



Installation and Initial Configuration

Skylar One version 12.5.3

Table of Contents

Introduction to Installing Skylar One	9
What is Skylar One?	10
Database Functions	10
User Interface	10
Data Collection	11
Message Collection	11
What is Skylar One Extended?	12
Computing	12
Load Balancing	12
Storage	13
Management	13
The Skylar One Agent	13
Third-Party Software	13
Preparing Hardware Appliances for Skylar One	14
Hardware Specifications	15
Prerequisites for Skylar One Hardware Appliances	15
Initial Configuration for Skylar One Hardware Appliances	15
Changing the Password for em7admin	15
Changing Network Settings	16
Ports for Skylar One Hardware Appliances	19
Preparing Virtual Machines for Skylar One	20
Virtual Machine Specifications	21
Ports for Virtual Appliances	21
Build Nodes or Appliances in This Order	21
Deploying a Node or Appliance on a VMware System	21
Deploying a Node or Appliance on a Microsoft Hyper-V System	22
Deploying a Node or Appliance on a Nutanix AHV System	23
Required Ports for Skylar One	24
Open Ports on the ScienceLogic All-In-One Appliance	25
Open Ports on the ScienceLogic Database Server Appliance	26
Open Ports on the ScienceLogic Administration Portal Appliance	27

Open Ports on the ScienceLogic Data Collector Appliance	28
Open Ports on the ScienceLogic Message Collector Appliance	29
Open Ports for ScienceLogic Subscription Billing	29
Open Ports for ScienceLogic PowerPacks	30
Apcon	30
Cisco: Cloud Services Platform	30
Cisco: Contact Center Enterprise	30
Cisco: CUCM	31
Cisco: ESA	31
Cisco: Meeting Server	31
Cisco: UC Ancillary	32
Cisco: UC VOS Applications	32
Cisco: UCS	32
Cisco: UCS Director	32
Cisco: UCS Standalone Rack Server	32
Cisco: Viptela	33
Citrix: Xen	33
Dell EMC: VMAX	33
Dell EMC: VNX	33
Dell EMC: XtremIO	33
ELK: AWS CloudTrail	33
ELK: Azure Activity Log	34
Hitachi Data Systems: VSP	34
HP 3PAR: SMI-S	34
IBM: AIX Monitoring	34
Kubernetes	34
Linux: Base Pack	35
Linux: SSH Automations	35
Microsoft: Azure	35
Microsoft: SQL Server Enhanced	35
Microsoft: Automation PowerPacks	35
Mongo DB	36

Monitoring Switches, Routers, and Firewalls with SNMP	36
Monitoring Windows Systems with PowerShell	36
Monitoring Windows Systems with WMI	37
MySQL	37
NetApp Base Pack	37
OpenStack	38
Oracle: Database	38
Palo Alto	38
Pure Storage: Flash Array	38
Restorepoint Automation PowerPack	38
Skylar One PowerFlow	38
SMI-S: Array	39
SoftLayer: Cloud	39
VMware: NSX	39
VMware: NSX-T	39
Installing Skylar One on Hardware Appliances and Virtual Appliances	40
Prerequisites	42
Workflow for Installing and Configuring an SL1 Collector	42
Downloading the ISO Image	43
Installing the Database Server	43
Installing an Administration Portal or SL1 Collector	44
Installing an Administration Portal or SL1 Collector	44
Licensing New Skylar One Appliances	45
Configuring a New Skylar One System for Traditional Communication	46
What is Traditional Communication?	46
Configuring a New SL1 Collector for Traditional Communication	47
Configuring Traditional Database Initiates System Accepts Communication	48
Managing the Nodes Page	51
Viewing the List of Registered Nodes	52
Viewing the Tokens on the Nodes Page	53
Recreating a Token	54
Licensing and Configuring an Appliance	55

Navigating the Classic Web Configuration Utility	57
Logging into the Classic Web Configuration Utility	57
Changing the Password for the Classic Web Configuration Utility	57
Licensing and Configuring a Database Server or All-In-One Appliance	58
Configuring an Administration Portal	58
Configuring a Data Collector or Message Collector	60
Other Initial Configuration Tasks	60
Configuring Logging for a Skylar One System	60
Defining the NTP Server	60
Creating a Bonded Interface from the Web Configuration Utility	61
Defining a Proxy Server from the Appliance Manager Page	62
Navigating the Node Configuration Utility	63
Logging into the Node Configuration Utility	63
Changing the Password for the Node Configuration Utility	63
Viewing the Collector Connection Status	64
Configuring the Proxy Server from the Node Configuration Utility	64
Adding a Bonded Interface from the Node Configuration Utility	65
Editing an Interface from the Node Configuration Utility	66
Configuring Skylar One for PhoneHome Communication	67
What is PhoneHome Communication?	69
Important Notes about PhoneHome Communication	70
Prerequisites for Configuring PhoneHome Communication	72
Overview of the PhoneHome Configuration	72
Configuring the Database Server for PhoneHome Communication	73
Before Configuring the Database Server for PhoneHome Communication	73
Understanding Database Server PhoneHome Configuration Options	74
Configuring a Single Database Server	74
Configuring a Database with a Non-default Address or Port	75
Configuring a Database with Multiple IP Addresses	76
Configuring PhoneHome Database Servers for High Availability and Disaster Recovery	76
Managing Proxy Connections for PhoneHome Communication	77
Adding a Proxy Configuration	78

Viewing a List of Proxy Connections	79
Deleting a Proxy Configuration	79
Configuring SL1 Collectors for PhoneHome Communication	80
Configuring Collector Initiates System Accepts Communication	81
Configuring Collector Initiates User Accepts Communication	83
Connecting a Skylar One Collector to the Skylar One Database Server using the Command-line Interface	85
System Accepted	86
User Accepted	86
Understanding PhoneHome Components	87
phd	87
phc	87
Using the Command-Line Interface for PhoneHome Collection	87
Viewing a List of PhoneHome Devices	90
Viewing Information about a Single PhoneHome Device	90
Renaming a PhoneHome Device	90
Checking the Status of a PhoneHome Collector	91
Syncing the Configuration of a PhoneHome System	91
Managing Port Forwarding for PhoneHome Communication	92
Viewing a List of Port Forwards	92
Adding a Port Forward	92
Removing a Port Forward	93
Managing Destinations	93
Viewing a List of Destinations	93
Adding a Destination Address	94
Removing a Destination	95
Logging PhoneHome Configuration Information	95
Tuning PhoneHome Settings	95
Viewing a List of Current PhoneHome Settings	96
Updating PhoneHome Settings	96
Clearing a PhoneHome Device	97
Deleting a PhoneHome Collector	98

Deleting a PhoneHome Database Server	99
Troubleshooting PhoneHome Configurations	100
Connectivity Issues from a Collector	100
ssh: handshake failed: ssh: unable to authenticate, attempted methods [none publickey], no supported methods remain	100
ssh: handshake failed: knownhosts: key mismatch	101
dial TCP <database_host_addr>:<port>: i/o timeout	101
dial TCP <database_host_addr>:<port>: connect: no route to host	101
dial TCP <database_host_addr>:<port>: connect: connection refused	101
Register Command Complains that the Token Has Expired	101
You Cannot See a Request You Sent on the Server and You Cannot Send Another Request ..	101
Status Shows Disconnected but the Check Succeeds	102
Installing Skylar One on AWS	103
AWS Instance Specifications	104
Deploying a Skylar One System on AWS	104
What are the ScienceLogic AMIs?	104
Getting the ScienceLogic AMI	105
Launching the New Instance	106
Security Rules for Each Appliance Type	110
All-In-One Appliance	110
Database Server	111
Administration Portal	113
Data Collector	115
Message Collector	116
Additional Configuration Steps	117
Assigning an EIP to the New Instance	118
Accessing the Appliance Using SSH	119
Gathering Information Required for Accessing the Appliance Using SSH	119
Configuring SSH	119
Connecting to Your Instance	119
Configuring the EC2 Instance	120
Web Configuration Tool	121

Rebooting Data Collectors and Message Collectors	122
Installing Skylar One in Azure	123
Azure Instance Specifications	125
Prerequisites	125
Configuring an Azure Resource Group and Storage Account	125
Generating a Source Blob SAS URL	125
Transferring Skylar One VHD Files to Your Azure Blob Container	126
Creating the Azure Image	127
Skylar One Appliance Deployment Order for Distributed Systems	127
Creating an Azure Virtual Machine	128
Setting the Public IP Address to Static	129
Configuring the Virtual Machine	129
Navigating the Setup and Config Page	131
What is the Setup and Config Page?	132
Setup and Config Journey Workflows	132
Taking a Tour of Skylar One	133
Discover and Monitor Hybrid Cloud Infrastructure	134
Onboarding your Hybrid Cloud Infrastructure	134
Resetting a Completed Journey Workflow	136
Setting a Journey Workflow as "Not Applicable"	136
Updating Skylar One	138

Chapter

1

Introduction to Installing Skylar One

Overview

This manual describes how to install and configure Skylar One (formerly SL1).

This chapter covers the following topics:

<i>What is Skylar One?</i>	10
<i>What is Skylar One Extended?</i>	12
<i>The Skylar One Agent</i>	13
<i>Third-Party Software</i>	13

What is Skylar One?

Skylar One offers you the capabilities to monitor your hybrid cloud infrastructure, improve service visibility, and automate your IT workflows.

In a **Distributed** system, there are four general functions that a Skylar One appliance can perform:

- Database functions
- User interface
- Data collection
- Message collection

In large Skylar One systems, dedicated nodes or appliances perform each function. In smaller systems, some nodes or appliances perform multiple functions. In an **All-In-One Appliance** system, a single Skylar One node or appliance performs all four functions.

Database Functions

The node or appliance that provides the database functions 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 nodes or appliances responsible for collecting data and collecting messages.
- Processing and normalizing collected data.
- Allocating tasks to the other nodes or appliances in the Skylar One 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:

- **Database Server.** A dedicated **Database Server** provides all database functions.
- **All-In-One Appliance.** An **All-In-One Appliance** performs all 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 node or appliance that provides the user interface also generates all scheduled reports and provides access to the ScienceLogic API. The following nodes or appliances provide the user interface:

- **Administration Portal.** A dedicated **Administration Portal** node or appliance can provide the user interface.

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

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

Data Collection

In a distributed system, nodes or appliances retrieve data from monitored devices and perform some pre-processing of collected data and execute automation actions.

The following appliances can perform the data collection function:

- **Data Collector.** One or more Data Collectors are 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.
- **All-In-One Appliance.** An **All-In-One Appliance** performs all functions.

NOTE: The Skylar One Agent can also be used to collect data from devices on which it can be installed. See the [System Requirements](#) page of the ScienceLogic 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 Skylar One Agent, or using a combination of both.

Message Collection

In a distributed system, nodes or appliances receive and process inbound, asynchronous syslog and trap messages from monitored devices.

The following nodes or appliances can perform the message collection function:

- **Message Collector.** A dedicated **Message Collector** receives and processes inbound, asynchronous syslog and trap messages from monitored devices.

NOTE: In distributed systems that use the Skylar One agent, the Message Collector passes agent data to the Database Server. On these distributed systems, the **Message Collector** must be a standalone node or appliance, not a combination **Data Collector/Message Collector**.

- **Data Collector.** A Data Collector can also perform the message collection function in addition to data collection.

- **All-In-One Appliance.** An **All-In-One Appliance** performs all functions.

What is Skylar One Extended?

The **Skylar One Extended Architecture** includes additional types of Skylar One nodes or appliances. The following Skylar One features require the Skylar One Extended Architecture:

- **Expanded Agent Capabilities.** You can configure the Skylar One agent to communicate with Skylar One via a dedicated Message Collector. However, this configuration limits the capabilities of the Skylar One agent. If you configure the Skylar One agent to communicate with Skylar One via a Compute Cluster, you expand the capabilities of the Skylar One 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 Skylar One. Data pipelines introduce message queues and communicate using encrypted web services.
- **Publisher.** Publisher enables the egress of data from Skylar One. Publisher can provide data for long-term storage or provide input to other applications to 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. Skylar One 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.

Skylar One Extended Architecture includes the following additional Skylar One functions:

Computing

Skylar One Extended includes a **Compute Cluster** that includes a minimum of three Compute Nodes. Compute nodes are the Skylar One appliances that transport, process, and consume the data from Data Collectors and the Skylar One Agent. Skylar One uses Docker and Kubernetes to deploy and manage these services. The compute node sends configuration data to the Database Server and performance data to the Storage Node cluster.

Load Balancing

A **load balancer** is the Skylar One node or 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

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

The Skylar One Agent

The **Skylar One agent** is a program that you can install on a device monitored by Skylar One (formerly SL1). There is a Windows agent, an AIX agent, a Solaris agent, and a Linux agent. The agent collects data from the device and pushes that data back to Skylar One.

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

You can configure an agent to communicate with either the Message Collector or the Compute Cluster.

NOTE: The following minimum agent versions are required for Skylar One 12.5.1 and later:

- **Windows** version 154
- **Linux** version 196
- **AIX** version 196
- **Solaris** version 196

Users who require agent-based log collection on a device with a Windows agent or a Linux agent must have the minimum Windows agent or Linux agent version. If you do not have the minimum required agent versions, ScienceLogic recommends that you upgrade using the **[Upgrade]** button on the **Agents** page (Devices > Agents), or by downloading and upgrading the agent manually. For more information, see the section on [Upgrading an Agent](#).

Third-Party Software

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

Chapter

2

Preparing Hardware Appliances for Skylar One

Overview

This chapter describes how to prepare hardware appliances before installing Skylar One (formerly SL1).

Use the following menu options to navigate the Skylar One user interface:

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

This chapter covers the following topics:

<i>Hardware Specifications</i>	15
<i>Prerequisites for Skylar One Hardware Appliances</i>	15
<i>Initial Configuration for Skylar One Hardware Appliances</i>	15
<i>Ports for Skylar One Hardware Appliances</i>	19

Hardware Specifications

For details about supported ScienceLogic hardware appliances, see the [System Requirements](#) page on the ScienceLogic Support Site.

Prerequisites for Skylar One Hardware Appliances

Perform the following steps to prepare a Skylar One appliance for configuration:

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

Initial Configuration for Skylar One Hardware Appliances

You must perform the following tasks during initial configuration of a Skylar One hardware appliance shipped by ScienceLogic:

- [Change the password for the administrative user `em7admin`.](#)
- [Change the network settings for the appliance.](#) This includes changing the following:
 - The IP address for the network gateway; you must have already allocated IP addresses for the Skylar One appliances
 - The primary IP address of the appliance
 - The Netmask for the primary IP address of the appliance
 - 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 Skylar One appliance or use SSH to access the server.
2. Log in as user `em7admin` with the current password.
3. At the shell prompt, type the following:

```
passwd
```

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

TIP: You can use the following special characters in the ***em7admin*** user account 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 Skylar One 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
```

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
```

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

6. 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-b3f3-32f61dccb8
```

```
DEVICE=ens32
```

```
ONBOOT=yes
```

```
IPADDR=10.64.68.20
```

```
PREFIX=24
```

```
GATEWAY=10.64.68.1
```

```
IPV6_PEERDNS=yes
```

```
IPV6_PEERROUTES=yes
```

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

```
sudo service network restart
```

Ports for Skylar One Hardware Appliances

See the chapter on [ports](#) to configure firewalls to allow traffic to and from the Skylar One appliances.

Chapter

3

Preparing Virtual Machines for Skylar One

Overview

This chapter describes how to prepare virtual nodes or appliances before installing Skylar One (formerly SL1).

Use the following menu options to navigate the Skylar One user interface:

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

This chapter covers the following topics:

<i>Virtual Machine Specifications</i>	21
<i>Ports for Virtual Appliances</i>	21
<i>Build Nodes or Appliances in This Order</i>	21
<i>Deploying a Node or Appliance on a VMware System</i>	21
<i>Deploying a Node or Appliance on a Microsoft Hyper-V System</i>	22
<i>Deploying a Node or Appliance on a Nutanix AHV System</i>	23

Virtual Machine Specifications

You can deploy Skylar One appliances on hypervisors from the following vendors:

- VMware
- Microsoft
- Nutanix
- Citrix
- RedHat

For details about the hypervisor versions that are supported and the requirements and specifications for each Skylar One appliance, see the [System Requirements](#) page on the ScienceLogic Support Site.

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

Ports for Virtual Appliances

See the chapter on [ports](#) to configure firewalls to allow traffic to and from the Skylar One appliances.

Build Nodes or Appliances in This Order

For ease of configuration, create nodes or appliances in this order:

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

Deploying a Node or Appliance on a VMware System

NOTE: The following instructions describe how to configure a ScienceLogic virtual machine in VMWare. If you are looking for resources and support for VMWare, see the VMWare Marketplace: <https://marketplace.cloud.vmware.com/>.

To deploy a Skylar One node or appliance on a VMware system:

1. If you have not already done so, [download the ISO file for Skylar One](#).
2. 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 [System Requirements](#) page on the ScienceLogic Support Site.
4. On the **Guest Operating System** page, select *Linux* as the **Guest Operating System**, and then select *Oracle Linux 8 (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 on the [System Requirements](#) page on the ScienceLogic Support Site.
 - Configure the CD/DVD drive to use the Skylar One ISO file.
7. Turn on the virtual machine and boot the VM from the Skylar One ISO file to begin [installing the Skylar One Database Server](#).
8. Repeat these steps for each node or appliance in your Skylar One system.

Deploying a Node or Appliance on a Microsoft Hyper-V System

To deploy a Skylar One node or appliance on a Microsoft Hyper-V system:

1. If you have not already done so, [download the ISO file for Skylar One](#).
2. Follow the instructions from Microsoft on creating a virtual machine in Hyper-V:
<https://learn.microsoft.com/en-us/windows-server/virtualization/hyper-v/get-started/create-a-virtual-machine-in-hyper-v>
3. When prompted to **Specify a Generation** for the VM, make one of the following selections:
 - **Generation 1**. Fully supports Oracle Linux and Skylar One.
 - **Generation 2**. To support Oracle Linux and Skylar One, you must disable the "secure boot" feature.
4. When prompted to **Assign Memory** and **Connect Virtual Hard Disk**, enter the hardware requirements as specified on the [System Requirements](#) page on the ScienceLogic Support Site.
5. On the **Installation Options** page, select **Install an operating system later**.
6. Complete the rest of the VM creation steps based on your specific needs. When you are done, click **[Finish]**.

7. If you selected a **Generation 2** virtual machine in step 3, 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
```

8. Follow the instructions from Microsoft to install the operating system (Oracle Linux 8 64 bit):
<https://learn.microsoft.com/en-us/windows-server/virtualization/hyper-v/get-started/create-a-virtual-machine-in-hyper-v>
9. In the Hyper-V Manager, right-click the VM and select *Connect*.
10. In the **Virtual Machine Connection** window, click **[Action]** and then select *Start*.
11. Repeat these steps for each node or appliance in your Skylar One system.
12. To install Skylar One on the Hyper-V virtual machine, see *Installing Skylar One on Hardware Appliances and Virtual Appliances*.

