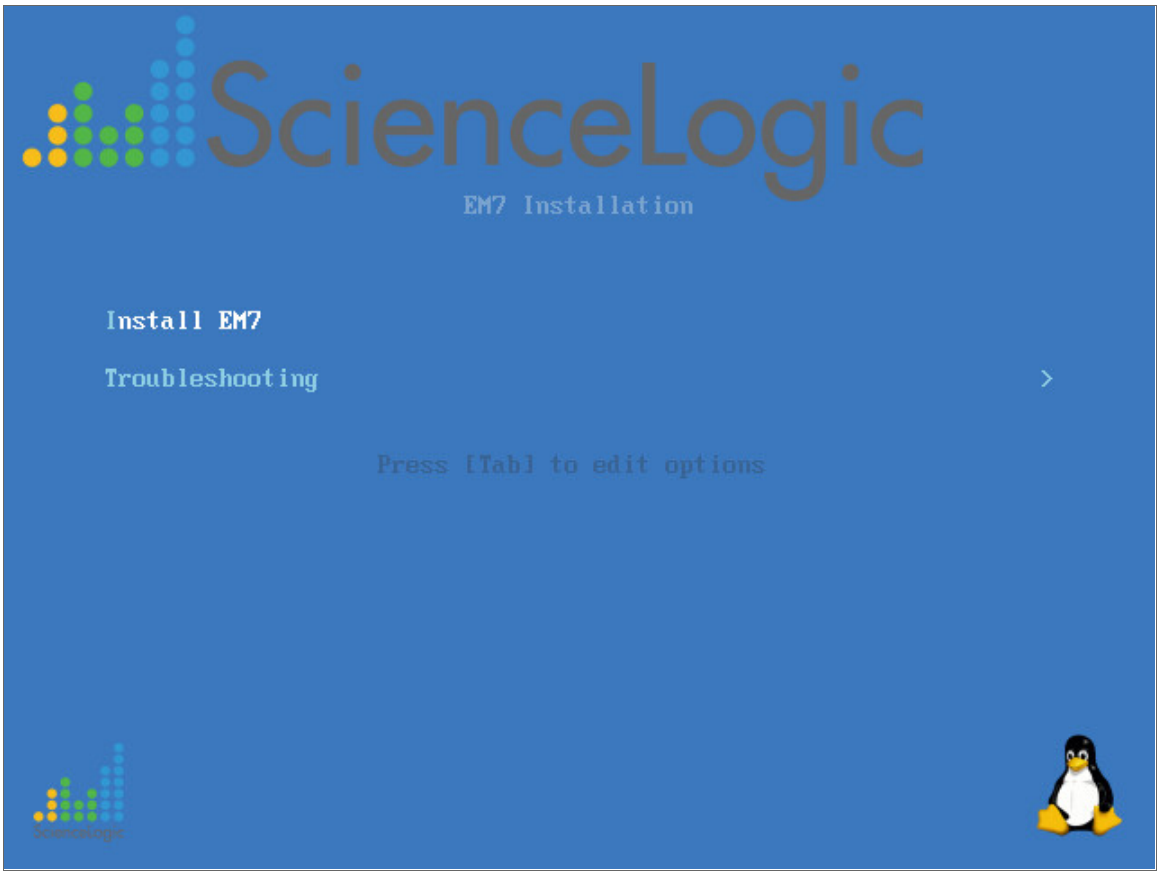

3. Select ***Install EM7***. The **Model Type** window appears.

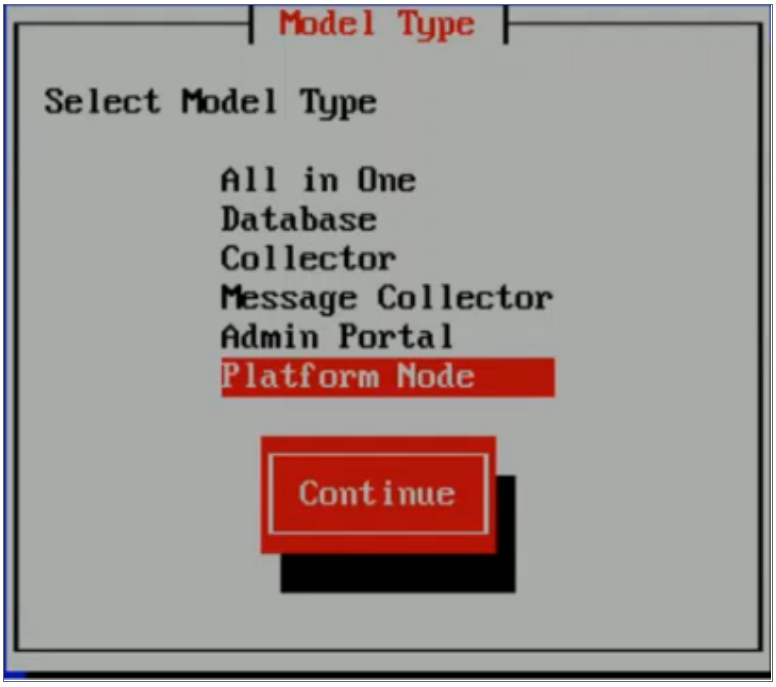

4. Select the ***Database***. Select **[Continue]**.
5. The Military Unique Deployment window appears. ***Do not select if you are not using a Military Unique Deployment.***

6. Select **[Continue]**. The **Database** window appears:

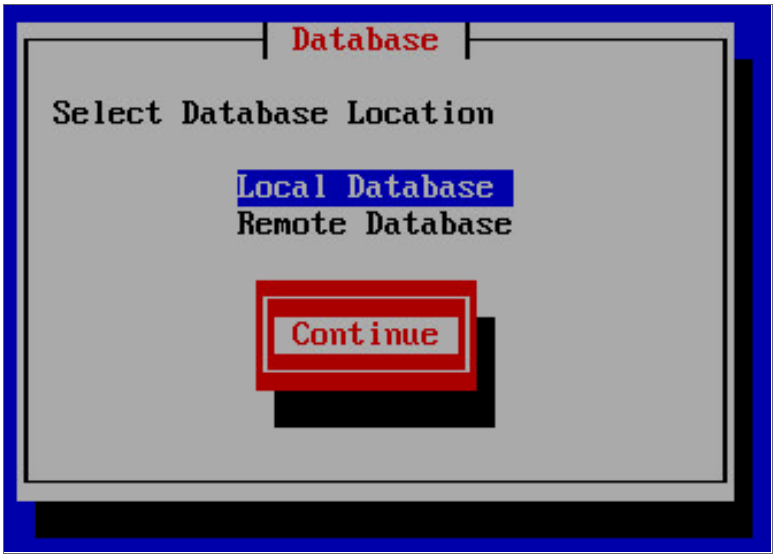

7. In the **Database** window, select ***Local Database*** and select **[Continue]**.
8. After the installer for the selected appliance type is loaded, the **Network Configuration** window appears.

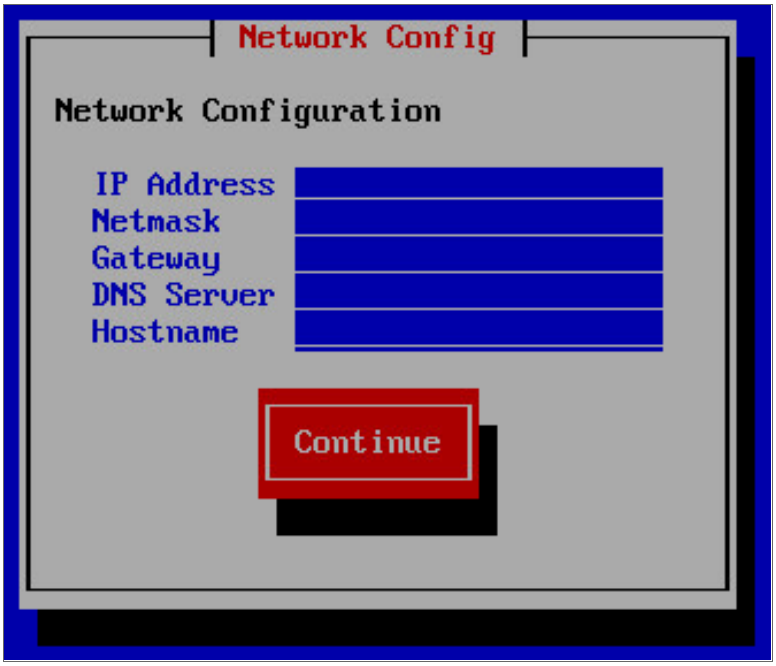

9. Enter the following information:
    - ***IP Address***. Type the primary IP address of the appliance.
    - ***Netmask***. Type the netmask for the primary IP address of the appliance.
    - ***Gateway***. Type the IP address for the network gateway.
    - ***DNS Server***. Type the IP address for the primary Nameserver.
    - ***Hostname***. Type the hostname for the appliance.

10. Select **[Continue]**.
11. The **System Password** window appears:

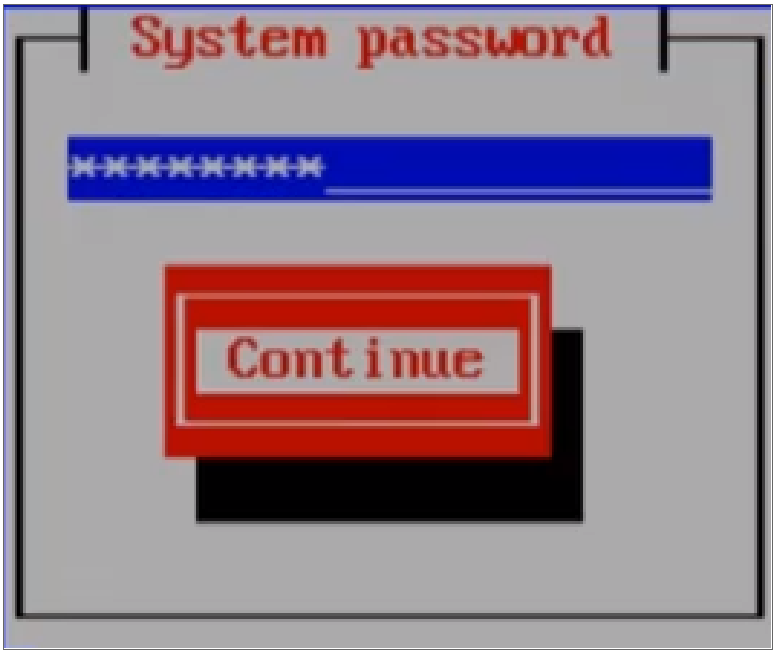

12. Type the password for the em7admin user on the operating system and select **[Continue]**.
13. Type the password for the em7admin user again and select **[Continue]**.
14. The appliance installer runs, and the virtual machine reboots automatically.
15. If you are using a VMware instance, after the appliance reboots, follow the instructions to *install VMware tools*.
16. *Follow the instructions to license the appliance*.

Installing the Administration Portal, Data Collector and/or Message Collector

After installing the Database Server, you can next install:

1. The Administration Portal (if applicable)
2. The Data Collectors
3. The Message Collectors (if applicable)

You can use the following instructions to build the Administration Portal, and one or more Data Collectors and Message Collectors.

1. Boot the appliance from the SL1 ISO.
2. Select ***Install EM7***. The **Model Type** window appears.

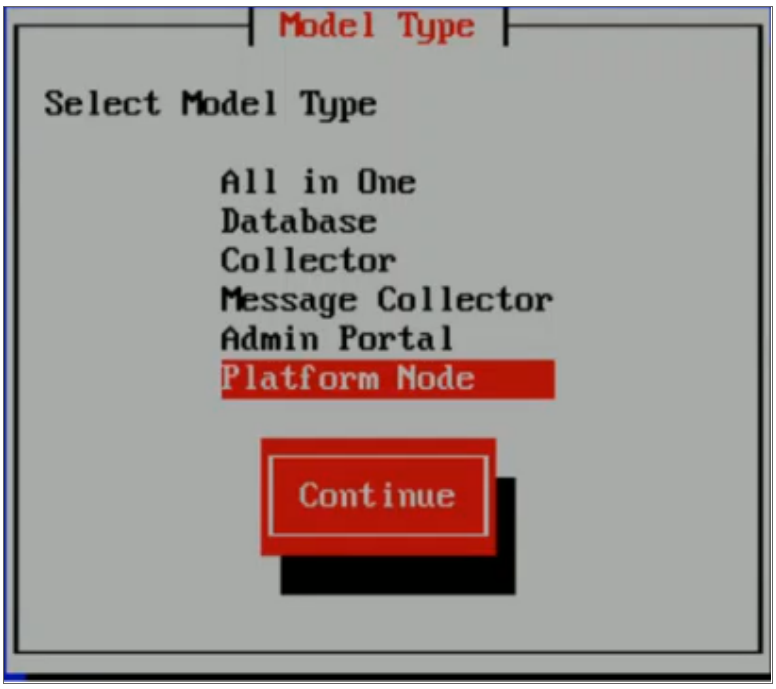

3. Select the appropriate appliance type. Select **[Continue]**.
4. The Military Unique Deployment window appears. ***Do not select if you are not using a Military Unique Deployment.***

5. After the installer for the selected appliance type is loaded, the **Network Configuration** window appears.

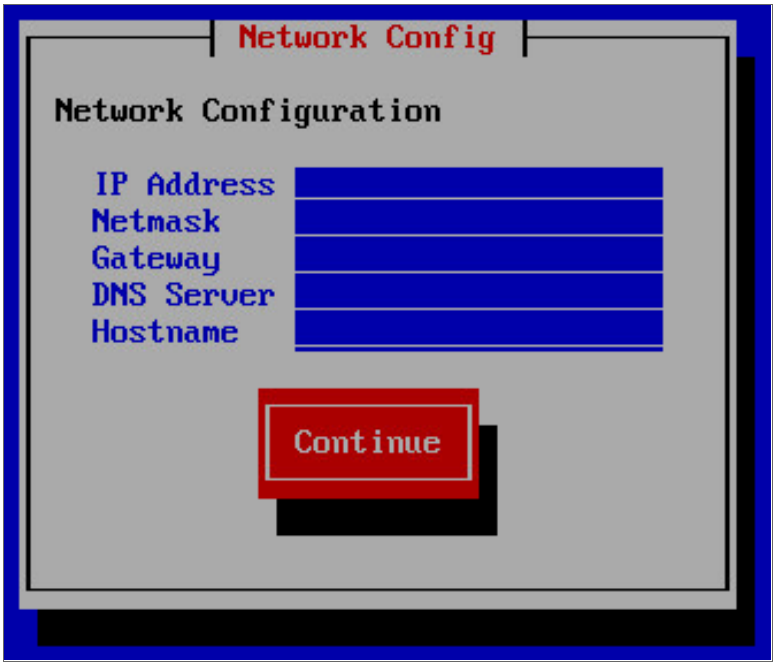

6. Enter the following information:
   - ***IP Address***. Type the primary IP address of the appliance.
   - ***Netmask***. Type the netmask for the primary IP address of the appliance.
   - ***Gateway***. Type the IP address for the network gateway.
   - ***DNS Server***. Type the IP address for the primary Nameserver.
   - ***Hostname***. Type the hostname for the appliance.

7. Select **[Continue]**.
8. The **System Password** window appears:

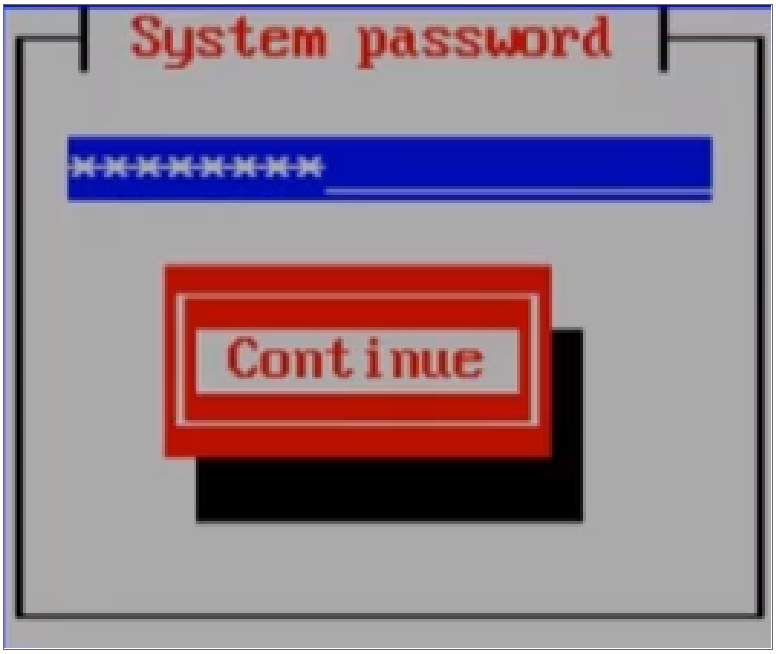

9. Type the password for the em7admin user on the operating system and select **[Continue]**.
10. Type the password for the em7admin user again and select **[Continue]**.
11. If you are using a VMware instance, after the appliance reboots, follow the instructions to *install VMware tools*.
12. *Follow the instructions to configure the appliance in the Web Configuration Tool*.

Licensing the SL1 Appliances

For details on licensing the SL1 Appliances, see the chapter *Licensing and Configuring an Appliance*.

Additional Steps for SL1 10.1

SL1 10.1.x includes an upgrade to MariaDB. The upgrade did not include a tool, jemalloc, that helps manage memory usage.

NOTE: This section applies only to the following releases:

- 10.1.0
- 10.1.1
- 10.1.2
- 10.1.3
- 10.1.4
- 10.1.4.1
- 10.1.4.2
- 10.1.5
- 10.1.5.1

For SL1 versions later than 10.1.5.1, jemalloc is included with the platform. For SL1 versions prior to 10.1.0, jemalloc is included with the platform.

To avoid problems with memory usage on Database Servers, perform the following steps after upgrading MariaDB for 10.1.x.

NOTE: Perform these steps first on the active Database Server and then on each additional Database Server in your SL1 System.

1. Open an SSH session to the Database Server.
2. To verify that the Database Server is not currently running jemalloc, enter the following at the shell prompt:

   ```
   silo_mysql -e 'show global variables like "version_malloc_library"'
   ```

   If the Database Server is not currently running jemalloc, the shell will display the following:

Variable Name	Value
version_malloc_library	system

3. Search for the file /usr/lib64/**libjemalloc.so.1**.

   If the file does not exist, contact ScienceLogic Customer Support to request the file jemalloc-3.6.0-1.el7.x86_64.rpm.

   To install the RPM, use a file-transfer utility, copy the file to a directory on the SL1 appliance. Then enter the following at the shell prompt:

   ```
   cd /usr/lib64
   sudo yum install jemalloc-3.6.0-1.el7.x86_64.rpm
   ```

4. Create the file /etc/systemd/system/mariadb.service.d/jemalloc.conf:

   ```
   vi /etc/systemd/system/mariadb.service.d/jemalloc.conf
   ```

5. Add the following lines to the file:

   ```
   [Service]
   ```

```
Environment="LD_PRELOAD=/usr/lib64/libjemalloc.so.1"
```

6. Save and close the file.
7. Reload the systemd config files:

```
sudo systemctl daemon-reload
```

8. Restart the Database Server:

   To restart the ***standalone Database Server*** or the ***primary Database Server in a cluster***, enter the following:

```
sudo systemctl restart mariadb
```

   To restart each ***secondary Database Server in a cluster***:

   a. Open an SSH session to the secondary Database Server. At the shell prompt, enter:

```
coro_config
```

   b. Select **1**.

   c. When prompted to put the Database Server into maintenance, select **y**.

   d. Open an SSH session to the primary Database Server. To pause SL1, enter the following at the shell prompt:

```
sudo touch /etm/.proc_mgr_pause
```

   e. In the SSH session for the secondary Database Server, restart MariaDB:

```
crm resource restart mysql
```

   f. After MariaDB has restarted successfully on the secondary Database Server, return to the SSH session on the primary Database Server. Remove the pause file for SL1:

```
sudo rm /tmp/.proc_mgr_pause
```

   g. In the SSH session on the secondary Database Server, take the Database Server out of maintenance. At the shell prompt, enter:

```
coro_config
```

   h. Select **1**.

   i. When prompted to take the Database Server out of maintenance, select **y**.

9. To verify that jemalloc is running on the Database Server, enter the following at the shell prompt:

```
silo_mysql -e 'show global variables like "version_malloc_library"'
```

   If the Database Server is currently running jemalloc, the shell will display something like the following:

Variable Name	Value
version_malloc_library	jemalloc 3.6.0-0-g46c0af68bd248b04df75e4f92d5fb804c3d75340

10. Perform these steps on each Database Server in your SL1 system.

Chapter 6

Licensing and Configuring an Appliance

Overview

This chapter describes how to license an SL1 appliance and add it to your SL1 system.

Upon installation, SL1 appliances are automatically licensed for 30 days. During these 30 days, you can perform the steps to obtain a permanent license from ScienceLogic.

SL1 appliances automatically generate a Registration Key file. This file is used by ScienceLogic to generate a unique License Key file. **You must not edit or alter the Registration Key file.** While performing the steps described in this chapter, you must obtain a License Key file by providing the Registration Key file to ScienceLogic.

For distributed SL1 systems, you must license the Database Server first. All other SL1 appliances in a distributed SL1 system depend on the Database Server for registration.

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:

Logging in to the Web Configuration Utility

Perform the following steps to log in to the Web Configuration Utility:

1. You can log in to the Web Configuration Utility using any web browser supported by SL1. The address of the Web Configuration Utility is in the following format:

   `https://ip-address-of-appliance:7700`

NOTE: For AWS instances, ***ip-address-of-appliance*** is the public IP for the AWS instance. To locate the public IP address for an AWS instance, go to AWS, go to the **Instances** page, and highlight an instance. The **Description** tab in the lower pane will display the public IP.

2. When prompted to enter your user name and password, log in as the "em7admin" user with either the default password of ***em7admin*** or the password you configured.

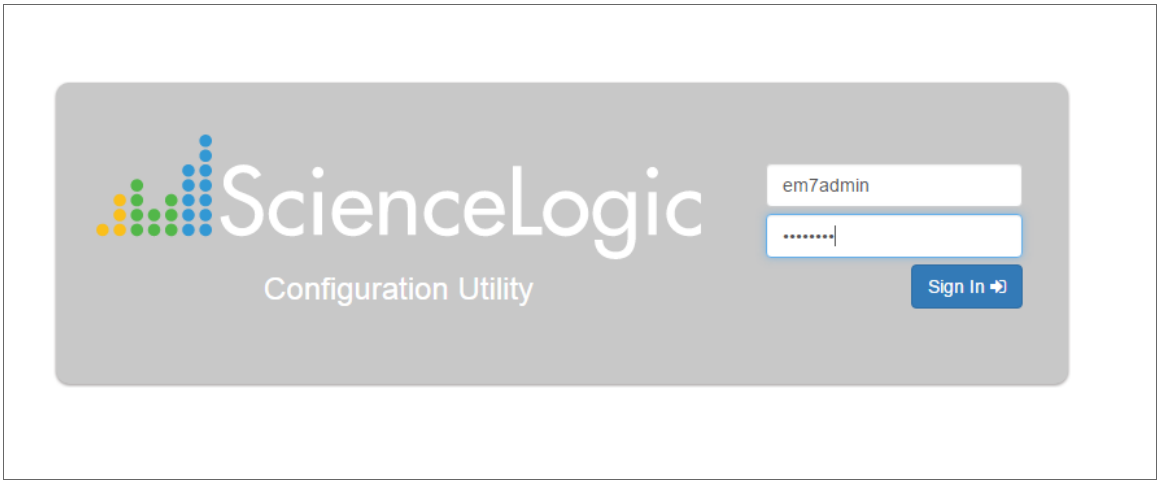

3. After logging in, the main **Configuration Utility** page appears:

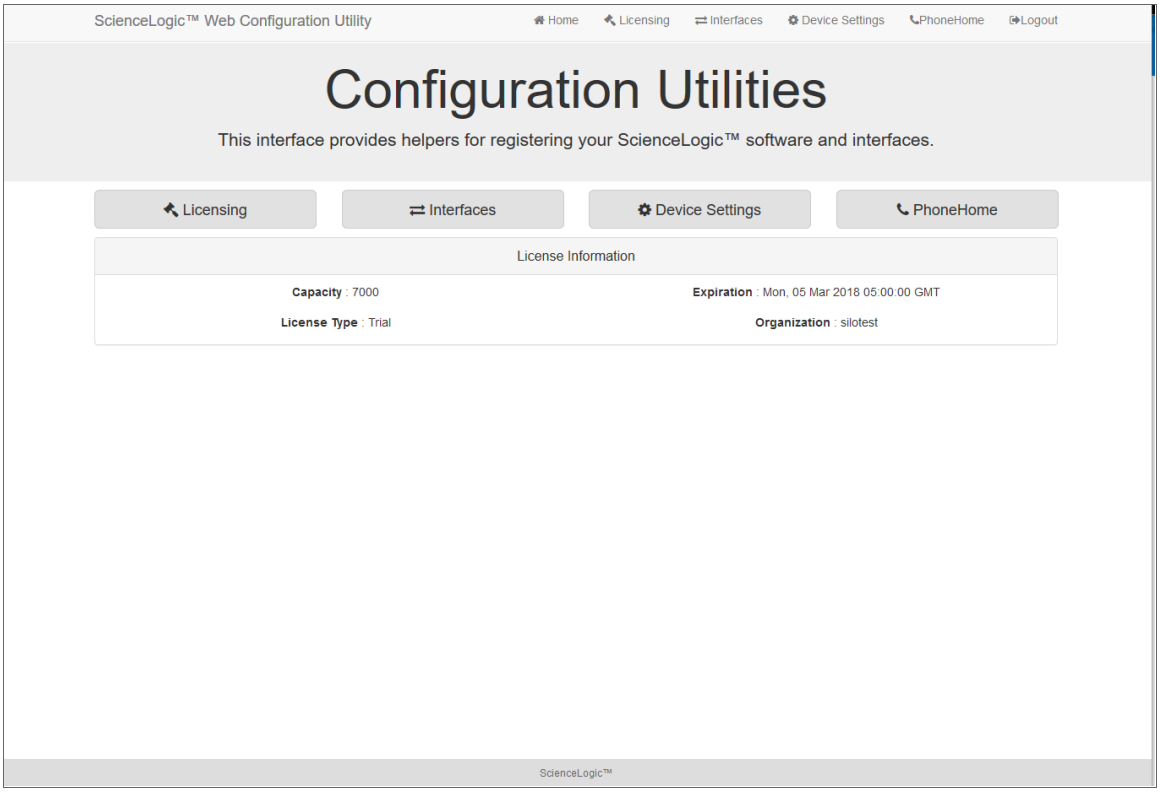

Changing the Password for the Web Configuration Utility

You can change the password for the Web Configuration Utility.

NOTE: If you want to change the password for the Web Configuration Utility on all SL1 appliances, you must log in to the Web Configuration Utility on each appliance and perform the steps in this section.

NOTE: You cannot change the username for the Web Configuration Utility. The username remains ***em7admin***.

To change the password for the Web Configuration Utility:

1. Log in to the Web Configuration Utility. The **Configuration Utilities** page appears.

2. Click the **[Device Settings]** button. The **Settings** page appears.

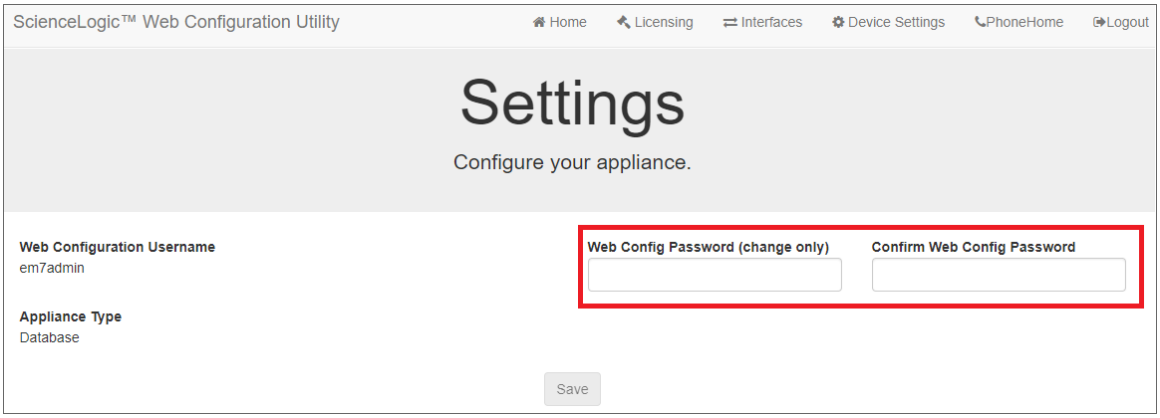

3. In the **Settings** page, type the following:
   - ***Web Config Password (change only)***. Type the new password.
   - ***Confirm Web Config Password***. Type the new password again.
4. Click **[Save]**
5. Perform steps 1-4 for each appliance for which you want to change the password for the Web Configuration Utility.