Deploying a Node or Appliance on a Nutanix AHV System

To deploy a Skylar One node or appliance on a Nutanix AHV virtual machine (VM):

1. If you have not already done so, *download the ISO file for Skylar One*.
2. Review Nutanix's best practices for deploying a Nutanix AHV VM: <https://portal.nutanix.com/page/documents/solutions/details?targetId=BP-2029-AHV:vm-deployment.html>

NOTE: The AHV VM must be configured to operate on Oracle Linux 8.

3. Follow the instructions from Nutanix for creating an AHV VM through the Prism Element web console: https://portal.nutanix.com/page/documents/details?targetId=AHV-Admin-Guide-v6_5:wc-vm-create-acropolis-wc-t.html
4. When creating the VM, in the **Compute Details** section, set the CPU and memory configuration allocations to the values recommended on the *System Requirements* page on the ScienceLogic Support Site for your environment and the appliance type you are installing.
5. In the **Disks** section, configure the disks to use the Skylar One ISO file.
6. In the **Network Adapters (NIC)** section, add a network adapter that enables the VM to connect to the internet.
7. When you are finished creating the AHV VM, click **[Save]**. Nutanix begins creating the VM.
8. After the VM is created, select the VM from the Nutanix console, then power it on and boot the VM from the Skylar One ISO file to begin *installing the Skylar One Database Server*.
9. Repeat these steps for each node or appliance in your Skylar One system.

Chapter

4

Required Ports for Skylar One

Overview

This chapter describes the ports that must be open on each Skylar One (formerly SL1) appliance. These open ports allow communication between appliances in a Skylar One system.

Some PowerPacks also require specific ports to be open for tasks such as monitoring, creating credentials, or gaining access through the firewall. Those ports are also described in this chapter.

Use the following menu options to navigate the Skylar One user interface:

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

This chapter covers the following topics:

Open Ports on the ScienceLogic All-In-One Appliance	25
Open Ports on the ScienceLogic Database Server Appliance	26
Open Ports on the ScienceLogic Administration Portal Appliance	27
Open Ports on the ScienceLogic Data Collector Appliance	28
Open Ports on the ScienceLogic Message Collector Appliance	29
Open Ports for ScienceLogic Subscription Billing	29
Open Ports for ScienceLogic PowerPacks	30

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
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. NOTE: For Military Unique Deployment (MUD) configurations, this utility and port are disabled by default. They can be enabled for initial configuration, but must be disabled again after the configuration process is complete.	TCP	7700
SNMP	Optional. SNMP information about the All-In-One Appliance can be collected by Skylar One.	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

Name	Description	Protocol	Port
SMTP	Optional. To receive inbound Email for tickets, events, and email round-trip monitoring.	TCP	25
NTP	Communication between the All-In-One Appliance and configured NTP server.	TCP	123

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. NOTE: For Military Unique Deployment (MUD) configurations, this utility and port are disabled by default. They can be enabled for initial configuration, but must be disabled again after the configuration process is complete.	TCP	7700
SNMP	Optional. SNMP information about the Database Server can be collected by Skylar One.	UDP	161

Name	Description	Protocol	Port
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
DRBD Replication	This port is open only if your Database Server appliance is configured for High Availability, Disaster Recovery, or both.	TCP	7789
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, but only for non-SaaS systems.	TCP	7705
EKMS Cluster Communication	If there is a firewall between the Database Server, Data Engine, and Administration Portal appliances, this port must be open to enable Enterprise Key Management Service (EKMS) cluster communication between those appliances.	TCP	8200

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

Name	Description	Protocol	Port
Web Configurator	Configuration Utility from browser session on user workstation. NOTE: For Military Unique Deployment (MUD) configurations, this utility and port are disabled by default. They can be enabled for initial configuration, but must be disabled again after the configuration process is complete.	TCP	7700
SNMP	Optional. SNMP information about the Administration Portal can be collected by Skylar One.	UDP	161
High Availability	Required when using Quorum with High Availability or High Availability and Disaster Recovery.	TCP	5403
EKMS Cluster Communication	If there is a firewall between the Database Server, Data Engine, and Administration Portal appliances, this port must be open to enable Enterprise Key Management Service (EKMS) cluster communication between those appliances.	TCP	8200

Open Ports on the ScienceLogic Data Collector Appliance

Name	Description	Protocol	Port
Data Pull	Requests from Database Servers to retrieve collected data. In a PhoneHome 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. NOTE: For Military Unique Deployment (MUD) configurations, this utility and port are disabled by default. They can be enabled for initial configuration, but must be disabled again after the configuration process is complete.	TCP	7700
SNMP	Optional. SNMP information about the Data Collector can be collected by Skylar One.	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

Name	Description	Protocol	Port
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 PhoneHome 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. NOTE: For Military Unique Deployment (MUD) configurations, this utility and port are disabled by default. They can be enabled for initial configuration, but must be disabled again after the configuration process is complete.	TCP	7700
SNMP	Optional. SNMP information about the Message Collector can be collected by Skylar One.	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 for ScienceLogic Subscription Billing

Name	Description	Protocol	Port
HTTPS Secure Interface	Required. Secure interface must be open for the Subscription Billing process to send information to ScienceLogic.	TCP	443

Open Ports for ScienceLogic PowerPacks

ScienceLogic recommends reviewing the required port data for each PowerPack below. Some PowerPacks require specific ports for tasks such as monitoring, creating credentials, or gaining access through the firewall.

NOTE: Many PowerPacks can be configured so that you can connect with the third-party platform via a proxy server. When you do so, you will need to open a port on that proxy server as well as in Skylar One to establish communication between the two platforms.

TIP: For more information about the configuration requirements for the PowerPacks below or other PowerPacks that are not included in this section, see the [Skylar One PowerPacks](#) documentation.

Apcon

Name	Description	Protocol	Port
SNMP	Required for SNMP credential.	UDP	161

Cisco: Cloud Services Platform

Name	Description	Protocol	Port
SNMP	Required for monitoring CSP clusters with SNMP community string read privileges, or if you have to create two SNMP credentials for CSP clusters.	UDP	161
SNMP	Required if you have to create two SNMP credentials for CSP clusters.	TCP	1610

Cisco: Contact Center Enterprise

Name	Description	Protocol	Port
REST API	Required for monitoring Contact Center Enterprise using REST API.	TCP	7890

Cisco: CUCM

Name	Description	Protocol	Port
PhoneHome Configuration	Potentially required based on your configuration. Requests from the PhoneHome Collector to the Database Server to retrieve collected data.	TCP	7707
PhoneHome Configuration	Potentially required based on your configuration. Requests from the Database Server to the Data Collector to retrieve collected data.	TCP	7705
SNMP	Potentially required based on your configuration. Enables communication between Skylar One Data Collector and the Cisco Unified CM cluster and CallManagers.	UDP	161
Cisco Unified Communications Manager	Potentially required based on your configuration. Enables communication between Skylar One Data Collector and the Cisco Unified CM cluster and CallManagers.	TCP	8443
<p>NOTE: The example credential included in older versions of the Cisco: CUCM Unified Communications Manager PowerPack used "80" as the default port number. If your Cisco Unified CM credential specifies port 80, Skylar One will automatically override that value and use port 8443 instead. If your Cisco Unified CM credential specifies any port other than 80, Skylar One will use that specified port.</p>			

Cisco: ESA

Name	Description	Protocol	Port
SNMP	Required for SNMP credential.	UDP	161

Cisco: Meeting Server

Name	Description	Protocol	Port
SNMP	Required for creating an SNMP credential for one IP address.	UDP	161

Name	Description	Protocol	Port
SSH	Required for creating a Basic/Snippet credential for one IP address or creating a Basic/Snippet credential on a system Mainboard Management Processor interface if monitoring more than one IP address.	TCP	22
HTTPS Secure Interface	Required for creating a Basic/Snippet credential for the API interface if monitoring more than one IP address.	TCP	443

Cisco: UC Ancillary

Name	Description	Protocol	Port
SSH	Required for SSH/Key credential.	TCP	22

Cisco: UC VOS Applications

Name	Description	Protocol	Port
Proxy Server	Used for proxy server port in SOAP/XML credential.	TCP	0
HTTPS Secure Interface	Required for creating a Basic/Snippet credential for REST API queries to Cisco Unity Connection servers and Cisco IM & Presence servers.	TCP	443

Cisco: UCS

Name	Description	Protocol	Port
HTTPS Secure Interface	Required for discovering UCS Manager over HTTPS.	TCP	443

Cisco: UCS Director

Name	Description	Protocol	Port
HTTP	Required for discovering UCS Director over HTTP.	TCP	80

Cisco: UCS Standalone Rack Server

Name	Description	Protocol	Port
HTTPS Secure Interface	Required for discovering UCS Rack Server over HTTPS.	TCP	443

Cisco: Viptela

Name	Description	Protocol	Port
HTTPS Secure Interface	Required for discovering Viptela over HTTPS.	TCP	443

Citrix: Xen

Name	Description	Protocol	Port
HTTPS Secure Interface	Required for the Citrix: Xen Basic/Snippet credential.	TCP	443

Dell EMC: VMAX

Name	Description	Protocol	Port
HTTP	Required for connecting to an SMI-S Provider over HTTP.	TCP	5988
HTTPS	Required for connecting to an SMI-S Provider over HTTPS.	TCP	5989

Dell EMC: VNX

Name	Description	Protocol	Port
HTTP	Required for connecting to an SMI-S Provider over HTTP.	TCP	5988
HTTPS	Required for connecting to an SMI-S Provider over HTTPS.	TCP	5989

Dell EMC: XtremIO

Name	Description	Protocol	Port
Proxy Server	Used for proxy server port in SOAP/XML credential.	TCP	0
HTTPS Secure Interface	Required for discovering Dell EMC XtremIO devices.	TCP	443

ELK: AWS CloudTrail

Name	Description	Protocol	Port
Elasticsearch	Required for the ELK: AWS Basic/Snippet credential.	TCP	9200

ELK: Azure Activity Log

Name	Description	Protocol	Port
Elasticsearch	Required for the ELK: Azure Activity Log Basic/Snippet credential.	TCP	9200

Hitachi Data Systems: VSP

Name	Description	Protocol	Port
HTTPS	Required for connecting to an SMI-S Provider over HTTPS.	TCP	5989

HP 3PAR: SMI-S

Name	Description	Protocol	Port
HTTPS	Required for connecting to an SMI-S Provider over HTTPS.	TCP	5989

IBM: AIX Monitoring

Name	Description	Protocol	Port
SSH	Required for SSH/Key credential.	TCP	22

Kubernetes

Name	Description	Protocol	Port
SSH	Typically used for connecting to Kubernetes nodes via SSH.	TCP	22
HTTPS	Can be used for connecting to Kubernetes cluster via HTTPS.	TCP	443
	Can be used for connecting to Kubernetes cluster via HTTPS. <div style="border: 1px solid black; padding: 10px; margin-top: 10px;">NOTE: If you would prefer to configure a customized IP port other than 443 or 8443, you can do so. For more information, see the section on Configuring Customized IP Ports in the Monitoring Kubernetes manual.</div>	TCP	8443

Linux: Base Pack

Name	Description	Protocol	Port
SSH	Required for SSH/Key credential.	TCP	22

Linux: SSH Automations

Name	Description	Protocol	Port
SSH	Required for SSH/Key credential.	TCP	22

Microsoft: Azure

Name	Description	Protocol	Port
SNMP	When using the run book automations included in the PowerPack to discover physical devices, allows the discovery session to use SNMP credentials.	UDP	161
PowerShell (HTTP)	When using the run book automations included in the PowerPack to discover physical devices, allows the discovery session to use PowerShell credentials over HTTP.	TCP	5985
PowerShell (HTTPS)	When using the run book automations included in the PowerPack to discover physical devices, allows the discovery session to use PowerShell credentials over HTTPS.	TCP	5986

Microsoft: SQL Server Enhanced

Name	Description	Protocol	Port
PowerShell (HTTP)	Required for users who want to connect to a SQL server using PowerShell credentials over HTTP.	TCP	5985

Microsoft: Automation PowerPacks

Name	Description	Protocol	Port
DNS Server	Required for forward and reverse DNS server availability for the Windows server.	TCP	53

Name	Description	Protocol	Port
Kerberos Authentication	Required for Kerberos authentication if using an Active Directory user account to access the Windows Domain Controller.	UDP	88
PowerShell (HTTP)	Required if connecting using PowerShell credentials over HTTP.	TCP	5985
PowerShell (HTTPS)	Required if connecting using PowerShell credentials over HTTPS.	TCP	5986

Mongo DB

Name	Description	Protocol	Port
MongoDB Server	Required when creating a MongoDB credential.	TCP	27017
SSH	Optional, but required if including SSH settings in the MongoDB credential.	TCP	22

Monitoring Switches, Routers, and Firewalls with SNMP

Name	Description	Protocol	Port
SNMP	Required for SNMP credential.	UDP	161

Monitoring Windows Systems with PowerShell

Name	Description	Protocol	Ports
SNMP	Required for SNMP credential	UDP	161
SNMP	At least one of the additional listed ports must be open on the device to discover SNMP-enabled Windows devices.	TCP	21, 22, 23, 25, or 80
DNS Server	Required for forward and reverse DNS server availability for the Windows server.	TCP	53
Kerberos Authentication	Required for Kerberos authentication if using an Active Directory user account to access the Windows Domain Controller.	UDP	88
PowerShell (HTTP)	Required if connecting using PowerShell credentials over HTTP.	TCP	5985
PowerShell (HTTPS)	Required if connecting using PowerShell credentials over HTTPS.	TCP	5986

Monitoring Windows Systems with WMI

Name	Description	Protocol	Ports
SNMP	Required for SNMP credential	UDP	161
SNMP	At least one of the additional listed ports must be open on the device to discover SNMP-enabled Windows devices.	TCP	21, 22, 23, 25, or 80
DNS Server	Required for forward and reverse DNS server availability for the Windows server.	TCP	53
WMI	Required for incoming network traffic to the remote machine.	TCP	135
WMI	Required for incoming network traffic to the remote machine.	TCP	445
WMI	In addition to ports 135 and 445, additional dynamically assigned ports must be open, typically in the listed ranges.	TCP	1025-5000, 49152-65535

MySQL

Name	Description	Protocol	Port
MySQL Server SSL Certificate	When configuring a SOAP/XML credential to support loading your SSL certificate on a database connection, you can specify one port or a range of ports. This will be based on your MySQL instance. For more information, see the section on Creating a SOAP/XML Credential for an SSL Certificate in the Monitoring MySQL manual.	N/A	N/A

NetApp Base Pack

Name	Description	Protocol	Port
HTTP (FIPS Mode)	Used for the NetAPP C-Mode appliance credential if Skylar One is running in FIPS-compliant mode.	TCP	80
SNMP	Required for SNMP credential.	UDP	161

OpenStack

Name	Description	Protocol	Port
Proxy Server	Used for proxy server port in SOAP/XML credential.	TCP	0

Oracle: Database

Name	Description	Protocol	Port
SSH	Required for SSH/Key credential for Linux users.	TCP	22
PowerShell (HTTP)	Required for Windows users who want to connect using PowerShell credentials over HTTP.	TCP	5985
PowerShell (HTTPS)	Required for Windows users who want to connect using PowerShell credentials over HTTPS.	TCP	5986

Palo Alto

Name	Description	Protocol	Port
SNMP	Required for SNMP credential.	UDP	161
HTTPS Secure Interface	Required for the Palo Alto Basic/Snippet credential.	TCP	443

Pure Storage: Flash Array

Name	Description	Protocol	Port
HTTPS Secure Interface	Required for discovering Pure Storage components over HTTPS or via API.	TCP	443

Restorepoint Automation PowerPack

Name	Description	Protocol	Port
SSH	Required for SSH/Key credential.	TCP	22

Skylar One PowerFlow

Name	Description	Protocol	Port
SSH	Required for SSH/Key credential.	TCP	22

SMI-S: Array

Name	Description	Protocol	Port
HTTPS	Required for connecting to an SMI-S Provider over HTTPS.	TCP	5989

SoftLayer: Cloud

Name	Description	Protocol	Port
HTTP	Required for discovering Softlayer: Cloud over HTTP.	TCP	80

VMware: NSX

Name	Description	Protocol	Port
HTTPS Secure Interface	Required for the VMware: NSX Basic/Snippet credential.	TCP	443

VMware: NSX-T

Name	Description	Protocol	Port
HTTPS Secure Interface	Required for the VMware: NSX-T Basic/Snippet credential.	TCP	443

Chapter

5

Installing Skylar One on Hardware Appliances and Virtual Appliances

Overview

This chapter describes how to install Skylar One (formerly SL1) on hardware appliances or virtual machines, including how to download the ISO image; install the Database Server, Administration Portal, and SL1 Collectors; and establish a connection between the new SL1 Collectors and the Database Server. An SL1 Collector can be either a Data Collector or a Message Collector.

NOTE: For detailed instructions on how to upgrade existing Skylar One deployments, see the section on [Updating Skylar One](#).

Use the following menu options to navigate the Skylar One user interface:

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

This chapter covers the following topics:

Prerequisites	42
Workflow for Installing and Configuring an SL1 Collector	42
Downloading the ISO Image	43
Installing the Database Server	43
Installing an Administration Portal or SL1 Collector	44

<i>Licensing New Skylar One Appliances</i>	45
<i>Configuring a New Skylar One System for Traditional Communication</i>	46
<i>Managing the Nodes Page</i>	51

Prerequisites

Before installing and configuring Skylar One, you must:

- Have already performed the prerequisites for all of the *ScienceLogic Hardware Appliances* or all of the *Virtual Appliances* in your Skylar One stack.
- Have a valid customer account that allows you to download the Skylar One ISO image. For details, contact your Customer Success Manager.
- Have access to the files for your SSL certificate.
- 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: For backup purposes, ScienceLogic supports *only* Skylar One backups to remote storage. Third-party backup types such as vmotion or VMware Snapshots can cause Skylar One outages. For details on properly configuring Skylar One backups, see the section on *Backup Management*.

Workflow for Installing and Configuring an SL1 Collector

The typical workflow for installing and configuring an SL1 Collector includes the following steps:

1. *Download the ISO image*. The ISO includes the Database Server, Administration Portal, and SL1 Collectors.
2. *Use the ISO to install the Database Server*.
3. *Use the ISO to install the Administration Portal and SL1 Collectors*.
4. *License the Skylar One appliances*.
5. Configure the new Skylar One system for one of the following communication types:
 - *Traditional communication*, in which the Database Server initiates a connection to the SL1 Collectors.
 - *PhoneHome communication*, in which the SL1 Collectors initiate an outbound connection to the Database Server.
6. *Use the Nodes page to manage nodes and tokens*.

Downloading the ISO Image

NOTE: The following ISO installation steps do not affect the performance of the Skylar One system. ScienceLogic recommends that you perform these steps at least 3 days before upgrading.

To download the ISO image:

1. Log in to the ScienceLogic Support Center at <https://support.scienceologic.com/s/> using your ScienceLogic customer account and password to access the site.
2. Select the **Skylar One** menu and choose *Downloads*. The **Skylar One (Skylar One) Platform Downloads** page appears.
3. Click the name of the Skylar One version you want to download. The **Release Version** page appears.
4. Click the link for the "Product Image" you want to download and scroll to the bottom of the page. The **Release File Details** page appears.
5. Click the **[Download File]** button for the ISO file to download the file to your local computer.