Licensing and Configuring a Database Server or All-In-One Appliance

You must perform the following steps in the Web Configuration Utility to license an All-In-One Appliance or Database Server:

1. *Log in to the Web Configuration Utility*. The **Configuration Utilities** page appears.
2. Click the **[Licensing]** button. The **Licensing Step 1** page appears.

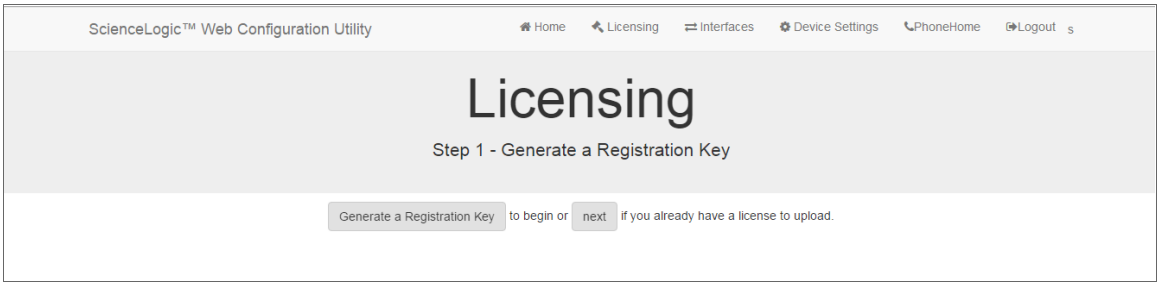

3. In the Licensing Step 1 page, click the **[Generate a Registration Key]** button.
4. When prompted, save the Registration Key file to your local disk.
5. Log in to the ScienceLogic Support Site (https://support.sciencelogic.com).

- Click your user name and from the menu select ***My Support and Customer Success***.
- On the next page, click the **[Submit a License Request]** button.
- Fill out the ***Appliance Information*** form and click the **[Submit License Request]** button.
- In the ***Upload Appliance Registration Key*** field, click the**[Upload Files]**button and navigate to the file where you saved the Registration Key file.
- ScienceLogic Customer Support will generate a license for the All-In-One Appliance or Database Server.

6. When you have the license for the All-In-One Appliance or Database Server, return to the Web Configuration Utility.

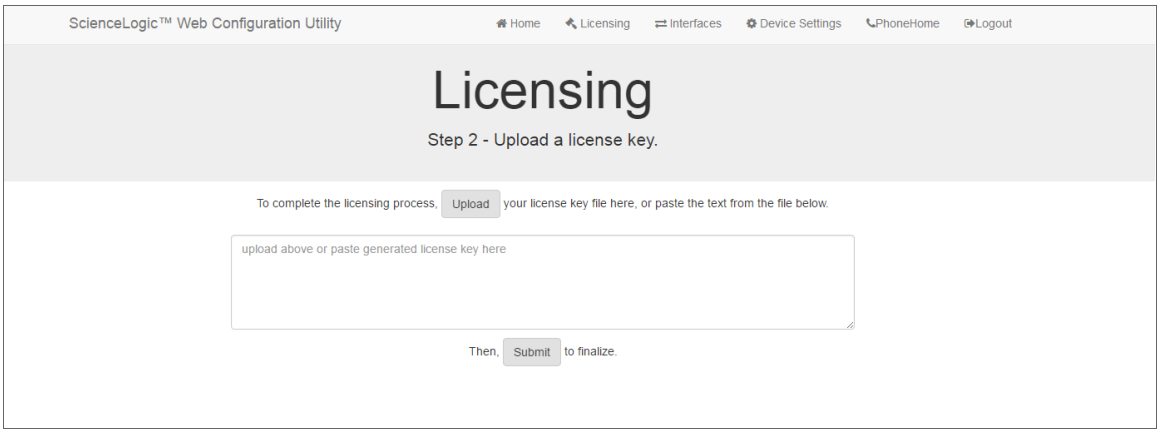

7. In the ***Licensing Step 2*** page, click the **[Upload]** button to upload the license file.
8. After navigating to and selecting the license file, click the **[Submit]** button to finalize the license. If the license key is correct and has been saved successfully, the message "Success: Thank you for licensing your ScienceLogic product!" appears.

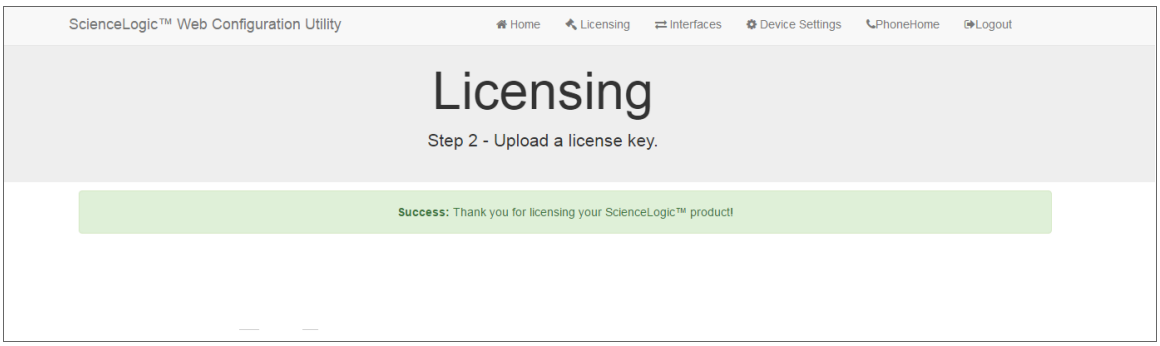

Configuring an Administration Portal

1. *Log in to the Web Configuration Utility*. The **Configuration Utilities** page appears.
2. Click the **[Device Settings]** button in the upper-right of the page. The **Settings** page appears.

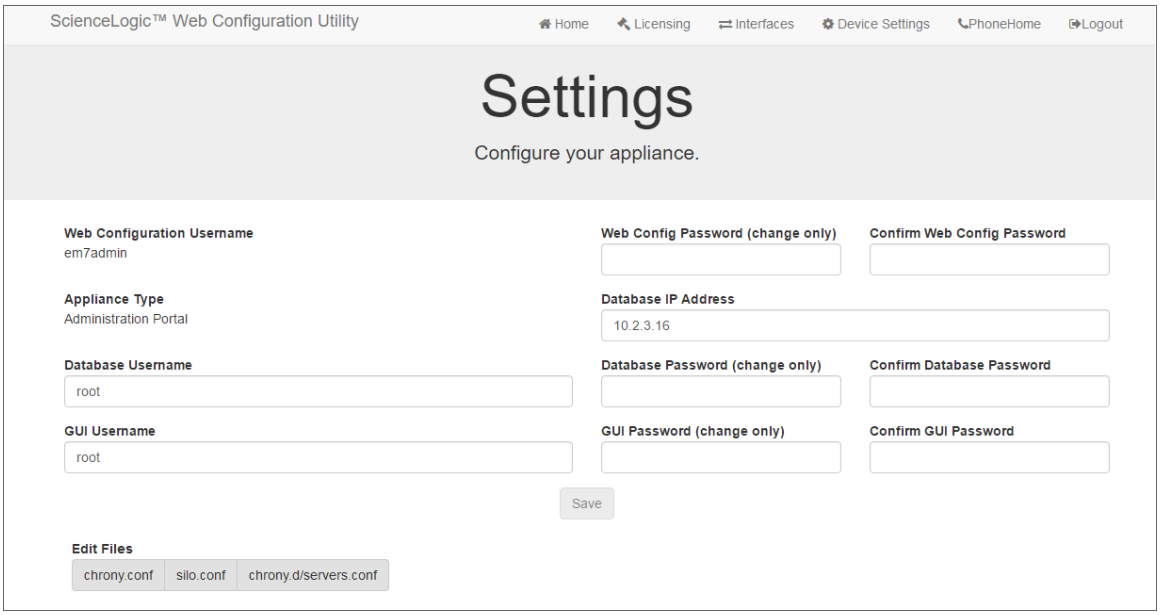

3. In the **Settings** page, enter the following:

   - ***Database IP Address***. The IP address of the primary ScienceLogic Database Server.
     - For an All-In-One Appliance with multiple Administration Portals, enter the IP address for the All-In-One Appliance.
     - If the Administration Portal and Database Server are AWS instances, supply the private IP address for the Database Server. To find the private IP of an AWS instance, go to AWS, go to the **Instances** page, and highlight an instance. The **Description** tab in the lower pane will display the private IP.
   - ***Database Username***. Username for the database account that the Administration Portal will use to communicate with the Database Server.
   - ***Accept the default values in all other fields***.

4. Click the **[Save]** button. You may now log out of the Web Configuration Utility.

Configuring a Data Collector or Message Collector

NOTE: The instructions for configuring a Data Collector or Message Collector for PhoneHome configuration differ from the instructions in this section. For details on configuring a Data Collector or Message Collector for PhoneHome configuration, see the chapter on *PhoneHome*.

You must perform the following steps in the Web Configuration Utility to configure a Data Collector or a Message Collector:

1. *Log in to the Web Configuration Utility* on the Data Collector or the Message Collector. The **Configuration Utilities** page appears.
2. Click the **[Device Settings]** button. The **Settings** page appears.

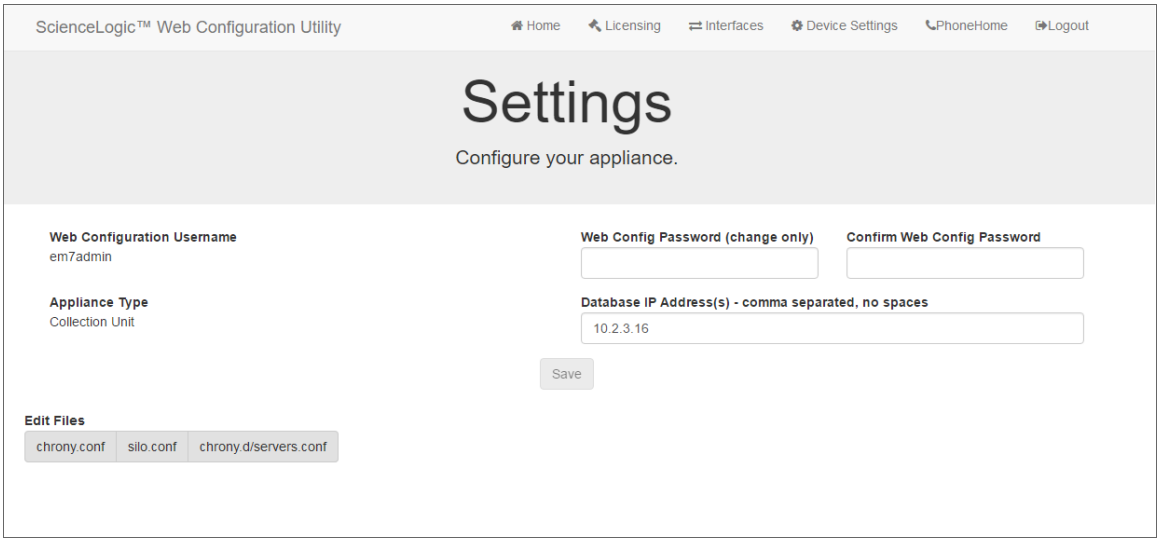

3. In the **Settings** page, update the following field:

    - ***Database IP Address***. The IP address of the ScienceLogic Database Server(s). If more than one Database Server will manage this appliance, type the IP addresses of the Database Servers, separated by commas. If the Data Collector or the Message Collector and the Database Server are AWS instances, supply the private IP address for the Database Server. To find the private IP of an AWS instance, go to AWS, go to the **Instances** page, and highlight an instance. The **Description** tab in the lower pane will display the private IP.

4. Click the **[Save]** button. You may now log out of the Web Configuration Utility.
5. Perform these steps for each Data Collector and Message Collector in your PhoneHome configuration.

Registering the Data Collector or Message Collector with the Database Server

After configuring a Data Collector or Message Collector in the Web Configuration Utility, you must register the appliance with the main Database Server in your SL1 system.

To register a Data Collector or Message Collector with the main Database Server, perform the following steps:

1. In the address bar of your browser, type the IP address of the SL1 appliance that provides the user interface for your SL1 system. The user interface is provided by either the Database Server or an Administration Portal. The login screen appears.
2. Log in as the "em7admin" user with the password "em7admin".
3. If this is your first successful login, you will be asked to agree to the End-user License Agreement. Read the End-user License Agreement then click the **[I Agree to The Terms Outlined Above]** button.

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

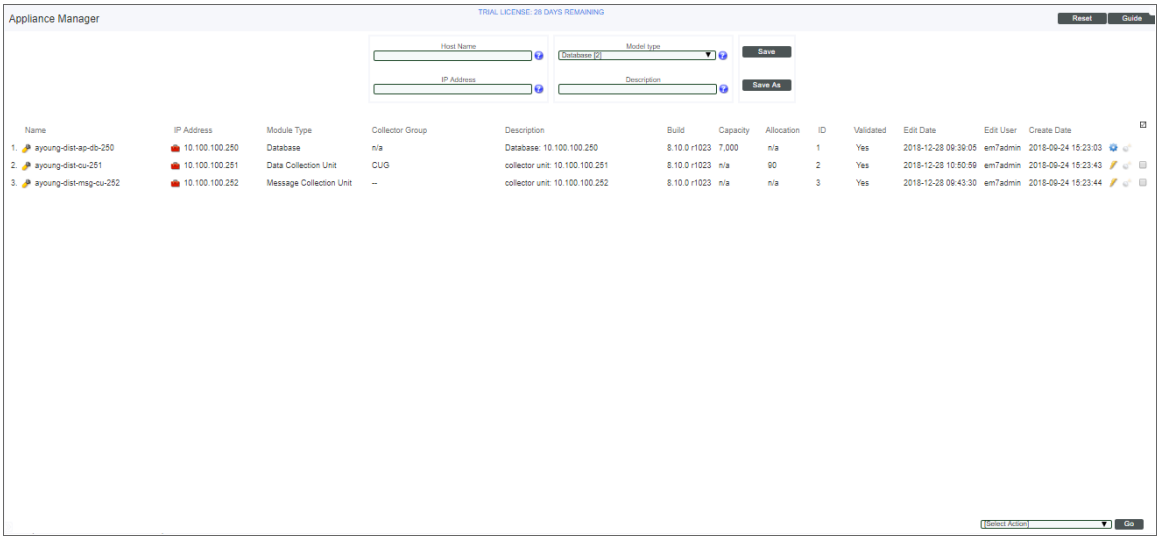

5. Supply values in the following fields:

   - ***Host Name***. Enter the hostname of the Data Collector or Message Collector.
   - ***IP Address***. Enter the IP address of the Data Collector or Message Collector. If the Data Collector or the Message Collector are AWS instances, supply the private IP address for the Data Collector or the Message Collector. To find the private IP of an AWS instance, go to AWS, go to the **Instances** page, and highlight an instance. The **Description** tab in the lower pane will display the private IP.
   - ***Model Type***. If you are configuring a Data Collector, select *Data Collection Unit [5]* from the drop-down list. If you are configuring a Message Collector, select *Message Collection Unit [6]* from the drop-down list.
   - ***Description***. Enter a description for the Data Collector or Message Collector. This field is optional.

6. Click the **[Save]** button. If the save is successful, the message "Appliance Registered" appears.

7. If you are using an AWS RDS system, select the wrench icon ( ) for the newly created Data Collector or Message Collector. Supply values in the ***DB User*** field and the ***DB Password*** field.

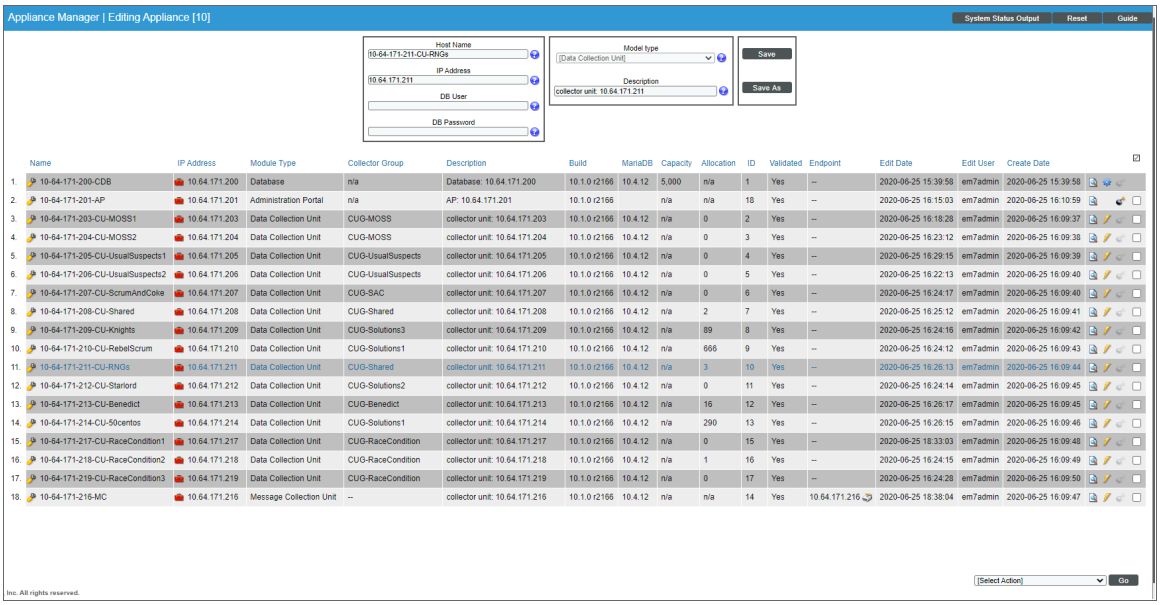

6. If all information is valid and the Database Server can communicate with the Data Collector or Message Collector, the appliance page will display "Yes" in the **Validated** column. If the **Validated** column displays "No" for longer than five minutes, double-check your settings and network connection.

Defining the Syslog Server

For each device except for Message Collectors and All-In-One Appliances, you must specify the IP address of the server to which the SL1 appliance will send syslog messages. Enter the IP address of the syslog server that will monitor this SL1 appliance. Usually, this is the IP address of a Message Collector, Data Collector, or All-In-One Appliance.

NOTE: A device configured with Transport Layer Security (TLS) support for an rsyslog server can successfully exchange messages with a SL1 appliance configured with TLS support for an rsyslog client.

To specify the syslog server:

1. Either go to the console of the SL1 appliance or use SSH to access the server.
2. Log in as user ***em7admin*** with the password you configured during setup.
3. Install the required Transport Layer Security (TLS) certificates by typing the following lines at the shell prompt:

```
mkdir -pv /etc/rsyslog.d/keys/ca.d
cd /etc/rsyslog.d/keys/ca.d/
```

NOTE: You might need to create a **ca.d** directory to contain the certificates needed for TLS encryption.

4. To define the syslog server, type the following at the shell prompt:

```
sudo vi /etc/rsyslog.d/siteconfig.conf
```

5. On a line of its own, add the following entry:

```
facility.priority@ip address of syslog server
```

where:

- *facility* specifies a valid facility value. These categories provide a general description of the originator of the message.
- *priority* specifies a valid priority value. These values specify severity.
- *ip address of syslog server* specifies the IP address of the syslog server that will monitor this SL1 appliance, usually a Data Collector or Message Collector.

NOTE: For details on valid facility and priority values, see https://docs.oracle.com/cd/E37670_01/E36387/html/ol_log_sec.html.

6. Save your changes and exit the file (:wq).
7. At the command line, type the following:

```
sudo service rsyslog restart
```

8. Repeat steps 1-7 on each SL1 appliance in your system.

Defining the NTP Server

By default, SL1 uses the time servers in the Red Hat Linux pool of time servers. If you want to use a different time server, you can edit the configuration files for the time server.

From the **Device Settings** page of the Web Configuration Utility, you can edit the following time server files:

- **chrony.d/servers.conf**. This configuration file contains additional settings for the various chrony time servers.
- **chrony.conf**. This configuration file contains settings related to the time server (chrony.d) used by SL1.

To configure a time server file:

1. *Log in to the Web Configuration Utility*. The **Configuration Utilities** page appears.

2. Click the **[Device Settings]** button. The **Settings** page appears.

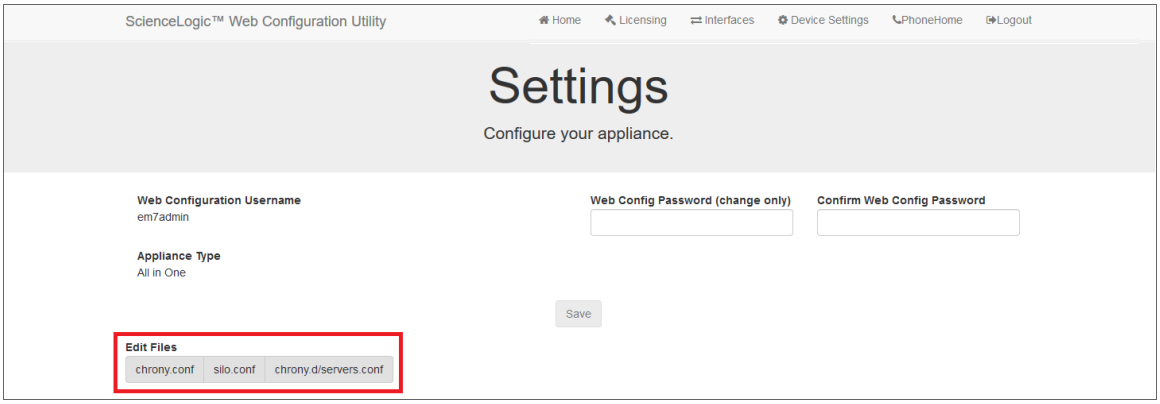

3. In the Edit Files section, click **chrony.d/servers.conf**. The Chrony.d/servers.conf Editor modal page appears:

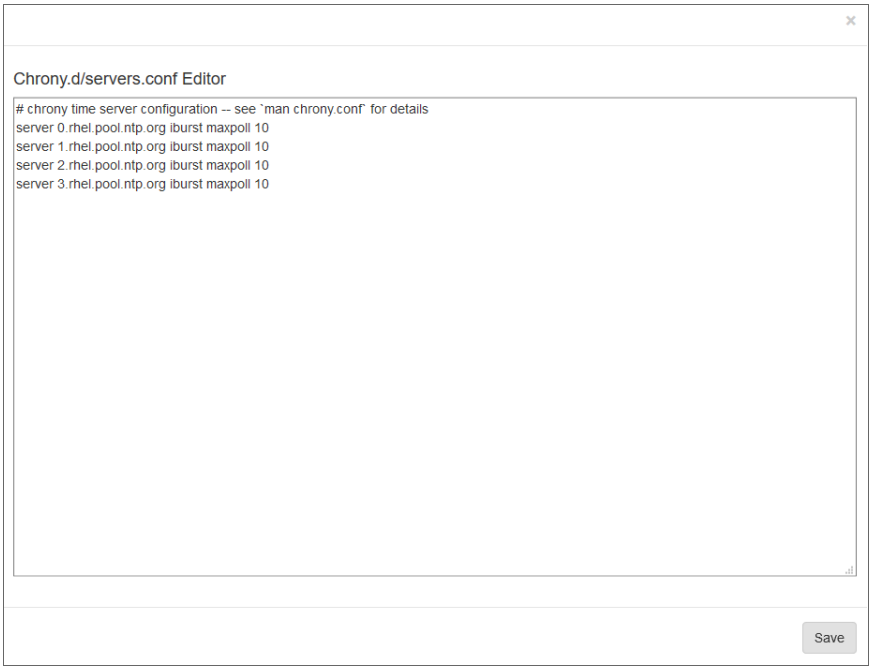

4. In the Chrony.d/servers.conf modal page, copy the first line that begins with **server**, such as **server 0.rhel.pool.ntp.org iburst maxpoll 10**.
5. Paste that line *above* the first line that begins with **server**.
6. Replace the hostname portion of the line with the IP address or fully qualified domain name of your preferred time server.
7. You can delete the additional lines or leave them as additional time servers.
8. To save your changes, click **Save** and then close the modal window.

9. If you need to configure the time server (chrony.d) used by SL1, click **chrony.conf** in the Edit Files section of the Settings page.

Creating a Bonded Interface

A bonded interface (also known as port trunking, channel bonding, link aggregation, and NIC teaming) allows you to combine multiple network interfaces (called "slave interfaces") into a single logical interface (called a "master interface"). A bonded interface can:

- increase available bandwidth
- provide redundancy

To the operating system, a bonded interface appears as a normal network interface. However, the bonded interface uses a round-robin protocol to assign network traffic to the slave interfaces that make up the bonded interface.

To create one or more bonded interfaces:

1. *Log in to the Web Configuration Utility*. The **Configuration Utilities** page appears.

2. Click the **[Interfaces]** button. The **Interfaces** page appears.

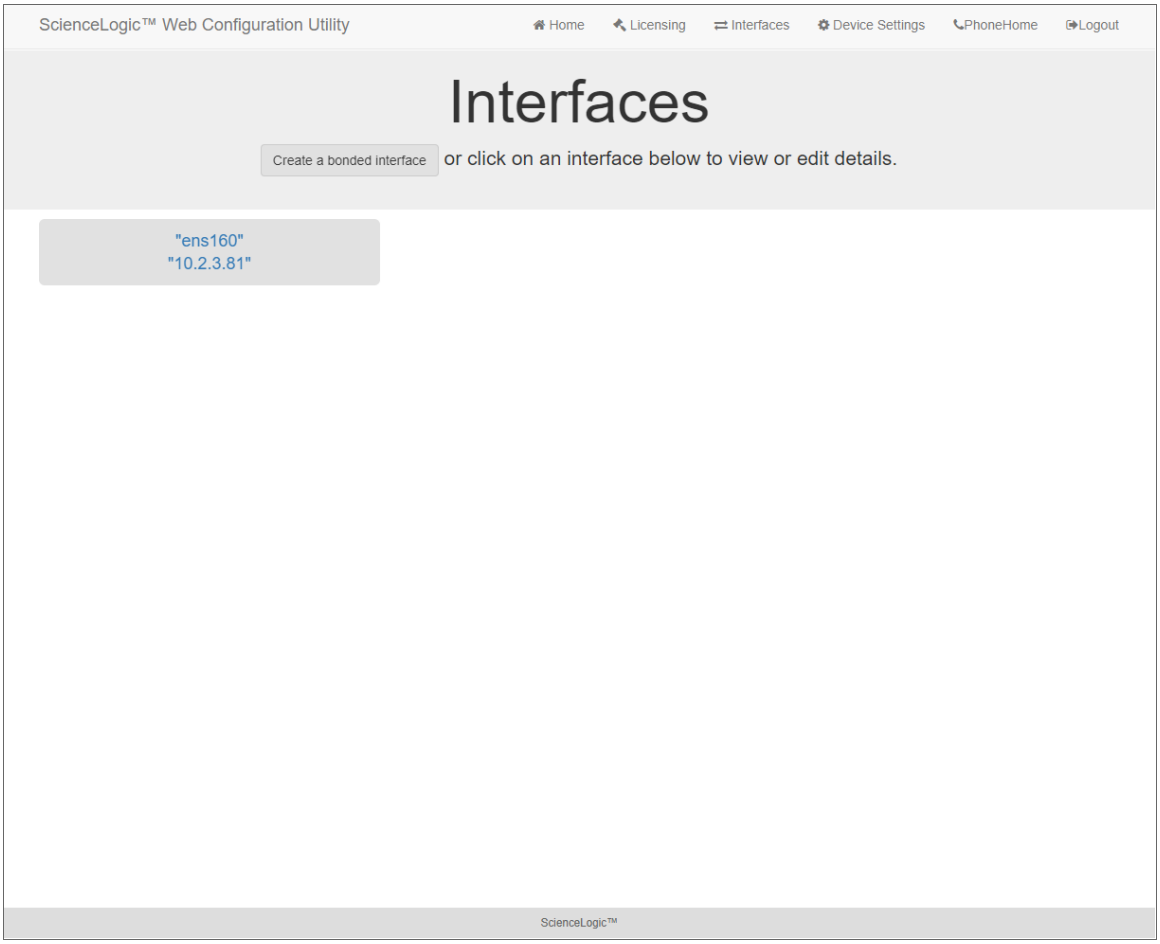

3. In the **Interfaces** page, click the **[Create a bonded interface]** button. The **Create a Bonded Interface** page appears.