Installing the Database Server

The Database Server should be the first node or appliance you install.

NOTE: The installation options were updated in Skylar One 12.2.0. The following steps are intended for use in Skylar One 12.2.0 and later.

NOTE: If you deploy the ISO version of Skylar One, you might get an "Appliance is not licensed" message on the login page. This situation occurs only if you use another tab or browser to log in to the Skylar One environment after deployment. If you use the same browser or tab that you used for the deployment, the user interface will be available.

To install the Database Server:

1. Boot the appliance from the Skylar One ISO. The **Installation** window appears.

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. Select **Install Skylar One (recommended)**. After the installer environment boots, the **Installation Type** menu appears.
3. Select **Typical (recommended)**, and then select **[Continue]**. The **Model Type** window appears.
4. Select **Database**. Select **[Continue]**.
5. In the **Database** window, select **Local Database** and select **[Continue]**. 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 node or appliance.
 - **Netmask**. Type the netmask for the primary IP address of the node or 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 node or appliance.
7. Select **[Continue]**. The **System Password** window appears.
8. Type the password for the em7admin user on the operating system and select **[Continue]**.
9. Type the password for the em7admin user again and select **[Continue]**.
10. The appliance installer runs, and the virtual machine reboots automatically, and you are returned to a login prompt.
11. *Follow the instructions to license the appliance.*
12. *Follow the instructions for installing the remaining nodes or appliances*: the Administration Portal, the Data Collectors, and the Message Collectors (if applicable).

Installing an Administration Portal or SL1 Collector

Before you can install a Skylar One Collector, you will need to [*use the ISO to install the Skylar One Database Server*](#), if it is not already installed.

After installing the Database Server, you can then install:

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

Installing an Administration Portal or SL1 Collector

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

NOTE: If you deploy the ISO version of Skylar One, you might get an "Appliance is not licensed" message on the login page. This situation occurs only if you use another tab or browser to log in to the Skylar One environment after deployment. If you use the same browser or tab that you used for the deployment, the user interface will be available.

To install an Administration Portal or an SL1 Collector in Skylar One:

1. Boot the collector from the Skylar One ISO. The **Installation** window appears.
2. Select **Install Skylar One (recommended)**. After the installer environment boots, the **Installation Type** menu appears.
3. Select **Typical (recommended)**, and then select **[Continue]**. The **Model Type** window appears.
4. Select the appropriate appliance type and then select **[Continue]**.
5. After the installer for the collector is loaded, the **Network Configuration** window appears.
6. Enter the following information:
 - **IP Address**. Type the primary IP address of the collector.
 - **Netmask**. Type the netmask for the primary IP address of the collector.
 - **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 collector.
7. Select **[Continue]**. The **System Password** window appears.
8. Type the password for the em7admin user on the operating system and select **[Continue]**.
9. Type the password for the em7admin user again and select **[Continue]**.
10. After you install the SL1 Collector, upgrade the collector if needed to make sure the collector is running the same version of Skylar One that the Database Server is running. Then you can connect the new collector with the Skylar One Database Server.

Licensing New Skylar One Appliances

After you have installed new Skylar One appliances, you must then license them. The method for doing so varies by appliance type.

For details on licensing the Skylar One appliance types, see the following sections:

- [*Licensing and Configuring a Database Server or All-In-One Appliance*](#)
- [*Configuring an Administration Portal*](#)
- [*Configuring a Data Collector or Message Collector*](#)

For additional details about licensing Skylar One appliances, including details about using the Classic Web Configuration Utility or Node Configuration Utility, defining syslog servers, defining proxy servers, and more, see the chapter on [*Licensing and Configuring an Appliance*](#).

Configuring a New Skylar One System for Traditional Communication

After you have installed your Skylar One appliances from the ISO image and licensed those appliances, you must configure the new Skylar One system for one of the following communication types:


- **Traditional communication**, in which the Database Server initiates communication with each SL1 Collector. This configuration method is described in the sections below.
- **PhoneHome communication**, in which the SL1 Collectors initiate communication with the Database Server, either through the use of tokens or with passwords and secret keys. This configuration method is described in the chapter on [Configuring Skylar One for PhoneHome Communication](#).

What is Traditional Communication?

Skylar One supports two methods for communication between a Database Server (a Skylar One Central Database or a Skylar One Data Engine) and the Collectors:

- Traditional
- PhoneHome


In the **Traditional** method, the Skylar One services on the Database Server initiate a new connection to the MariaDB port on the collector to read and write data. The connection request traverses the network, including the Internet if necessary, eventually reaching the collector. For this approach to work, the collector administrator must allow ingress communication from the Database Server on TCP port 7707, which is the MariaDB port on the collector. The communication is encrypted using SSL whenever possible.



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

In the **PhoneHome** method, the collectors initiate an outbound connection to the Database Server over SSH. The connection requests originate from edge to core via TCP, using port 7705 by default.

After authenticating, the client forwards the local MariaDB port onto the Database Server using a loopback remote IP address. A corresponding Skylar One appliance is added using the loopback IP. When the Skylar One services on the database try to make a connection to the collector's MariaDB, they connect locally to the loopback IP address, in contrast to reaching out to the collector's IP or DNS name. The communication is encrypted.



The benefits of this method are that no ingress firewall rules need to be added, as the collector initiates an outbound connection, and no new TCP ports are opened on the network that contains the Data Collectors.

NOTE: While you do not need to add any ingress firewall rules, a best practice is to add an egress firewall rule that allows SSH traffic from the collector on the server's port to either all available destination addresses on the DB or to the specific address on the DB that you know the collector will be able to reach. Custom firewall rules must use the rich rules syntax and added to `/etc/siteconfig/firewalld-rich-rules.siteconfig`.

The PhoneHome configuration uses public key/private key authentication to maintain the security of the Database Server. You can use RSA256 and RSA512 algorithms for key authentication when configuring PhoneHome communication.

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.

Configuring a New SL1 Collector for Traditional Communication

After you install an SL1 Collector, use the **Add Node** wizard on the **Nodes** page (Manage > Nodes > Add Nodes) to configure your new SL1 Collector. This configuration process:

- Registers the SL1 Collector in Skylar One
- Connects the SL1 Collector to the Database Server so it can share its collected data
- Aligns the SL1 Collector to a new or existing Collector Group.

While navigating through the **Add Node** wizard, the **Choose Connection Type** window appears. This window enables you to determine the method in which the SL1 Collector and Database Server will communicate. The options are:

Connection Type	Used For
<i>Collector Initiates System Accepts</i>	Token-based PhoneHome Communication
<i>Collector Initiates User Accepts</i>	Password/secret-based PhoneHome Communication
<i>Database Initiates System Accepts</i>	Traditional Communication

NOTE: Part of the setup for SL1 Collectors takes place in the **Node Configuration Utility**, which has its own user interface separate from the Skylar One user interface. The **Nodes** page and the **Node Configuration Utility** replace some of the functionality previously found in the **Web Configuration Utility** in earlier versions of Skylar One.

All connection types require a token that Skylar One generates as part of the wizard. A **token** is a JSON web token (JWT) that contains a set of secure data that Skylar One uses to establish communication between the SL1 Collector and the Database Server. This token expires after a predefined time from the time of generation; by default, this expiration time is 30 minutes, but it can be extended to a maximum of 2 hours. The token encodes all destination addresses.

Configuring Traditional Database Initiates | System Accepts Communication

This section describes how to register and connect an SL1 Collector to the Database Server using the **Database Initiates | System Accepts** option. This is a "traditional" or non-PhoneHome collector connection type.

To connect an SL1 Collector to the Database Server for traditional communication:

- On the **[Registered]** tab on the **Nodes** page (Manage > Nodes), click **[Add Nodes]**. The **Choose Connection Type** window of the Add Node wizard appears.
- Select **Database Initiates | System Accepts** and click **Next**. The **Define Collector Properties** window appears.
- Complete the following fields as needed:
 - Collector Name.** Type the name the collector used when registering the collector. Skylar One will update this value with the collector hostname.
 - Collector IP Address or Hostname.** Type the IP address in this field so the Database Server can connect to the collector. Required.
 - Collector Description.** Type a description of the collector. This field is optional.

- **Collector Group.** The new collector must be aligned to a Skylar One Collector Group. You have the following options for this field:
 - Select an existing Collector Group from the drop-down.
 - Create a new Collector Group for the collector by clicking the plus icon (+). On the **Add Collector Group** modal, you can name the new group and choose to make that Collector Group available to all current and future organizations. You can also limit the Collector Group to specific organizations.

NOTE: The **All current and future organizations** toggle is enabled by default. If you want to limit Organization access to the new Collector Group, disable this toggle and select the organization or organizations from the drop-down.

- **Collector Type.** Your options include:
 - *Data Collector.* This is the most commonly used type. A Data Collector retrieves a specific set of information from monitored devices. A Data Collector can also work as a Message Collector.
 - *Message Collector.* A Message Collector receives and processes inbound, asynchronous syslog and trap messages from monitored devices.

4. Click **[Generate Token]**. The **Configure Collector** window appears.

NOTE: You can go back to a previous step at any point in the wizard, but when you click the **[Generate Token]** button, Skylar One always generates a new token. You cannot retrieve this particular token if you close the Add Node wizard. The generated token expires after 30 minutes.

5. Click the Copy icon (📋) to copy the token in the **Token** field.
6. Open the Node Configuration Utility by clicking the Open icon (🔗) in the **Node Configuration Utility** field. The login page for the Node Configuration Utility opens in a new browser window.

TIP: If you did not specify an IP address or a hostname in step 2 of this wizard, you will need to open a new browser window and type the IP address or hostname for the collector, followed by ":7700/node-config", such as "https://10.1.1.100:7700/node-config".

NOTE: If the node type is not a collector, the Node Configuration Utility will display the following message: "This page will only be visible if you are on a collector."

7. Log in to the Node Configuration Utility using the same username and password that you used when you installed the collector. After you log in, the collector and the Skylar One Database Server attempt to connect. The connection will fail, which is expected. The **Connect Collector** page appears with an empty **Paste token** text field.
8. Paste the token you copied in step 5 in the **Paste token** field.

NOTE: If the collector and Database Server are not able to connect, make sure that port 7707 is open between the Database Server and the collector.

9. Alternately, instead of pasting the token here, you can save time with additional configuration that you will need to do later by clicking **[Manual Entry]**, selecting **Database Initiated Connection**, and adding the IP addresses for the Database Servers (CMDBs) in the text box.

TIP: Using this option lets you add all IP addresses for your Database Servers (CMDBs), including primary, High Availability (HA) and Disaster Recovery (DR) servers.

10. After pasting the token or manually adding the IP addresses, click **[Register]** or **[Register Database]**, based on your choices in the two previous steps. When the connection is made, a **Success** dialog states that the collector was configured to accept a connection from the Database Server, and that you need to register the collector in Skylar One if you have not already done so. Click the link in the **Status** dialog to get more information about registering a collector.
11. Click **[OK]** on the **Success** dialog. The **Connect Collector** page appears, with a message stating that the collector can receive inbound connection requests.
12. After you connect the new collector, you will need to manually register the collector in Skylar One by navigating to the **Appliance Manager** page (System > Settings > Appliances).
13. At the top of the **Appliance Manager** page, complete the following fields:
 - **Host Name.** Type the host name of the collector.
 - **IP Address.** Type the IP address of the collector.
 - **Model Type.** Select the type of appliance (Data Collector or Message Collector) you are registering.

NOTE: When you select either type of collector, the **DB User** and **DB Password** fields appear. If the Database Server has different credentials from the collector, type the credentials for MariaDB on the Collector. This credential was entered when the ISO was deployed.

- **Description.** Type a description for the Data Collector or Message Collector. This field is optional.

- **Sharing Permissions.** Indicates if the appliance is shared or private. Choices are:
 - *Shared.* The appliance can be viewed by users across all organizations.
 - *Private.* The appliance can be viewed only by user accounts assigned to the System organization.

NOTE: The **Sharing Permissions** field displays only for Administrator user accounts and user accounts assigned to the System organization.

- **DB User.** Type a user name that can access the MariaDB database on the Data Collector or Message Collector.

NOTE: This user is the default database user for MariaDB. This user has the same password as the admin and root user, and the password is set during the initial installation. If you installed Skylar One from the ISO, the DB user name will be "clientdbuser".

- **DB Password.** Type a password that allows access to the MariaDB database on the Data Collector or Message Collector.

14. Click **[Save].** If the save is successful, the message "Appliance Registered" displays.
15. If all information is valid and the Database Server can communicate with the Data Collector or Message Collector, the **Appliance Manager** page displays the Skylar One 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.
16. Perform steps 13-15 for each Data Collector and Message Collector in your configuration.
17. Finally, align the new collector with the relevant Collector Group by going to the **Collector Groups** page (Manage > Collector Groups, or System > Settings > Collector Groups in the classic user interface).
18. Select the Collector Group you want to use, select the new collector from the **Message Collector Selection** field or the **Message Collector Selection** field, and click **[Save].** (If you are using the classic user interface, click the edit icon (✎) next to the Collector Group you want to use, select the new collector from the **Collector Selection** field, and click **[Save].**)
19. Go to the **[Registered]** tab on the **Nodes** page (Manage > Nodes), where you can now see the new collector in the list, aligned with the Collector Group you specified.

Managing the Nodes Page

The following topics describe how to use and add information on the **Nodes** page.

Viewing the List of Registered Nodes

The **[Registered]** tab of the **Nodes** page lets you manage the nodes used for installing Skylar One collectors, Skylar One instances, and other related appliances. You can also click the **[Add Node]** button to connect a Skylar One collector to a Skylar One Database Server.

The **[Pending]** tab displays a list of pending requests for establishing a connection between a collector and a Skylar One Database Server. The **[Tokens]** tab displays a list of existing and expired tokens used for connecting collectors.

The **[Pending]** tab and the **[Tokens]** tab do not display on an All-In-One Skylar One system.

NOTE: The **Nodes** page replaces some of the functionality previously found in the Web Configuration utility and the **Appliance Manager** page.

TIP: You can filter the items on this inventory page by typing filter text or selecting filter options in one or more of the filters found above the columns on the page. For more information, see "Filtering Inventory Pages" in the *Introduction to Skylar One* manual.

TIP: You can adjust the size of the rows and the size of the row text on this inventory page. For more information, see the section on "Adjusting the Row Density" in the *Introduction to Skylar One* manual.

By default, the **Nodes** page displays the following about each node:

- **Name.** Name of the node.
- **IP.** Primary IP address for the node.
- **Status.** The node status types include:
 - Available
 - Unavailable
 - Failed Over
 - Available Failed Over
 - Unconfigured
 - Unlicensed
- **Node Type.** The node types include:
 - All-In-One Appliance
 - Application Server (Administration Portal)

- Compute Node
- Collector Unit (Data Collector)
- Database Server
- Message Collector
- Storage Node
- **Database Version.** Version number of the Database Server for an All-In-One Appliance or a Database Server node.
- **Collector Groups.** For Data Collectors and All-In-One Appliances, specifies the Collector Group name associated with the node.

In addition, you can click the **[Grid Settings]** button and select *Column Preferences* to add the following columns to the **Nodes** page:

- **Node ID.** Unique numeric ID, automatically assigned by Skylar One to each node on the **Nodes** page.
- **Capacity.** For Database Servers, specifies the licensed capacity of the node.
- **Description.** Description of the node.
- **Patch Level.** Most recent patch version number for the node, where applicable.
- **Release Version.** Skylar One version running on the node.
- **Version ID.** Unique numeric ID, automatically assigned by the platform to each Skylar One version.
- **Created.** Date and time the node was registered and licensed.
- **Edit User.** User who last edited the node's information.
- **Last Edited.** Date the node's information was discovered or last edited.
- **Task Manager Paused.** Specifies whether the task manager service is paused. This value is updated every two minutes.
- **Needs Reboot.** Specifies whether the node requires reboot to add latest kernel or security updates.
- **Allocation.** For Data Collectors, specifies the number of devices aligned with the node.
- **Endpoint.** Skylar One Agent endpoint for the Gen 1 Agent.
- **Collector Group ID.** For Data Collectors and All-In-One Appliances, specifies the Collector Group ID associated with the node.

Viewing the Tokens on the Nodes Page

The **[Tokens]** tab on the **Nodes** page lists the existing and expired tokens that get used when connecting a collector. A **token** is a JSON web token (JWT) that contains a set of secure data that Skylar One uses to establish communication between the new SL1 Collector and the Skylar One Database Server.

By default, tokens for a "Collector Initiates | System Accepts" connection type have a 30-minute expiration period.

The **[Tokens]** tab lists the following:

- Collector registration details entered by the user at the time of token creation (collector hostname, description)
- Collector type (Data Collector or Message Collector) and aligned Collector Group
- Details about the token (including its type, date of creation, and expiration date)

A token inherits organization membership from the Collector Group to which it is aligned to allow multi-tenancy.

Recreating a Token

Expired tokens cannot be recovered on the **[Tokens]** tab, but you can recreate an expired token, which lets you generate a new token with the same collector details. Recreating the token actually deletes the existing token, but retains the user-supplied collector registration details to use in the new token.

To recreate an expired token:

1. Go to the **[Tokens]** tab on the **Nodes** page (Manage > Nodes).
2. Click the **Actions** menu (⋮) and select *Recreate* for the expired token. The **Recreated Token** window appears.
3. Click the **[Copy]** button to copy the token, and then paste the copied token into the Node Configuration Utility.

Chapter

6

Licensing and Configuring an Appliance

Overview

This chapter describes how to license a Skylar One (formerly SL1) appliance and add a Skylar One appliance to your Skylar One system.

There are two utilities you can use to perform various functions for the setup and editing of your Skylar One appliance.

These two utilities include:

- The Classic Web Configuration Utility (default)
- The Node Configuration Application

Upon installation, Skylar One appliances are automatically licensed for 90 days, with a capacity of 1,000 devices. During these 90 days, you can perform the steps to obtain a permanent license from ScienceLogic.

Skylar One 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 Skylar One systems, you **must** license the Database Server first. All other Skylar One appliances in a distributed Skylar One system depend on the Database Server for registration. Be sure to license your appliances before using the latest release of Skylar One.

Use the following menu options to navigate the Skylar One user interface:

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

This chapter covers the following topics:

<i>Navigating the Classic Web Configuration Utility</i>	57
<i>Licensing and Configuring a Database Server or All-In-One Appliance</i>	58
<i>Other Initial Configuration Tasks</i>	60
<i>Navigating the Node Configuration Utility</i>	63

Navigating the Classic Web Configuration Utility

The default utility application for configuring your appliance is the Classic Web Configuration Utility. This utility allows you to perform many different functions surrounding the configuration of your appliance.

In the Classic Web Configuration Utility, you can:

- Configure an Administration Portal
- Configure a Data Collector or Message Collector
- Register the Data Collector or Message Collector with the Database Server
- Define a Syslog, NTP, and/or Proxy Server(s)
- Create a Bonded Interface

Logging into the Classic 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 Skylar One. 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 the appropriate password.
3. After logging in, the main **Configuration Utility** page appears.