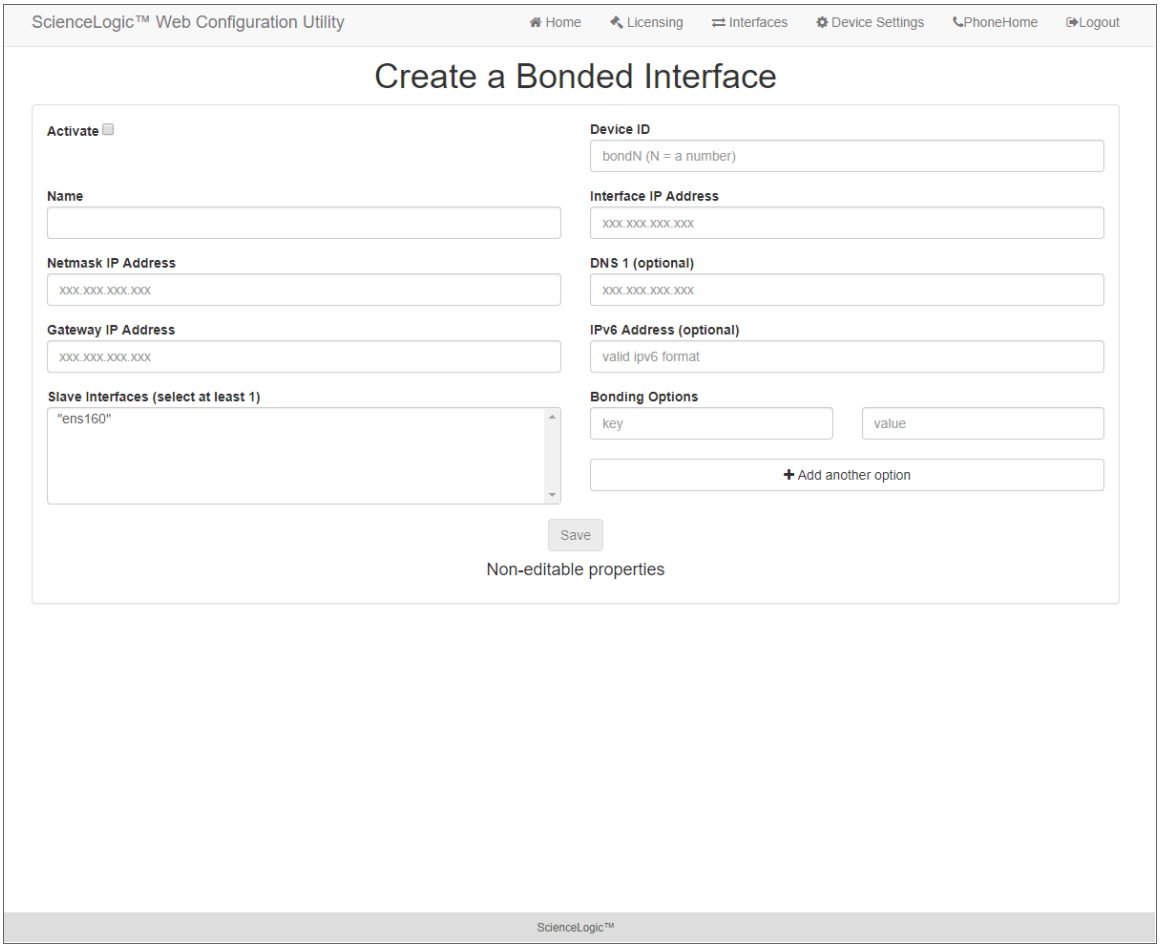

4. In the **Create a Bonded Interface** page, enter the following:

   - ***Device ID***. Required. ID for the bonded interface. Enter a string with the format:

     `bondN`

     where *N* is a number. For example, you could enter ***bond0***, ***bond1***, or ***bond64***.

     If the device ID already exists in the SL1 System, the SL1 system will display an error message.

   - ***Name***. Required. Enter a user name for the bonded interface.
   - ***Interface IP Address***. Required. Enter the IP address for the bonded interface in standard IPv4, dotted-octet format.
   - ***Netmask IP Address***. Required. Enter the netmask for the bonded interface in standard IPv4, dotted-octet format.

- ***Slave Interfaces***. Required. Select one or more interfaces from the list of available interfaces. The selected interfaces will be used by the new bonded interface.
- ***DNS1***. Optional. Enter the IP address of the DNS server that the bonded interface will use. Enter the IP address in standard IPv4, dotted-octet format.
- ***Gateway IP Address***. Optional. Enter the IP address of the gateway device or router that the bonded interface will use. Enter the IP address in standard IPv4, dotted-octet format.
- ***IPv6 Address***. Optional. Enter the IP address for the bonded interface, in IPv6 format.
- ***Bonding Options***. Optional. You can enter one or more bonding options. For each option, enter the name of the option in the *key* field and the value in the *value* field.

For details on bonding options, see the Red Hat documentation on Bonding Interface Parameters: https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Deployment_Guide/sec-Specific_Kernel_Module_Capabilities.html#s3-modules-bonding-directives

Defining a Proxy Server

A proxy server enables SL1 appliances to get system updates when the appliance does not have a direct connection to the internet. A proxy server also enables ScienceLogic Database Servers to send subscription licensing data to ScienceLogic.

Each SL1 appliance can define its own proxy server.

To define a proxy server:

1. Go to the **Appliance Manager** page (System > Settings > Appliances).
2. Find the appliance for which you want to define a proxy server. Click its toolbox icon ( ).
3. When prompted to enter your username and password, log in as the "em7admin" user with either the default password of ***em7admin*** or the password you configured.

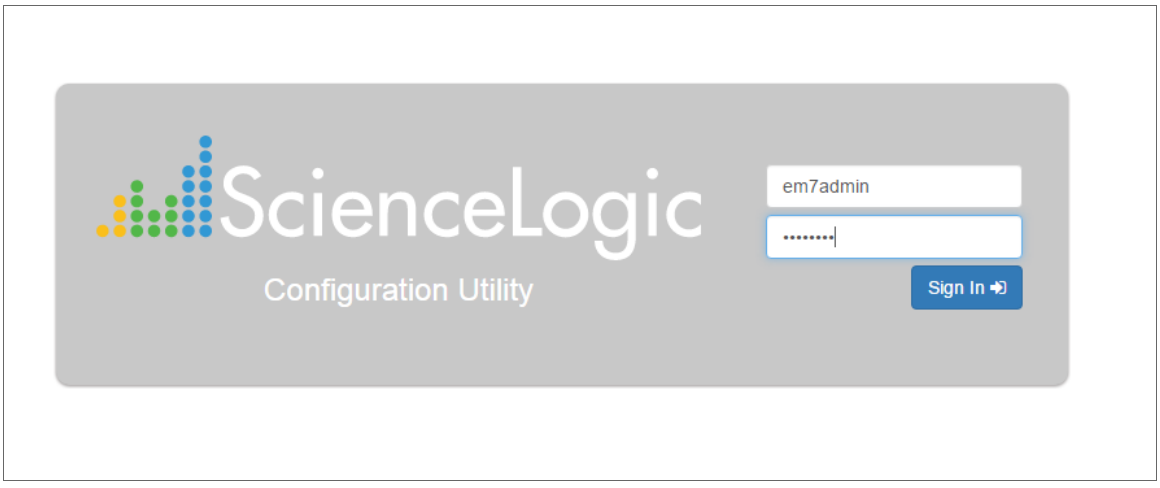

4. After logging in, the main **Configuration Utility** page appears:

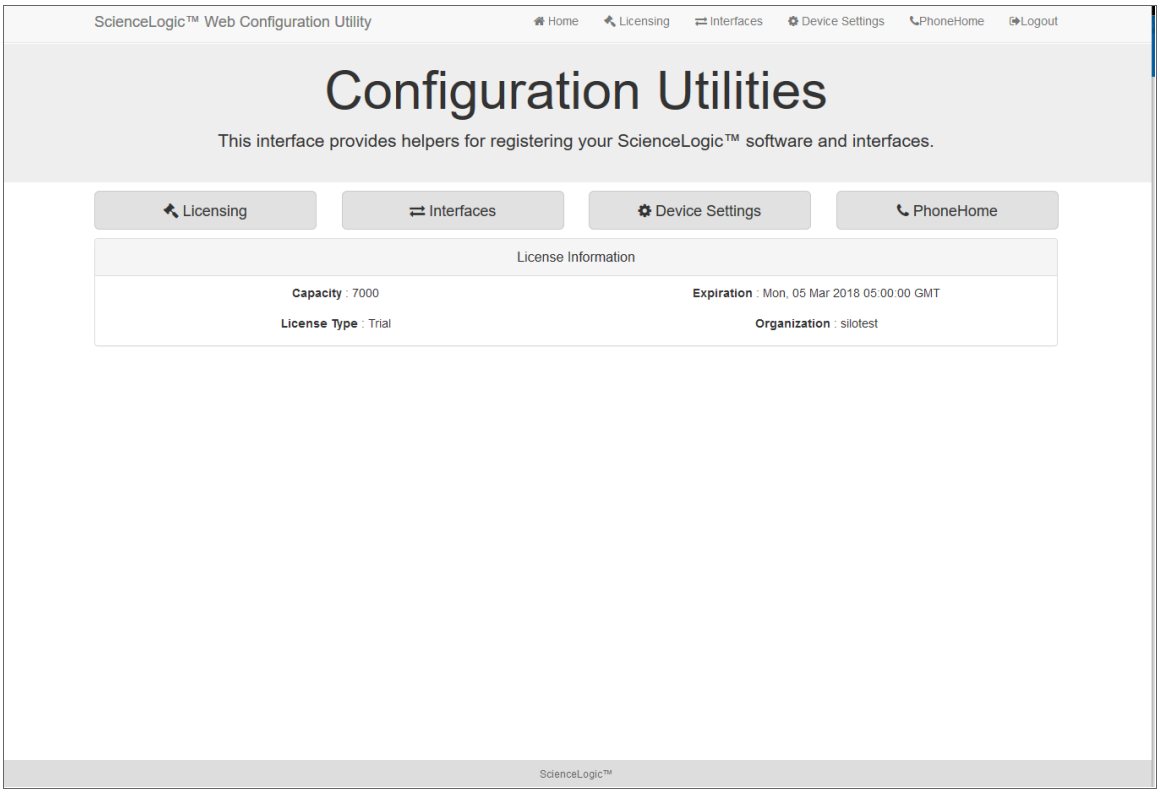

5. Click the **[Device Settings]** button. The **Settings** page appears.

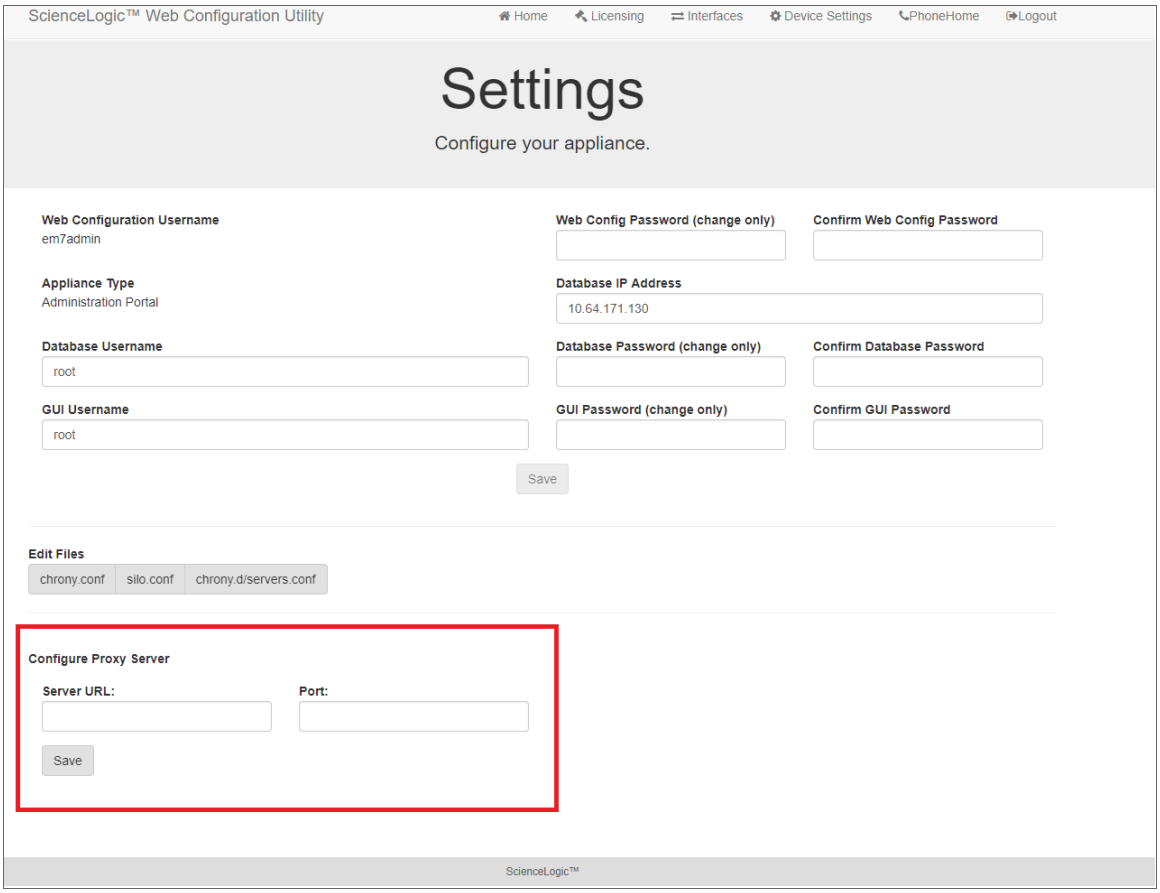

6. Enter values in the following fields:
    - ***Server URL***. Type the URL of the proxy server. For example, "http://10.2.12.51".
    - ***Port***. Type the port on the proxy server to which the SL1 appliance will talk.
7. Click **[Save]**.

Chapter 7

Configuring SL1 for PhoneHome Communication

Overview

This chapter explains how to configure SL1 to use PhoneHome Communication.

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 (•••).

The following topics cover the details of configuring PhoneHome Communication:

What is PhoneHome Communication?

SL1 supports two methods for communication between Database Servers and the Data Collectors and Message Collectors in a system:

The traditional method, where the Database Server initiates communication with each Data Collector and Message Collector. The Database Server periodically pushes configuration data to the Data Collectors and Message Collectors and retrieves data from the Data Collectors and Message Collectors.

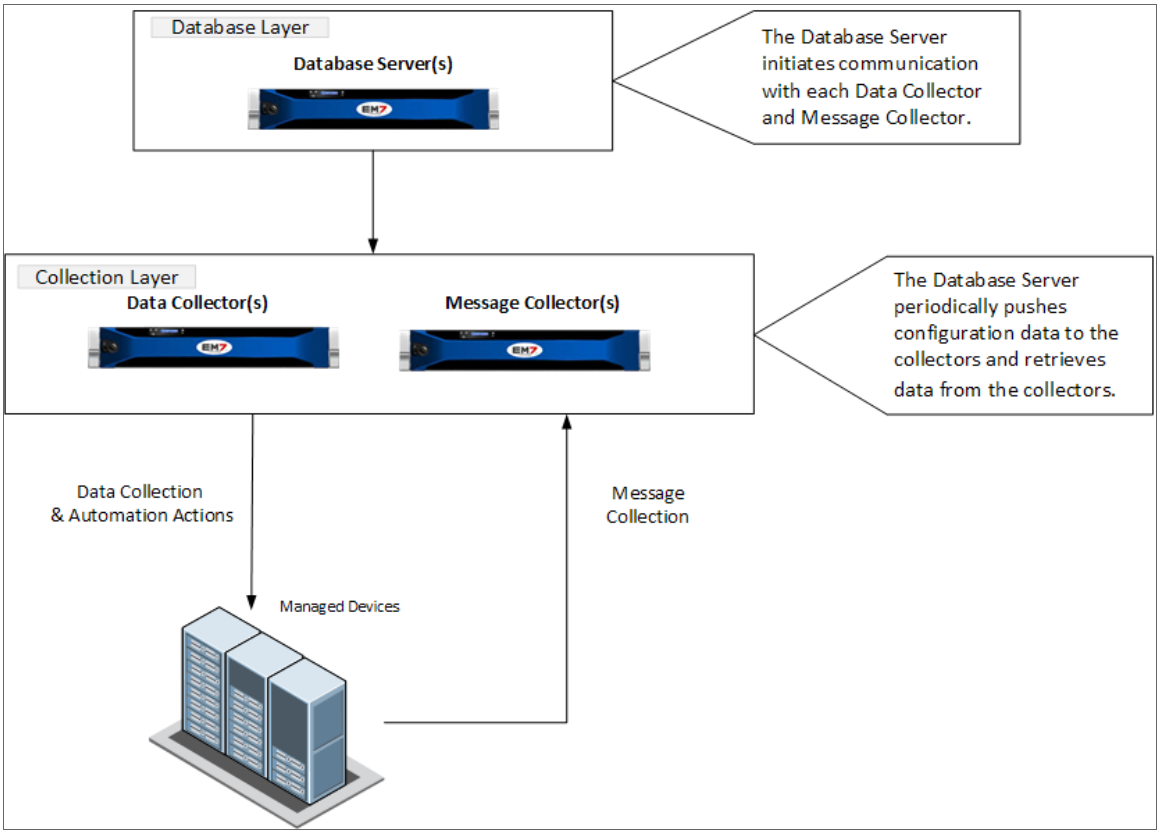

The benefit of this method is that communication to the Database Server is extremely limited, so the Database Server remains as secure as possible.

The PhoneHome method, where the Data Collectors and Message Collectors initiate communication with the Database Server. The Data Collectors and Message Collectors create an SSH tunnel. The Database Server uses the SSH tunnel to periodically push configuration data to the Data Collectors and Message Collectors and retrieve data from the Data Collectors and Message Collectors.

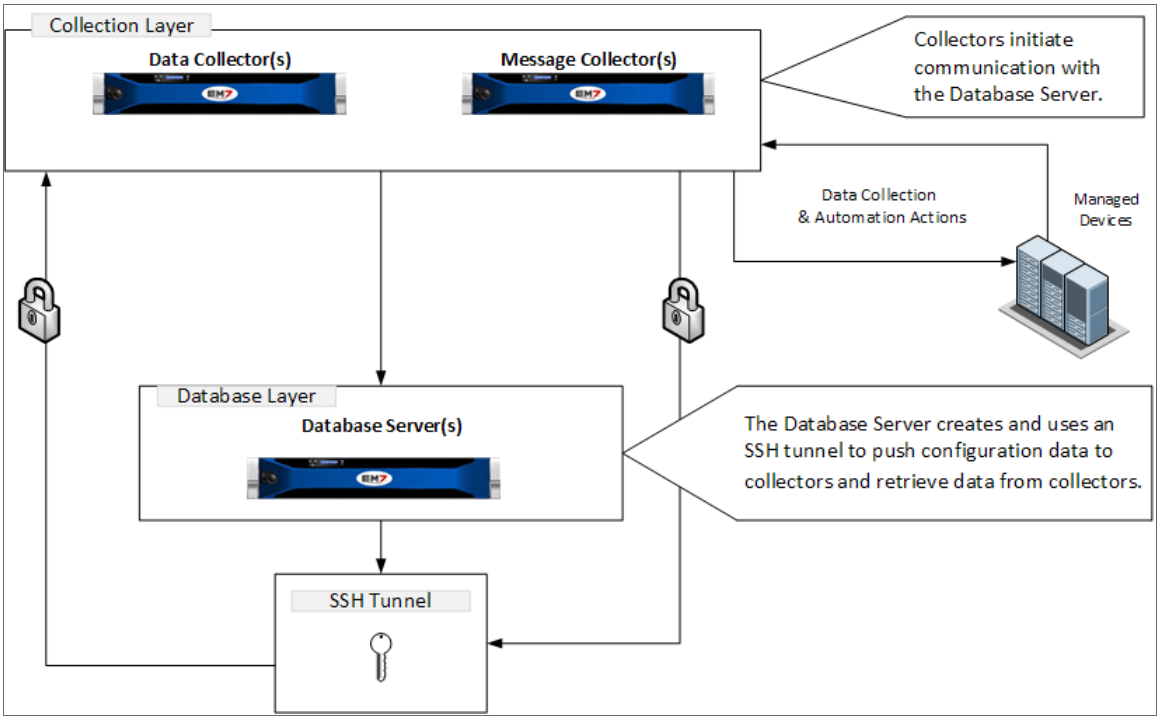

The benefits of this method are that no firewall rules must be added on the network that contains the Data Collectors, and no new TCP ports are opened on the network that contains the Data Collectors.

The PhoneHome configuration uses public key/private key authentication to maintain the security of the Database Server. Each Data Collector is aligned with an SSH account on the Database Server and uses SSH to communicate with the Database Server. Each SSH account on the Database Server is highly restricted, has no login access, and cannot access a shell or execute commands on the Database Server.

Prerequisites

Before configuring PhoneHome communication in your ScienceLogic environment, you must:

- Have SSH access or console access to each database.
- On each ScienceLogic appliance, know the username and password for access to the console as **em7admin**.
- Ensure that the Database Server opens a port for PhoneHome communication. The default port used by the Configuration Utility is **7705**, but you can use other ports besides 7705.

NOTE: If you use a proxy in your PhoneHome configuration, *perform the steps in the proxy section* before configuring the other steps in the PhoneHome configuration. The remaining configuration steps require the proxy for communication.

Overview of the PhoneHome Configuration

For a configuration that includes one or more Database Servers, perform the following steps to use PhoneHome communications:

1. *Configure one or more Database Servers for PhoneHome*. Each Database Server must have SL1 installed, have an IP address, and be licensed with ScienceLogic.

NOTE: If you are using a High Availability and Disaster Recovery configuration, see *Configuring PhoneHome for High Availability and Disaster Recovery* to configure Database Servers.

2. *Configure the Data Collectors and Message Collectors for PhoneHome*. Each Data Collector or Message Collector must have SL1 installed and have an IP address.
3. *Define the Database Server* associated with each Data Collector or Message Collector appliance.
4. *Register the Data Collectors and Message Collectors* in SL1.
5. As needed, *define port forwarding* for each collector to use SSH from the Database Server to access that Data Collector or Message Collector.
6. See the *Troubleshooting section* for additional help.

Configuring the Database Server for PhoneHome

In PhoneHome communication, the Database Server communicates with the Data Collectors and the optional Message Collectors. The Database Server stores all the configuration information for the PhoneHome configuration.

In 8.14.0 or later releases, PhoneHome configuration is stored in tables on the Database Server. The information is accessible to all Database Servers in the SL1 system. Any Database Server in the SL1 system can provide network access.

To configure a Database Server for PhoneHome communication, you must first perform the following required steps:

- Install the SL1 on the Database Server.
- Assign an IP address to the Database Server.
- License the Database Server with ScienceLogic. For more information, see *Licensing and Configuring a Database Server*.

Next, configure the Database Server for PhoneHome communication:

1. Go to the console of the Database Server or use SSH to access the server.
2. Log in as user ***em7admin*** with the password you configured during setup.
3. For the Database Server, you must open a port to accept incoming connection requests. To do this, type the following at the shell prompt:

   ```
   sudo phonehome open-control-port port_number
   ```

   where *port_number* is an unused port number greater than 1000. The Configuration Utility uses port 7705 as the default port. If you want to use the default port, specify **7705** in this command. To use a *different* port, specify that port number in this command.

4. To define the Database Server (to itself), type the following at the shell prompt:

   ```
   sudo phonehome add database
   ```

5. Review the output, which should look like the following:

   ```
   Created local: #0
   Reloading sshd configurations
   Created database: #11
   Changing password for user: "phonehome11"
   Created Device Id: "11"
   Created token: "phonehome://11@71.97.6.197/ee4sdRRK8yNu"
   ```

6. Note the ID number for the database (***11*** in our example).

7. If the database is behind a firewall, you need to define the public-facing IP address of the Database Server and also define the port to use for SSH communication from PhoneHome servers to the Database Server. To do this, type the following at the shell prompt:

   ```
   sudo phonehome set appliance_ID -ip=IP_address -port=port_number
   ```

   where *appliance_ID* is the value you noted in step 6, *IP_address* is the public-facing IP address, and *port_number* is the port you want to use for SSH communication to and from the Database Server.

   For example:

   ```
   sudo phonehome set 11 -ip=71.197.6.197 -port=7705
   ```

8. You must now configure the Data Collectors and Message Collectors in your network. To do this, go the next section.

Configuring the Data Collectors and Message Collectors for PhoneHome

This section describes how to configure a Data Collector and a Message Collector for use in a PhoneHome configuration.

Before configuring a Data Collector and a Message Collector for use in a PhoneHome configuration, you must first perform the following required steps:

- Install SL1 on each Data Collector and Message Collector
- Assign an IP address to each Data Collector and Message Collector

NOTE: If your PhoneHome configuration uses proxy servers, do not use this section to configure a Data Collector or Message Collector. See the *section on proxy servers* instead.

To configure a Data Collector or Message Collector as part of a PhoneHome configuration:

1. On the Data Collector or Message Collector, log in to the Web Configuration Utility using any web browser supported by SL1. The address of the Web Configuration Utility is in the following format:

   `https://ip_address_of_collector:7700`

2. Type the address of the Web Configuration Utility in the Address bar of your browser, replacing "*ip-address-of-appliance*" with the IP address of the Data Collector or Message Collector.
3. When prompted to enter your user name and password, log in as the "em7admin" user with either the default password of ***em7admin*** or the password you configured.