Changing the Password for the Classic Web Configuration Utility

If you want to change the password for the Web Configuration Utility on all Skylar One appliances, you must log in to the Web Configuration Utility on each node or appliance and perform the steps in this section.

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 by navigating to `https://<ip-address-of-appliance>:7700` and entering your credentials. 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 node or 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 by navigating to `https://<ip-address-of-appliance>:7700` and entering your credentials. 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 Center](#) and go to the **ScienceLogic Product Licensing** page (Support > License & Image Requests).
6. Under the **Skylar One** heading, click **[Request License]**.
7. Fill out the **Appliance Information** form and click the **[Submit License Request]** button.
8. 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

You must perform the following steps in the Web Configuration Utility to configure an Administration Portal:

1. Log in to the Web Configuration Utility by navigating to `https://<ip-address-of-appliance>:7700` and entering your credentials. The **Configuration Utilities** page appears.

2. Click the **[Device Settings]** button in the upper-right of the page. The **Settings** page appears.
3. On the **Settings** page, enter the following:
 - **Database IP Address.** The IP address(es) of the primary ScienceLogic Database Server(s). If this is a High Availability or Disaster Recovery (HA/DR) system, use the Virtual IP address in this field.
 - 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.
5. In Skylar One, go to the **Appliance Manager** page (System > Settings > Appliances).
6. Supply values in the following fields:
 - **Host Name.** Enter the hostname of the Administration Portal, where relevant.
 - **IP Address.** Enter the IP address of the Administration Portal. If this is a High Availability or Disaster Recovery (HA/DR) system, use the Virtual IP address in this field. If the Administration Portal is an AWS instance, supply the private IP address for the Administration Portal. 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.** Select *Administration Portal [3]* from the drop-down list.
 - **Description.** Enter a description of the Administration Portal. This field is optional.
 - **Sharing Permissions.** Indicates if the appliance is shared or private. Choices are:
 - *Shared.* The appliance can be viewed by users across all organizations.
 - *Private.* The appliance can be viewed only by user accounts assigned to the System organization.

NOTE: The **Sharing Permissions** field displays only for Administrator user accounts and user accounts assigned to the System organization.

7. Click the **[Save]** button. If the save is successful, the message "Appliance Registered" appears.
8. If you are using an AWS RDS system, select the wrench icon (🔧) for the newly created Administration Portal. Supply values in the **DB User** field and the **DB Password** field.
9. If all information is valid and the Database Server can communicate with the Administration Portal, 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.

Configuring a Data Collector or Message Collector

You do not need to use the Web Configuration Utility to configure and register a Data Collector or Message Collector. Instead, configuration and registration for SL1 Collectors occurs during their initial setup. The exact process varies based on which of the following connection types you are using:

Connection Type	Used For
Collector Initiates / System Accepts	Token-based PhoneHome Communication
Collector Initiates / User Accepts	Password/secret-based PhoneHome Communication
Database Initiates / System Accepts	Traditional Communication

Click the links in the table above to see instructions for configuring each connection type.

Other Initial Configuration Tasks

This section describes other initial configuration tasks you might need to complete when setting up a new Skylar One system.

Configuring Logging for a Skylar One System

For each device except for Message Collectors and All-In-One Appliances, you must specify the IP address of the server to which the Skylar One appliance will send syslog messages.

For full instructions on configuring logging in your Skylar One system, see the section on [Logging in Skylar One Version 11.3.0 and Later](#).

Defining the NTP Server

By default, Skylar One 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 Skylar One.

To configure a time server file:

1. Log in to the Web Configuration Utility by navigating to `https://<ip-address-of-appliance>:7700` and entering your credentials. 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 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 Skylar One, click **chrony.conf** in the Edit Files section of the Settings page.

Creating a Bonded Interface from the Web Configuration Utility

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.

NOTE: This section describes how to create bonded interfaces from the Web Configuration Utility. You can also do so using the [Node Configuration Utility](#).

To create one or more bonded interfaces:

1. Log in to the Web Configuration Utility by navigating to `https://<ip-address-of-appliance>:7700` and entering your credentials. 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 Skylar One System, the Skylar One 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 from the Appliance Manager Page

A proxy server enables Skylar One 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 Skylar One appliance can define its own proxy server.

NOTE: This section describes how to create define a proxy server from the **Appliance Manager** page. You can also do so using the [Node Configuration Utility](#).

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 the appropriate password.
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 Skylar One appliance will talk.
7. Click **[Save]**.

Navigating the Node Configuration Utility

The Node Configuration Utility is the other utility application you can use to configure and edit your Skylar One appliance.

In the Node Configuration Utility, you can:

- View the Collector connection status
- Configure a proxy server
- Add or edit a bonded interface

Logging into the Node Configuration Utility

Perform the following steps to access the Node Configuration Utility:

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

```
https://<ip-address-of-appliance>:7700/node-config
```

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 the appropriate password.
3. After logging in, the main **Node Configuration Utility** home page appears.

Changing the Password for the Node Configuration Utility

If you want to change the password for the Node Configuration Utility on all Skylar One appliances, you must log in to the Node Configuration Utility on each node or appliance and perform the steps in this section.

You cannot change the username for the Node Configuration Utility. The username remains **em7admin**.

To change the password for the Node Configuration Utility:

1. Log in to the Node Configuration Utility by navigating to <https://<ip-address-of-appliance>:7700/node-config> and entering your credentials. The **Configuration Utilities** page appears.
2. Click the drop-down arrow icon next to the username credential in the top-right corner and select **[Change Password]**.
3. On the **Change Password** page, type the following:

- **Current Password.** Type the current password.
- **New Password.** Type the new password.
- **Confirm New Password.** Type the new password again.

4. Click **[Change Password]**.

Viewing the Collector Connection Status

You can view the connection status for a Data Collector from multiple places in the Node Configuration Utility. You can view connection details for both the Collector and the Database Server in the utility.

To view the collector connection status:

1. Log in to the Node Configuration Utility by navigating to <https://<ip-address-of-appliance>:7700/node-config> and entering your credentials. The **Configuration Utilities** page appears.
2. Click the **[Connection]** icon located in the left-side navigation menu of the Node Configuration Utility. The **Collector Connection Status** page appears.
3. From this page, you can perform a few functions. You can:
 - Click **[Refresh Status]** to troubleshoot your collector's connection.
 - Click **[Disconnect & Clear Configuration]** to close the outgoing connection from this collector to all configured destinations. It will also clear all local configurations. A warning prompt will appear that asks you to confirm your action.

You can also access the **Collector Connection Status** page from the Node Configuration home page by clicking *View Connection Details* on the home page.

Configuring the Proxy Server from the Node Configuration Utility

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

Each Skylar One appliance can configure its own proxy server.

NOTE: This section describes how to configure a proxy server from the Node Configuration Utility. You can also do so using the [Appliance Manager Page](#).

To configure a proxy server:

1. Log in to the Node Configuration Utility by navigating to `https://<ip-address-of-appliance>:7700/node-config` and entering your credentials. The **Configuration Utilities** page appears.
2. Click the **[Settings]** icon located in the left-side navigation menu of the Node Configuration Utility.
3. 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 Skylar One appliance will talk.
4. Click **[Save]**.

Adding a Bonded Interface from the Node Configuration Utility

A bonded interface, which is 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").

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.

NOTE: This section describes how to create bonded interfaces from the Node Configuration Utility. You can also do so using the [Web Configuration Utility](#).

To add one or more bonded interfaces:

1. Log in to the Node Configuration Utility by navigating to `https://<ip-address-of-appliance>:7700/node-config` and entering your credentials. The **Configuration Utilities** page appears.
2. Click the **[Interfaces]** icon located in the left-side navigation menu of the Node Configuration Utility. The **Interfaces** page appears.
3. Click **[Add Bonding Interface]**. The **Add Bonding Interface** page appears.
4. Select the **[Activate]** button if you want this interface to be activated after you add it.

5. Complete the following fields:
 - **Name**. Required. Enter a user name for the bonded interface.
 - **Device ID**. Required. ID 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.
 - **DNS**. 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.
 - **Choose Bonded Interface**. Select your bonded interface from the drop-down list.
 - **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. Click **Add Another Option** for the addition of multiple bonding options.
6. Click **[Save]**.

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

Editing an Interface from the Node Configuration Utility

You can also edit an already existing bonded interface from the Node Configuration Utility.

1. Log in to the Node Configuration Utility by navigating to `https://<ip-address-of-appliance>:7700/node-config` and entering your credentials. The **Configuration Utilities** page appears.
2. From the **[Interfaces]** page, click the ellipses icon (⋮) located in the column to the right of your bonded interface.
3. Click **[Edit]**. The Interface Configuration window appears for editing.
4. Complete the **Interface**, **IP Address**, **Configuration**, and **Network** fields as needed for your interface.
5. Click **[Save]**.

Chapter

7

Configuring Skylar One for PhoneHome Communication

Overview

This chapter explains how to configure Skylar One (formerly SL1) to use PhoneHome communication.

If you are using a new Skylar One system or a system that has not previously used PhoneHome communication for collectors, you or your Skylar One administrator will need to configure each Database Server in the Skylar One system to accept these connections.

NOTE: PhoneHome communication is not available for All-In-One Appliances.

Use the following menu options to navigate the Skylar One user interface:

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

This chapter covers the following topics:

What is PhoneHome Communication?	69
Important Notes about PhoneHome Communication	70
Prerequisites for Configuring PhoneHome Communication	72
Overview of the PhoneHome Configuration	72
Configuring the Database Server for PhoneHome Communication	73


<i>Managing Proxy Connections for PhoneHome Communication</i>	77
<i>Configuring SL1 Collectors for PhoneHome Communication</i>	80
<i>Understanding PhoneHome Components</i>	87
<i>Using the Command-Line Interface for PhoneHome Collection</i>	87
<i>Viewing a List of PhoneHome Devices</i>	90
<i>Viewing Information about a Single PhoneHome Device</i>	90
<i>Renaming a PhoneHome Device</i>	90
<i>Checking the Status of a PhoneHome Collector</i>	91
<i>Syncing the Configuration of a PhoneHome System</i>	91
<i>Managing Port Forwarding for PhoneHome Communication</i>	92
<i>Managing Destinations</i>	93
<i>Logging PhoneHome Configuration Information</i>	95
<i>Tuning PhoneHome Settings</i>	95
<i>Clearing a PhoneHome Device</i>	97
<i>Deleting a PhoneHome Collector</i>	98
<i>Deleting a PhoneHome Database Server</i>	99
<i>Troubleshooting PhoneHome Configurations</i>	100

What is PhoneHome Communication?

Skylar One supports two methods for communication between a Database Server (a Skylar One Central Database or a Skylar One Data Engine) and the Collectors:

- Traditional
- PhoneHome


In the *Traditional* method, the Skylar One services on the Database Server initiate a new connection to the MariaDB port on the collector to read and write data. The connection request traverses the network, including the Internet if necessary, eventually reaching the collector. For this approach to work, the collector administrator must allow ingress communication from the Database Server on TCP port 7707, which is the MariaDB port on the collector. The communication is encrypted using SSL whenever possible.



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

In the **PhoneHome** method, the collectors initiate an outbound connection to the Database Server over SSH. The connection requests originate from edge to core via TCP, using port 7705 by default.

After authenticating, the client forwards the local MariaDB port onto the Database Server using a loopback remote IP address. A corresponding Skylar One appliance is added using the loopback IP. When the Skylar One services on the database try to make a connection to the collector's MariaDB, they connect locally to the loopback IP address, in contrast to reaching out to the collector's IP or DNS name. The communication is encrypted.



The benefits of this method are that no ingress firewall rules need to be added, as the collector initiates an outbound connection, and no new TCP ports are opened on the network that contains the Data Collectors.

NOTE: While you do not need to add any ingress firewall rules, a best practice is to add an egress firewall rule that allows SSH traffic from the collector on the server's port to either all available destination addresses on the DB or to the specific address on the DB that you know the collector will be able to reach. Custom firewall rules must use the rich rules syntax and added to `/etc/siteconfig/firewalld-rich-rules.siteconfig`.

The PhoneHome configuration uses public key/private key authentication to maintain the security of the Database Server. You can use RSA256 and RSA512 algorithms for key authentication when configuring PhoneHome communication.

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.

Important Notes about PhoneHome Communication

Before attempting to configure PhoneHome communication for your Skylar One system, be advised of the following:

- PhoneHome communication is not available for All-In-One Appliances.
- If you are using a proxy in your PhoneHome configuration, you should configure the proxy before you configure SL1 Collectors. For more information, see the section on [Adding a Proxy Configuration](#).
- If you are using a high-availability (HA) or disaster recovery (DR) setup, you can configure up to three PhoneHome Database Servers.
- PhoneHome communication uses secure shell (SSH). You cannot use PhoneHome over HTTP(S) or an HTTP(S) proxy.
- ScienceLogic does not recommend putting a PhoneHome Database Server behind a load balancer or a NAT gateway, as PhoneHome communication is designed to enable active connections to all Database Servers at any given time. If you must use a load balancer or NAT gateway, make sure each Database Server is behind a separate load balancer or NAT gateway.
- For destination addresses, use IP addresses whenever possible. Use a DNS name only if it uniquely identifies one host and does not point to a load balancer or is a round-robin for multiple hosts.
- If you have an AWS configuration, set up AWS hosts for the Database Server using an Elastic IP. In the event of a disaster recovery, this will make it easier to rebuild the Database Server without needing to change the IP address.
- Most intrusion detection/prevention systems will flag and drop SSH traffic on ports other than 22, which is the default SSH port. Since the PhoneHome server listens on ports other than 22, this often causes issues with onboarding PhoneHome collectors. You should ensure that your intrusion detection/prevention systems are configured to allow SSH traffic on the server's port.

Prerequisites for Configuring PhoneHome Communication

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

- Have *installed* and *licensed* the Database Server and SL1 Collectors.
- Have SSH access or console access to each Database Server.
- On each ScienceLogic appliance, know the username and password for access to the console. Note that the MySQL password matches the login password for Skylar One unless one or both of the passwords were manually changed after installation.
- Ensure that all Skylar One appliances are running the same version of Skylar One that the Database Server is running.
- Ensure that the Database Server opens a port for PhoneHome communication. The default port used by the Configuration Utility is **7705**. If you are on a SaaS Skylar One system, you must use port 7705. If you are on a non-SaaS system, you can use other ports besides 7705, but make sure those ports are not already being used.

CAUTION: Allow-listing port 7705 in the firewall is not enough. If the firewall does a layer 7 (application layer) filtering, you must create an exception rule to allow any outgoing traffic from the Data Collector to all the Database Servers on the control port, which is port 7705 by default. Some firewalls enable this by default and, as a result, those firewalls will drop SSH traffic on a non-standard port like 7705 in this situation.

IMPORTANT: If you use a proxy in your PhoneHome configuration, you must *add the proxy configuration* before configuring the SL1 Collectors for PhoneHome communication.

Overview of the PhoneHome Configuration

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

1. *Configure one or more Database Servers for PhoneHome.*
2. *Add a proxy connection*, if applicable to your configuration. Otherwise, you can skip this step.
3. *Configure the SL1 Collectors for PhoneHome.* If needed, update the collector to the same version of Skylar One that the Database Server is running.

After you have configured PhoneHome communications for your Skylar One system, you can also:

1. Familiarize yourself with the `phd` and `phc` *PhoneHome components*.
2. Learn how to use the *command-line interface* for PhoneHome communications.
3. As needed, perform administrative functions on your PhoneHome system, such as:
 - *View a list of PhoneHome devices*
 - *View information about a single PhoneHome device*
 - *Rename a PhoneHome device*
 - *Check the status of a PhoneHome collector*
 - *Check the connection between PhoneHome devices*
 - *Sync the configuration of a PhoneHome system*
 - *Define port forwarding* for each collector to use SSH from the Database Server to access that collector
 - *Associate a new destination address* with a PhoneHome Database Server
 - *View logs relating to PhoneHome configuration*
 - *Tune various PhoneHome settings*
 - *Clear a PhoneHome device*
 - Delete a *PhoneHome collector* or *Database Server*
4. See the *Troubleshooting section* for additional help.

Configuring the Database Server for PhoneHome Communication

The first step in establishing PhoneHome communication is to configure a PhoneHome Database Server. This can be either a Central Database (CDB) appliance or a Data Engine (DE) appliance.

In PhoneHome communication, the Database Server communicates with the SL1 Collectors. The Database Server stores all the configuration information for the PhoneHome configuration. Server-client authentication happens over the configuration store in MariaDB on the Database Server.

Setting up a Database Server prepares the server to listen to incoming connections from a PhoneHome collector. This process also opens the firewall rules on the configured port and labels the configured port for SSH traffic in the SE Linux subsystem.

PhoneHome configuration is stored in tables on the Database Server. The information is accessible to all Database Servers in the Skylar One system. Any Database Server in the Skylar One system can provide network access.

Before Configuring the Database Server for PhoneHome Communication

Make sure you have answers to the following questions before setting up the Database Server for PhoneHome communication:

- Is the Database Server a single CDB or DE, or is there a High Availability (HA) or Disaster Recovery (DR) pair?
- Is the CDB or DE behind a NAT gateway?
- Do you want the PhoneHome server to listen on to the default port 7705, or do you want to customize the port?

NOTE: SaaS Skylar One systems must use port 7705.

- Does the Database Server have multiple routable IP addresses to it, and do you plan to have PhoneHome collectors from different subnets connect to the Database Server?

IMPORTANT: Each Database Server must have Skylar One installed, have an IP address assigned to it, and be licensed with ScienceLogic. For more information about licensing, see [Licensing and Configuring a Database Server](#).

Understanding Database Server PhoneHome Configuration Options

The following sections explain how to configure the Database Server based on your Skylar One environment.

NOTE: After you configure the Database Server for PhoneHome communication, you must [add a proxy host](#) (if necessary for your setup) and then configure the Data Collectors and Message Collectors in your network. For more information, see [Configuring SL1 Collectors for PhoneHome Communication](#).

Configuring a Single Database Server

The most basic Skylar One environment contains a single Database Server. This setup makes the following assumptions:

- The Database Server has a public IP address assigned to one of its network interfaces or has a private IP address.
- All the PhoneHome collectors will be on the same network and will be able to reach the private IP address of the PhoneHome Database Server.
- The PhoneHome Database Server will be configured to listen on port 7705.
- The PhoneHome Database Server will be named "ph-db-1". Naming the PhoneHome collector is optional, but recommended.

To configure a single Database Server for PhoneHome communication:

1. Go to the console of the Database Server or use SSH to access the server and log in as user **em7admin** with the password you configured during setup.
2. Run the following command:

```
sudo phonehome setup -n ph-db-1
```

The `setup` command creates a PhoneHome device in the config store along with its corresponding RSA host key. You can use RSA256 and RSA512 algorithms for key authentication. The command also adds the default non-loopback IP address, corresponding to the hostname, as the default destination address. However, you can define a custom destination address if required. The command also adds a firewall rule to allow incoming connections on the specified port and labels it as SSH port (`ssh_port_t`) in the SELinux subsystem.

Configuring a Database with a Non-default Address or Port

CAUTION: If you are configuring a Skylar One system in a SaaS environment, you must use 7705 as the port for PhoneHome communication. Custom ports are not supported for PhoneHome communication on SaaS systems.

You can configure a PhoneHome Database Server to use a non-default address or port in the following situations:

- You want the PhoneHome server to listen on a non-default port, or on an address that is different than the output of the `getaddrbyhostname` syscall.
- The database appliance is behind a NAT gateway
- The database appliance is set up on a cloud host, like AWS, where the public IP is not assigned directly to the network interface of the virtual host.

To configure a Database Server with a non-default address or port:

1. Go to the console of the Database Server or use SSH to access the server and log in as user **em7admin** with the password you configured during setup.
2. Run the following command:

```
sudo phonehome setup -n ph-db-1 -a <addr>
```

where `<addr>` is an IPv4 address or DNS name in "host:port" format, such as `203.65.33.22:7809` or `ph-db1.example.com:8899`.

NOTE: The PhoneHome server process runs as an unprivileged user that will not be able to bind to a privileged port (1-1023). Therefore, when you choose a custom port, you must choose port 1024 or higher.

Configuring a Database with Multiple IP Addresses

You can assign multiple addresses to a destination if required. The list of addresses can be a mix of IPv4 addresses and DNS names.

To configure a Database Server with multiple IP addresses:

1. Go to the console of the Database Server or use SSH to access the server and log in as user **em7admin** with the password you configured during setup.
2. Run the following command:

```
sudo phonehome setup -n ph-db-1
```

3. Run the following command:

```
sudo phonehome destination add <id> <addr>
```

where **<id>** is the resulting device ID for the PhoneHome Database Server and **<addr>** is an actual address string in "host:port" format.

NOTE: The port must be the same for all addresses, because a PhoneHome server is not capable of listening on multiple ports.

4. Repeat this command for every address that you want to add to the destination.

Configuring PhoneHome Database Servers for High Availability and Disaster Recovery

If you are using a high-availability (HA) or disaster recovery (DR) setup, you can configure up to three PhoneHome Database Servers.

In an HA/DR PhoneHome configuration, there is no notion of a control node. Every Database Server in an HA/DR setup can participate in all operations.

ScienceLogic does not recommend putting a PhoneHome Database Server behind a load balancer or a NAT gateway, as PhoneHome communication is designed to enable active connections to all Database Servers at any given time. If you are configuring PhoneHome communication for HA/DR and you must use a load balancer or NAT gateway, make sure each Database Server is behind a separate load balancer or NAT gateway.

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

To configure PhoneHome Database Servers for HA/DR:

1. To configure the primary Database Server for PhoneHome communication, follow the instructions in the section [Configuring a Single Database Server](#).
2. To add a secondary Database Server, run the following command on each Database Server appliance in HA:

```
sudo phonehome setup -n <name>
```

where `<name>` is a customized name other than `ph-db-1`.

NOTE: If the IP or port that the collector would connect to is different from the IP of the network interface IP on the appliance, add `-a <addr>` to the above command, where `<addr>` is an IPv4 address or DNS name in "host:port" format.

3. Repeat step 2 if you are adding a third Database Server. Otherwise, proceed to step 4.
4. To add SL1 Collectors to your PhoneHome setup, follow the instructions in the section [Configuring SL1 Collectors for PhoneHome Communication](#).

NOTE: Alternatively, you can configure the SL1 Collectors for PhoneHome communication [using the command line](#).

Managing Proxy Connections for PhoneHome Communication

If your organization requires that you use a proxy for outbound requests, you can configure one or more proxy connections between the SL1 Collectors and the Database Server.

IMPORTANT: If you use a proxy in your PhoneHome configuration, you must perform the steps in the section about [Adding a Proxy Connection before you configure SL1 Collectors](#). The other steps in the PhoneHome configuration setup will require the proxy for communication.

Otherwise, if you are configuring PhoneHome communication and do not require a proxy connection, you can skip ahead to the section on [Configuring SL1 Collectors for PhoneHome Communication](#).

For example, you might use a proxy connection if your SL1 Collector does not have a direct outbound internet connection to reach the Database Server. A PhoneHome proxy configuration includes the destination address—either the address of the Database Server or that of the next proxy host—and the address of the proxy server to which the client must connect to reach the destination.

There can be one or more proxy hosts in between an SL1 Collector and a Database Server, thus forming a proxy chain.

Adding a Proxy Configuration

If you are using a proxy in your PhoneHome configuration, you should configure the proxy before you configure SL1 Collectors. The other steps in the PhoneHome configuration will require the proxy for communication.

NOTE: Only SSH proxies are supported for PhoneHome configurations. SOCKS over SSH is not supported.

To add a proxy connection between an SL1 Collector to the Database Server:

1. Go to the console of the SL1 Collector.
2. Run the following command on the SL1 Collector, replacing *<Destination Address>* with the address and port of the Database Server appliance to which you want to connect, *<Proxy Address>* with the proxy host address and port you want to use as a tunnel, and *<Proxy User>* with the username used to log in to the proxy host:

```
phonehome proxy new <Destination Address> <Proxy Address> <Proxy User>
```

NOTE: Addresses should be in the format *<host>:<port>*. The host can be either an IP address or a DNS name.

For example, if you want to configure the SL1 Collector to connect to the Database Server with an address of 202.35.52.71 through a proxy host with the address 10.1.17.68 with the user em7admin, you would run the following command:

```
phonehome proxy new 202.35.52.71:7705 10.1.17.68:22 em7admin
```

If you are connecting to the Database Server through a chain consisting of multiple proxies, you should add the proxy configurations in reverse order, starting with the destination address and last proxy host address, then the last proxy host address and previous proxy host address, and so forth, until you get to the first proxy host.

For example, if you want to connect to the Database Server with an address of 202.42.63.79 through proxy host A with an address of 192.168.0.3 with the user proxyuser, and also proxy host B with an address of 10.2.13.79 with the user em7admin, then you would run the following commands:

```
phonehome proxy new 202.42.63.79:7705 10.2.13.79:22 em7admin
```

```
phonehome proxy new 10.2.13.79:22 192.168.0.3 proxyuser
```

NOTE: New proxy configurations do not take effect until the PhoneHome client is restarted or the next watchdog cycle occurs.

NOTE: When you run the command, the system prompts you for a password for the proxy host. The system uses this password to automatically configure and validate SSH key-based authentication to the host; the next time you need to run anything via the proxy host, it will use the collector's private key for authentication rather than prompting you for the password. Optionally, you can disable this behavior by adding "-n" to the end of the command. If you do so, you must then manually configure the proxy's SSH key-based authentication.

If you get a "handshake failed: ssh..." error message when adding a new proxy:

1. In Skylar One, go to the **Appliance Manager** page (System > Settings > Appliances) and click the edit button (edit icon) for that appliance.
2. Review the collector's MariaDB credentials. This error can occur if the collector and the Database Server (CDB) use different credentials.

For example, if the Database Server has been updated and the ISO for the Database Server is before Skylar One version 11.3.0, while the collector was deployed with Skylar One version 11.3.0 or later, the Database Server will be using `root/<password>`, and the collector would be using `clientdbuser/<password>`.

Viewing a List of Proxy Connections

To view a list of proxy connections from an SL1 Collector to the Database Server:

1. Go to the console of the SL1 Collector.
2. Run the following command on the SL1 Collector:

```
phonehome proxy list
```

Deleting a Proxy Configuration

To add a proxy configuration between an SL1 Collector to the Database Server:

1. Go to the console of the SL1 Collector.
2. Run the following command on the SL1 Collector, replacing `<Destination Address>` with the address and port of the Database Server appliance to which the proxy is connecting:

```
phonehome proxy delete <Destination Address>
```

NOTE: Addresses should be in the format `<host>:<port>`. The host can be either an IP address or a DNS name.

NOTE: Removed proxy configurations do not take effect until the PhoneHome client is restarted or the next watchdog cycle occurs.

Configuring SL1 Collectors for PhoneHome Communication

After you install an SL1 Collector, use the **Add Node** wizard on the **Nodes** page (Manage > Nodes > Add Nodes) to configure your new SL1 Collector. This configuration process:

- Registers the SL1 Collector in Skylar One
- Connects the SL1 Collector to the Database Server so it can share its collected data
- Aligns the SL1 Collector to a new or existing Collector Group.

While navigating through the **Add Node** wizard, the **Choose Connection Type** window appears. This window enables you to determine the method in which the SL1 Collector and Database Server will communicate. The options are:

Connection Type	Used For
<i>Collector Initiates System Accepts</i>	Token-based PhoneHome Communication
<i>Collector Initiates User Accepts</i>	Password/secret-based PhoneHome Communication
<i>Database Initiates System Accepts</i>	Traditional Communication

NOTE: Part of the setup for SL1 Collectors takes place in the **Node Configuration Utility**, which has its own user interface separate from the Skylar One user interface. The **Nodes** page and the **Node Configuration Utility** replace some of the functionality previously found in the **Web Configuration Utility** in earlier versions of Skylar One.

All connection types require a token that Skylar One generates as part of the wizard. A **token** is a JSON web token (JWT) that contains a set of secure data that Skylar One uses to establish communication between the SL1 Collector and the Database Server. This token expires after a predefined time from the time of generation; by default, this expiration time is 30 minutes, but it can be extended to a maximum of 2 hours. The token encodes all destination addresses.

The processes for setting up the two PhoneHome communication types—"Collector Initiates | System Accepts" and "Collector Initiates | User Accepts"—through the Skylar One user interface and the Node Configuration Utility are described below. Alternatively, you can also [configure these communication types using the command line](#).

Configuring Collector Initiates | System Accepts Communication

This section describes how to register and connect an SL1 Collector to the Database Server using the **Collector Initiates | System Accepts** option. This is a token-based PhoneHome collector connection type.

To connect an SL1 Collector to the Database Server for token-based PhoneHome communication:

1. On the **[Registered]** tab on the **Nodes** page (Manage > Nodes), click **[Add Nodes]**. The **Choose Connection Type** window of the Add Node wizard appears.
2. Select **Collector Initiates | System Accepts** and click **Next**. The **Define Collector Properties** window appears.
3. Complete the following fields as needed:
 - **Collector Name**. Type the name the collector used when registering the collector. Skylar One will update this value with the collector hostname.
 - **Collector IP Address or Hostname**. Type the IP address or the hostname of the collector. This information is optional but recommended, as it is used in Step 3 of the wizard to create a link to the collector's Node Configuration Utility, where you will input the token you generate.
 - **Collector Description**. Type a description of the collector. This field is optional.
 - **Collector Group**. The new collector must be aligned to a Skylar One Collector Group. You have the following options for this field:
 - Select an existing Collector Group from the drop-down.
 - Create a new Collector Group for the collector by clicking the plus icon (+). On the **Add Collector Group** modal, you can name the new group and choose to make that Collector Group available to all current and future organizations. You can also limit the Collector Group to specific organizations.

The **All current and future organizations** toggle is enabled by default. If you want to limit Organization access to the new Collector Group, disable this toggle and select the organization or organizations from the drop-down.

- **Collector Type**. Your options include:
 - **Data Collector**. This is the most commonly used type. A Data Collector retrieves a specific set of information from monitored devices. A Data Collector can also work as a Message Collector.
 - **Message Collector**. A Message Collector receives and processes inbound, asynchronous syslog and trap messages from monitored devices.

4. Click [Generate Token]. The **Configure Collector** window appears.

NOTE: You can go back to a previous step at any point in the wizard, but when you click the **[Generate Token]** button, Skylar One always generates a new token. You cannot retrieve this particular token if you close the Add Node wizard. The generated token expires after 30 minutes.

5. In the **Token** field, click the Copy icon (📋) to copy the token .
6. Open the Node Configuration Utility by clicking the Open icon (🌐) in the **Node Configuration Utility** field. The login page for the Node Configuration Utility opens in a new browser window.

TIP: If you did not specify an IP address or a hostname in step 2 of this wizard, you will need to open a new browser window and type the IP address or hostname for the collector, followed by ":7700/node-config", such as "https://10.1.1.100:7700/node-config".

NOTE: If the node type is not a collector, the Node Configuration Utility will display the following message: "This page will only be visible if you are on a collector."

7. Log in to the Node Configuration Utility using the same username and password that you used when you installed the collector. After you log in, the collector and the Skylar One Database Server attempt to connect. The connection will fail, which is expected. The **Connect Collector** page appears with an empty **Paste token** text field.
8. Paste the token you copied in step 5 in the **Paste token** field.

NOTE: If you did not generate a token, you can click **[Manual Entry]**, select *User Accepted Connection Request*, and add the IP addresses for the Database Servers (CMDBs) in the text box.

9. After pasting the token, click **[Register]** or **[Register Database]**, based on your choices in the previous step. When the connection is made, a **Success** dialog states that the collector was registered and the connection to the database was initiated.
10. Click **[OK]** on the **Success** dialog. The **Collector Connection Status** page displays details about the collector and the Database Server, along with the connection state, which can be "Connected", "Not Connected", or "Unknown". "Unknown" indicates that Skylar One has not yet completed its first check of the connection state; click **[Refresh Status]** after a few moments and the status should update to "Connected".
11. On the **Collector Connection Status** page, click the expand icon (⌄) to view the connection path. The health of each hop in the connection is reported separately, but hops after an unresponsive hop will not be checked. This "Connection Path" information can be useful in diagnosing collector-database connection issues.

12. To view any changes to the connection status, click **[Refresh Status]**.

NOTE: If you want to disconnect the collector and close the SSH tunnel between the collector and the Database Server, click *Disconnect & Clear Configuration*. This action will close the outgoing connection from the collector to all configured destinations, and it will also clear all local configuration. This action cannot be undone.

13. Close the Node Configuration Utility.
14. In Skylar One, go to the **[Registered]** tab on the **Nodes** page, where you can now see the new collector in the list, aligned with the Collector Group you specified in the Add Node wizard.
15. Go to the **Appliance Manager** page (System > Settings > Appliances), click the edit icon (✎) for the new collector, and update the **DB User** and **DB Password** fields with the MariaDB credentials. The **DB User** is the default database user for MariaDB. This user has the same password as the admin and root user, and the password is set during the initial installation. If you installed Skylar One from the ISO, the **DB User** name will be *clientdbuser*.

Configuring Collector Initiates | User Accepts Communication

This section describes how to register and connect an SL1 Collector to the Database Server using the **Collector Initiates | User Accepts** option. This is a password/secret key PhoneHome collector connection type.

To connect an SL1 Collector to the Database Server for password/secret key PhoneHome communication:

1. On the **[Registered]** tab on the **Nodes** page (Manage > Nodes), click **[Add Nodes]**. The **Choose Connection Type** window of the Add Node wizard appears.
2. Select **Collector Initiates | User Accepts** and click **Next**. The **Define Collector Properties** window appears.
3. Complete the following fields as needed:
 - **Collector Name**. Type the name the collector used when registering the collector. Skylar One will update this value with the collector hostname.
 - **Collector IP Address or Hostname**. Type the IP address or the hostname of the collector. This information is optional but recommended, as it is used to create a link to the collector's Node Configuration Utility, where you will input the token you generate.
 - **Collector Description**. Type a description of the collector. This field is optional.

- **Collector Group.** The new collector must be aligned to a Skylar One Collector Group. You have the following options for this field:
 - Select an existing Collector Group from the drop-down.
 - Create a new Collector Group for the collector by clicking the plus icon (+). On the **Add Collector Group** modal, you can name the new group and choose to make that Collector Group available to all current and future organizations. You can also limit the Collector Group to specific organizations.

NOTE: The **All current and future organizations** toggle is enabled by default. If you want to limit Organization access to the new Collector Group, disable this toggle and select the organization or organizations from the drop-down.

- **Collector Type.** Your options include:
 - **Data Collector.** This is the most commonly used type. A Data Collector retrieves a specific set of information from monitored devices. A Data Collector can also work as a Message Collector.
 - **Message Collector.** A Message Collector receives and processes inbound, asynchronous syslog and trap messages from monitored devices.

4. Click **[Generate Token]**. The **Configure Collector** window appears.

NOTE: You can go back to a previous step at any point in the wizard, but when you click the **[Generate Token]** button, Skylar One always generates a new token. You cannot retrieve this particular token if you close the Add Node wizard. The generated token expires after 30 minutes.

5. Click the Copy icon (📋) to copy the token in the **Token** field.
6. Open the Node Configuration Utility by clicking the Open icon (🔗) in the **Node Configuration Utility** field. The login page for the Node Configuration Utility opens in a new browser window.

TIP: If you did not specify an IP address or a hostname in step 2 of this wizard, you will need to open a new browser window and type the IP address or hostname for the collector, followed by ":7700/node-config", such as "https://10.1.1.100:7700/node-config".

NOTE: If the node type is not a collector, the Node Configuration Utility will display the following message: "This page will only be visible if you are on a collector."

7. Log in to the Node Configuration Utility using the same username and password that you used when you installed the collector. After you log in, the collector and the Skylar One Database Server attempt to connect. The connection will fail, which is expected. The **Connect Collector** page appears with an empty **Paste token** text field.
8. Paste the token you copied in step 5 in the **Paste token** field.

NOTE: If you did not generate a token, you can click **[Manual Entry]**, select *User Accepted Connection Request*, and add the IP addresses for the Database Servers (CMDBs) in the text box.

9. After pasting the token, click **[Register]** or **[Register Database]**, based on your choices in the previous step. When the connection is made, the **Success** dialog contains a six-digit confirmation code. Click the Copy icon (copy icon) to copy the confirmation code.
10. Click **[OK]** on the **Success** dialog. The **Collector Connection Status** page displays details about the connection request and the same six-digit confirmation code.
11. In Skylar One, click **[See Pending Requests]** on Step 3 of the Add Node wizard. The **[Pending]** tab on the **Nodes** page appears with the pending request.
12. Select the Actions icon (Actions icon) next to the pending request for the new collector and select *Accept*. The **Accept Request** dialog appears.
13. Paste the six-digit confirmation code you copied in step 9 from the **Connect Collector** page of the Node Configuration Utility and click **[Validate]**. The **Configure Collector** dialog displays a summary of the collector information you entered in the Add Node wizard.
14. Edit the collector information and collector group as needed, and then click **[Save]**. The **Configure Collector** dialog displays a summary of your information.
15. Click **[OK]**. The **[Registered]** tab on the **Nodes** page displays the new collector, aligned with the collector group you specified.
16. Go to the **Appliance Manager** page (System > Settings > Appliances), click the edit icon (edit icon) for the new collector, and update the **DB User** and **DB Password** fields with the MariaDB credentials. The **DB User** is the default database user for MariaDB. This user has the same password as the admin and root user, and the password is set during the initial installation. If you installed Skylar One from the ISO, the **DB User** name will be *clientdbuser*.

Connecting a Skylar One Collector to the Skylar One Database Server using the Command-line Interface

As an alternative to onboarding SL1 Collectors via the user interface, you can instead choose to onboard SL1 Collectors using the command-line interface if you prefer to do so. This section describes how to onboard SL1 Collectors based on whether you want a "system accepted" connection type or a "user accepted" connection type.

System Accepted

In this connection method, the database administrator creates a new token on the database appliance.