4. After you log in, the **Configuration Utilities** page appears:

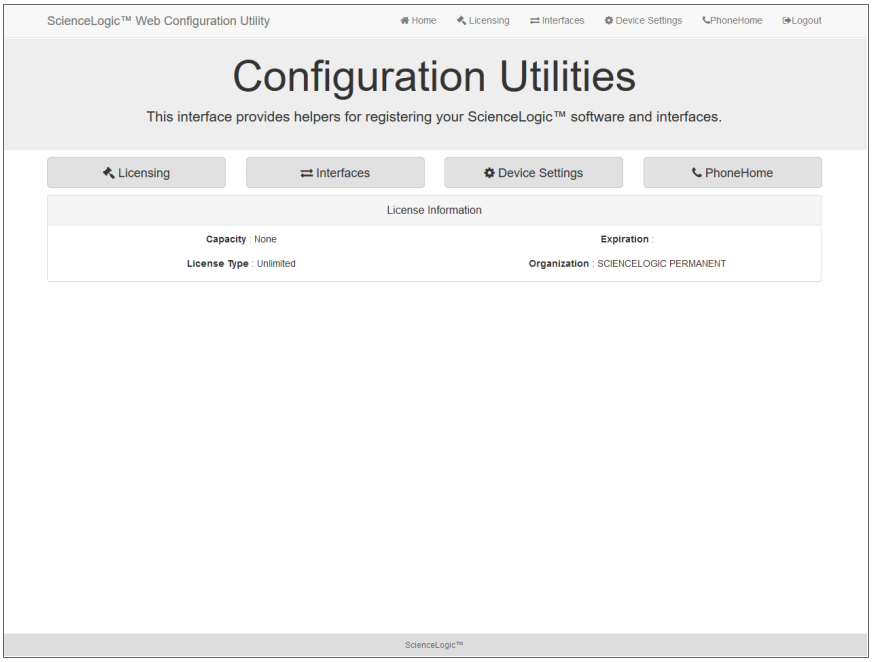

5. Click the **[PhoneHome]** button. The **PhoneHome - Collector** page appears.

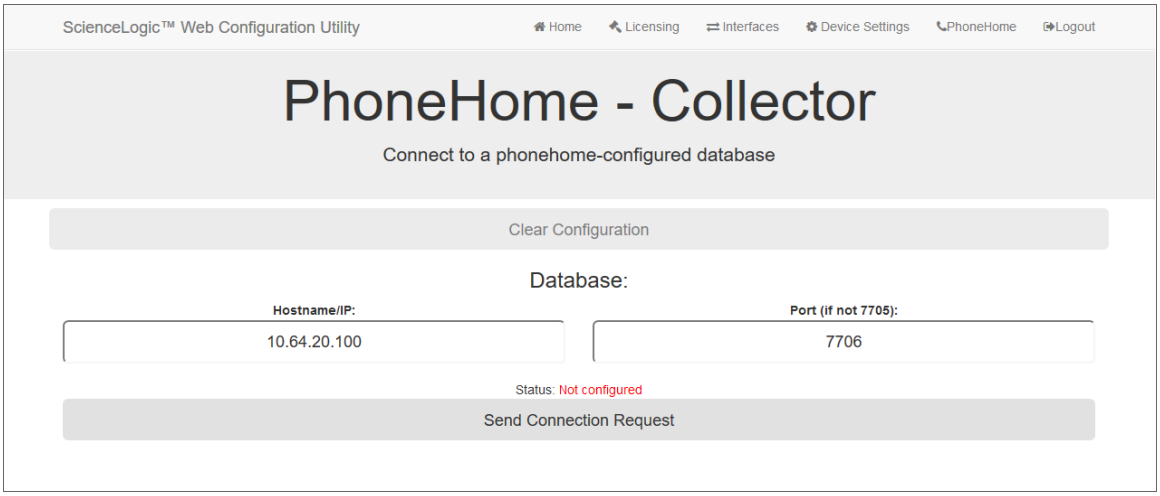

6. Complete the following fields:

    - ***Hostname/IP***. Type the Hostname or IP address of the Database Server that is configured for PhoneHome.
    - ***Port (if not 7705)***. Optional. Port number for SSH communications with the Database Server that is configured for PhoneHome. If you are using a port *other* than 7705 on the Database Server, type the port number in this field. Otherwise, leave this field blank.

NOTE: The option ***Make request with the Web Configuration Utility via HTTPS*** was removed in 8.14.0. In versions of SL1 of 8.14.0 or later, this option does not appear in the Web Configuration Utility page for PhoneHome.

7. Click the **[Send Connection Request]** button to send the request from the Data Collector to the Database Server that is configured for PhoneHome. After clicking the **[Send Connection Request]** button, the **PhoneHome - Collector** page displays the status *Pending database approval*.

NOTE: Do not close the browser window or navigate away from this page while the connection request is being processed.

8. In a new browser window, open the ScienceLogic Web Configuration Utility for the Database Server. To do this, type the following, replacing "*ip-address-of-database*" with the IP address of the Database Server:

    `https://ip_address_of_database:7700`

9. When prompted to enter your username and password, log in as the "em7admin" user with either the default password of ***em7admin*** or the password you configured. The main **Configuration Utility** page appears.

10. Click the **[PhoneHome]** button. The **PhoneHome Database - Master** page appears.

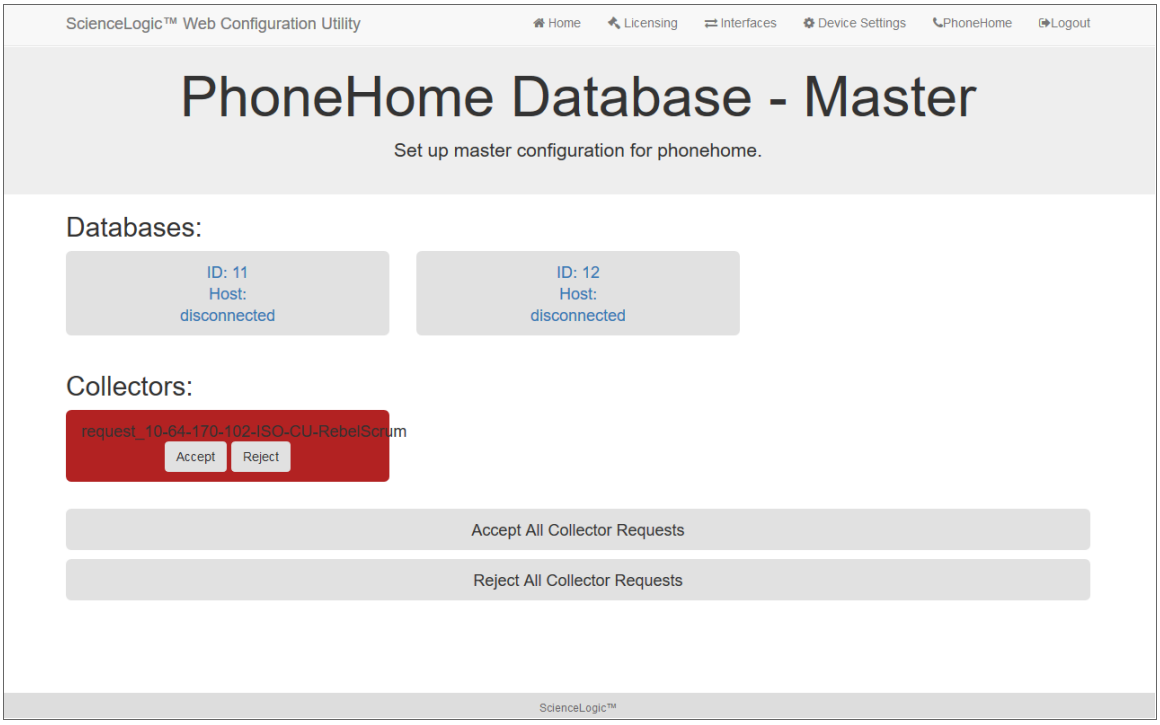

11. Note that the list of Collectors includes a request. Click the **[Accept]** button for that collector. The Status for the Collector now displays as *Approved*.
12. On the Data Collector or Message Collector, open the ScienceLogic Web Configuration Utility and click the **[PhoneHome]** button. The **PhoneHome - Collector** page appears:

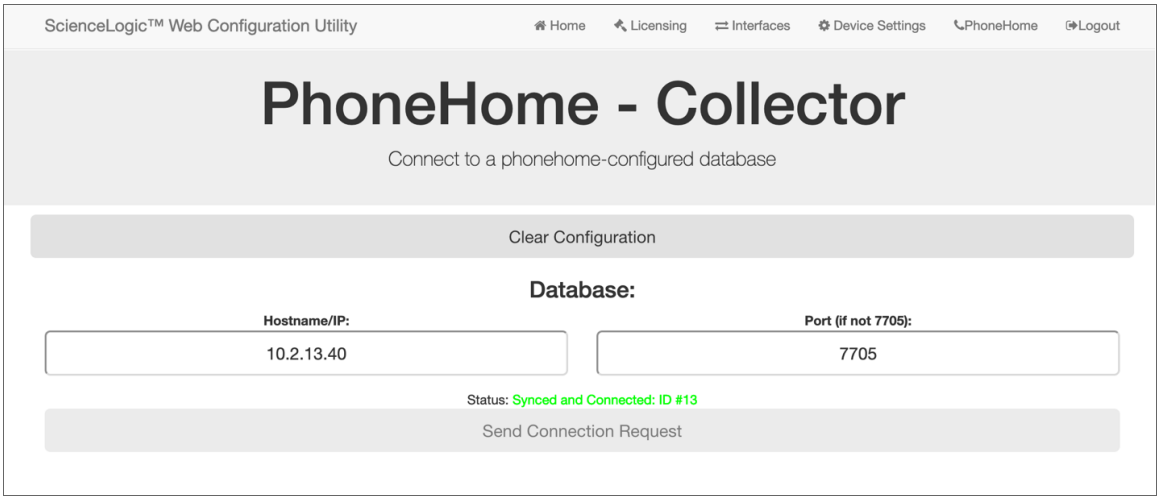

13. Click the **[Check Approval]** button. Note that the Status message is now *Configured - ID [phonehome_user_number]*.
14. If you refresh the page, the status field displays the message *Synced and Connected*.

If you have a large number of collectors, you can perform the following steps to approve multiple collectors at the same time:

1. On each Data Collector or Message Collector, follow steps 1-7 in the *previous procedure* to send the connection request for each collector.
2. Open the ScienceLogic Web Configuration Utility for the Database Server and click the **[PhoneHome]**button.
3. Click the **[Accept All Collector Requests]** button.
4. Open the ScienceLogic Web Configuration Utility for each collector, click the **[PhoneHome]**button, and then click the **[Check Approval]** button.
5. Repeat step 4 until you have approved all of your collectors.

Registering the Data Collectors and Message Collectors

> **NOTE**: Perform the steps in this section after you have successfully established a PhoneHome connection between Data Collectors or Message Collectors and the main Database Server. The steps in this section ensure that the SL1 system uses the loopback address that is assigned to each Data Collector and Message Collector upon *successful completion of a phonehome connection request*.

In this section:

- In the Web Configuration Utility, you must define the Database Server associated with each Data Collector or Message Collector appliance.
- In SL1, you must register the Data Collectors or Message Collectors with the main Database Server.

Define the Database Server for Each Collector

You must perform the following steps in the Web Configuration Utility to configure a Data Collector or a Message Collector:

1. Log in to the Web Configuration Utility on the Data Collector or the Message Collector. The **Configuration Utilities** page appears.

2. Click the **[Device Settings]** button. The **Settings** page appears.

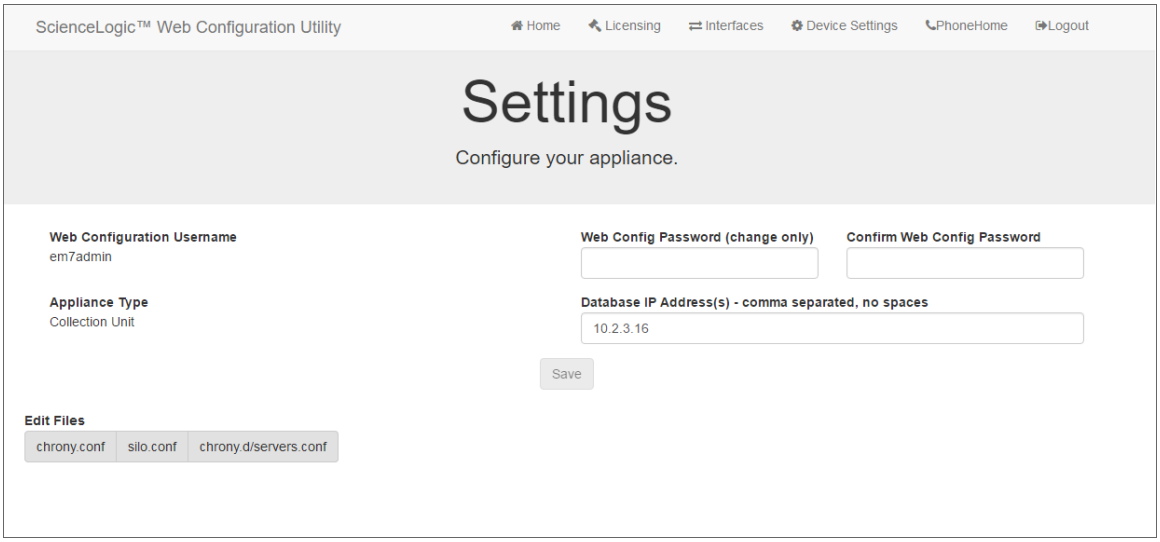

3. On the **Settings** page, update the following field:
    - ***Database IP Address***. The IP address of the ScienceLogic Database Server(s). If more than one Database Server will manage this appliance, type the IP addresses of the Database Servers, separated by commas. You *cannot* paste an IP address into this field.
4. Click the **[Save]** button.
5. Click the **[Logout]** button to log out of the Web Configuration Utility for this collector.
6. Perform steps 1-5 for each Data Collector and Message Collector in your PhoneHome configuration.

Register the Collectors with the Main Database Server

In SL1, to register a Data Collector or Message Collector with the main Database Server:

1. In the address bar of your browser, type the IP address of the SL1 appliance that provides the user interface for your SL1 system. The user interface is provided by either the Database Server or an Administration Portal. The login page appears:

2. Log in as the "em7admin" user with the password "em7admin" (or the password you configured).
3. If this is your first successful login, you will be asked to agree to the End-user License Agreement. Read the End-user License Agreement and then click the **[I Agree to The Terms Outlined Above]** button.
4. Go to the **Appliance Manager** page (System > Settings > Appliances):

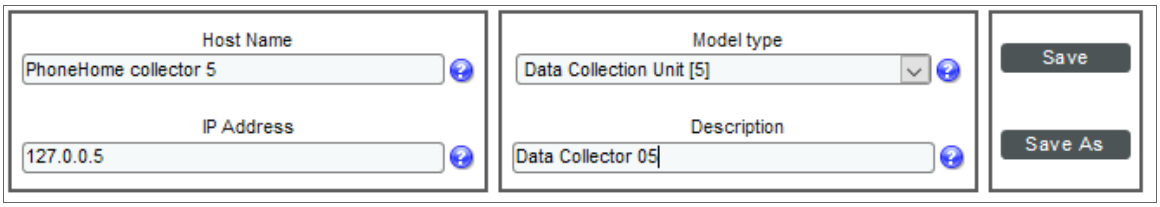

5. Complete the following fields:
   - ***Host Name***. Type the host name of the Data Collector or Message Collector.
   - ***IP Address***. Type the **loopback IP address** of the Data Collector or Message Collector. The loopback IP address is a special, virtual network interface that your computer uses to communicate with itself. This address also allows you to view content on a server in the same way a client would. In most cases, the loopback address is 127.0.0.1, although the loopback address can be any IP address in the 127.0.0.0/8 block.

TIP: If you go to the Web Configuration Utility of the Database Server and click the **PhoneHome** button, you can view a list of all the connected collectors, along with their IDs. The ID indicates the loopback address. For example, if the ID of a given collector is 12, then its loopback address is 127.0.0.12.

   - ***Model Type***. Select the type of appliance (Data Collector or Message Collector) you are registering.
   - ***Description***. Type a description for the Data Collector or Message Collector. This field is optional.
6. Click the **[Save]** button. If the save is successful, the message "Appliance Registered" displays.
7. If all information is valid and the Database Server can communicate with the Data Collector or Message Collector, the **Appliance Manager** page displays the ScienceLogic version installed on the collector in the **Build** column. If the **Build** column remains blank for longer than five minutes, double-check your settings and network connection.
8. Perform these steps for each Data Collector and Message Collector in your PhoneHome configuration.

Defining Port Forwarding

If you define port forwarding for each Data Collector or Message Collector in the PhoneHome configuration, you can use SSH from the Database Server to access the Data Collector or Message Collector.

To define port forwarding:

1. Either go to the console of the Database Server or use SSH to access the server.

2. Log in as user ***em7admin*** with the password you configured during setup.
3. For each Data Collector and/or Message Collector, type the following at the shell prompt:

   ```
   sudo phonehome set appliance_ID -forwards=port_number
   ```

   where:

   - *appliance_ID* is device ID for the Data Collector or Message Collector.
   - *port_number* is the port you want to use for SSH communication from the Database Server to the Data Collector or Message Collector.

   For example:

   ```
   sudo phonehome set 12 -forwards=22
   ```

4. For every other server in the PhoneHome configuration, go to the console of the Database Server or use SSH to access the server.
5. Log in as user ***em7admin*** with the password you configured during setup.
6. Type the following at the shell prompt:

   ```
   sudo phonehome sync
   ```

7. Now, whenever you are SSHed in to the Database Server, you can SSH to the Data Collector or Message Collector.
8. To use the forward port, append "100" to the front of the port you defined in step #3 and use the loopback IP of the Data Collector or Message Collector using port 10022:

   ```
   ssh -p 10022 root@127.0.0.12
   ```

Using Custom Options for AutoSSH

By default, SL1 stores settings for AutoSSH for PhoneHome configurations in the file /opt/em7/backend/phonehome/em7_ph_tunnels on each Data Collector and Message Collector in your configuration.

If you want to use custom AutoSSH settings for a specific Data Collector or Message Collector in your configuration, you can create the file /etc/phonehome/local.conf and define custom AutoSSH options for that server.

To define custom SSH options for a server:

1. Log in to the console of the Data Collector or Message Collector as the root user.
2. Open the file /etc/phonehome/local.confg with a text editor like vi:

   ```
   vi /etc/phonehome/local.conf
   ```

3. Add one or more custom settings for AutoSSH. You can define:

- TCPKeepAlive = "*yes or no*". Specifies whether the client will send a null packet to the server (to keep the connection alive). Uses the TCP layer to send the packet. The default value is "no". If you set this value to zero (0), this feature is disabled. Your connection will drop if idle for too long.
- ServerAliveInterval = "*number of seconds*". The number of seconds the client will wait before sending a null packet to the server (to keep the connection alive). Uses the SSH layer to send the packet. The default value is "10". If you set this value to zero (0), this feature is disabled.
- StrictHostKeyChecking = "*yes or no*". If this flag is set to "yes", AutoSSH will never automatically add host keys to AutoSSH configuration and will refuse to connect to hosts whose host key has changed. This option forces the user to manually add all new hosts. If this flag is set to "no", ssh will automatically add new host keys to the known hosts files. The default value is "no".
- ServerAliveCountMax = "*number of messages*". The maximum number of unacknowledged null packets the client will send to the server (to keep the connection alive). After the maximum number of unacknowledged null packets, the client will drop the SSH connection to the server. The default value is "2". If you set this value to zero (0), this feature is disabled. Your connection will drop if idle for too long.
- CUSTOM_PARAMS = "-o *parameter = argument*". Any additional SSH parameters can be configured with this option. For example:

```
CUSTOM_PARAMS="-o ExitOnForwardFailure=yes"
```

NOTE: To determine the format for entries in the /etc/phonehome/local.confg file, see the file /opt/em7/backend/phonehome/em7_ph_tunnels.

4. Save your changes and exit the file.

Configuring PhoneHome for High Availability and Disaster Recovery

This section describes how to configure the Database Servers in your system for use in a PhoneHome configuration.

NOTE: You can use the same Database Servers in both a PhoneHome configuration and a traditional configuration.

After performing the steps in this section, go the section on *Configuring the Data Collectors and Message Collectors* to complete the configuration.

Configuring the Primary Database Server for High Availability and Disaster Recovery

To configure the primary Database Server for PhoneHome communication:

1. Go to the console of the primary Database Server or use SSH to access the server.
2. Log in as user ***em7admin*** with the password you configured during setup.
3. For the primary Database Server, you must first open a port to accept incoming connection requests. To do this, type the following at the shell prompt:

   ```
   sudo phonehome open-control-port port_number
   ```

   where *port_number* is an unused port number greater than 1000. The default value in the Configuration Utility is 7705. If you want to use the default port later in the Configuration Utility, specify "7705" in this command.

4. To define the primary Database Server (to itself), type the following at the shell prompt:

   ```
   sudo phonehome add database
   ```

5. Review the output, which should look like the following:

   ```
   Created local: #0
   Reloading sshd configurations
   Created database: #11
   Changing password for user: "phonehome11"
   Created Device Id: "11"
   Created token: "phonehome://11@71.97.6.197/ee4sdRRK8yNu"
   ```

6. Note the ID number for the primary Database Server (***11*** in our example).
7. To define the public-facing IP address of the primary Database Server and the port to use for SSH communication from PhoneHome servers to the primary Database Server, type the following at the shell prompt:

   ```
   sudo phonehome set appliance_ID -ip=IP_address -port=port_number
   ```

   where *port_number* is an unused port number greater than 1000. The Configuration Utility uses port 7705 as the default port. If you want to use the default port, specify **7705** in this command. To use a *different* port, specify that port number in this command.

   For example:

   ```
   sudo phonehome set 11 -ip=71.197.6.197 -port=7705
   ```

8. Start the PhoneHome watchdog. To do so, type the following at the shell prompt:

   ```
   sudo systemctl enable em7_ph_watchdog
   sudo systemctl start em7_ph_watchdog
   ```

9. If your SL1 System uses multiple databases for high availability and/or disaster recovery, you must create a record for the secondary Database Server on the primary Database Server. To do so, type the following at the shell prompt:

   ```
   sudo phonehome add database
   ```

10. The output will look like this:

```
Reloading sshd configurations
Created database: #13
Changing password for user: "phonehome13"
Created Device Id: "13"
Created token: "phonehome://13@10.64.68.31:22/GmHtYvDd9O0V"
```

11. Note the ID number for the secondary Database Server. You will need this value later in the configuration.
12. Copy and save the token for the secondary Database Server. You will need this value later in the configuration.
13. To define the public-facing IP address of the secondary Database Server and the port to use for SSH communications from PhoneHome servers to the secondary Database Server, type the following at the shell prompt:

```
sudo phonehome set appliance_ID -ip=IP_address -port=port_number
```

where:

- *appliance_ID* is the value you noted in step 5.
- *IP_address* is the public-facing IP address.
- *port_number* is the port you want to use for SSH communication to and from the primary Database Server.

For example, we could enter:

```
sudo phonehome set 13 -ip=71.197.6.198 -port=7705
```

Configuring the Secondary Database Server for High Availability and Disaster Recovery

To configure a secondary Database Server as part of a PhoneHome configuration:

1. Either go to the console of the secondary Database Server or use SSH to access the server.
2. Log in as user ***em7admin*** with the password you configured during setup.
3. For the secondary Database Server, you must first open a port to accept incoming connection requests. To do this, type the following at the shell prompt:

```
sudo phonehome open-control-port port_number
```

where:

- *port_number* is an unused port number greater than 1000. The default value in the Configuration Utility is 7705. If you want to use the default value, specify "7705".

NOTE: ScienceLogic recommends that you use the same port number on each database in your PhoneHome configuration.

4. To register the secondary Database Server, type the following at the shell prompt:

```
sudo phonehome register appliance_token
```

where:

- *appliance_token* is the URL you saved during step 11 in the section ***Configuring the Primary Database for High Availability and Disaster Recovery***.

5. The output looks like this:

```
Registered device successfully
```

6. Type the following at the shell prompt:

```
sudo phonehome sync
```

7. The output looks like this:

```
Started synchronization
Synchronized: collectors
Synchronized: databases
Reloading sshd configurations
Finished synchronizations
```

Configuring Data Collectors and Message Collectors for High Availability and Disaster Recovery

You must now configure the Data Collectors and Message Collectors in your network. To do this, see ***Configuring the Data Collectors and Message Collectors***.

NOTE: If your PhoneHome configuration uses proxy servers, do not use this section to configure a Data Collector or Message Collector. See the ***section on proxy servers*** instead.

Syncing the High Availability and Disaster Recovery System

After adding Data Collector(s) or Message Collector(s) to your PhoneHome configuration, you must once again execute the sync command on all ***Database Servers*** and then on the newly configured Collectors in the PhoneHome configuration.

To sync your PhoneHome configuration:

1. Either go to the console of the Database Server (or the new Collectors) or use SSH to access the server.
2. Log in as user ***em7admin*** with the password you configured during setup.
3. At the shell prompt, type the following:

```
sudo phonehome sync
```

4. Perform these steps on each Database Server, Data Collector, and Message Collector in your PhoneHome configuration.

Adding a New Secondary Database Server

To add a new secondary Database Server to an existing PhoneHome configuration:

1. On the primary Database Server, perform the steps from the section *Configuring the Primary Database Server*. These are the steps that define the secondary Database Server, including saving the token and saving the new configuration.
2. On the new secondary Database Server, perform the steps from the section *Configuring the Secondary Database Server for High Availability and Disaster Recovery*.
2. Either go to the console of the SL1 appliance or use SSH to access the new secondary Database Server. Log in as "root".
4. At the shell prompt, type the following:

```
phonehome status
```

5. The new secondary Database Server should be connected to each Data Collector in the PhoneHome configuration.

Configuring One or More Proxy Servers

If your organization requires that you use a proxy for outbound requests, you can configure one or more Data Collectors to act as proxy servers. These proxy servers will sit between the Data Collectors in your PhoneHome configuration and the Database Server in your PhoneHome configuration.

To use one or more Data Collectors as proxy servers in a PhoneHome configuration:

- Ensure that the SSH port is open on each Data Collector that acts as a proxy server.
- Ensure that the SSH port is open on each Database Server in the PhoneHome configuration.

> **NOTE**: If you use a proxy in your PhoneHome configuration, ***perform the steps in this section before you configure the other steps in this chapter***. The other steps in the PhoneHome configuration will require the proxy for communication.

To configure your PhoneHome configuration to use a proxy server, you must:

1. Configure a Database Server for PhoneHome configuration as either a stand-alone Database Server (*Configuring the Database Servers*) or a High Availability Database Server (*Configuring the Database Servers for High Availability*)
2. *Edit the ssh_config file*.
3. Use the command line to configure Data Collectors *that connect via proxy.*
4. *Copy the SSH key to each proxy*.
5. *Synchronize the* Data Collectors with the Database Server.

Editing ssh_config

1. Either go to the console of the Data Collector that will be part of the PhoneHome configuration or use SSH to access the server.

NOTE: Perform these steps on the Data Collector that will be part of the PhoneHome configuration, ***not on the Data Collector that will serve as a proxy server***.

2. Log in as user ***em7admin*** with the password you configured during setup.
3. Open the file /etc/ssh/ssh_config with vi or another text editor:

```
sudo vi /etc/ssh/ssh_config
```

4. Add the following lines to the file:

```
Host hostname_of_primary_Database_Server

ProxyCommand ssh -q em7admin@proxy_hostname nc %h %p
```

where:

hostname_of_primary_Database_Server is the hostname for the primary Database Server.

proxy_hostname is the hostname of the proxy server that directly communicates with the Database Server. If you have a chain of proxy servers, this value is the hostname of the last proxy server in that chain (the proxy server that connects to the Database Server).

NOTE: If you use hostnames to configure proxy servers, you must use hostnames for all PhoneHome configuration. If you use IP addresses to configure proxy servers, you must use IP addresses for all PhoneHome configuration. You cannot mix hostnames and IP addresses in ssh_config and in PhoneHome configuration.

5. If applicable, for all secondary databases, add the following lines to the file:

```
Host hostname_of_secondary_Database_Server
ProxyCommand ssh -q em7admin@proxy_hostname nc %h %p
```

where:

hostname_of_secondary_Database_Server is the hostname for the secondary Database Server.

proxy_hostname is the hostname of the proxy server that directly communicates with the secondary Database Server. If you have a chain of proxy servers, this value is the hostname of the last proxy server in that chain (the proxy server that connects to the Database Server).

6. If you have more than one proxy server, add the following lines to the file:

```
Host hostname_of_proxy_server
ProxyCommand ssh -q em7admin@proxy_hostname nc %h %p
```

where:

hostname_of_proxy_server is the hostname of the current proxy server (the proxy server you are creating an entry for). For example, you could create an entry for "ProxyServer2".

proxy_hostname is the hostname of the proxy server that is previous in the chain and communicates with the current proxy server. If your entry is for "ProxyServer2", you would specify "ProxyServer1" in this field.