To connect a collector using the System Accepted method with the command-line interface:

1. Go to the console of the Database Server or use SSH to access the server and log in as user **em7admin** with the password you configured during setup.
2. Run the following command:

```
phonehome token new <model_type> <name> <CUG_ID> <description>
```

where:

- **<model_type>** is either a **cu** for a Data Collector or **mc** for a Message Collector.
- **<name>** is the name of the PhoneHome collector. You can use any name you want.
- **<CUG_ID>** is the numeric ID of a collector group from Skylar One.
- **<description>** is the descriptive text about the collector.

3. Make a note of the resulting token and share it with the collector administrator.
4. The collector administrator registers the collector using the token value by running the following command on the Skylar One collector:

```
sudo phonehome register <token>
```

User Accepted

In this connection method, the collector administrator sends a registration request from the collector.

To connect a collector using the User Accepted method with the command-line interface:

1. Go to the console of the Skylar One collector or use SSH to access the collector and log in as user **em7admin** with the password you configured during setup.
2. Run the following command on the collector:

```
sudo phonehome request send <address_1> [<address_2> <address_3> ...  
<address_n>] [-l <label>]
```

where:

- **<address>** is the destination address of the database server, in "host:port" format. You can include multiple addresses to one or multiple databases. Separate multiple addresses with a space.
- **<label>** is an optional field you can use to associate a human-friendly identifier with the request. Every request is identified by a random string on the server side, and it might be confusing for the database administrator to find a specific request if numerous requests are coming from other collectors.

3. Make a note of the one-time secret and share it with the database administrator.

4. The Database administrator accepts the incoming request using the one-time secret by running the following command on the Database Server:

```
phonehome request accept <uuid> <model_type> <name> <CUG_ID>  
<description> <one_time_secret>
```

where:

- <uuid> is the unique ID of the request.
- <model_type> is either a `cu` for a Data Collector or `mc` for a Message Collector.
- <name> is the name of the PhoneHome collector. You can use any name you want.
- <CUG_ID> is the numeric ID of a collector group from Skylar One to which you want to assign this collector.
- <description> is the descriptive text about the collector.
- <one_time_secret> is the secret generated when sending a request from the collector that you made note of in step 3.

Understanding PhoneHome Components

This section describes two important PhoneHome components, `phd` and `phc`.

phd

The `phd` PhoneHome server daemon is installed and managed as a systemd service that is enabled on PhoneHome Database Servers. The server daemon listens to a port (7705 by default) and accepts incoming SSH connections from the PhoneHome client (`phc`) as well as OpenSSH clients. This service supports public key authentication for registered PhoneHome clients and collectors, as well as challenge-response authentication for the initial registration. The authentication-related configuration is stored in MariaDB; as such, it does not require creating local (Linux) users on the Database Server. Some aspects of the `phd` configuration will be stored on the local filesystem.

phc

The `phc` PhoneHome client runs as a service in systemd on PhoneHome SL1 Collectors. It is responsible for establishing a tunnel with the `phd` that is running on the Database Server and forwarding the local MariaDB port from the SL1 Collector to the Database Server.

Using the Command-Line Interface for PhoneHome Collection

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 the `phonehome` command to configure and troubleshoot your PhoneHome configuration.

To use the `phonehome` command:

1. Either go to the console of the Skylar One appliance or use SSH to access the server. Log in as "root".
2. At the command prompt, type the following:

```
phonehome <command>
```

where `<command>` is one of the following commands:

Command	Used For	See Also
<code>clear</code>	Clears the PhoneHome configuration on a PhoneHome device. The <code>clear</code> command will also <i>disable</i> the PhoneHome <code>phd</code> service. You can use the <code>clear</code> command on a Database Server to block future connection requests from Data Collectors and secondary Database Servers in an HA/DR configuration.	Clearing a PhoneHome Device
<code>check</code>	Checks the state of the connection from an SL1 Collector to the Database Server, visualizing the network path from the SL1 Collector to the Database Server as well as any proxy hops in between, if applicable. The output indicates any failures connecting to any hop.	Checking the Connection Between PhoneHome Devices
<code>client</code>	Runs the PhoneHome client (installed as a systemd service <code>phc</code>).	Understanding PhoneHome Components
<code>config</code>	Displays and enables you to edit PhoneHome configuration related to the server and client.	Tuning PhoneHome Settings
<code>delete</code>	Deletes a PhoneHome SL1 Collector. This argument prevents you from deleting any SL1 Collector with an associated Skylar One appliance.	Deleting a PhoneHome Collector
<code>destination</code>	Enables you to add, remove, or view addresses to a PhoneHome Database Server.	Managing Destinations
<code>forwards</code>	Enables you to add, remove, or view ports forwarded from an SL1 Collector to the Database Server.	Managing Port Forwarding for PhoneHome Communication
<code>list</code>	Displays a list of PhoneHome devices (Database Servers and Collectors).	Viewing a List of PhoneHome Devices
<code>migrate</code>	Migrates the configuration from the classic PhoneHome setup to the new PhoneHome setup. This is done automatically during upgrade, if you are upgrading from a version of Skylar One prior to 11.2.0.	Running the Pre-upgrade Test for Existing PhoneHome Connections
<code>proxy</code>	Enables you to add, remove, or view proxy configurations along the network path from an SL1 Collector to the Database Server.	Managing Proxy Connections for PhoneHome Communication
<code>register</code>	Registers a new SL1 Collector as a PhoneHome	Connecting a Skylar One

Command	Used For	See Also
	collector with a token.	Collector to the Skylar One Database Server using the Command-line Interface
rename	Renames an existing Phone Home device: <code>phonehome rename <id> <new_name></code> .	Renaming a PhoneHome Device
request	Enables you to send, view, accept, or reject an SL1 Collector registration request.	Connecting a Skylar One Collector to the Skylar One Database Server using the Command-line Interface
server	Runs the PhoneHome server (installed as a systemd service <code>phd</code>).	Understanding PhoneHome Components
setup	Configures a new PhoneHome Database Server.	Configuring the Database Server for PhoneHome Communication
status	Displays the status of the PhoneHome SL1 Collectors. The output is tabular by default but supports JSON output as well. The output does not contain the remote loopback IP address of PhoneHome SL1 Collectors, nor does it list PhoneHome Database Servers.	Checking the Status of a PhoneHome Collector
sync	Syncs the configuration from the Database Server.	Syncing the Configuration of a PhoneHome System
token	Enables you to create, view, or delete registration tokens.	Connecting a Skylar One Collector to the Skylar One Database Server using the Command-line Interface
view	Displays the state of an SL1 Collector. This argument must be run on a Database Server.	Viewing Information about a Single PhoneHome Device

Additionally, after configuring communication between a Database Server and an SL1 Collector, you can go to the console of the SL1 Collector or Database Server and run the following commands to view more information about your servers and collectors:

- To ensure that the PhoneHome service is active on the Database Server and view additional configuration information about the server:

```
systemctl status phd.service
```

- If the PhoneHome service is disconnected on a Database Server or SL1 Collector and you want to start it:

```
systemctl start phc
```

Viewing a List of PhoneHome Devices

The `phonehome list` command lists all of the PhoneHome devices in your Skylar One system, including the Database Server and SL1 Collector, including the addresses for the Database Server and the remote IP address corresponding to the collectors.

To view a list of PhoneHome devices:

1. Go to the console of the SL1 Collector or Database Server.
2. Run the following command on the SL1 Collector or Database Server:

```
sudo phonehome list
```

To view a list of only the PhoneHome Database Servers, run the following command:

```
sudo phonehome destination list
```

To view information about a specific PhoneHome Database Servers, run the following command:

```
sudo phonehome destination list --id <id>
```

where `<id>` is the PhoneHome device ID for the Database Server.

Viewing Information about a Single PhoneHome Device

The `phonehome view` command displays the state of a single PhoneHome device. This command must be run on a Database Server.

To view information about the PhoneHome configuration of a specific PhoneHome device:

1. Go to the console of the Database Server.
2. Run the following command on the SL1 Collector or Database Server:

```
sudo phonehome view <id>
```

where `<id>` is the PhoneHome device ID for the Database Server or SL1 Collector.

Renaming a PhoneHome Device

The `phonehome rename` command enables you to rename a PhoneHome device. You can run this command only from a Database Server, and you must know the PhoneHome device ID of the device that you want to rename.

To rename a PhoneHome device:

1. Go to the console of the Database Server.
2. Run the following command on the SL1 Collector or Database Server:

```
sudo phonehome rename <id><new_name>
```

where `<id>` is the PhoneHome device ID for the Database Server or SL1 Collector that you want to rename and `<new_name>` is the new name that you want to apply to the device.

Checking the Status of a PhoneHome Collector

The `phonehome status` command displays the status of the PhoneHome SL1 Collectors against all available databases. The output is tabular by default but supports JSON output as well. In the color output mode, the command will print the status of disconnected collectors in red.

The output does not contain the remote loopback IP address of PhoneHome SL1 Collectors, nor does it list PhoneHome Database Servers.

To check the status of a PhoneHome SL1 Collectors:

1. Go to the console of the SL1 Collector.
2. Run the following command:

```
sudo phonehome status
```

where you can optionally add the following parameters to the command:

- `-n` to disable live probing to the collector and instead use the periodic server check results, which happens every minute by default
- `-x` to enable extended output that includes a column indicating the last change timestamp
- `-c` to disable color output
- `-j` to output the data in JSON instead of a table

Syncing the Configuration of a PhoneHome System

The `phonehome sync` command syncs the configuration from the Database Server. This command can be run on the SL1 Collector.

To sync the configuration of a PhoneHome system:

1. Go to the console of the SL1 Collector.
2. Run the following command:

```
sudo phonehome sync
```

Managing Port Forwarding for PhoneHome Communication

A port forward is a configuration that allows a PhoneHome client to "copy" a local port from the SL1 Collector to the Database Server, essentially making the local port available on the Database Server appliance as if it was physically present on that appliance itself.

NOTE: The local MariaDB port 7707 on the collector is forwarded to the Database Server by default.

Viewing a List of Port Forwards

To view a list of ports forwarded from an SL1 Collector to the Database Server:

1. Go to the console of the SL1 Collector.
2. Run the following command on the SL1 Collector:

```
sudo phonehome forwards list
```

NOTE: This list will not include the MariaDB port 7707, which is forwarded by default.

Adding a Port Forward

To add a port forward:

1. Go to the console of the SL1 Collector.
2. Run the following command on the SL1 Collector, replacing `<Remote Port>` with the port on the Database Server onto which the local port will be forwarded and `<Local Port>` with the local port to forward from the SL1 Collector:

```
sudo phonehome forwards add <Remote Port> <Local Port>
```

NOTE: Ports should be in the format :`<port>`.

NOTE: The remote port should be an unprivileged port greater than 1023.

For example, if you want to forward SSH port 22 from the SL1 Collector to the Database Server appliance as port 10022 to enable a Database Server administrator to SSH into the SL1 Collector from the Database Server appliance, you would run the following command:

```
sudo phonehome forwards add :10022 :22
```

NOTE: New forwards do not take effect until the PhoneHome client is restarted or the next watchdog cycle occurs.

Removing a Port Forward

To remove a port forward:

1. Go to the console of the SL1 Collector.
2. Run the following command on the SL1 Collector, replacing `<Remote Address>` with the port on the Database Server appliance onto which the local port was forwarded and `<Local Address>` with the local port that was forwarded from the SL1 Collector:

```
sudo phonehome forwards remove <Remote Address> <Local Address>
```

NOTE: Addresses should be in the format `:<port>`.

For example, if you want to no longer forward SSH port 22 from the SL1 Collector to the Database Server appliance as port 10022, run the following command:

```
sudo phonehome forwards remove :10022 :22
```

NOTE: Deleted forwards do not take effect until the PhoneHome client is restarted or the next watchdog cycle occurs.

Managing Destinations

A destination is a list of addresses associated with a Database Server. A PhoneHome Database Server can have one or more destination addresses associated with it.

NOTE: Destination addresses can be IPv4 addresses or DNS names, or a combination of both.

Viewing a List of Destinations

To view a list of all destinations in your stack:

1. Go to the console of the SL1 Collector.
2. Run the following command on the SL1 Collector:

```
sudo phonehome destination list
```

This command provides a list of all Database Servers with their device IDs, addresses, and priorities. Priorities define the order in which an SL1 Collector will try to connect to the individual addresses. It will keep trying until it can connect to one of them.

NOTE: To view a list of destinations on a specific PhoneHome Database Server, run the following command, where *<Device ID>* is the ID of the PhoneHome Database Server:

```
phonehome destination list --id <Device ID>
```

Adding a Destination Address

To add a new destination address:

1. Go to the console of the SL1 Collector.
2. Run the following command on the SL1 Collector, where *<Device ID>* is the ID of the device to which you want to add a new address and *<Address>* is the destination address:

```
sudo phonehome destination add <Device ID> <Address>
```

NOTE: Addresses should be in the format *<host>:<port>*.

NOTE: Host addresses can be IPv4 addresses or DNS names.

If successful, you will get a message confirming that the new address was successfully added to the destination.

For example, if you wanted to add the destination address 192.168.0.13, with port 7705 open, to the device with the device ID 2, run the following command:

```
phonehome destination add 2 192.168.0.13:7705
```

NOTE: The port you open must match the port that is open for the original device. Otherwise, you will receive an error.

NOTE: Optionally, you can add the suffix `--priority <Priority>` to establish the destination's priority, or use the suffix `--force` to force add a destination address, even if the port does not match with the device's listed port.

Removing a Destination

To remove an existing address from a destination:

1. Go to the console of the SL1 Collector.
2. Run the following command on the SL1 Collector, where `<Device ID>` is the ID of the device from which you want to remove a destination address and `<Address>` is the destination address:

```
sudo phonehome destination remove <Device ID> <Address>
```

NOTE: Addresses should be in the format `<host>:<port>`.

NOTE: Host addresses can be IPv4 addresses or DNS names.

NOTE: You cannot remove an address from a destination if it is the destination's only address.

Logging PhoneHome Configuration Information

In Skylar One, the server hosts are stored in the `journald` log for the `phd` service on the Database Server and in the `journald` log for the `phc` service on the Collector.

To view those logs, run the following commands on the Database Server or Collector:

```
sudo journalctl -u phd.service
```

```
sudo journalctl -u phc.service
```

Tuning PhoneHome Settings

A PhoneHome setting is a customizable configuration that impacts how a PhoneHome server or client behaves. Some settings impact both the server and client; others are localized to either just the server or just the client.

NOTE: Updated PhoneHome settings do not take effect until the PhoneHome server or client is restarted or the next watchdog cycle occurs.

Viewing a List of Current PhoneHome Settings

To view a list of current PhoneHome settings:

1. Go to the console of the SL1 Collector or Database Server.
2. Run the following command on the SL1 Collector or Database Server:

```
sudo phonehome config list
```

When you run the command, the system returns a list that includes each configuration setting, its value, a description, and an indication of whether the setting affects the client, the server, or both.

Updating PhoneHome Settings

To set a new value for an existing PhoneHome setting:

1. Go to the console of the SL1 Collector or Database Server.
2. Run the following command on the SL1 Collector or Database Server:

```
sudo phonehome config set <setting_name> <new_value>
```

For example, if you want to change the client timeout value to 30 seconds, you would run the following command:

```
sudo phonehome config set client_timeout 30s
```

You can update the following settings:

Configuration	Setting	Description	Default Value	Affects
Client Timeout	client_timeout	Maximum amount of time allowed for the client to connect to a Database Server, after which the connection times out. The value is an actual time value, such as 30s, 5m, or 2h.	30s	Client
Exit on Forward Failure	exit_on_forward_failure	Indicates whether to close the connection to the Database Server if any custom ports fail to forward. This is not applicable to MariaDB port forwarding (port 7707). If the MariaDB port fails to forward, the client closes the connection regardless of this setting. The value is either true or false.	false	Client
Watchdog Frequency Duration	watchdog_freq	Amount of time between watchdog service cycles. The value is an actual time value, such 30s, 5m, or 2h.	1m0s	Both

Configuration	Setting	Description	Default Value	Affects
Fail Watchdog on Additional Forwards	<code>fail_wd_add_forward</code>	Indicates whether to close the connection from an SL1 Collector and mark it as disconnected if additional forwards fail. The value is either true or false.	false	Server
Port Ping Timeout	<code>port_ping_timeout</code>	Maximum allowed time for a Database Server's watchdog to connect to the forwarded port before it marks the SL1 Collector as disconnected and closes the incoming client connection. The value is an actual time value, such as 30s, 5m, or 2h.	10s	Server
Token Time to Live (TTL)	<code>token_ttl</code>	Default amount of time a token is valid before it expires. The value is an actual time value, such as 30s, 5m, or 2h. The maximum value is 2h.	30m0s	Server
Expired Token Cleanup Frequency	<code>expired_token_cleanup_freq</code>	Amount of time after which an expired token is deleted by the server. The value is an actual time value, such as 30s, 5m, or 2h.	48h0m0s	Server
Keepalive Timeout	<code>keepalive_timeout</code>	The timeout value for sending keepalive requests to the server. Adjusting this value can be helpful for PhoneHome collectors with high network latency. The value is an actual time value between 10s and 10m.	20s	Client

Clearing a PhoneHome Device

The `phonehome clear` command clears the PhoneHome configuration on a PhoneHome device. This command will also *disable* the PhoneHome `phd` service if it is run on the Database Server.

For PhoneHome SL1 Collectors, the `phonehome clear` command clears the PhoneHome configuration, stops the PhoneHome client, and deletes the client keys. However, it will not delete the collector's configuration that is stored on the Database Server. To delete the Database Server's configuration related to the client, you must use the `phonehome clear` command on the SL1 Collector and then execute the `phonehome delete` command on the Database Server.

For PhoneHome Database Servers, the `phonehome clear` command clears the PhoneHome configuration and stops the PhoneHome server. You can also use the `phonehome clear` command on a Database Server to block future connection requests from Data Collectors and secondary Database Servers in an HA/DR configuration.

To clear a PhoneHome device:

1. Go to the console of the SL1 Collector or Database Server.
2. Run the following command on the SL1 Collector or Database Server:

```
sudo phonehome clear
```

NOTE: For PhoneHome Database Servers, you can alternatively use the command `phonehome clear -d`. This deletes the device record associated with the Database Server, including the host key. For more information, see the section on [Deleting a PhoneHome Database Server](#).

Deleting a PhoneHome Collector

You can delete a PhoneHome SL1 Collector only if it has no corresponding Skylar One appliance.

Therefore, to delete a PhoneHome SL1 Collector, you must also perform the following steps, if applicable:

- If the SL1 Collector has a corresponding Skylar One appliance, you must delete that appliance before you can delete the SL1 Collector.
- If the corresponding Skylar One appliance is included in a collector group, you must delete that collector group before you can delete the appliance and then the Skylar One Collector. If there are more than one collectors in the collector group, you can edit the collector group to deselect that collector instead of deleting the collector group.
- If the Skylar One appliance's collector group includes other devices, you must move those devices to a different collector group before you can delete the appliance's collector group, then the appliance, and finally the Skylar One Collector.

WARNING: Once you delete a PhoneHome SL1 Collector, you cannot undelete it. Deleting an SL1 Collector will delete all configuration related to the device and cause all Database Servers to close incoming connections from the device.