For example, if you had the following configuration:

PhoneHome Data Collector -> ProxyServer1 -> ProxyServer2 -> ProxyServer3 -> Database Server

You would add the following entries to /etc/ssh/ssh_config:

```
Host ProxyServer2
ProxyCommand ssh -q em7admin@ProxyServer1 nc %h %p

Host ProxyServer3
ProxyCommand ssh -q em7admin@ProxyServer2 nc %h %p

Host EM7_DB1
ProxyCommand ssh -q em7admin@ProxyServer3 nc %h %p
```

For another example, if you had the following configuration:

PhoneHome Data Collector -> ProxyServer1 -> ProxyServer2 -> ProxyServer3 -> primary Database Server

PhoneHome Data Collector -> ProxyServer1 -> ProxyServer2 -> ProxyServer3 -> secondary Database Server

You would add the following entries to /etc/ssh/ssh_config:

```
Host ProxyServer2
ProxyCommand ssh -q em7admin@ProxyServer1 nc %h %p

Host ProxyServer3
ProxyCommand ssh -q em7admin@ProxyServer2 nc %h %p

Host EM7_DB1
ProxyCommand ssh -q em7admin@ProxyServer3 nc %h %p

Host EM7_DB2
ProxyCommand ssh -q em7admin@ProxyServer3 nc %h %p
```

7. Save (:wq) your changes to the /etc/ssh/ssh_config file.

Configuring Data Collectors that Connect to the PhoneHome Database Server Through a Proxy

To configure a Data Collector that connect to the PhoneHome Database Server through a Proxy:

1. Either go to the console of the Database Server or use SSH to access the server.
2. Log in as user ***em7admin*** with the password you configured during setup.
3. For the Database Server, you must first open a port to accept incoming connection requests. To do this, type the following at the shell prompt:

   ```
   sudo phonehome open-control-port port_number
   ```

   where:

   - *port_number* is an unused port number greater than 1000.

4. To define the Data Collector (to the Database Server), type the following at the shell prompt:

   ```
   sudo phonehome add collector
   ```

5. The output will look like this:

   ```
   Created collector: #12
   Changing password for user: "phonehome12"
   Created Device Id: "12"
   Created token: "phonehome://12@10.64.68.31:22/om3Odt3iPEAD
   ```

6. Note the token for the Data Collector.
7. Either go to the console of the Data Collector or use SSH to access the server.

> **NOTE**: Perform these steps on the Data Collector that will be part of the PhoneHome configuration, ***not on the Data Collector that will serve as a proxy server***.

8. Register the Data Collector with the Database Server . To do this, type the following at the shell prompt:

   ```
   sudo phonehome register token
   ```

   where:

   - *token* is the value you noted in step 5.

Copying the SSH key to Each Proxy

You must now copy the SSH key to each proxy server. To do this:

1. Either go to the console of the Data Collector that will be part of the PhoneHome configuration (not the proxy server) or use SSH to access the server.

2. Log in as user ***em7admin*** with the password you configured during setup.
3. At the shell prompt, type the following:

```
ssh-copy-id -i /home/phonehome0/.ssh/id_rsa.pub em7admin@[IP_address_of_proxy_server]
```

4. Perform step 3 for each proxy server in your PhoneHome configuration.

Synchronize the Data Collectors with the Database Server

After adding Data Collector(s) or Message Collector(s) to your PhoneHome configuration, you must execute the sync command on all ***Database Servers*** in the PhoneHome configuration.

To execute the sync command on all Database Servers:

1. Either go to the console of the Database Server or use SSH to access the server.
2. Log in as user ***em7admin*** with the password you configured during setup.
3. At the shell prompt, type the following:

```
sudo phonehome sync
```

4. Perform these steps on each Database Server in your PhoneHome configuration.

The Watchdog Service

Each Database Server, Data Collector and Message Collector in a PhoneHome configuration runs a service called watchdog. The watchdog process automatically checks the connection between the Data Collector or Message Collector and the Database Server. If the connection is stale, the watchdog service automatically forces the Database Server to reconnect to the Data Collector or Message Collector.

The watchdog service can also detect configuration changes. If the PhoneHome configuration changes, the watchdog service will detect the changes and synchronize the configuration data on each device in the PhoneHome configuration.

The watchdog service is started automatically on each Data Collector, Message Collector, and secondary Database Server.

To view information about the watchdog service:

1. Log in to the console of the Data Collector, Message Collector, and secondary Database Server as the root user.
2. At the command line, type the following:

```
phonehome watchdog view
```

3. You should see something like the following:

```
Current settings:

autosync: yes
interval: 20
```

```
state: enabled
autoreconnect: yes
timeoutcount: 2
check: default
```

4. You can change any of these settings by typing the following at the command line:

   ```
   phonehome watchdog set -settingvalue
   ```

   where *setting* is one of the settings displayed with the view command and *value* is the value to assign to that setting.

5. For details about the watchdog service, type the following at the command line:

   ```
   phonehome watchdog help
   ```

For details about the arguments and settings for watchdog, see the section on *Using the Command Line Interface*.

Logging PhoneHome Configuration Information

The PhoneHome configuration logs information to the following files:

- ***/var/log/phonehome/phonehome0.log***. Resides on each device in the PhoneHome configuration. This log file stores the date and time that devices are added to or removed from the PhoneHome configuration and each configuration action, including token generation, device registration, and configuration data synchronization, performed for each device. This log is rotated daily.
- ***/var/log/phonehome/phonehome<device ID>.log***. Resides on the Database Server. This log file stores an entry for each action requested by or performed on a specific device (specified by device ID). This log is rotated daily.
- ***/home/$USER/logs/shell.log***. Resides in the home directory of each PhoneHome Database Server and each PhoneHome Data Collector. This log file stores information about registration operations and periodic checks performed by the Data Collectors. SL1 auto-rotates these log files when they exceed 512MB, to prevent the log file from filling the /home partition.
- ***/var/log/phonehome/ph_watchdog.log***. Resides on the Database Server. This log file stores information about the watchdog service. This log is rotated daily.

Using the Command-Line Interface

If you have access to the console for each appliance in the PhoneHome configuration, or if you have SSH access to each appliance in the PhoneHome configuration, you can use a shell session and the **phonehome** command to configure and troubleshoot your PhoneHome configuration.

The phonehome Command

To use the **phonehome** command:

1. Either go to the console of the SL1 appliance or use SSH to access the server. Log in as "root".

> **NOTE:** For details on enabling and using SSH with SL1, see the manual ***Security***.

2. At the command prompt, you can type the following:

   phonehome *argument*

   where *argument* is one of the following:

   - ***add*** *appliance_type* or *request_file*. Run this command on the primary Database Server. Adds an appliance to the current PhoneHome configuration.
     - *appliance_type*. Type one of the following:
       - collector. Adds a Data Collector or Message Collector to the PhoneHome configuration.
       - database. Adds a primary Database Server or a secondary Database Server to the PhoneHome configuration.
     - *request_file*. When the Data Collector or Message Collector sends a request to the Database Server, the Database Server creates a request file in the directory **/home/phonehomerequest/requests**. You can specify the full pathname of a request file to accept a request and add a new Data Collector or Message Collector to the PhoneHome configuration.

> **NOTE**: The **phonehome *add*** *request_file* command performs the same operations as selecting the **[Accept]** button for a request in ScienceLogic Web Configuration Utility.

   - ***check*** **-json** *yes*. Run this command on any appliance in the PhoneHome configuration. Executes diagnostic steps to aid in troubleshooting.

     The **phonehome** command first tries to connect to the primary Database Server.

     If you issue this command from a Database Server, the command checks the status of the database port, the SSH port, and port for the web configuration tool for each Data Collector and Message Collector.

     If you issue this command from a Data Collector or Message Collector, the command checks the status of the database port, the SSH port, and port for the web configuration tool for each Database Server.

- **-json** *yes*. Displays output in json format.

- ***clear*** *clear_type*. Clears the PhoneHome configuration, as specified in the *clear_type* argument.
  - *clear_type*. Specifies which configuration to remove. Can be one of the following:
    - *client*. Run this command on the secondary Database Server, Data Collector, or Message Collector. Removes the PhoneHome connection (SSH tunnel). The appliance can then no longer connect to the primary Database Server.
    - *users*. Run this command on the primary Database Server. Removes the PhoneHome configuration for all appliances except the primary Database Server.
    - *all*. Run this command on the primary Database Server. Removes the PhoneHome configuration for each Data Collector, Message Collector, secondary Database Server, and the primary Database Server.

- ***close-control-port*** *port_number*. Run this command on Database Servers (primary and secondary). Blocks future connection requests from Data Collectors and secondary Database Servers.
- ***connect***. Run this command from the Data Collectors, Message Collectors, or secondary Database Server. Starts communication between the primary Database Server and the Data Collector, Message Collector, or secondary Database Server.
- ***delete*** *appliance_ID*. Run this command on the primary Database Server. Deletes an appliance from the current PhoneHome configuration.
  - *appliance_ID*. Enter the numeric ID of the appliance. You can find this ID with the **phonehome status** command.

- ***disconnect***. Run this command from the Data Collector(s), Message Collector(s), or secondary Database Server. Stops communication between the primary Database Server and the Data Collector, Message Collector, or secondary Database Server.

- ***help***. Run this command from any appliance in the PhoneHome configuration. Displays information about each parameter for the phonehome command.
- ***help extra***. Run this command from any appliance in the PhoneHome configuration. Displays information about the basic steps to configure a PhoneHome configuration.
- ***mysql*** *appliance_id*. Run this command on the primary Database Server. Tests the connection to the MySQL database. If the *appliance_ID* specifies a Data Collector or Message Collector, the **phonehome** command will test the MySQL connection using the loopback address of the Data Collector or Message Collector and port 7707. If the *appliance_ID* specifies a Database Server, the **phonehome** command will test the MySQL connection using the public IP address of the Database Server and port 7706.
  - *appliance_ID*. Enter the numeric ID of the appliance. You can find this ID with the **phonehome status** command.

- ***open-control-port*** *port_number*. Run this command on Database Servers (primary and secondary). Adds an entry for the specified SSH port to the /etc/sysconfig/iptables file on the current server.

- ***reconnect***. Run this command from the Data Collector(s), Message Collector(s), or secondary Database Server. Stops and then restarts communication between the primary Database Server and the Data Collector(s),Message Collector(s), or secondary Database Server.
- ***register*** *device_token*. Run this command from the Data Collector(s),Message Collector(s), or secondary Database Server. Registers the appliance with the primary Database Server.

  After you generate a ***token*** for a Data Collector or Message Collector (either with **phonehome token** or **phonehome add**), go to the Data Collector or Message Collector and use the **phonehome register** command to register the Data Collector or Message Collector with the primary Database Server. The Data Collector or Message Collector will then upload its public key to the primary Database Server and download its configuration for PhoneHome from the primary Database Server. After executing this command, the Data Collector or Message Collector will automatically connect to the Database Server.

  In configurations that have multiple Database Servers: After you generate a ***token*** for a secondary Database Server (either with **phonehome token** or **phonehome add**), go to the secondary Database Server and use the **phonehome register** command to register the secondary Database Server with the primary Database Server. The secondary Database Server will then upload its public key to the primary Database Server and download its configuration for PhoneHome from the primary Database Server.

  - *device_token*. Enter the token you generated for the Data Collector, Message Collector, or secondary Database Server.

- ***reload***. Can be run on any appliance in the PhoneHome configuration. Stops the em7_sshd and em7_ph_service processes, finds and applies any configuration changes, and restarts the service.
- ***request*** [*protocol*]://[*database_IP]* [*no_verify*]. Run this command from the Data Collector or Message Collector to send a request to the Database Server.
  - *protocol*. Enter the protocol to use to send the request to the Database Server. Choices are *phonehome* or *https*.
  - *database_IP*. The IP address of the Database Server in the PhoneHome configuration.
  - *no_verify*. Optional. If you specified *https* in the **protocol** option, you can specify ***no_verify*** to disable SSL verification.

NOTE: The **phonehome *request*** command performs the same operations as sending a request to the Database Server from the ScienceLogic Web Configuration Utility. Specifying ***no_verify*** performs the same operation as not selecting the ***Verify SSL Cert*** checkbox.

You can use the **phonehome request** command and the **phonehome add *request_file*** command to add a Data Collector or Message Collector to a PhoneHome Configuration. Go to the Data Collector or Message Collector and use the **phonehome request** command to send a request to join the PhoneHome configuration. Go to the Database Server and use the **phonehome add *request_file*** command to accept the request from the Data Collector or Message Collector. Go to the Data Collector or Message Collector again and execute the **phonehome request** command a second time to retrieve the request approval and set up the connection.

- ***set*** *appliance_ID -parameter=value*. Run on the primary Database Server. For a specific device, assigns a value to a parameter:
    - *appliance_ID*. Enter the numeric ID of the appliance. You can find this ID with the **phonehome status** command.
    - *parameter*. Can be one of the following parameters, preceded by a dash:
        - *name*. Specifies the name of the device in the ***Name*** field in the Web Configuration Utility.
        - *ssh*. Specifies whether or not to enable port forwarding for the SSH port for this device. Possible values are "yes" or "no".
        - *ip*. Specifies the IP address of the device in the ***IP Address*** field in the Web Configuration Utility.
        - *forwards*. Enables port forwarding for one or more ports. Specify one or more port numbers, separated by a space.
    - *value*. Value to assign to the parameter, surrounded by double quotes.

    For example:

    ```
    phonehome set 11 -ssh yes -name "Reston"
    ```

    - This example affects the device with an appliance ID of "11".
    - The example enables port forwarding for SSH.
    - The example enables port forwarding for the Web Configuration Utility.
    - The example sets the device's device name to "Reston".

- ***ssh*** *appliance_id*. Run on the primary Database Server. Tests the SSH connection to the specified appliance. If the *appliance_ID* specifies a Data Collector or Message Collector, the **phonehome** command will test the SSH connection using the loopback address of the Data Collector or Message Collector and port 10022. If the appliance ID specifies a Database Server, the **phonehome** command will test the SSH connection using the public IP address of the Database Server and defined SSH port.

    - *appliance_ID*. Enter the numeric ID of the appliance. You can find this ID with the **phonehome status** command.

- ***status***. Can be run on any appliance in the PhoneHome configuration. Displays the name and status of each currently defined PhoneHome appliance.
- ***sync***. Run this command from the Data Collectors or Message Collectors. Downloads the current configuration for PhoneHome from the primary Database Server to the Data Collector or Message Collector.
- ***token*** *appliance_ID*. Run this command from the primary Database Server. This command creates a URL that allows the Data Collector(s), Message Collector(s), or secondary Database Server to log in to the primary Database Server, upload a public key to the primary Database Server, and download the configuration for PhoneHome from the primary Database Server.

    - *appliance_ID*. Enter the numeric ID of the Data Collector, Message Collector, or secondary Database Server. You can find this ID with the **phonehome status** command.

- ***view*** *appliance_id* **-json***yes*. Run this command from the primary Database Server. Displays the name, type, loopback IP, port status, revision number, and SSH status of the Data Collector, Message Collector, or secondary Database Server specified in *appliance_ID*.

    - *appliance_ID*. Enter the numeric ID of the appliance that you want. You can find this ID with the **phonehome status** command.
    - **-json** *yes*. Displays output in json format.

- ***wake*** *appliance_id*. Run this command from the primary Database Server. Depending on the specified *appliance_ID*, stops and then restarts communication between the Database Server and the Data Collector, Message Collector, or secondary Database Server.

    - *appliance_ID*. Enter the numeric ID of the appliance that you want. You can find this ID with the **phonehome status** command.

- ***watchdog*** *option*. Run this command from the Data Collector, Message Collector, or secondary Database Server. The watchdog service runs automatically on each Data Collector, Message Collector, or secondary Database Server and checks the connection to the primary Database Server. If the connection is stale, the watchdog service automatically forces the primary Database Server to reconnect to the Data Collector, Message Collector, or secondary Database Server. The watchdog service can also detect configuration changes. If the PhoneHome configuration changes, the watchdog service will detect the changes and synchronize the configuration data on each device in the PhoneHome configuration.

NOTE: The ***watchdog*** and ***autosync*** services are not available for versions of SL1 ***earlier than the 7.5.3 ISO***.

You can use this command to control the watchdog service. The options are:

- *start*. Starts the PhoneHome watchdog service.
- *stop*. Stops the PhoneHome watchdog service.
- *status*. Gets the status of the PhoneHome watchdog service.
- *view*. Displays the current settings for the watchdog service.
- *set* -**parameter** *value*. Sets the value of a parameter for the watchdog service. Parameters are:
- interval seconds. Specify the interval, in seconds, at which to execute the watchdog service. The default value is "50".
- autosync (yes, no). Specifies whether or not you want the watchdog service to cause configuration data to be synchronized automatically at regular intervals.
- autoreconnect (yes, no). Specifies whether or not you want the watchdog service to reconnect stale connections automatically.
- state (enabled, disabled). Specifies whether or not the watchdog service is running.
- timeoutcount number. Specifies the number of failed calls to the watchdog service before stopping and restarting the watchdog. The default value is "3".
- check (ssh, db, default). Specifies which port the watchdog service checks. The default value is "db".
- run -verbose (yes, no). Manually starts the watchdog service if it is not already running.
- -verbose (yes, no). Specifies whether or not to display verbose logging to standard output.

The monitor_phonehome Command

The `monitor_phonehome` utility is available to use when debugging or diagnosing PhoneHome related issues. It is especially useful when escalating a PhoneHome issue to ScienceLogic Support.

Troubleshooting the PhoneHome Configuration

Using the PhoneHome Troubleshooting Script

The troubleshooting tool for PhoneHome provides diagnostics about the connection between a Data Collector and a Database Server in the PhoneHome system.

To use the troubleshooting tool:

1. Either go to the console of a Data Collector or use SSH to access the Data Collector.
2. Enter the following at the shell prompt:

```
sudo phtb -h
```

Available Commands

- This command checks the SSH connectivity between the Data Collector and the PhoneHome primary Database Server

```
sudo phtb destination
```

  - If the connection is healthy, the output looks like this:

```
Executing phonehome destination check
Running check for destination with host 52.70.238.48, port 7705, user
phonehome13
================================================================================
======
Remote server 52.70.238.48 reported back the status of this device (check
command): [Status: forwarded, Summary: db (Alive), Phonehome status:
operational]
SSH test to destination (Name: Phone Home database 11, IP: 52.70.238.48,
Port: 7705) is successful

Destination check summary: [Total: 1, Skipped: 0, Success: 1, failed: 0]
```

  - If the connection is not healthy, the output looks like this:

```
Executing phonehome destination check
Running check for destination with host 10.152.2.100, port 7705, user
phonehome13
================================================================================
======
Failed to check destination 10.152.2.100 on port 7705. Error: dial tcp
10.152.2.100:7705: i/o timeout

Destination check summary: [Total: 1, Skipped: 0, Success: 0, failed: 1]
```

- This command checks the SSH credentials.

```
sudo phtb probe-host <hostname or IP address of server> <port number> <username>
<-p password or -k path to private key> <options>
```

  - Here is an example with output for correct credentials:

```
phtb probe-host 52.70.238.48 7705 phonehome13 -k /home/phonehome0/.ssh/id_
rsa.pub

Executing probe for host 52.70.238.48 on port 7705
Successfully established SSH connection to host 52.70.238.48 on port 7705
```

  - Here is an example with output for incorrect credentials:

```
phtb probe-host 52.70.238.48 7705 em7admin -p em7admin

Executing probe for host 52.70.238.48 on port 7705
Failed to obtain SSH connection to host 52.70.238.48 on port 7705. Error:
ssh: handshake failed: ssh: unable to authenticate, attempted methods
[none], no supported methods remain
```

- This command verifies the connection to a proxy server.

```
sudo phtb proxy <hostname or IP address of server> <port number> <optional -u
username>
```

  - Here is an example with output for "no proxy found":

```
sudo phtb proxy 52.70.238.48 7705

Executing proxy check for host 52.70.238.48, port 7705
There is not any proxy configured for host 52.70.238.48. No further action
needed.
```

Basic Troubleshooting

Problem	Possible Cause	Diagnostics

Can't register a server or sync a server	Confirm that the Data Collector can "see" the phonehome port on the Database Server.	At the command line of the server that can't sync or register:
		`nmap -p port_number IP_address_of_database_server`
		For example:
		`nmap -p 7705 71.197.6.197`

Can't register a server or sync a server	Confirm that the PhoneHome port is open on the firewall on the Database Server.	At the command line of the Database Server:
		`iptables -nL`
		You should see output that specifies that the port accepts connections.
		To open the port, run this command:
		`sudo phonehome open-control-port`

Can't register a server or sync a server	Ensure that the Data Collector has line-of-sight with the Database Server	`sudo tcptraceroute <IP_address_of_database>`
Can't register a server or sync a server	Ensure the Data Collector can initiate an SSH session to the phonehome port on the Database Server	`sudo -u phonehome0 -s ssh -vvv -o StrictHostKeyChecking=no -p <control_port_(usually_7705)> phonehome<device_id_of_collector>@<IP_address_of_database>` Output should include: • Remote server string which specifies a version of OpenSSH • SL1 trying various authentication methods until one succeeds or fails • If authentication is attempted, this means that the Data Collector can establish a TCP connection on the control port to the Database Server

Problem	Possible Cause	Diagnostics

Can't register a server	Ensure that tunnel service is running on the Data Collector.	`sudo systemctl status em7_ph_tunnels.service`
		If the services is not running, force a sync using the following command:
		`sudo phonehome sync`

Problem with authentication	Ensure that the Data Collector is registered with the Database Server.	
Problem with authentication	Check that that the public keys of the Data Collector (/home/phonehome0/.ssh/id_rsa.pub) are configured correctly.	Ensure that the public keys of the Data Collector(/home/phonehome0/.ssh/id_rsa.pub) are present in the following places: On the Database Server, in /home/phonehome<id>/.ssh/authorized_keys On the Database Server, in /home/phonehome<id>/remote_key file On the Database Server, in the database table sysinfo..phonehome_devices, find the row for the Data Collector. The **pub** column should contain the public keys
Forgot to copy the token for a server		`phonehome token appliance_ID`
Confirm that Data Collector (s) and/or Message Collector(s) are successfully configured for PhoneHome		At the command line of the Database Server: `netstat –an |grep –i listen |grep "127.0.0." |grep 7707` Output displays each Data Collector and Message Collector that is listening for PhoneHome communications.

Chapter 8

Installing SL1 on AWS

Overview

This chapter describes how to install SL1 on an Amazon Web Services EC2 instances. An instance is a virtual server that resides in the AWS cloud.

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:

NOTE: For more information about monitoring Amazon Web Services in SL1, see the ***Monitoring Amazon Web Services*** manual.

AWS Instance Specifications

For details about AWS and the requirements and specifications for each SL1 appliance, see the ScienceLogic Support Site. https://support.sciencelogic.com/s/system-requirements?tabset-e65a2=f5872

Deploying an SL1 System on AWS

For ease of configuration, create appliances in this order:

1. Database Server
2. Administration Portal (if applicable)
3. Data Collectors
4. Message Collectors (if applicable)

What Are the ScienceLogic AMIs?

An instance is a virtual server that resides in the AWS cloud. An Amazon Machine Image (AMI) is the collection of files and information that AWS uses to create an instance. A single AMI can launch multiple instances.

For details on AMIs, see http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIs.html.

The ScienceLogic AMIs are defined by ScienceLogic. ScienceLogic has created an AMI for each type of ScienceLogic appliance. You can use a ScienceLogic AMI to create Elastic Compute Cloud (EC2) instances for each type of ScienceLogic appliance.

NOTE: Elastic Compute Cloud (EC2) instances are virtual servers that come in a variety of configurations and can be easily changed as your computing needs change. For more information on EC2, see http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts.html.

The ScienceLogic AMIs are private and are for ScienceLogic customers only. After you collect specific information about your AWS account, you can send a request (and the collected information) to ScienceLogic, and ScienceLogic will share the ScienceLogic AMIs with you.

NOTE: As of 8.10.0 and later releases, ScienceLogic AMIs support Enhanced Network Adapters (ENAs).

Getting the ScienceLogic AMI

To get access to the ScienceLogic AMIs:

1. Log in to the *ScienceLogic Support Site*.
2. Click your user name and from the menu select ***My Support and Customer Success***.
3. On the next page, click the **[Submit an AMI Request]** button.
4. Fill out the ***Request Amazon AMI*** form and click the **[Submit AMI Request]** button.
5. Repeat steps 1-4 for each type of SL1 appliance you want to install on AWS.
6. ScienceLogic Customer Support will send you an email confirming that they have shared the ScienceLogic AMI with your AWS account.
7. To view the ScienceLogic AMIs in your AWS account, go to the **AWS Management Console** page. Under the heading ***Compute***, click **[EC2]**.

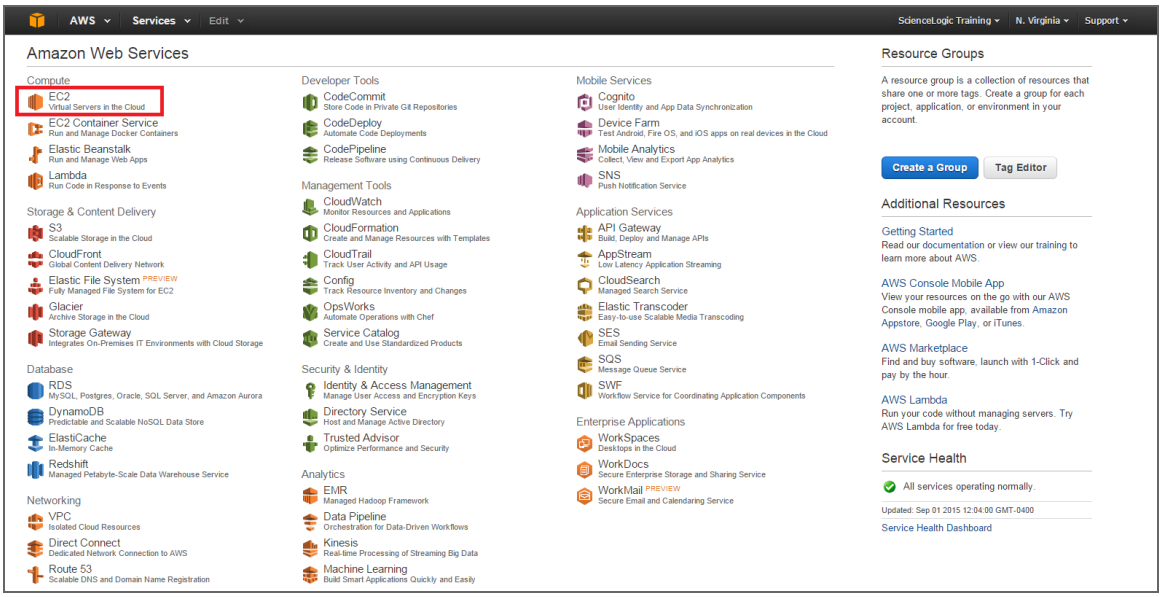

8. In the **EC2 Dashboard** page, go to the left navigation bar. Under the heading ***Images***, click **[AMIs]**.
9. In the main pane, under ***Filters***, click **[Owned by me]** and then select *Private images*.
10. You should see AMIs with names that begin with "EM7" and end with the current release number for SL1. You should see an AMI for each type of SL1 appliance.

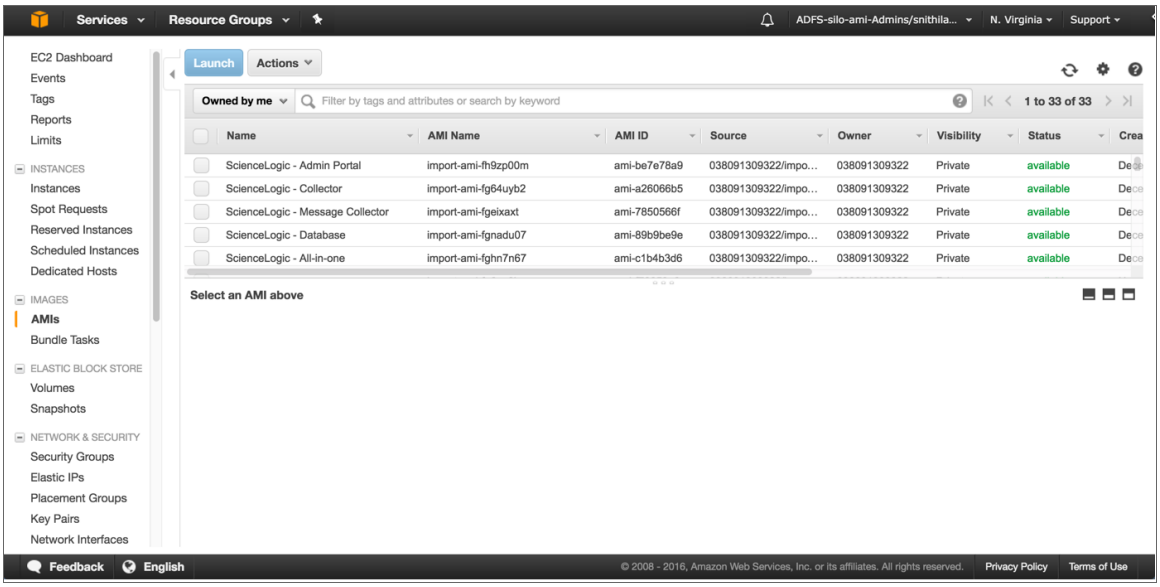

11. If you do not see AMIs with names that begin with "EM7", your EC2 Dashboard might have a default region that does not match the region for the ScienceLogic AMIs. To change the current region in the EC2 dashboard, click the region pull-down in the upper right and choose another region. Do this until you find the ScienceLogic AMIs.

NOTE: A region is a geographic location. AWS has data centers that include multiple regions. You can specify that an instance reside in a specific region. For more details on regions, see http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html.

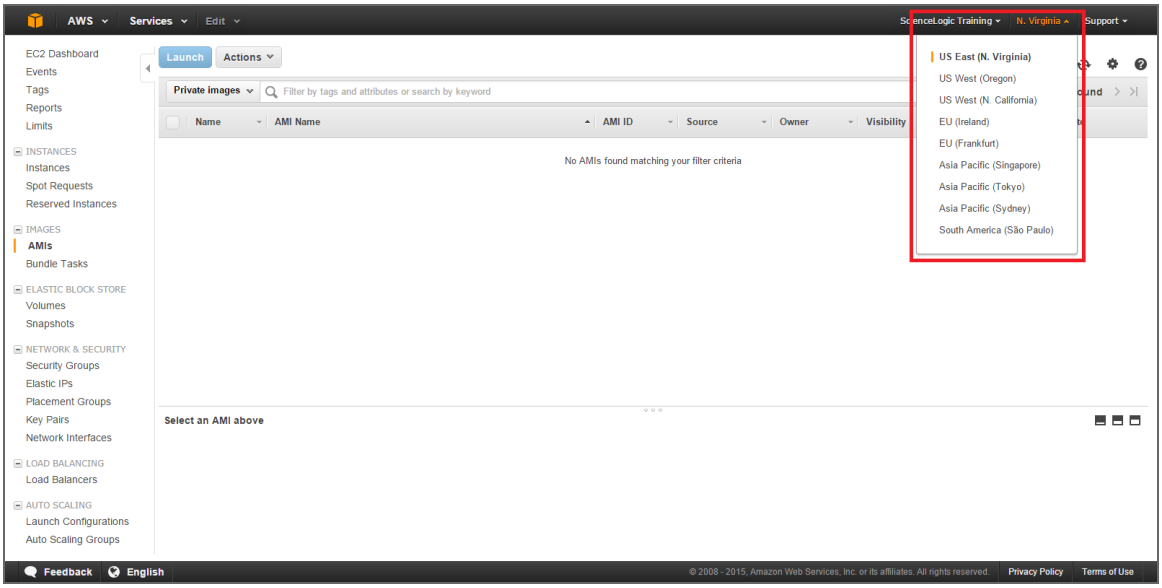

Launching the New Instance

To complete the steps listed in this chapter, you must have *received the ScienceLogic AMIs*.

This chapter assumes that you will launch each new EC2 instance into a VPC subnet with a primary IP address that is static and private.

NOTE: For more information on VPCs and VPC subnets, see http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Introduction.html.

For details about the recommended instance type for each ScienceLogic appliance, see System Requirements page on the *ScienceLogic Support Site*.

You can use multiple AWS instances to create a distributed SL1 System. For each instance, you must specify the correct instance type, storage size, and security rules. All these parameters are described in this chapter.

To launch the new EC2 instance from the ScienceLogic AMI:

1. Go to the ***EC2 Dashboard***:

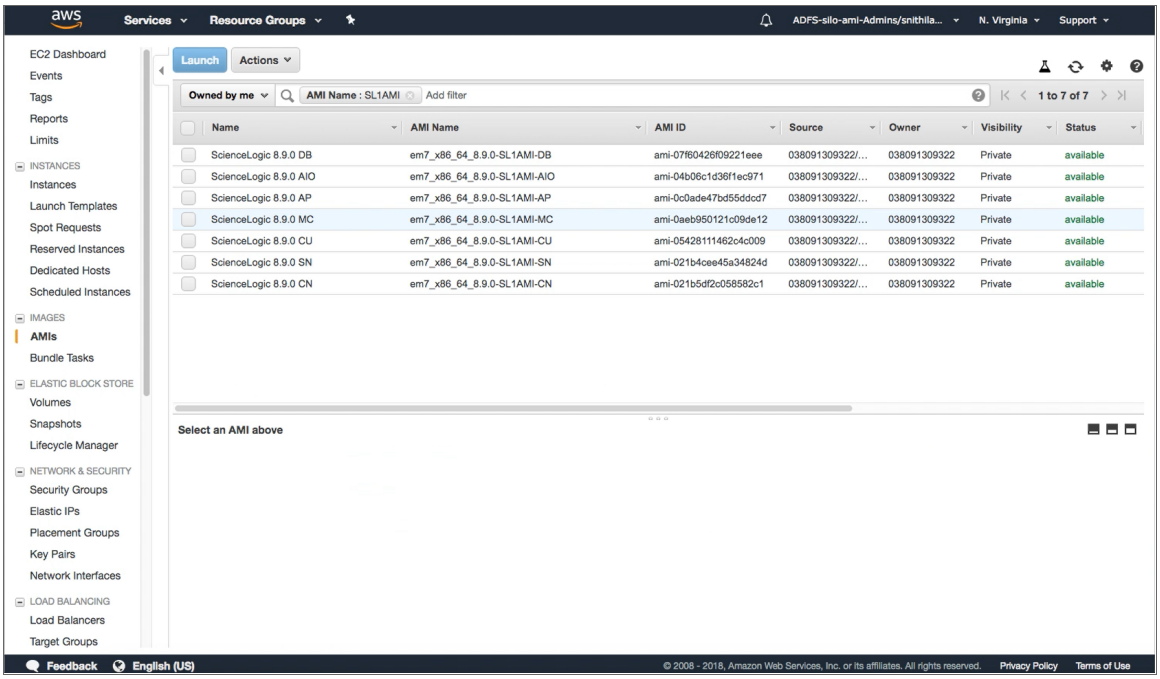

2. Select the ScienceLogic AMI that matches the ScienceLogic appliance you want to create. Click the **[Launch]** button.

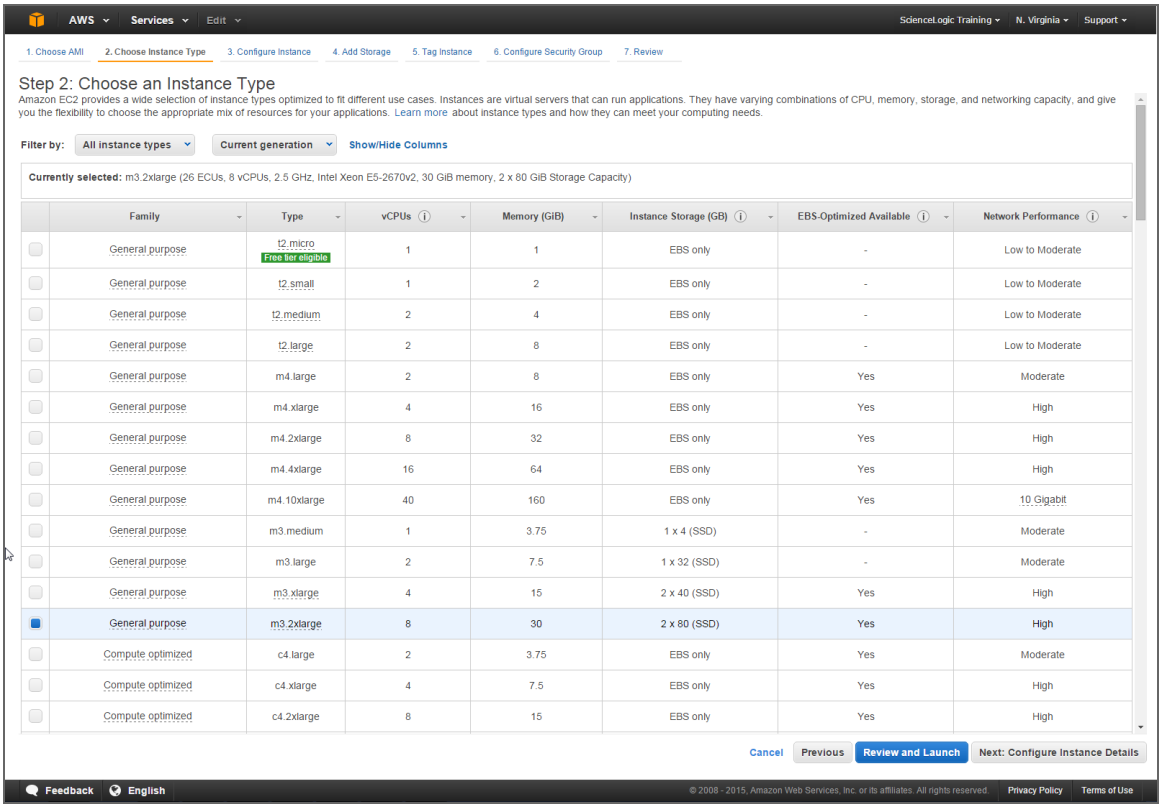

3. In the **Choose Instance Type** page, choose the instance type recommended for the AMI. Choose the size and type that fulfills your needs.

NOTE: For details about the recommended instance type for each ScienceLogic appliance, see the ScienceLogic Support Site. https://support.sciencelogic.com/s/system-requirements?tabset-e65a2=f5872

4. Click the **[Next: Configure Instance Details]** button.
5. In the **Configure Instance Details** page, define the following:
   - ***Number of Instances***. Enter "1".
   - ***Request Spot Instances***. Do not select.
   - ***Network***. For VPC-enabled accounts, specify the network where the instance will reside. If you are unsure of the network, accept the default.
   - ***Subnet***. For VPC-enabled accounts, specify the subnet where the instance will reside. If you are unsure of the subnet, accept the default.
   - ***Auto-assign Public IP***. If you select *Enable*, AWS will assign an IP address from the public pool to this instance. If you select *Disable*, you must assign an ***Elastic IP Address (EIP)*** to the instance.

NOTE: If you select *Enable* in the ***Auto-assign Public IP*** field, the IP address will change each time the instance is stopped or terminated. For All-In-One Appliances and for Administration Portals, you might want to use an Elastic IP address (EIP), which is a persistent IP address. See the section on *Elastic IP Addresses (EIP)* for details.

NOTE: For more information on Elastic IP Addresses, see http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html.

- ***IAM role***. If your organization uses IAM roles, select the appropriate role.
- ***Shutdown behavior***. Select *Stop*.
- ***Enable termination protection***. Selecting this checkbox is not required. Configure this checkbox according to your organization's procedures.
- ***Monitoring***. Do not select this checkbox.
- ***EBS-optimized instance***. Do not select this checkbox.
- ***Tenancy***. Select *Shared tenancy (multi-tenant hardware)*.

6. Click the **[Next: Add Storage]** button.

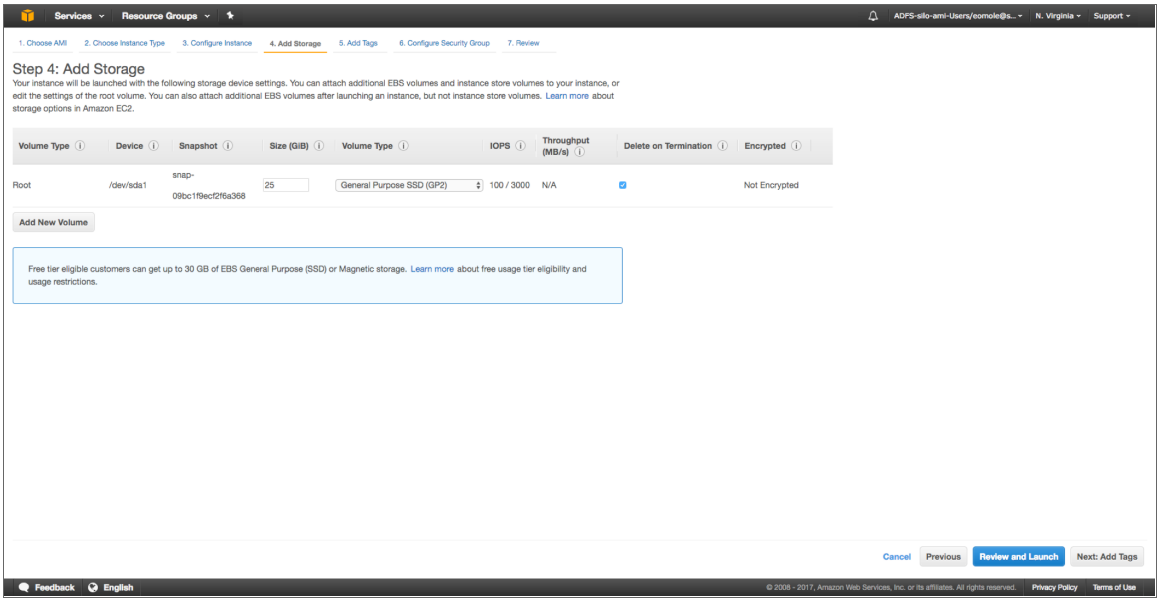

7. In the **Add Storage** page, select the checkbox in the ***Delete on Termination*** column.
8. In the **Add Storage** page, increase the size of the /dev/sda1 partition as follows:

SL1 Appliance	Type	>Device	Size in GB
Administration Portal	Instance Store	/dev/sda1	85
Message Collector without ScienceLogic Agent	Instance Store	/dev/sda1	85
Message Collector with ScienceLogic Agent	Instance Store	/dev/sda1	85
Database Server	EBS	/dev/sda1	105
All-In-One Appliance	EBSNVMe SSD	/dev/sda1	105
Data Collector	Instance Store	/dev/sda1	85

NOTE: The /dev/sda1 partition will contain the database.

9. Click the **[Next: Tag Instance]** button.

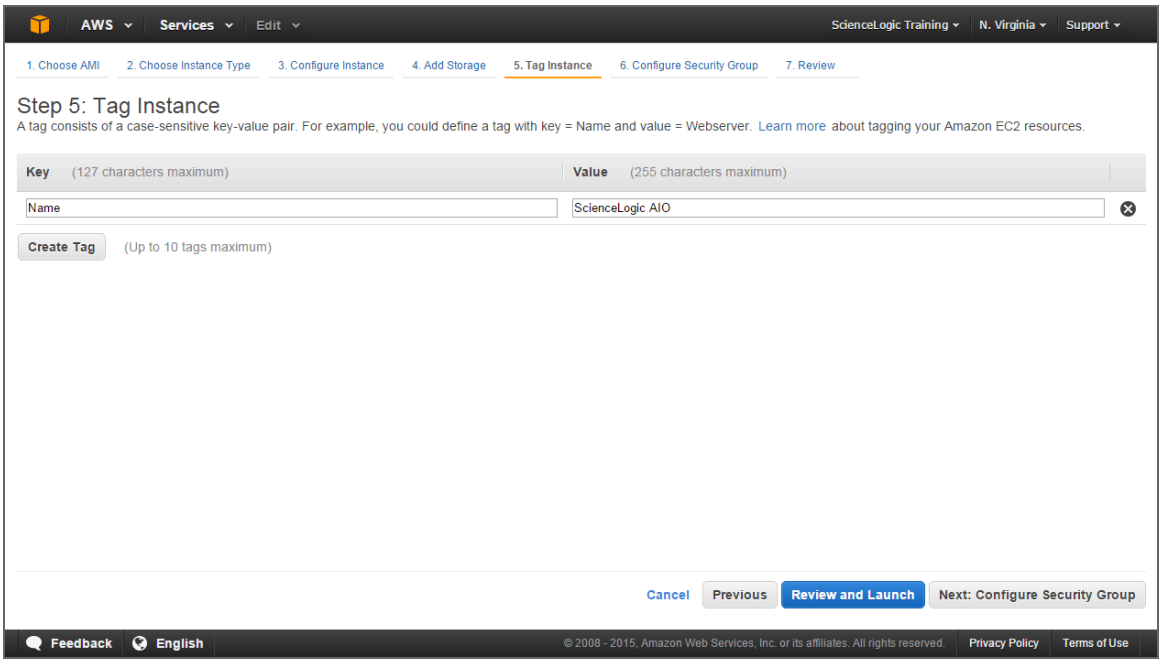

10. In the **Tag Instance** page, assign a descriptive tag to this instance. For example, you could enter "Name" in the ***Key*** field and "ScienceLogic AIO" in the ***Value*** field. This is optional.

NOTE: For more information on tags, see http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html.

11. Click the **[Next: Configure Security Group]** button.

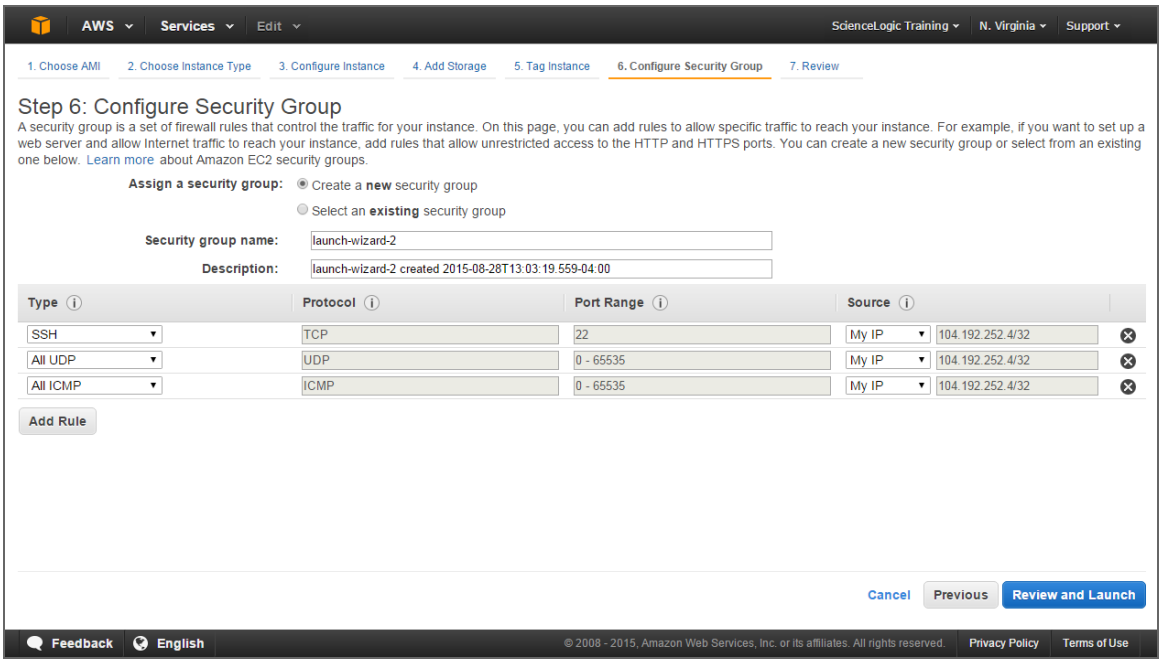

12. A security group is a reusable set of firewall rules. In the **Configure Security Group** page, do the following:

    - ***Assign a security group***. Select *Create a new security group*.
    - ***Security group name***. Enter a name or accept the default name.
    - ***Description***. Accept the default value in this field.

13. Use the *following tables* to create security rules for each type of SL1 appliance. After completing each row, click the **[Add Rule]** button.

14. Click the **[Review and Launch]** button and review the details of the new instance. Fix any problems to meet the requirements of your organization.

15. Click the **[Launch]** button.

16. Amazon EC2 instances use public-key cryptography for authentication. Select ***create a new key pair now***. You can enter a name for the private key. AWS will store the public key on its servers and automatically download the file that contains the private key to your browser. The private key is stored in a file that ends in .pem. You will need this file again when you *configure SSH* access to your AWS instances.

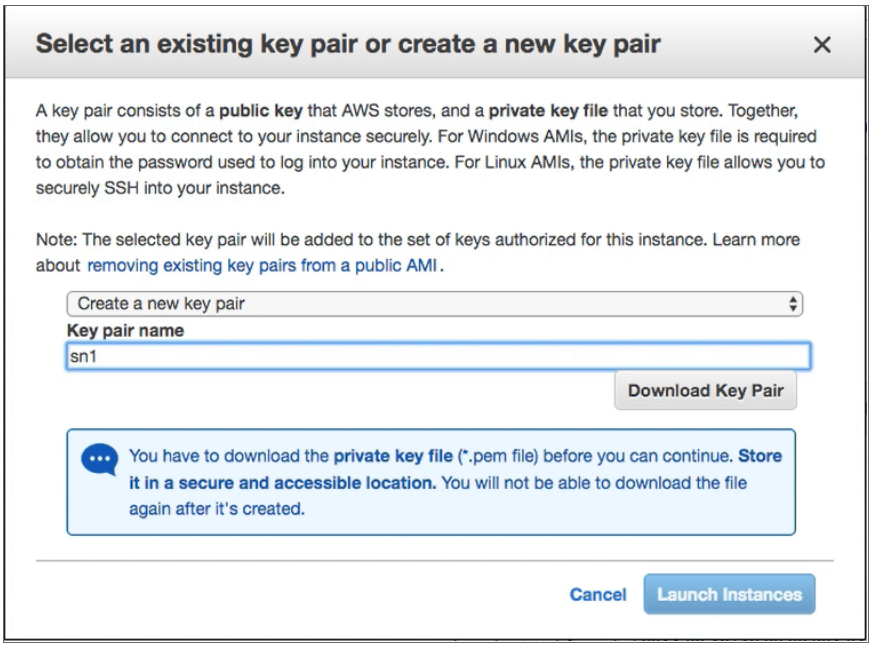

17. Amazon EC2 instances use public-key cryptography for authentication.

    - Select ***create a new key pair now***.
    - ***Key pair name***. Enter a name for the private key.
    - ***Download Key Pair***. AWS will store the public key on its servers and automatically download the file that contains the private key to your browser. The private key is stored in a file that ends in .pem. You will need this file again when you *configure SSH* access to your AWS instances.

> **NOTE**: Do not select an existing key unless you have previously downloaded and saved the key. You cannot retrieve an existing key a second time.

18. Click the **[Launch Instances]** button.
19. The **Launch Status** page displays the status of the new instance.
20. While the Launch runs in the background, go to the **Instances** page and provide a value in the ***Name*** field.
21. When the instance launch has completed, click the **[View Instances]** button to see your new instance.

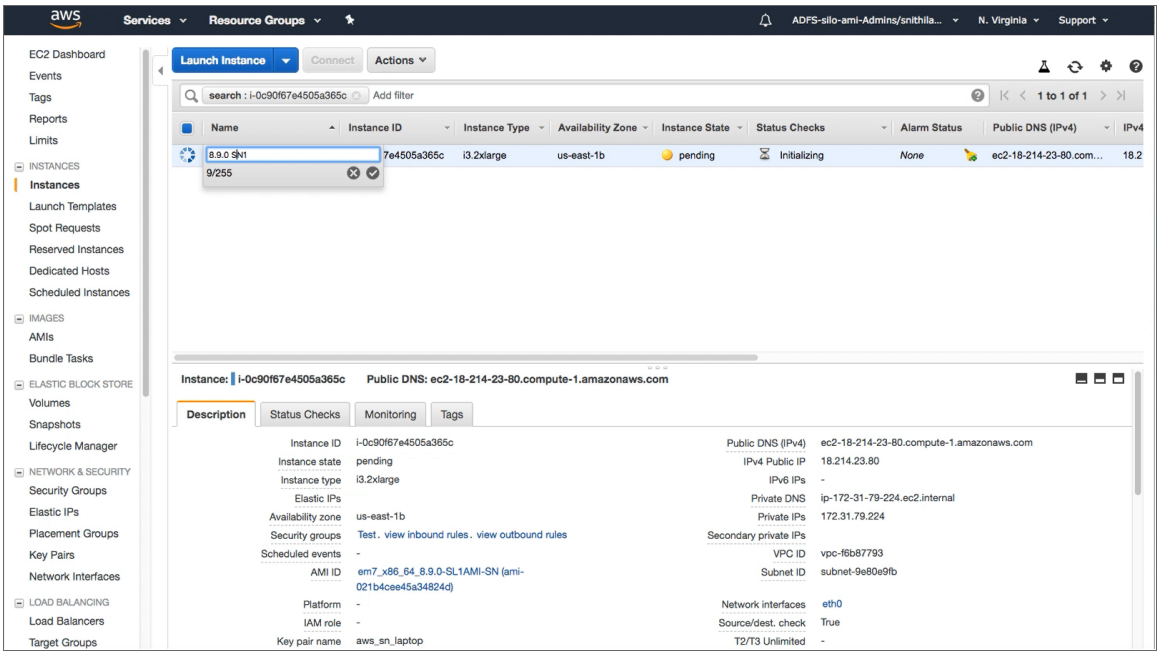

22. When the instance launch has completed, click the **[View Instances]** button to see your new instance.
23. For all nodes, continue to the steps listed in *Additional Configuration Steps*.

Security Rules for Each Appliance Type

All-In-One Appliance

Inbound

Type	Protocol	Port Range	Source	Description
SSH (edit the default SSH rule)	TCP	22	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	SSH. For SSH sessions from the user workstation to the appliance. This is necessary to start the installation wizard.

Type	Protocol	Port Range	Source	Description
HTTP	TCP	80	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	HTTP from browser session on user workstation.
HTTPS	TCP	443	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	HTTPS from browser session on user workstation.
Custom TCP Rule	TCP	7700	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	ScienceLogic Web Configurator. Configuration Utility from browser session on user workstation. This is necessary to license the appliance.
Custom UDP Rule	UDP	162	Specify a list of IP addresses for all managed devices from which you want to receive SNMP traps. Configure this list according to your requirements, your AWS configuration, and your security rules.	SNMP Traps. Necessary to receive SNMP traps from managed devices.

Type	Protocol	Port Range	Source	Description
Custom UDP Rule	UDP	514	Specify a list of IP addresses for all managed devices from which you want to receive Syslog messages. Configure this list according to your requirements, your AWS configuration, and your security rules.	Syslog messages. Necessary to receive syslog messages from managed devices.
SMTP	TCP	25	Specify a list of IP addresses for all managed devices from which you want to receive email messages. Configure this list according to your requirements, your AWS configuration, and your security rules.	Necessary to receive inbound email for tickets, events, and email round-trip monitoring.
Custom TCP Rule	TCP	123	Enter the IP address of the NTP server. Configure this list according to your requirements, your AWS configuration, and your security rules.	NTP. Communication between the All-In-One Appliance and configured NTP server.

Database Server

Inbound

Type	Protocol	Port Range	Source	Description
SSH (edit the default SSH rule)	TCP	22	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	SSH. For ssh sessions from user workstation to the appliance. This is necessary to start the installation wizard.

Type	Protocol	Port Range	Source	Description
SMTP	TCP	25	Specify a list of IP addresses for all managed devices from which you want to receive email messages. Configure this list according to your requirements, your AWS configuration, and your security rules.	Necessary to receive inbound email for tickets, events, and email round-trip monitoring.
HTTP **NOTE**: Required only if you are using the Administration Portal on the Database	TCP	80	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	HTTP from browser session on user workstation.
Custom TCP Rule	TCP	123	Enter the IP address of the NTP server. Configure this list according to your requirements, your AWS configuration, and your security rules.	NTP. Communication between the Database Server and configured NTP server.
Custom UDP Rule	UDP	161	Specify an IP address for each Data Collector that you will allow to can collect SNMP information about the Database Server. Configure this list according to your requirements, your AWS configuration, and your security rules.	SNMP Agent. Allows SNMP information about the Database Server to be collected by SL1.