To delete a PhoneHome SL1 Collector:

1. Go to the console of the SL1 Collector.
2. Run the following command on the SL1 Collector, replacing `<id>` with the PhoneHome device ID of the device you want to delete:

```
sudo phonehome delete <id>
```

NOTE: This command only works for deleting a collector. You cannot use this command to delete a Database Server.

One of the following will occur:

- If the device does not have a corresponding Skylar One appliance on the stack, a confirmation prompt appears, asking you to confirm that you want to delete the device. Type "Y" and press Enter. The device is deleted and you can skip the rest of this section.
- If the device does have a corresponding Skylar One appliance, a message similar to the following appears:

```
Error: Cannot delete a phonehome device that has a corresponding
appliance: [Module ID: 10, Name: example-device-cu1, CUG(s):
cug-dc09]
```

If you receive an error message, proceed to the next step.

3. Go to the **Appliance Manager** page (System > Settings > Appliances).
4. Locate the device with the *ID* that matches the *Module ID* value that was returned in the error message in step 2, and then do one of the following:
 - If the appliance is not part of a collector group, click its delete icon (>Delete) to delete it. You can then repeat steps 1 and 2 to delete the SL1 Collector.
 - If the appliance is part of a collector group, the delete icon is disabled. Proceed to the next step.
5. Go to the **Collector Group Management** page (System > Settings > Collector Groups).
6. Locate the collector group with the name that matches the *CUG* value that was returned in the error message in step 2, and do one of the following:
 - If the collector group does not contain any devices, click its delete icon (>Delete) to delete it. You can then repeat steps 3 and 4 to delete the appliance.
 - If the collector group contains devices, the delete icon is disabled. Proceed to the next step.
7. Go to the **Device Manager** page (Registry > Devices > Device Manager).
8. Select the checkbox for each device that you want to move to a different collector group.
9. In the **Select Action** field (in the lower right), select *Change Collector Group* and then select a collector group.
10. Click the **[Go]** button. The selected devices will now be aligned with the selected collector group.
11. Repeat steps 5 and 6, and then work your way backwards as needed, completing steps 3 and 4, followed by steps 1 and 2. Repeat these steps as needed until the device is deleted successfully in step 2.

Deleting a PhoneHome Database Server

To delete a PhoneHome Database Server:

1. Go to the console of the Database Server that you want to delete.
2. Run the following command:

```
sudo phonehome clear -d
```

A confirmation prompt appears, asking you to confirm that you want to delete the device. Type "delete" and press Enter.

NOTE: You must run this command from the Database Server that you want to delete. You cannot run it from any other Database Server or the Administration Portal.

WARNING: Once you delete a PhoneHome Database Server, you cannot undelete it. Deleting a Database Server will delete all configuration related to the device and close all incoming connections from PhoneHome SL1 Collectors.

Troubleshooting PhoneHome Configurations

This section describes how to troubleshoot issues some users experience when configuring PhoneHome communications.

Connectivity Issues from a Collector

You can run the following command on the SL1 Collector or Database Server to check connectivity issues:

```
sudo phonehome check -x
```

This command visualizes the network path from the SL1 Collector to the Database Server as well as any proxy hops in between, if applicable. The output reports back any failures connecting to any hop.

These are some of the common error messages seen with the disconnected host:

ssh: handshake failed: ssh: unable to authenticate, attempted methods [none publickey], no supported methods remain

There are two possible causes if the disconnected error is shown on the database host:

- Client keys have been reconfigured on the collector.
- The server does not have a valid record of the client. This would happen if a database administrator would delete the device record, but would not run clear on the collector itself.

If this happens on an intermediary proxy host, this means that the SSH key-based authentication has not been set properly with the proxy host.

ssh: handshake failed: knownhosts: key mismatch

This means there is an old entry for the given destination (or proxy) in `/etc/phonehome/known_hosts` that needs to be deleted from the file.

dial TCP <database_host_addr>:<port>: i/o timeout

This issue can be caused due to any of the following reasons:

- The Database Server is inaccessible or shut down.
- The Database Server is up but the **phd** service is down.
- A firewall rule has been added that prevents a connection from the SL1 Collector to the Database Server.

dial TCP <database_host_addr>:<port>: connect: no route to host

This error means that either the Database Server is shut down or it is experiencing a network connectivity issue.

dial TCP <database_host_addr>:<port>: connect: connection refused

This error means that the **phd** service on the Database Server host is not active/running.

Register Command Complains that the Token Has Expired

A PhoneHome token has a default time to live of 30 minutes, although this can be extended up to two hours using the command-line interface to generate the token. After this set time, the token expires. The `register` command lets you know that the token is expired and the Database Server will reject the request if you attempt to use it.

If this happens, you have two options:

- Ask the database administrator to issue you a new token since the old one has expired.
- Send a request from the SL1 Collector instead and let the database administrator know the one-time secret so they can accept the request on the Database Server.

You Cannot See a Request You Sent on the Server and You Cannot Send Another Request

When you send a request, the request is stored on the Database Server for an administrator to accept or reject. A request never expires.

If there is any failure with storing the request, the `phonehome request send` command will fail and display an error. This can happen if a database administrator deletes or rejects the request by mistake.

The SL1 Collector does not get any feedback when an administrator rejects a request on the Database Server, and the tool prevents you from sending duplicate requests because it thinks that there is already a queued request.

You can override this by using the `-f | --force` flag with the `phonehome request send` command.

Status Shows Disconnected but the Check Succeeds

This means that the SL1 Collector is able to connect to the Database Server successfully but is failing to forward the ports.

Status changes are not immediate. To determine a collector's status, the Database Server needs to run a watchdog cycle, which happens every minute by default. Therefore, if you have very recently registered an SL1 Collector or restarted the `phc` service, wait for another watchdog cycle to see if the status changes from disconnected to forwarded. If this does not happen, you can check the logs for more details on the forwarding issue. To do so, use the following commands:

- On the Database Server: `journalctl -u phd.service -f -n`
- Client SL1 Collector: `journalctl -u phc.service -f -n`

Chapter

8

Installing Skylar One on AWS

Overview

This chapter describes how to install Skylar One (formerly SL1) on an Amazon Web Services EC2 instance, which is a virtual server that resides in the AWS cloud.

Use the following menu options to navigate the Skylar One user interface:

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

This chapter covers the following topics:

<i>AWS Instance Specifications</i>	104
<i>Deploying a Skylar One System on AWS</i>	104
<i>What are the ScienceLogic AMIs?</i>	104
<i>Getting the ScienceLogic AMI</i>	105
<i>Launching the New Instance</i>	106
<i>Security Rules for Each Appliance Type</i>	110
<i>Database Server</i>	111
<i>Administration Portal</i>	113
<i>Data Collector</i>	115
<i>Message Collector</i>	116
<i>Additional Configuration Steps</i>	117

<i>Assigning an EIP to the New Instance</i>	118
<i>Accessing the Appliance Using SSH</i>	119
<i>Connecting to Your Instance</i>	119
<i>Configuring the EC2 Instance</i>	120
<i>Web Configuration Tool</i>	121
<i>Rebooting Data Collectors and Message Collectors</i>	122

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

AWS Instance Specifications

For details about AWS and the requirements and specifications for each Skylar One appliance, see the ScienceLogic Support Site: <https://support.scienceologic.com/s/system-requirements?tabset-3429b=db66f>.

Deploying a Skylar One System on AWS

For ease of configuration, create nodes or appliances in this order:

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

NOTE: The following instructions describe how to configure a ScienceLogic virtual machine in AWS. If you are looking for resources and support for AWS Cloud, see the Amazon AWS Marketplace: <https://aws.amazon.com/marketplace/>.

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 Center](#) and go to the **ScienceLogic Product Licensing** page (Support > License & Image Requests).
2. Under the **Amazon AWS AMI Request** heading, click **[Submit AMI Request]**. The **Request Amazon AMI** page appears.

NOTE: If you are an Amazon Web Service GovCloud user, you will need to contact ScienceLogic Support to get the ScienceLogic AMI.

3. Fill out the **Request Amazon AMI** form and click the **[Submit AMI Request]** button.
4. Repeat steps 2-4 for each type of Skylar One appliance you want to install on AWS.
5. ScienceLogic Customer Support will send you an email confirming that they have shared the ScienceLogic AMI with your AWS account.
6. To view the ScienceLogic AMIs in your AWS account, go to the **AWS Management Console** page. Under the heading **Compute**, click **[EC2]**.
7. In the **EC2 Dashboard** page, go to the left navigation bar. Under the heading **Images**, click **[AMIs]**.
8. In the main pane, under **Filters**, click **[Owned by me]** and then select *Private images*.
9. You should see AMIs with names that begin with "EM7" and end with the current release number for Skylar One. You should see an AMI for each type of Skylar One appliance.

10. 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

This chapter describes how to launch a new EC2 instance from the ScienceLogic AMI.

Before you begin, be advised of the following:

- To complete the steps listed in this chapter, you must have already *received the ScienceLogic AMIs*. If you have just completed the steps in that section, you can start at step 4 in this section.
- 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. For more information on VPCs and VPC subnets, see http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Introduction.html.
- You can use multiple AWS instances to create a distributed Skylar One system. For each instance, you must specify the correct instance type, storage size, and security rules. All these parameters are described in this chapter.
- For details about the recommended instance type for each ScienceLogic appliance, see the *System Requirements* page on the ScienceLogic Support site.

To launch the new EC2 instance from the ScienceLogic AMI:

1. Go to the *EC2 Dashboard*.
2. In the left navigation bar, under the heading *Images*, click **[AMIs]**.
3. In the main pane, under **Filters**, click **[Owned by me]** and then select *Private images*.
4. From the list, select the checkbox of the ScienceLogic AMI that matches the ScienceLogic appliance you want to create, then click the **[Launch instance from AMI]** button.
5. On the **Launch an instance** page, complete the following fields:
 - **Name and tags**. Add a descriptive name and one or more tags for this instance.

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

- **Application and OS Images (Amazon Machine Image)**. This field is prepopulated with your ScienceLogic AMI.

- **Instance Type.** Select the instance type recommended for the AMI that meets the [system requirements](#) for the instance you are launching.
- **Key pair (login).** Create a new key pair to connect to your instance. Alternatively, you can select an existing key pair, but only if have previously downloaded and saved the private key, as you cannot retrieve an existing private key a second time.

NOTE: Amazon EC2 instances use public-key cryptography for authentication. When you create a new key pair, AWS stores the public key on its servers and automatically downloads the file that contains the private key to your browser in a file that ends in ".pem". You will need this .pem file again when you [configure SSH](#) access to your AWS instances.

- **Network settings.** Expand this section, click [**Edit**], and update the fields as needed based on your environments needs. Options include:

- **VPC.** For accounts enabled for virtual private clouds, select the network where the instance will reside. If you are unsure of the network, use the default, which is based on your region.
- **Subnet.** For VPC-enabled accounts, select or create the subnet where the instance will reside. If you are unsure of the subnet, use the default.
- **Auto-assign Public IP.** If you select *Enable*, AWS will assign an IPv4 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, hibernated, 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>.

- **Auto-assign IPv6 IP.** Select whether you want to *Enable* or *Disable* the ability for AWS to automatically assign an IPv6 address to this instance.

NOTE: If you select *Enable* in the **Auto-assign IPv6 IP** field, the IP address will change each time the instance is terminated, but not when it is stopped or hibernated. You cannot assign an elastic IP address for IPv6.

- **Firewall (security groups).** Select an existing security group or create a new security group for your instance. You must ensure that your security group has rules that allow traffic to and from your AWS instances, as all other traffic will be ignored. If you create a new security group, add a name and description of the security group as well as inbound security group rules. Use the *following tables* to create security rules for each type of Skylar One appliance. After completing each row, click the [**Add security group rule**] button.

- **Configure storage.** Add the amount of storage you need that meets the *system requirements* for the instance you are launching. Using the **Advanced** view, increase the size of the /dev/sda1 partition as follows:

Skylar One 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

In addition, make the following update in this section:

- **Delete on Termination.** Select *Yes*.
- **Advanced details.** Expand this section and update the fields as needed based on your environment's needs. At a minimum, update the following fields:
 - **IAM instance profile.** If your organization uses IAM roles, select the appropriate role.
 - **Shutdown behavior.** Select *Stop*.
 - **Termination protection.** Configure this setting according to your organization's procedures.
 - **Detailed CloudWatch monitoring.** Select *Disable*.
 - **EBS-optimized instance.** Select *Disable*.
 - **Tenancy.** Select *Shared - Run a shared hardware instance*.
 - **Metadata accessible.** Select *Enabled*.
 - **Metadata version.** Select *V1 and V2 (token optional)*.

NOTE: For more information about all of your options when launching a new instance, see <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-launch-parameters.html>.

6. In the **Summary** panel, enter the number of instances you need to launch in the **Number of instances** field and then click **[Launch instance]**.

NOTE: It might take several minutes for your instance to launch.

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

Security Rules for Each Appliance Type

NOTE: Configure this list according to your requirements, your AWS configuration, and your security rules.

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 <i>My IP</i> . If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.	SSH. For SSH sessions from the 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 <i>My IP</i> . If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.	HTTP from browser session on user workstation.

Type	Protocol	Port Range	Source	Description
HTTPS	TCP	443	If you will always log in from a single IP address, select <i>My IP</i> . If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.	HTTPS from browser session on user workstation.
Custom TCP Rule	TCP	7700	If you will always log in from a single IP address, select <i>My IP</i> . If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.	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.	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.	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.	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.	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 <i>My IP</i> . If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.	SSH. For ssh sessions from user workstation to the appliance. This is necessary to start the installation wizard.
SMTP	TCP	25	Specify a list of IP addresses for all managed devices from which you want to receive email messages.	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 <i>My IP</i> . If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.	HTTP from browser session on user workstation.
Custom TCP Rule	TCP	123	Enter the IP address of the NTP server.	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.	SNMP Agent. Allows SNMP information about the Database Server to be collected by Skylar One.
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 <i>My IP</i> . If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.	HTTPS from browser session on user workstation.

Type	Protocol	Port Range	Source	Description
Custom TCP Rule	TCP	7700	If you will always log in from a single IP address, select <i>My IP</i> . If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.	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.	MySQL. Communication from Administration Portal
Custom TCP Rule	TCP	8008	If you will always log in from a single IP address, select <i>My IP</i> . If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.	Administrative Web Interface (PHPMyAdmin) from browser session on user workstation
Custom TCP Rule	TCP	8200	If there is a firewall between the Database Server, Data Engine, and Administration Portal appliances, this port must be open to enable Enterprise Key Management Service (EKMS) cluster communication between those appliances.	EKMS Cluster Communication

Administration Portal

Inbound

Type	Protocol	Port Range	Source	Description
SSH (edit the default SSH rule)	TCP	22	<p>If you will always log in from a single IP address, select <i>My IP</i>.</p> <p>If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.</p>	SSH. For ssh sessions from user workstation to the appliance. This is necessary to start the installation wizard.
HTTP	TCP	80	<p>If you will always log in from a single IP address, select <i>My IP</i>.</p> <p>If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.</p>	HTTP from browser session on user workstation.
HTTPS	TCP	443	<p>If you will always log in from a single IP address, select <i>My IP</i>.</p> <p>If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.</p>	HTTPS from browser session on user workstation.
Custom TCP Rule	TCP	123	Enter the IP address of the NTP server.	NTP. Communication between the Administration Portal and configured NTP server.
Custom TCP Rule	TCP	7700	<p>If you will always log in from a single IP address, select <i>My IP</i>.</p> <p>If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.</p>	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 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.	SNMP Agent. Allows SNMP information about the Administration Portal to be collected by Skylar One.
Custom TCP Rule	TCP	8200	If there is a firewall between the Database Server, Data Engine, and Administration Portal appliances, this port must be open to enable Enterprise Key Management Service (EKMS) cluster communication between those appliances.	EKMS Cluster Communication

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 <i>My IP</i> . If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.	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.	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.	SNMP Agent. Allows SNMP information about the Data Collector to be collected by Skylar One.

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.	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.	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 <i>My IP</i> . If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.	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 Data Collector.	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 <i>My IP</i> . If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.	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.	NTP. Communication between the Message Collector and configured NTP server.

Type	Protocol	Port Range	Source	Description
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.	SNMP Agent. Allows SNMP information about the Message Collector to be collected by Skylar One.
Custom UDP Rule	UDP	162	Specify a list of IP addresses for all managed devices from which you want to receive SNMP traps.	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.	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 <i>My IP</i> . If you will log in to the instance from multiple IP addresses, enter those IP addresses, separated by commas, in this field.	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.	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](#)
 - [Configuring the EC2 Instance](#)
 - [Licensing the Appliance](#)
- For instances of the **Administration Portal**:

- *Assigning an EIP to the instance* (optional step)
- *Accessing the Appliance Using SSH*
- *Configuring the EC2 Instance*
- *Configuring the Appliance*
- For instances of the **Data Collector and Message Collector**:
 - *Assigning an EIP to the instance* (optional step)
 - *Accessing the Appliance Using SSH*
 - *Configuring the EC2 Instance*
 - *Configuring the Appliance*
 - *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, meaning that 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*.
2. In the left navigation pane, under the **Network & Security** heading, click **[Elastic IPs]**.
3. Click **[Allocate Elastic IP address]**.

4. On the **Allocate Elastic IP address** page, update the EIP settings and tags based on your needs for your Skylar One instance. When you are finished, click **[Allocate]**.
5. From the EC2 Dashboard, in the left navigation pane, under the **Network & Security** heading, click **[Elastic IPs]**.
6. Select the EIP you allocated, then click the **[Actions]** menu and select *Associate Elastic IP address*.
7. In the **Resource type** field, select *Instance*, then select the Skylar One appliance instance you want to associate with the EIP.
8. Click **[Associate]**. The Skylar One appliance instance is now associated with the new EIP.

Accessing the Appliance Using SSH

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

This chapter assumes that you have access to SSH on the command line.

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 Skylar One 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 Skylar One 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 Skylar One 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](#).

Connecting to Your Instance

On Unix and Linux systems, you can connect to your Skylar One 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 Skylar One appliance instance.

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

Configuring the EC2 Instance

To configure each new EC2 instance, perform the following steps:

1. Use SSH to access the EC2 instance using its public IP address, username, and the SSH key defined in the section [Accessing the Appliance Using SSH](#):

```
ssh -i <private key path> em7admin@<vm-ip-address>
```

2. If you are performing a fresh installation, you will be prompted by the Message of the Day to set up the MariaDB password.
3. If this is a new STIG installation, the Message of the Day will also contain instructions for setting a new password. Follow these instructions for setting a password on the **em7admin** account.

If you are updating a STIG system, the Message of the Day will display a security warning stating that the **em7admin** account does not have a password, along with instructions for setting the password.

4. Use the following command to edit the `/etc/silo.conf` file:

```
sudo visilo --no-validation
```

5. In the `/etc/silo.conf` file, update the following section or sections:

- For the **clientdbuser** account:

```
[LOCAL]
```

```
dbpasswd = <NEW_PASSWORD>
```

```
[CENTRAL]
```

```
dbpasswd = <NEW_PASSWORD>
```

NOTE: The `CENTRAL` section does not appear for all appliance types. If it does, then the `dbpasswd` values should match in both sections.

- For the **ap_user** account:

```
[CENTRAL]
```

```
ap_user = apuser
```

```
ap_pass = <NEW_PASSWORD>
```

NOTE: The `CENTRAL` section does not appear for all appliance types.

6. Save the file (`:wq`) and enter `y` to move the changes to the `/etc/siteconfig/siloconf.siteconfig` file automatically.
7. If you are upgrading Skylar One, run the following command:

```
sudo systemctl restart nextui php-fpm nginx
```

NOTE: If this is a new STIG installation, you can skip this step.

8. Repeat these steps on the other Skylar One appliances in your stack as needed to update the passwords for those appliances as well.

Web Configuration Tool

- For instances of the **Database Server** or **All-In-One Appliance**, see the section on [Licensing and Configuring a Database Server or All-In-One Appliance](#).
- For instances of the **Administration Portal**, see the section on [Configuring an Administration Portal](#).

- For instances of the **Data Collector and Message Collector**, see the section on [Configuring a Data Collector or Message Collector](#).

Rebooting Data Collectors and Message Collectors

After installing a Skylar One appliance 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
```

Chapter

9

Installing Skylar One in Azure

Overview

This chapter describes how to deploy a Skylar One (formerly SL1) virtual machine in Azure from a virtual hard disk (VHD) file. ScienceLogic supports the following subscription types for deploying Skylar One in Azure:

- Azure Global
- Azure Government

NOTE: For Azure resources and support, see the Microsoft Azure Marketplace:
<https://azuremarketplace.microsoft.com/>

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

NOTE: High-availability for Azure deployments is supported for installations of 12.1.x and later that are running on Oracle Linux 8 (OL8). ScienceLogic recommends that customers running Skylar One versions prior to 12.1.x upgrade to 12.1.x or later, convert to OL8, and then complete the high-availability setup and configuration. For more information about upgrading, see the section on "Updating Skylar One" in the ***System Administration*** manual.

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

Use the following menu options to navigate the Skylar One user interface:

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

This chapter covers the following topics:

<i>Azure Instance Specifications</i>	125
<i>Prerequisites</i>	125
<i>Generating a Source Blob SAS URL</i>	125
<i>Transferring Skylar One VHD Files to Your Azure Blob Container</i>	126
<i>Creating the Azure Image</i>	127
<i>Skylar One Appliance Deployment Order for Distributed Systems</i>	127
<i>Creating an Azure Virtual Machine</i>	128
<i>Setting the Public IP Address to Static</i>	129
<i>Configuring the Virtual Machine</i>	129

Azure Instance Specifications

For details about Azure and the requirements and specifications for each Skylar One appliance, see the ScienceLogic Support Site: <https://support.scienceologic.com/s/system-requirements?tabset-3429b=f4ac1>.

Prerequisites

You must have *an Azure Resource group and a storage account that includes at least one blob container* before attempting to deploy Skylar One.

In addition, before attempting to deploy Skylar One in Azure, you should confirm that the following prerequisites are all true:

- Your Azure policies allow you to create a new virtual machine using the ScienceLogic virtual hard disk (VHD).
- Your virtual machine naming convention does not conflict with any existing policies in Azure.
- You are using virtual networks and subnets that allow access and the creation of new virtual machines.
- Your resource group allows you to create new virtual networks.

Configuring an Azure Resource Group and Storage Account

To create and configure an Azure storage account:

1. Log in to the Azure Portal and follow Microsoft's instructions for [Creating resource groups](#) to create a new Azure resource group.
2. After creating the resource group, follow Microsoft's instructions for [Creating an Azure storage account](#) to create a new storage account for your resource group. While configuring the storage account, make sure to associate it with the resource group you created in step 1.
3. After creating the storage account, follow Microsoft's instructions for [Creating a blob container](#) to add a new container to the storage account you created in step 2.

Generating a Source Blob SAS URL

To upload the Skylar One VHD file to your Azure storage account, you must generate a source URL in the Azure portal. You will need this URL when you *transfer the Skylar One VHDs to your Azure blob container*.

To generate a source blob SAS URL:

1. Log in to the Azure portal and navigate to your storage account, then to your container.
2. Right-click the container name and then select *Generate SAS*.

3. In the **Generate SAS** pane, in the *Signing method* field, select *User delegation key*.
4. In the **Permissions** field, select *Read* and *Write*.
5. In the **Start** and **Expiry** fields, specify the start and end dates for the token's validity. The token will be valid only within this time frame.
6. Click **[Generate SAS Token URL]**, then copy the source blob SAS URL.

Transferring Skylar One VHD Files to Your Azure Blob Container

You can transfer Skylar One VHD files directly to your Azure blob container from the ScienceLogic Support Center. This enables you to transfer all appliance types for Skylar One in parallel in just a few minutes.

To transfer Skylar One VHD files to your Azure blob container:

1. Log in to the [ScienceLogic Support Center](#) and go to the **ScienceLogic Product Licensing** page (Support > License & Image Requests).
2. Click the **[Request Azure VHD]** button. The **Azure VHD Request** page appears.
3. Review and, if necessary, complete the prerequisites listed on the page , then complete the following fields:
 - **Please Choose a Product.** Select *Skylar One*.
 - **Product Type or Version.** Select the version of Skylar One you want to deploy.
 - **Appliance Types.** Select the Skylar One appliances you want to deploy.
 - **Blob SAS URL.** Type or paste [the Blob SAS URL you generated](#).
4. Click the **[Submit Azure VHD Request]** button. The VHDs are transferred to your desired container.

NOTE: This transfer might take several minutes.

Creating the Azure Image

After transferring the Skylar One VHD files to your Azure blob container, you must create the Azure image.

You can do so in one of two ways:

- Using Azure Compute Gallery. For instructions, see Microsoft's documentation on [Creating an image definition using Azure Compute Gallery](#).
- Using Azure's legacy managed images feature. Instructions for this method are included below. To use this method, you will need the following information:
 - [The Skylar One VHD file](#) and the local file path to the file
 - [Resource Group name](#)
 - [Container URI](#)

To create the Azure image using Azure's legacy managed images feature:

1. Open Microsoft Azure PowerShell and log in to your Azure account:

```
Login-AzureRmAccount
```

2. Run the following command:

```
az image create --name <image_name> -g <resource_group_name> --os-disk-caching ReadWrite --os-type Linux --source <container_URI>
```

where:

- `<image_name>` specifies what you want to call the Azure image (such as "dbimage123").
- `<resource_group_name>` specifies the resource group you created in Azure.
- `<container_URI>` specifies the destination value you provided when you uploaded the Skylar One VHD file. When entering the container URI, you must include the .vhd file name. For example: <https://azuretest.blob.core.windows.net/vhds/skylaroneinazure.vhd>

NOTE: This command might return a large amount of JSON output. This is normal.

3. Repeat these steps for each Skylar One appliance you want to build.

Skylar One Appliance Deployment Order for Distributed Systems

To deploy a distributed Skylar One 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)

Creating an Azure Virtual Machine

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

To create an Azure virtual machine (VM) using the Azure portal:

1. Log in to your Azure account, then go to the [Microsoft Marketplace](#).
2. In the search bar, type "ScienceLogic," then select the Skylar One appliance type that you want to deploy.
3. From the image's **Overview** page, click **[Get It Now]**.
4. In the modal that appears, review the software plan details and then click **[Continue]** to confirm the agreement. Your Azure portal opens with the image download already selected.
5. Click **[Create]** to start the VM creation process.
6. During the VM creation process, do the following:
 - Under **Project Details**, select the appropriate **Resource group**.
 - Under **Instance Details**, follow the [System Requirements](#) for sizing the VM based on the appliance type you are deploying.
 - Under **Administrator Account**:
 - Set the **Authentication type** to *SSH public key*.
 - Enter the **Username** as "em7admin".
 - For the **SSH public key source**, use an existing key or generate a new one.
 - Under **Inbound Port Rules**, select the necessary *inbound ports*.
7. Complete the rest of the VM creation steps based on your specific needs. (For more information, see Microsoft's instructions for [Creating a Linux VM](#).) When you are finished, click **[Review + create]**.
8. Review the details for the VM you are about to create. When you are ready, click **[Create.]**
9. When the **Generate new key pair** window appears, click **[Download private key and create resource]**. When you do so, your key is downloaded as the file **myKey.pem**. Make note of the file's download location.

10. When the deployment has completed, click **[Go to resource]**.
11. On the page for your new VM, make note of the VM's **Public IP address**.
12. Repeat these steps for each Skylar One appliance you want to build.

Setting the Public IP Address to Static

To ensure the IP address for the virtual machine remains the same after reboot, which ensures that your Skylar One appliances work properly, you must configure the public IP address to be static. To do so:

1. In the Azure Portal, enter "network interfaces" in the search box at the top of the portal. Select **Network interfaces** in the search results.
2. From the list of network interfaces, select the network interface you want to modify.
3. Click **Settings > IP configurations**, and then select the IP configuration that you want to modify.
4. In the **Edit IP configuration** window, click **[Disassociate]** in the **Public IP Address** field, and then select **[Associate public IP address]**.
5. In the **Public IP address** field, click **[Create a public IP address]**.
6. Under **Add a public IP address**, type a name for your IP address in the **Name** field and select **Static** in the **Assignment** field.
7. Click **[OK]** and then click **[Save]**.
8. Repeat these steps for each Skylar One appliance you want to build.

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, username, and the SSH key defined when *creating your Azure virtual machine*.

```
ssh -i <private key path> em7admin@<vm-ip-address>
```

2. If you are performing a fresh installation, you will be prompted by the Message of the Day to set up the MariaDB password.
3. If this is a new STIG installation, the Message of the Day will also contain instructions for setting a new password. Follow these instructions for setting a password on the **em7admin** account.

If you are updating a STIG system, the Message of the Day will display a security warning stating that the **em7admin** account does not have a password, along with instructions for setting the password.

4. Use the following command to edit the **/etc/silo.conf** file:

```
sudo visilo --no-validation
```

5. In the `/etc/silo.conf` file, update the following section or sections:

- For the `clientdbuser` account:

```
[LOCAL]
```

```
dbpasswd = <NEW_PASSWORD>
```

```
[CENTRAL]
```

```
dbpasswd = <NEW_PASSWORD>
```

NOTE: The `CENTRAL` section does not appear for all appliance types. If it does, then the `dbpasswd` values should match in both sections.

- For the `ap_user` account:

```
[CENTRAL]
```

```
ap_user = apuser
```

```
ap_pass = <NEW_PASSWORD>
```

NOTE: The `CENTRAL` section does not appear for all appliance types.

6. Save the file (`:wq`) and enter `y` to move the changes to the `/etc/siteconfig/siloconf.siteconfig` file automatically.
7. If you are upgrading to this release of Skylar One, run the following command:

```
sudo systemctl restart nextui php-fpm nginx
```

NOTE: If this is a new STIG installation, you can skip this step.

8. Repeat these steps on the other Skylar One appliances in your stack as needed to update the passwords for those appliances as well.

Chapter

10

Navigating the Setup and Config Page

Overview

This chapter describes how to navigate the **Setup and Config** page in Skylar One (formerly SL1) to help you get started with Skylar One.

Use the following menu options to navigate the Skylar One user interface:

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

This chapter covers the following topics:

<i>What is the Setup and Config Page?</i>	132
<i>Setup and Config Journey Workflows</i>	132

What is the Setup and Config Page?

The **Setup and Config** page (🔗) displays all information relevant to getting started in Skylar One for administrator-level users. Included on this page are a number of *journeys*, intuitive self-service workflows that will guide you through the most common Skylar One system tasks. Click the name of a workflow to get started.

This page also contains informational cards that provide you with the proper resources for Skylar One setup and configuration.

The informational cards on this page include:

- **Get Started.** Displays a list of available user journeys and their journey status. Click the name of the journey to get started. The journeys include:
 - Take a Tour of Skylar One
 - Discover and Monitor Hybrid Cloud Infrastructure
- **Resources.** Hosts additional external resources to help you with setup and configuration; these links include:
 - Training Portal
 - ScienceLogic Support
- **Overview.** Provides links to the user journeys. These journeys include guided tours and interactive wizards that help you set up and refine your Skylar One environment:
- **Next Steps.** Contains links to other pages in Skylar One where you can continue working after completing some or all of a journey:
 - Manage Devices
 - Manage Collector Groups
 - Manage Organizations
 - Manage Users
 - Manage Access Hooks

Setup and Config Journey Workflows

This section provides the information you need to follow the Setup and Config user journeys. You can use the Setup and Config page's journey cards as a guide to the overall Skylar One setup and configuration process.

For the best experience in following the Setup and Config journeys, it is recommended that you:

1. Familiarize yourself with the Skylar One product as whole by clicking through the **[Take a Tour of Skylar One]** journey and tracking your knowledge process with the journey's status buttons.

2. Follow the steps in the **[Discover and Monitor Hybrid Cloud Infrastructure]** journey card for a proper understanding of the setup and configuration process for your specific infrastructure. This space provides all of the information required for a successful setup in Skylar One; track your progress throughout with the journey's status buttons.

The status buttons on a card let you apply a specific status for an activity.

You can apply the following statuses for each journey card's individual activities:

- **[Not Started]**. This status serves as a "to-do" status for an activity that hasn't been attempted yet.
- **[In Progress]**. This status allows you to track and re-enter activities that have been started, but not completed.
- **[Complete]**. This status should be applied to any activity that is finished. You can also reset a completed workflow if you need to be guided through a workflow again. To do this, see *Resetting a Journey Workflow*.

IMPORTANT: A workflow can be set as *Not Applicable* if that workflow and its activities do not apply to you. This status removes that infrastructure's workflow from your "to-do" list and the workflow will not be tracked. To do this, see *Setting a Journey Workflow as "Not Applicable"*.

Taking a Tour of Skylar One


The first journey you can view in the **Overview** information card is the **[Take a Tour of Skylar One]** journey. This journey provides you a space to access product videos and track your progress as you learn the Skylar One system.

To access Skylar One educational product videos:

1. In the **Get Started** section, click the arrow (>) button next to the status in the **[Take a Tour of Skylar One]** journey card. A new **[Take a Tour of Skylar One]** card appears.
2. On the new card, click the pop-out redirect button (↗) to be redirected to a Skylar One product video landing page. The videos located on this page contain informational walk-throughs for Skylar One's key features and use cases.

To update your **[Take a Tour of Skylar One]** status:

1. In the **Get Started** section, click the arrow (>) button next to the status in the **[Take a Tour of Skylar One]** journey card.
2. As you become more familiar with Skylar One, click to update the status button drop-down. You can select **[Not Started]**, **[In Progress]**, or **[Complete]**. Your selected status then updates and appears across the entire workflow and **Setup and Config** pages:



Discover and Monitor Hybrid Cloud Infrastructure

The second journey card available allows you to onboard AWS, Azure, or VMware applications in order to begin data collection. This process is called "guided discovery". The workflow for each application provides a checklist of onboarding workflow activities.


NOTE: If you want to discover one of the third-party products that are available as an option when using guided discovery, you must have the corresponding PowerPack installed on your Skylar One system. For example, if you want to discover an Amazon Web Services account, you must have the "Amazon Web Services" PowerPack installed.

Onboarding your Hybrid Cloud Infrastructure

The onboard workflow for this Setup and Config user journey guides and points you to the correct pages for your infrastructure's setup.

To onboard your hybrid cloud infrastructure:


1. In the **Overview** section of the **Setup and Config** page, click the arrow (>) next to the status button in the **[Discover and Monitor Hybrid Cloud Infrastructure]** journey card. A new **[Discover and Monitor Hybrid Cloud Infrastructure]** card appears.
2. On the new card, click the arrow (>) to select your service infrastructure:



3. A checklist of activities appears. Click the pop-out redirect button (↗) to review how to complete each individual activity for your specific infrastructure (such as **AWS**). This button redirects you to a Skylar One Product Documentation page with the relevant information to walk-through the on-boarding activity.
4. The different workflow activities point you to the relevant documentation or the corresponding on-boarding wizard for each activity:
 - **Provide Access:** See the corresponding documentation to set up Skylar One credentials.
 - **Deploy Collector:** See the *Installing and Configuring an SL1 Collector* section in the Installation manual.
 - **Discover Infrastructure:**
 - For Amazon Web Services: See the **AWS Guided Discovery** section in the *Monitoring Amazon Web Services* manual.
 - For Azure: See the **Microsoft Azure Guided Discovery** section in the *Monitoring Microsoft Azure* manual.

- For VMware: See the **VMware Guided Discovery** section in the **Monitoring VMware** manual.

5. After you have completed the activities, you can return to the workflow pictured below and mark your progress as **[Complete]**. The workflow status updates in both the **[Discover and Monitor Hybrid Cloud Infrastructure]** journey card and the **[Get Started]** information card.



TIP: When you update the activity status to **[In Progress]**, the workflow status updates in both the **[Discover and Monitor Hybrid Cloud Infrastructure]** journey card and the **[Get Started]** information card as well.

Resetting a Completed Journey Workflow

The onboard workflow for this Setup and Config user journey allows you to reset an already completed infrastructure setup if needed. You can also select more than one workflow if multiple are needed to reset.

To reset a journey workflow:

1. From the **Discover and Monitor Hybrid Cloud Infrastructure** page, click the ellipses icon (⋮) and select **[Reset]**.
2. Select your desired infrastructure in the **Reset Workflows** model and click **[Confirm]**. That infrastructure will then appear with the **Not Started** status on your **Discover and Monitor Hybrid Cloud Infrastructure** page.

Setting a Journey Workflow as "Not Applicable"

You have the option to set an individual journey's workflow as **Not Applicable** on the **Discover and Monitor Hybrid Cloud Infrastructure** page. By setting a workflow as **Not Applicable**, the journey's workflow page moves that activity's workflow to the bottom of the page along with any other **Not Applicable** workflows. This keeps your activity workflows organized and allows for easy tracking.

To set an activity workflow as **Not Applicable**:

1. From the **Discover and Monitor Hybrid Cloud Infrastructure** page, click the ellipses icon next to your infrastructure and select **[Not Applicable]**.
2. A **Dismiss Workflows** modal appears. Select the workflow(s) that are not applicable to you. Skylar One will then organize that selection as ***Not Applicable*** and remove it from your immediate view on the page.
3. Click **[Confirm]**. The infrastructure(s) workflow will then appear as ***Not Applicable*** on the **Discover and Monitor Hybrid Cloud Infrastructure** page.
4. If you click the arrow button (>) next to the ***Not Applicable*** workflow, that workflow and its subsequent workflow activities will appear greyed out. To undo the ***Not Applicable*** status from this page and display the workflow again, click **[Display]** and confirm your changes.

Chapter

11

Updating Skylar One

Overview

For information on updating an existing Skylar One (formerly SL1) system, see the *Updating Skylar One* chapter of the ***System Administration*** manual, which describes how to update the software on your Skylar One appliances.

© 2003 - 2026, ScienceLogic, Inc.

All rights reserved.

ScienceLogic™, the ScienceLogic logo, and ScienceLogic's product and service names are trademarks or service marks of ScienceLogic, Inc. and its affiliates. Use of ScienceLogic's trademarks or service marks without permission is prohibited.

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 herein, the information provided in this document 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 herein at any time without notice.



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

International: +1-703-354-1010