Type	Protocol	Port Range	Source	Description
HTTPS **NOTE**: Required only if you are using the Administration Portal on the Database	TCP	443	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	HTTPS from browser session on user workstation.
Custom TCP Rule	TCP	7700	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	ScienceLogic Web Configurator. Configuration Utility from browser session on user workstation. This is necessary to license the appliance.
Custom TCP Rule	TCP	7706	Specify an IP address for each Data Collector that you will allow to collect SNMP information about the Database Server. Configure this list according to your requirements, your AWS configuration, and your security rules.	MySQL. Communication from Administration Portal
Custom TCP Rule	TCP	8008	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	Administrative Web Interface (PHPMyAdmin) from browser session on user workstation

Administration Portal

Inbound

Type	Protocol	Port Range	Source	Description
SSH (edit the default SSH rule)	TCP	22	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	SSH. For ssh sessions from user workstation to the appliance. This is necessary to start the installation wizard.
HTTP	TCP	80	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	HTTP from browser session on user workstation.
HTTPS	TCP	443	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	HTTPS from browser session on user workstation.
Custom TCP Rule	TCP	123	Enter the IP address of the NTP server. Configure this list according to your requirements, your AWS configuration, and your security rules.	NTP. Communication between the Administration Portal and configured NTP server.

Type	Protocol	Port Range	Source	Description
Custom TCP Rule	TCP	7700	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	ScienceLogic Web Configurator. Configuration Utility from browser session on user workstation. This is necessary to license the appliance.
Custom UDP Rule	UDP	161	Specify an IP address for each Data Collector that you will allow to can collect SNMP information about the Administration Portal. Configure this list according to your requirements, your AWS configuration, and your security rules.	SNMP Agent. Allows SNMP information about the Administration Portal to be collected by SL1.

Data Collector

Inbound

Type	Protocol	Port Range	Source	Description
SSH (edit the default SSH rule)	TCP	22	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	SSH. For ssh sessions from user workstation to the appliance. This is necessary to start the installation wizard.

Type	Protocol	Port Range	Source	Description
Custom TCP Rule	TCP	123	Enter the IP address of the NTP server. Configure this list according to your requirements, your AWS configuration, and your security rules.	NTP. Communication between the Data Collector and configured NTP server.
Custom UDP Rule	UDP	161	Specify an IP address for each Data Collector that you will allow to collect SNMP information about this Data Collector. Configure this list according to your requirements, your AWS configuration, and your security rules.	SNMP Agent. Allows SNMP information about the Data Collector to be collected by SL1.
Custom UDP Rule	UDP	162	Specify a list of IP addresses for all managed devices from which you want to receive SNMP traps. Configure this list according to your requirements, your AWS configuration, and your security rules.	SNMP Traps. Necessary to receive SNMP traps from managed devices.
Custom UDP Rule	UDP	514	Specify a list of IP addresses for all managed devices from which you want to receive Syslog messages. Configure this list according to your requirements, your AWS configuration, and your security rules.	Syslog messages. Necessary to receive syslog messages from managed devices.
Custom TCP Rule	TCP	7700	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	ScienceLogic Web Configurator. Configuration Utility from browser session on user workstation. This is necessary to license the appliance.

Type	Protocol	Port Range	Source	Description
Custom TCP Rule	TCP	7707	Specify the IP address of the Database Server that you want to retrieve data from the Data Collector. Configure this list according to your requirements, your AWS configuration, and your security rules.	Data Pull. Allows the Database Server to retrieve data from the Data Collector

Message Collector

Inbound

Type	Protocol	Port Range	Source	Description
SSH (edit the default SSH rule)	TCP	22	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	SSH. For ssh sessions from user workstation to the appliance. This is necessary to start the installation wizard.
Custom TCP Rule	TCP	123	Enter the IP address of the NTP server. Configure this list according to your requirements, your AWS configuration, and your security rules.	NTP. Communication between the Message Collector and configured NTP server.
Custom UDP Rule	UDP	161	Specify an IP address for each Data Collector that you will allow to collect SNMP information about this Message Collector. Configure this list according to your requirements, your AWS configuration, and your security rules.	SNMP Agent. Allows SNMP information about the Message Collector to be collected by SL1.

Type	Protocol	Port Range	Source	Description
Custom UDP Rule	UDP	162	Specify a list of IP addresses for all managed devices from which you want to receive SNMP traps. Configure this list according to your requirements, your AWS configuration, and your security rules.	SNMP Traps. Necessary to receive SNMP traps from managed devices.
Custom UDP Rule	UDP	514	Specify a list of IP addresses for all managed devices from which you want to receive Syslog messages. Configure this list according to your requirements, your AWS configuration, and your security rules.	Syslog messages. Necessary to receive syslog messages from managed devices.
Custom TCP Rule	TCP	7700	If you will always log in from a single IP address, select *My IP*. If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field. Configure this list according to your requirements, your AWS configuration, and your security rules.	ScienceLogic Web Configurator. Configuration Utility from browser session on user workstation. This is necessary to license the appliance.
Custom TCP Rule	TCP	7707	Specify the IP address of the Database Server that you want to retrieve data from the Message Collector. Configure this list according to your requirements, your AWS configuration, and your security rules.	Data Pull. Allows the Database Server to retrieve data from the Message Collector.

Additional Configuration Steps

After the instance is successfully launched, perform these additional steps to complete configuration:

- For instances of the ***Database Server*** or ***All-In-One Appliance***:

- *Assigning an EIP to the instance* (optional step)
- *Accessing the Appliance Using SSH*
- *Licensing the instance in the Web Configuration Tool*

- For instances of the ***Administration Portal***:
  - *Assigning an EIP to the instance* (optional step)
  - *Accessing the Appliance Using SSH*
  - *Configuring the instance in the Web Configuration Tool*
- For instances of the ***Data Collector and Message Collector***:
  - *Assigning an EIP to the instance* (optional step)
  - *Accessing the Appliance Using SSH*
  - *Configuring the instance in the Web Configuration Tool*
  - *Rebooting Data Collectors and Message Collectors*

Assigning an EIP to the New Instance

This chapter assumes you have already *received the ScienceLogic AMI* and *created an EC2 instance* based on the ScienceLogic AMI.

AWS can assign a public-facing IP address to your new instance. However, the IP address will change each time the instance is stopped or terminated. If you will be accessing an All-In-One Appliance or an Administration Portal appliance from the internet, ScienceLogic recommends you use an Elastic IP address (EIP).

An EIP is a permanent static address that belongs to an account (not an instance) and can be reused. An EIP address is required only if you want the public IP address to remain constant. When you assign an EIP to an instance, the instance still retains its private IP address in its VPC.

If you use an AWS VPN to access the All-In-One Appliance or Administration Portal appliance, that is you can access the All-In-One Appliance or Administration Portal appliance only through your corporate network, you do not have to assign an EIP to the All-In-One Appliance or Administration Portal appliance .

NOTE: For more information on Elastic IP, see http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html

NOTE: AWS accounts are limited five Elastic IP addresses.

To assign an EIP to your new instance:

1. Go to the *EC2 Dashboard*:

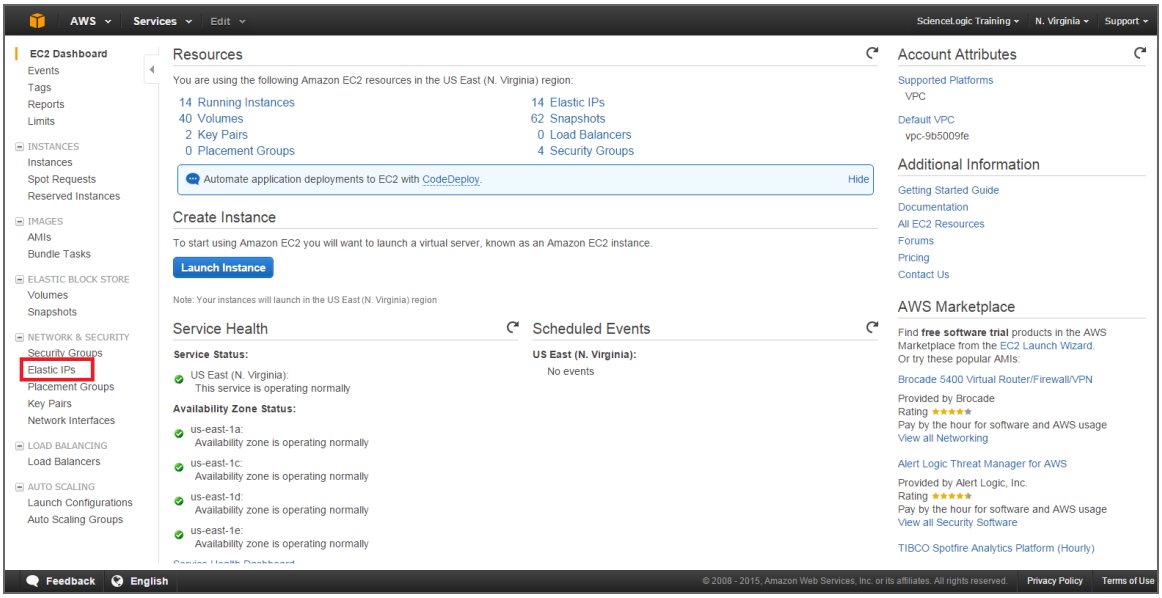

- In the left navigation pane, under the **Network & Security** heading, click **[Elastic IPs]**.

2. In the **Allocate New Address** page:

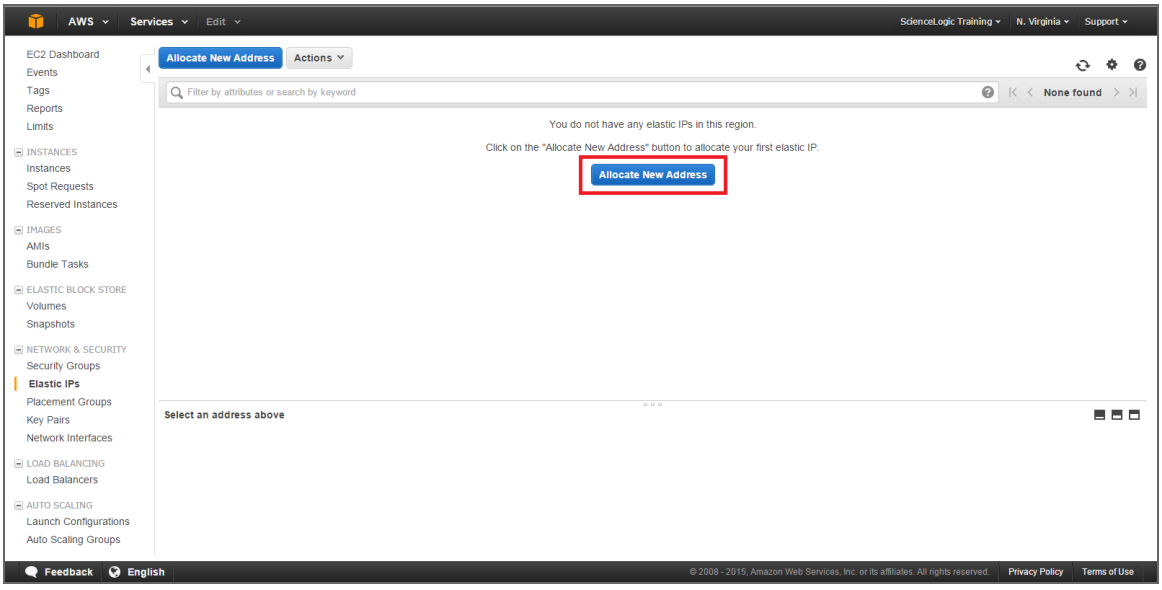

- Click the **[Allocate New Address]** button and then click the **[Yes, Allocate]** button.

3. In the **Allocate New Address** page:

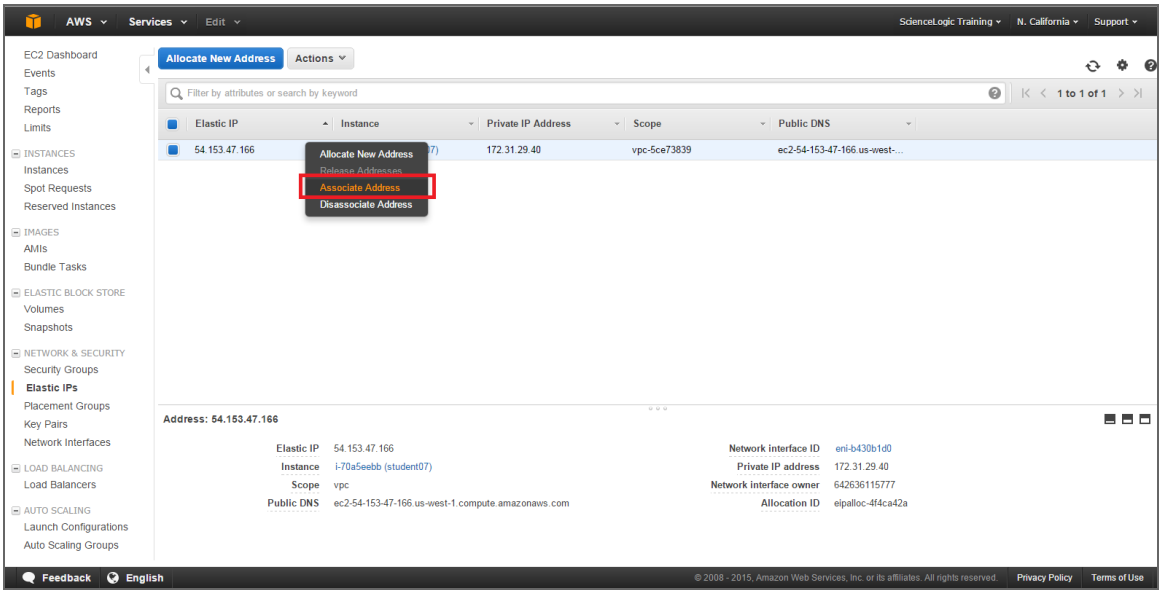

- Right-click the new address and select ***Associate Address*** from the drop-down menu.

4. In the **Associate Address** modal page:

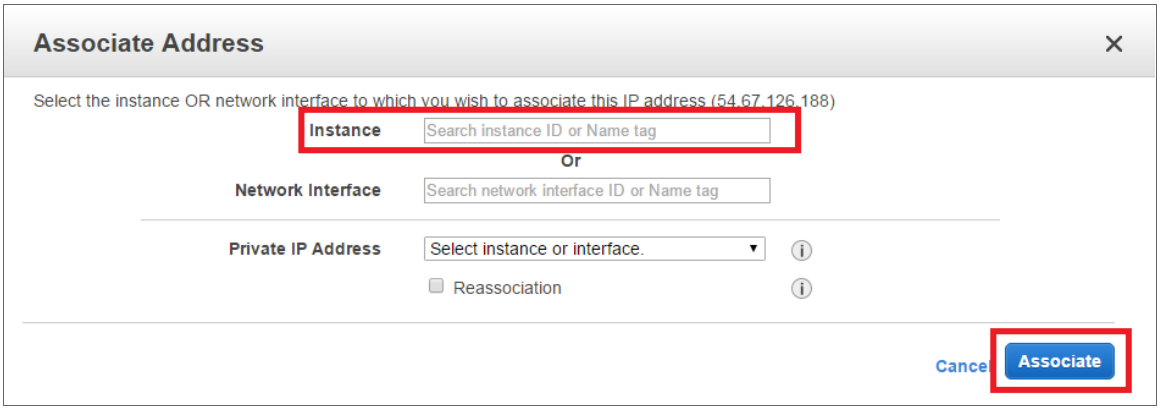

- Select the new SL1 appliance instance in the ***Instance*** field, then click the **[Associate]** button. The SL1 appliance instance is now associated with the new EIP.

Accessing the Appliance Using SSH

This chapter assumes you have already *received the ScienceLogic AMI*s and *created an EC2 instance* based on the ScienceLogic AMI.

This chapter assumes that you have access to SSH on the command line (for UNIX users) or have installed PuTTY (for Windows users).

Gathering Information Required for Accessing the Appliance Using SSH

To gather the required information:

1. Go to the *EC2 Dashboard*:
2. In the left navigation pane, under the **Instances** heading, select ***Instances***.

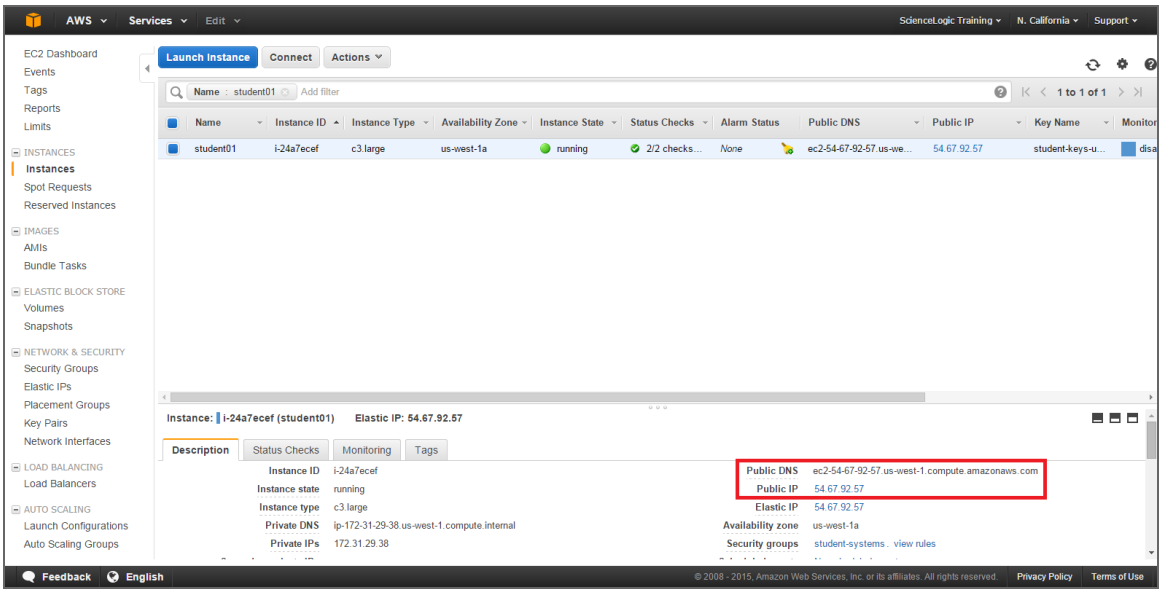

3. Click in the row that contains the SL1 appliance instance.
4. The lower pane contains information about the instance. Write down the ***Public DNS*** and ***Public IP***.
5. If you are using AWS instances to create a distributed SL1 system, perform this step for each AWS instance you want to include in the distributed system.

Configuring SSH

Before you can use SSH with the SL1 appliance instance, you must ensure that SSH can use the .pem file downloaded earlier during the configuration. For details on downloading the .pem file, see the last few steps in the section on *Launching the EC2 Instance*.

UNIX and LINUX Users

You can connect to your SL1 appliance instance using the SSH command.

NOTE: You should store the .pem file in a secure location. ScienceLogic recommends you store the .pem file in $HOME/.ssh. ScienceLogic also recommends you change the permissions on the .pem file to allow only read-only access by the owner of the .pem file.

To connect using the .pem file generated by AWS, enter the following at the shell prompt:

```
ssh -i ~/.ssh/my-aws-key.pem em7admin@[hostname or IP address]
```

where:

- ***~/.ssh/my-aws-key.pem***. Replace with the name and full path to your .pem file.
- ***hostname or IP address***. Replace with the hostname or public-facing IP address of the SL1 appliance instance.

You can also configure your SSH client to automatically select the correct key file when accessing the SL1 appliance instance. For details, see the man page for ssh_config for your flavor of UNIX.

Windows Users

You can connect with your SL1 appliance instance using PuTTY and SSH as the em7admin user. However, you must first convert the private key for your instance into a format that PuTTY can use. See the following for detailed instructions on using PuTTY SSH and converting your private key:

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/putty.html

Web Configuration Tool

- For instances of the ***Database Server*** or ***All-In-One Appliance***, see the section on *Licensing the instance in the Web Configuration Tool*
- For instances of the ***Administration Portal***, see the section on *Configuring the instance in the Web Configuration Tool*
- For instances of the ***Data Collector and Message Collector*** , see the section on *Configuring the instance in the Web Configuration Tool*

Rebooting Data Collectors and Message Collectors

After installing a Data Collector or a Message Collector as an AWS instance, you must reboot the instance.

To reboot the AWS instance:

1. Connect to the command-line interface of the appliance as the em7admin user using SSH. See the *Accessing the Appliance Using SSH* section for more information.
2. Execute the following command:

```
sudo reboot
```

Additional Steps for SL1 10.1

SL1 10.1.x includes an upgrade to MariaDB. The upgrade did not include a tool, jemalloc, that helps manage memory usage.

NOTE: This section applies only to the following releases:

- 10.1.0
- 10.1.1
- 10.1.2
- 10.1.3
- 10.1.4
- 10.1.4.1
- 10.1.4.2
- 10.1.5
- 10.1.5.1

For SL1 versions later than 10.1.5.1, jemalloc is included with the platform. For SL1 versions prior to 10.1.0, jemalloc is included with the platform.

To avoid problems with memory usage on Database Servers, perform the following steps after upgrading MariaDB for 10.1.x.

NOTE: Perform these steps first on the active Database Server and then on each additional Database Server in your SL1 System.

1. Open an SSH session to the Database Server.
2. To verify that the Database Server is not currently running jemalloc, enter the following at the shell prompt:

```
silo_mysql -e 'show global variables like "version_malloc_library"'
```

   If the Database Server is not currently running jemalloc, the shell will display the following:

Variable Name	Value
version_malloc_library	system

3. Search for the file /usr/lib64/**libjemalloc.so.1**.

If the file does not exist, contact ScienceLogic Customer Support to request the file jemalloc-3.6.0-1.el7.x86_64.rpm.

To install the RPM, use a file-transfer utility, copy the file to a directory on the SL1 appliance. Then enter the following at the shell prompt:

```
cd /usr/lib64
sudo yum install jemalloc-3.6.0-1.el7.x86_64.rpm
```

4. Create the file /etc/systemd/system/mariadb.service.d/jemalloc.conf:

```
vi /etc/systemd/system/mariadb.service.d/jemalloc.conf
```

5. Add the following lines to the file:

```
[Service]
Environment="LD_PRELOAD=/usr/lib64/libjemalloc.so.1"
```

6. Save and close the file.
7. Reload the systemd config files:

```
sudo systemctl daemon-reload
```

8. Restart the Database Server:

To restart the ***standalone Database Server*** or the ***primary Database Server in a cluster***, enter the following:

```
sudo systemctl restart mariadb
```

To restart each ***secondary Database Server in a cluster***:

a. Open an SSH session to the secondary Database Server. At the shell prompt, enter:

```
coro_config
```

b. Select **1**.
c. When prompted to put the Database Server into maintenance, select **y**.
d. Open an SSH session to the primary Database Server. To pause SL1, enter the following at the shell prompt:

```
sudo touch /etm/.proc_mgr_pause
```

e. In the SSH session for the secondary Database Server, restart MariaDB:

```
crm resource restart mysql
```

f. After MariaDB has restarted successfully on the secondary Database Server, return to the SSH session on the primary Database Server. Remove the pause file for SL1:

```
sudo rm /tmp/.proc_mgr_pause
```

g. In the SSH session on the secondary Database Server, take the Database Server out of maintenance. At the shell prompt, enter:

```
coro_config
```

h. Select **1**.
i. When prompted to take the Database Server out of maintenance, select **y**.

9. To verify that jemalloc is running on the Database Server, enter the following at the shell prompt:

```
silo_mysql -e 'show global variables like "version_malloc_library"'
```

If the Database Server is currently running jemalloc, the shell will display something like the following:

Variable Name	Value
version_malloc_library	jemalloc 3.6.0-0-g46c0af68bd248b04df75e4f92d5fb804c3d75340

10. Perform these steps on each Database Server in your SL1 system.

Chapter 9

Installing SL1 in Azure

Overview

This chapter describes how to deploy a ScienceLogic virtual machine in Azure from a VHD image file.

To deploy an SL1 appliance in Azure, you need the following components:

- Windows PowerShell version 5 or later
- PowerShellGet and Azure RM modules installed
- Azure CLI tool
- An Azure Resource group
- An Azure storage account that includes at least one blob container
- An Azure Network Security Group (NSG)

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 (•••).

The steps to create these components in Azure are described throughout this chapter. This chapter includes the following topics:

NOTE: If you are configuring a Database, ScienceLogic recommends allocating four times the memory for the Database as compared to the memory for the Data Collectors.

TIP: A single Azure image file can be used to create multiple virtual machines. For example, you can use the same Azure VHD file for the Database Server to create multiple Database Servers.

System Requirements

For details about supported hypervisors and the requirements and specifications for each SL1 appliance, see the ScienceLogic Support Site: https://support.sciencelogic.com/s/system-requirements

Deploying an SL1 System in Azure

This section provides the information you need to deploy SL1 in Azure and includes procedures for each step.

SL1 Appliance Deployment Order for Distributed Systems

To deploy a distributed SL1 system on Azure instances, create appliances in this order:

1. Database Server
2. Administration Portal (if applicable)
3. Data Collectors
4. Message Collectors (if applicable)

Installing and Configuring the Azure CLI

Azure CLI is a tool that lets you manage resources in Azure. To complete the SL1 installation on Azure using the procedures in this section, you must have the Azure CLI tool installed.

To install and configure the Azure CLI tool:

1. Download and install the Azure CLI tool from the Microsoft website: https://docs.microsoft.com/en-us/cli/azure/install-azure-cli-windows?view=azure-cli-latest
2. After installation completes, search for and click on "PowerShell" in Windows to start the program.
3. In PowerShell, type "az login". You will be prompted to sign into your Microsoft Azure account. After you log in, you will see information about your Azure subscription within the PowerShell window.

Configuring an Azure Resource Group and Storage Account

To create and configure an Azure storage account:

1. Log in to the Azure Portal, and then click **Resource groups** on the left menu.

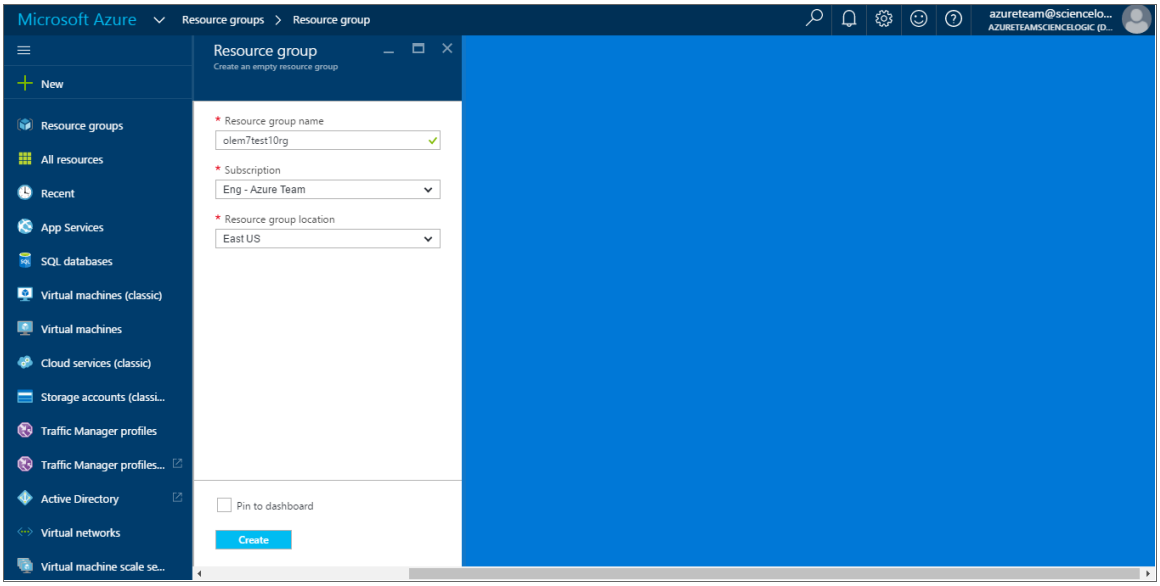

2. Click the **[Add]** button and add information for a new Resource group. Click the **[Create resource group]** button to create the Resource group.
3. After creating the Resource group, click **Storage accounts** on the left menu.

4. Click the **[Add]** button, and then click the **[Create Storage accounts]** button to create a new Storage account. When creating the Storage account, complete the following fields:

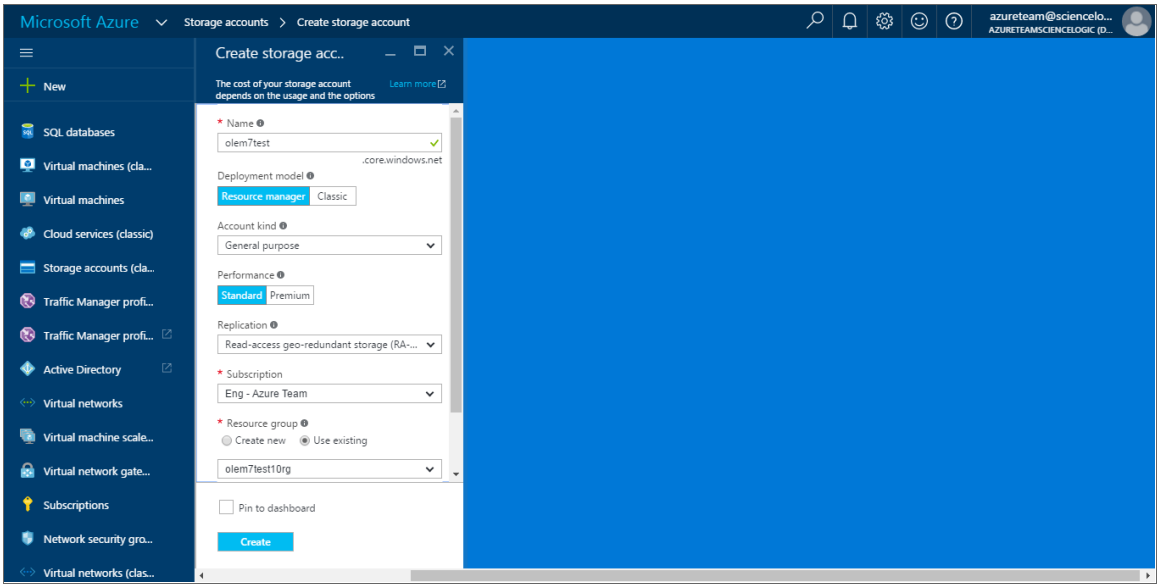

- ***Deployment model***. Select *Resource manager*.
- ***Account kind***. Select *General purpose*.
- ***Resource group***. Select *Use existing*, and then select the Resource group you created in step 2.

5. After creating the Storage account, click **Storage accounts** on the left menu, and then select the newly created Storage account.
6. Under the Services section of the Storage account pane, click **Blobs**. The Blob service blade information appears.

7. In the Blob service pane, click the **Plus** icon to add a new container. Type a name for the container and select *Blob* as the ***Access Type***. When you are finished, click the **[Create]** button to create the Blob container.

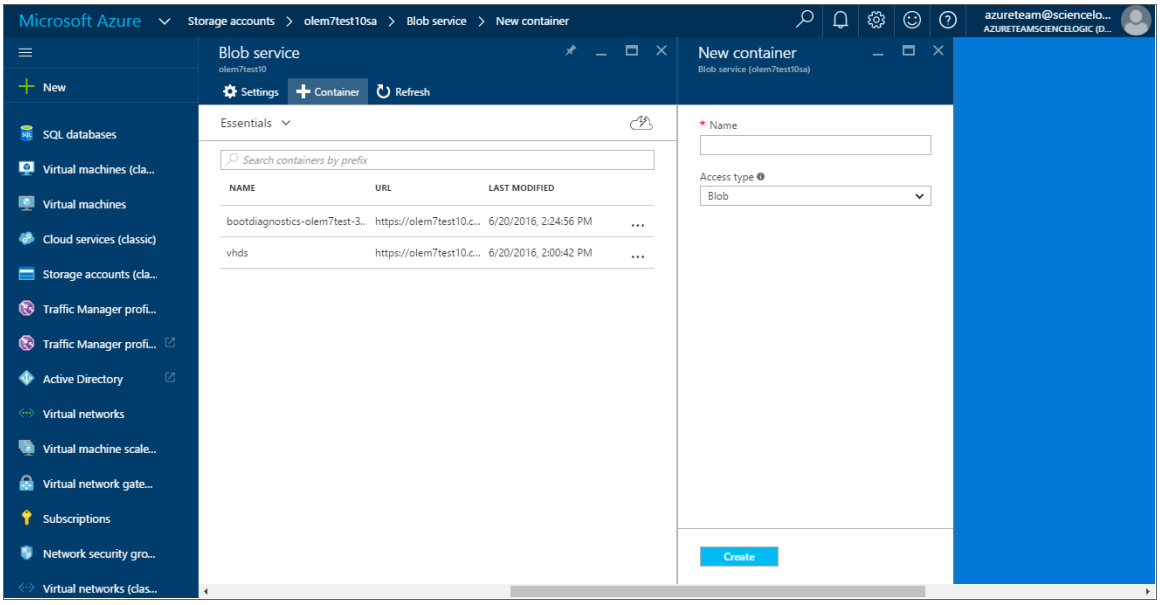

Obtaining the Container URL for an Azure Storage Account

Before you can upload the VHD image file, you will need to determine the BLOB_URI value from your Azure storage account. The container URL, along with the VHD filename you append, is used for the BLOB_URI value you will use when you upload the VHD file. This BLOB_URI tells Azure where to put the VHD file and what to call it.

To get the container URL:

1. Log in to the Azure portal.
2. Select **Storage Accounts**.
3. Select the **Containers** section.
4. Click the ellipsis (...) button to the right of the container name to open the pop-up menu.
5. Click **Container Properties** in the pop-up menu. You will see a URL displayed in the Properties.
6. Click the blue Copy icon on the Properties page to copy the URL for the container.

To create the BLOB_URI value, append the destination (that is, where you want the VHD to be stored in the container) to the container URL.

Example:

If the container URL is:

```
https://azuretest.blob.core.windows.net
```

Your BLOB_URI value might be:

```
https://azuretest.blob.core.windows.net/vhds/em7inazure.vhd
```

Where "/vhds" is the directory on the container, and "em7inazure.vhd" is the name of the VHD image file you will be uploading.

Uploading a VHD Image File to an Azure Storage Account

After creating the Resource group, Storage account, and Blob container, you must upload the ScienceLogic VHD image file to the Blob container. To do so, you will need the following information:

- The ScienceLogic VHD file
- Resource Group name
- Blob container URI
- Local file path to the VHD file

Downloading the ScienceLogic VHD File

To download the ScienceLogic VHD file:

1. Open a browser session and go to:

   ```
   support.sciencelogic.com
   ```

2. Go to the **Product Downloads** menu and choose ***Platform***.
3. Find the platform version that you want to download. Click on its name.
4. Expand the ***Release Files*** list and find an entry of Record Type *Product Image*.
5. In the **Release File Downloads** pane to the right, download the .vhd files for each SL1 appliance.

Uploading the VHD File to an Azure Container

To upload the ScienceLogic VHD file to your Blob container, perform the following steps:

1. Open Microsoft Azure PowerShell and log in to your Azure account using the cmdlet Add-AzureRmAccount:

   ```
   Login-AzureRmAccount
   ```

2. You created a resource group and storage container blob to which you will upload your VHD image file in *Configuring an Azure Resource Group and Storage Account*. You identified the BLOB_URI in *Obtaining the Container URL for an Azure Storage Account*. Now you are ready to upload your VHD image file.

> **NOTE:** The virtual machine that you create in *Creating an Azure Virtual Machine* must be in the same resource group as the storage account.

3. Add your VHD file to the storage account with the following cmdlet:

```
Add-AzureRmVhd -Destination '<BLOB_URI>' -LocalFilePath '<VHD_LOCAL_FILE_
PATH>' -ResourceGroupName '<RESOURCE_GROUP>'
```

   Where:

   - **BLOB_URI** specifies the BLOB_URI where you will upload your VHD file.
   - **VHD_LOCAL_FILE_PATH** specifies the path on your machine for the VHD file you want to upload
   - **RESOURCE_GROUP** specifies the resource group you created and that will be used when you create the Azure VM.

4. Repeat step 3 for each VHD file.

Creating the Disk Image

> **NOTE:** The following steps require that you have an ARM resource group and storage account with the VHD file uploaded.

To create the disk image:

1. Open Microsoft Azure PowerShell and log in to your Azure account using the cmdlet Add-AzureRmAccount:

```
Login-AzureRmAccount
```

2. Run the following command:

```
az disk create -g <resource-group-name> -n <disk-name> --source <BLOB_URI>
```

   Where:

   - **resource-group-name** specifies the resource group you created in Azure.
   - **disk-name** specifies what you want to call the disk (such as "dbdisk123").
   - **BLOB_URI** specifies the destination value you provided when you uploaded the VHD file. This is also the BLOB_URI you created in *Obtaining the Container URL for an Azure Storage Account*.

Creating an Azure Virtual Machine

> **NOTE:** The following steps require that you have an ARM resource group and storage account with the VHD file uploaded.

To create an Azure virtual machine:

1. Open Microsoft Azure PowerShell and log in to your Azure account using the cmdlet Add-AzureRmAccount:

   ```
   Login-AzureRmAccount
   ```

2. Run the following command:

   NOTE: The virtual machine that you create must be in the same resource group as the storage account.

   ```
   az vm create -g <resource_group_name> -n <VM_name> --os-type linux --attach-
   os-disk <disk_name> --vnet-name <virtual_network_name> --subnet <subnet_name>
   --size <VM_type>
   ```

   Where:

   - **resource_group_name** specifies the resource group you created in Azure.
   - **VM_name** specifies what you want to call your VM.
   - **disk_name** specifies the name you gave to the disk when you created it.
   - **virtual_network_name** specifies the virtual network name you want to use within Azure.
   - **subnet_name** specifies the subnet name you want to use within Azure.
   - **VM_type** specifies the Azure VM type you want to use, such as "Standard_DS2_VS". You can find a list of Azure machine types and attributes here: https://docs.microsoft.com/en-us/azure/virtual-machines/sizes-general

Setting the Virtual Machine Allocation Method to Static

To ensure the IP address for the virtual machine remains the same after reboot, you must set the allocation method to ***static***. To do so:

1. In the Azure Portal, navigate to the Virtual machine pane and verify that the virtual machine has a public IP address and a virtual network/subnet set.

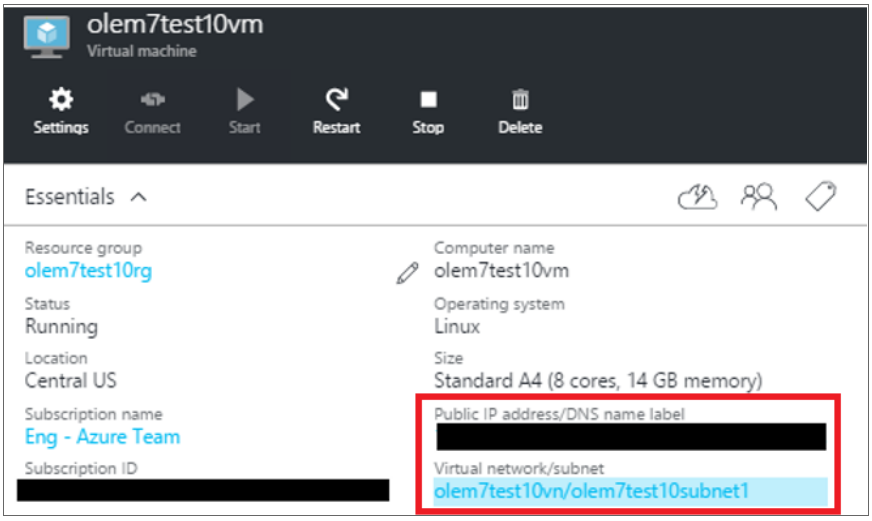

2. Click the name of the Virtual network/subnet. The Virtual network pane appears.

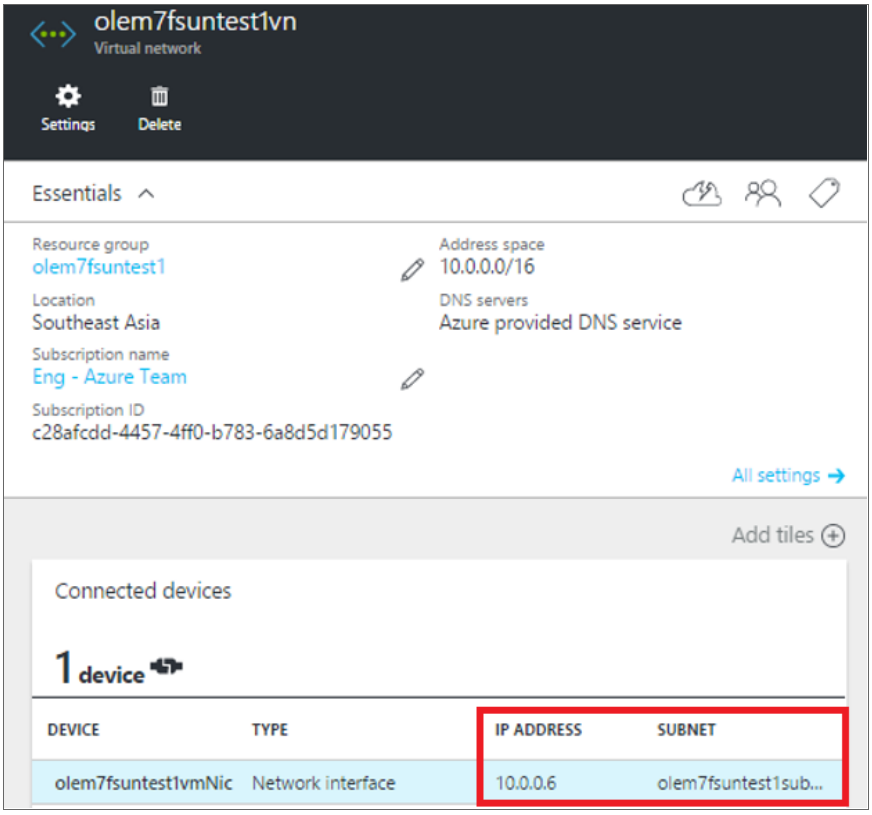

3. In the "Connected devices" section of the Virtual network pane, click the Network interface. The Network interface pane appears.
4. Click the Public IP address value, and then click the **[Dissociate]** button.

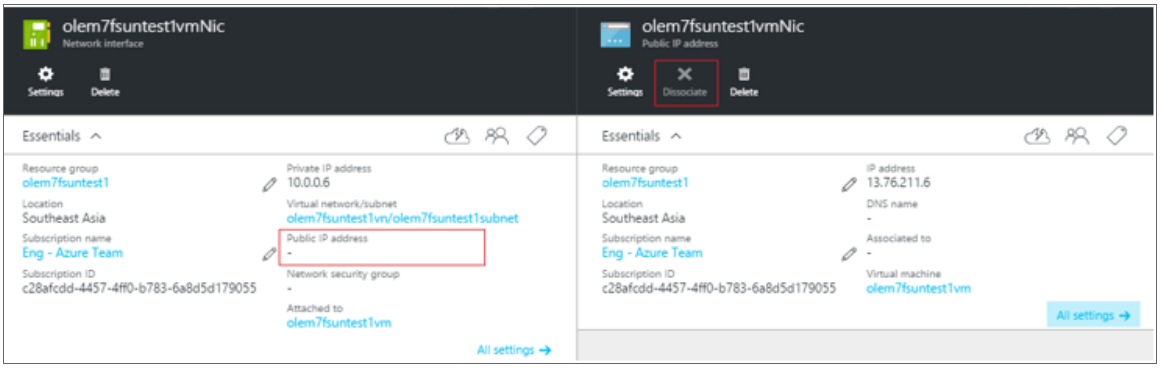

5. In the Network interface pane, click on **Settings > IP addresses**. Then, do one of the following:
   - If you **are not** using a VPN, complete steps 6 through 9. Ignore step 10.
   - If you **are** using a VPN, skip ahead to step 10.

6. If you **are not** using a VPN, then in the IP addresses pane, select *Enabled* in the ***Public IP address*** field and then click on the ***IP address*** field.
7. Click the **[Create new]** button.
8. In the Create public IP address pane, type a name for your IP address in the ***Name*** field and select *Static* in the ***Assignment*** field.
9. Click **[OK]** and then click **[Save]**.
10. If you **are** using a VPN, then in the IP addresses pane, select *Disabled* in the ***Public IP address*** field and then select a ***Subnet***. You can use the default values for all other fields.

Configuring Ports on SL1 Appliances

You must next create a Network security group that will specify the ports required for communication between the SL1 appliances and that will specify the ports required for communication between the SL1 appliances and the monitored devices in your network.

To configure the ports for communication:

1. In the Azure Portal, navigate to the Network security groups pane, and then click the **[Add]** button. The Create network security group pane appears.

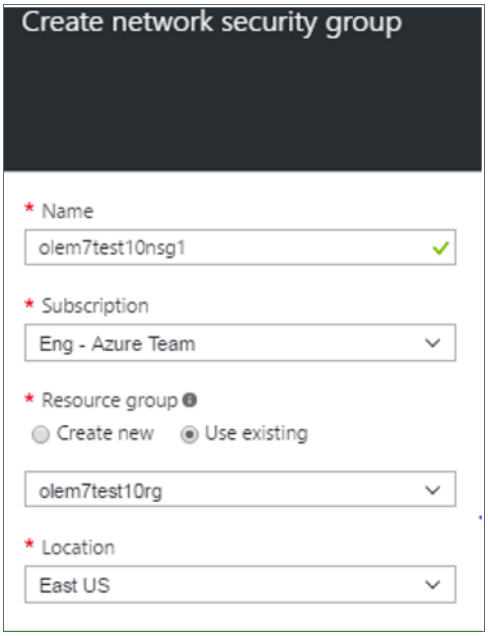

2. Type the information for the Network security group (name, subscription, resource group, and location), then click **[Create]**.
3. In the Network security groups pane, click the newly created Network security group, and then click the **[Settings]** button.
4. In the Settings pane, click **Inbound security rules**.

5. In the Inbound security rules pane, click the **[Add]** button. The Add inbound security rule pane appears.

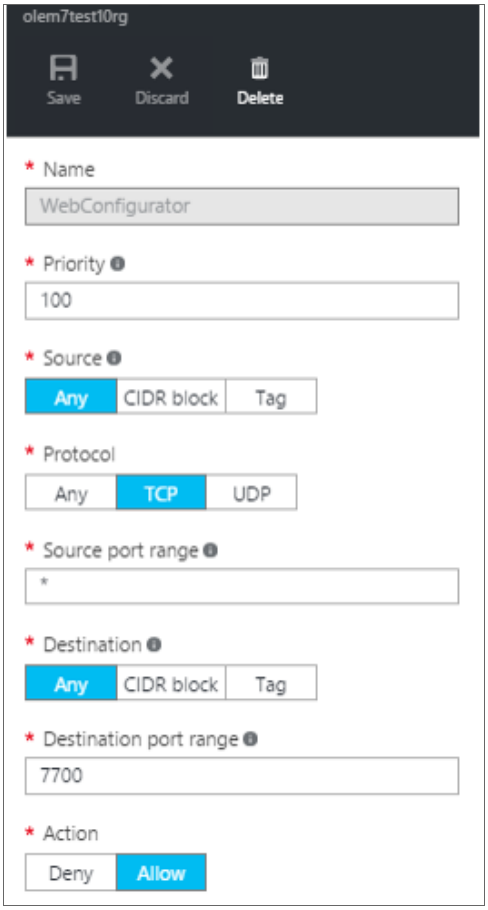

6. Use the table below to create security rules.
7. Repeat steps 5 and 6 to create an inbound security rule for each of the ports listed in the table below.
8. After creating all of the inbound security rules, navigate to the Virtual machine pane and click the **[Settings]** button.
9. In the Settings pane, click **Network interfaces** and then click on the name of the Network interface.
10. In the Network interface pane, click the **[Settings]** button and then click **Network security group**.
11. Select the newly created network security group to associate it with the Network interface.
12. Perform steps 9-11 for each network interface in your SL1 system.

NOTE: ScienceLogic recommends that you limit the Source port range for security reasons.

Type	Protocol	Port	Description
SSH	TCP	22	SSH. This is necessary to start the installation wizard.
SMTP	TCP	25	Necessary to receive inbound email for tickets, events, and email round-trip

Type	Protocol	Port	Description
			monitoring.
HTTP	TCP	80	HTTP from browser session or user workstation.
Custom TCP Rule	TCP	123	NTP. Communication between the All-In-One Appliance and configured NTP server.
Custom UDP Rule	UDP	161	SNMP Agent. Allows SNMP information about the SL1 appliance to be collected by SL1.
Custom UDP Rule	UDP	162	SNMP Traps. Necessary to receive SNMP traps from managed devices.
HTTPS	TCP	443	HTTPS from browser session or user workstation.
Custom UDP Rule	UDP	514	Syslog messages. Necessary to receive syslog messages from managed devices.
Custom TCP Rule	TCP	7700	ScienceLogic Web Configurator. Configuration Utility from browser session or user workstation. This is necessary to license the appliance.
Custom TCP Rule	TCP	7705	ScienceLogic PhoneHome. See *Configuring SL1 for PhoneHome Communications*.
Custom TCP Rule	TCP	7706	MySQL. Communication from Administration Portal.
Custom TCP Rule	TCP	7707	Data Pull. Allows the Database Server to retrieve data from the SL1 appliance.
Custom TCP Rule	TCP	8008	Administrative Web Interface (PHPMyAdmin) from browser session on user workstation.

Configuring the Virtual Machine

To configure each virtual machine, perform the following steps:

1. Use SSH to access the virtual machine using its public IP address and the username and password that were defined in step 2 of the section *Creating an Azure Virtual Machine*.
2. Run em7_install.sh in a special operational mode:

```
sudo /opt/em7/share/scripts/em7_install.sh --instance-init-only
```

3. On the Administration Portal (and the Database Server only if you are using the Administration Portal on the Database Server), run the following command to start the web server:

```
sudo service nginx start
```

4. Using vi or another text editor, edit the /etc/silo.conf file

```
sudo vi /etc/silo.conf
```

5. In the LOCAL section, set ipaddress to the Azure virtual machine's public IP address. If a VPN is used, however, set the ipaddress field to the VM's private IP address. For example, see the bolded text below:

```
[LOCAL]
rootdir = /opt/em7
vardir = /var/lib/em7
logdir = /var/log/em7
rundir = /run/em7
ipaddress = 172.16.10.10
dbdir = /data/db
dbserver = 172.16.10.11
dbport = 7706
dbuser = root
dbpasswd = em7admin
portcheck = /usr/bin/nmap
model_type = 1
eventmanager = internal,email,syslog,trap,dynamic
```

6. Perform the required steps in the Web Configuration Tool.

   - For instances of the ***Database Server*** or ***All-In-One Appliance***:
     - *Licensing the instance in the Web Configuration Tool*

NOTE: Upon installation, SL1 appliances are automatically licensed for 30 days. During these 30 days, you can perform the steps to obtain a permanent license from ScienceLogic.

   - For instances of the ***Administration Portal***:
     - *Configuring the instance in the Web Configuration Tool*
     - When prompted for the IP address of the Database Server,
       - If you are not using a VPN, enter the public IP address of the Database Server.
       - If you are using a VPN, use the private IP address of the Database Server.
   - For instances of the ***Data Collector and Message Collector***:
     - *Configuring the instance in the Web Configuration Tool*
     - When prompted for the IP address of the Database Server,
       - If you are not using a VPN, enter the public IP address of the Database Server.
       - If you are using a VPN, use the private IP address of the Database Server. .

7. Open a browser session to SL1 (to the Administration Portal). Go to the **Appliance Manager** page (System > Settings > Appliances).
8. If you are using an All-In-One Appliance, you will see two entries for the All-In-One Appliance.

- Select the bomb icon for the All-In-One Appliance for which the bomb icon () is enabled.
- In the remaining entry, select the wrench icon (). In the top pane, enter the IP Address specified in Azure for the All-In-One Appliance. Click **[Save]**.

9. If you are using a distributed system, you will see two entries for the Database Server.
   - Select the bomb icon for the Database Server for which the bomb icon () is enabled.

Additional Steps for SL1 10.1

SL1 10.1.x includes an upgrade to MariaDB. The upgrade did not include a tool, jemalloc, that helps manage memory usage.

NOTE: This section applies only to the following releases:

- 10.1.0
- 10.1.1
- 10.1.2
- 10.1.3
- 10.1.4
- 10.1.4.1
- 10.1.4.2
- 10.1.5
- 10.1.5.1

For SL1 versions later than 10.1.5.1, jemalloc is included with the platform. For SL1 versions prior to 10.1.0, jemalloc is included with the platform.

To avoid problems with memory usage on Database Servers, perform the following steps after upgrading MariaDB for 10.1.x.

NOTE: Perform these steps first on the active Database Server and then on each additional Database Server in your SL1 System.

1. Open an SSH session to the Database Server.
2. To verify that the Database Server is not currently running jemalloc, enter the following at the shell prompt:

```
silo_mysql -e 'show global variables like "version_malloc_library"'
```

If the Database Server is not currently running jemalloc, the shell will display the following:

Variable Name	Value
version_malloc_library	system

3. Search for the file /usr/lib64/**libjemalloc.so.1**.

   If the file does not exist, contact ScienceLogic Customer Support to request the file jemalloc-3.6.0-1.el7.x86_64.rpm.

   To install the RPM, use a file-transfer utility, copy the file to a directory on the SL1 appliance. Then enter the following at the shell prompt:

   ```
   cd /usr/lib64
   sudo yum install jemalloc-3.6.0-1.el7.x86_64.rpm
   ```

4. Create the file /etc/systemd/system/mariadb.service.d/jemalloc.conf:

   ```
   vi /etc/systemd/system/mariadb.service.d/jemalloc.conf
   ```

5. Add the following lines to the file:

   ```
   [Service]
   Environment="LD_PRELOAD=/usr/lib64/libjemalloc.so.1"
   ```

6. Save and close the file.
7. Reload the systemd config files:

   ```
   sudo systemctl daemon-reload
   ```

8. Restart the Database Server:

   To restart the ***standalone Database Server*** or the ***primary Database Server in a cluster***, enter the following:

   ```
   sudo systemctl restart mariadb
   ```

   To restart each ***secondary Database Server in a cluster***:

   a. Open an SSH session to the secondary Database Server. At the shell prompt, enter:

      ```
      coro_config
      ```

   b. Select **1**.
   c. When prompted to put the Database Server into maintenance, select **y**.
   d. Open an SSH session to the primary Database Server. To pause SL1, enter the following at the shell prompt:

      ```
      sudo touch /etm/.proc_mgr_pause
      ```

   e. In the SSH session for the secondary Database Server, restart MariaDB:

      ```
      crm resource restart mysql
      ```

   f. After MariaDB has restarted successfully on the secondary Database Server, return to the SSH session on the primary Database Server. Remove the pause file for SL1:

      ```
      sudo rm /tmp/.proc_mgr_pause
      ```

   g. In the SSH session on the secondary Database Server, take the Database Server out of maintenance. At the shell prompt, enter:

      ```
      coro_config
      ```

   h. Select **1**.
   i. When prompted to take the Database Server out of maintenance, select **y**.

9. To verify that jemalloc is running on the Database Server, enter the following at the shell prompt:

```
silo_mysql -e 'show global variables like "version_malloc_library"'
```

If the Database Server is currently running jemalloc, the shell will display something like the following:

Variable Name	Value
version_malloc_library	jemalloc 3.6.0-0-g46c0af68bd248b04df75e4f92d5fb804c3d75340

10. Perform these steps on each Database Server in your SL1 system.

Troubleshooting

If the Data Collector continuously displays a message saying the collector is working when running a Dynamic Application, **DO NOT** restart the Azure virtual machine, as doing so could cause you to lose SSH access to the machine.

Instead, do the following:

1. Using the command line interface, verify whether you can run the Dynamic Application in debug mode by typing the following command:

```
sudo /usr/local/silo/proc/dynamic_single.py <did> <app_id>
```

2. Restart the data pull processes (em7_hfpulld, em7_lfpulld, em7_mfpulld) by typing the following command:

```
sudo service <service_name> restart
```

Chapter 10

Updating SL1

Overview

For information on updating an existing SL1 system, see the manual ***Updating, Monitoring, and Maintaining SL1***. The *Updating, Monitoring, and Maintaining SL1* manual describes how to update the software on your SL1 appliances.

Contact ScienceLogic to get access to the ***Updating, Monitoring, and Maintaining SL1*** manual.

© 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

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

International: +1-703-354-1010