



System Administration

Skylar One version 12.5.4

Table of Contents

Introduction	14
Who Should Read This Manual?	15
What's In This Manual?	15
Requirements	15
Global Settings	16
Global Settings for API	17
Global Settings for Appliances	18
The Web Configuration Utility	21
Global Settings for Asset Automation	22
Global Settings for System Behavior	23
Configuring Single Instance Login for AP2	39
Global Settings for Data Retention	40
Normalization and Roll-Up of Performance Data	43
Collection of Raw Data	44
Data Normalization and Rollup	45
Example	47
Storage of Raw and Rolled Up Data	47
Global Settings for Inbound Email and Outbound Email	48
Global Settings for Login Alert Messages	51
Global Settings for Password Reset Emails	51
Defining the Email Message for "I forgot my password"	52
Global Settings for Security	53
Global Settings for System Thresholds	54
Global Settings for Interface Thresholds	59
Settings in Silo.conf	66
Disabling the User Interface on a Database Server	76
Collector Groups	78
What is a Collector Group?	80
Installing, Configuring, and Licensing Data Collectors	80
Technical Information About Data Collectors	81
Duplicate IP Addresses	81

Open Ports	81
Viewing the List of Collector Groups	82
Viewing the List of Collector Groups in the Classic Skylar One User Interface	83
Creating a Collector Group	84
Pre-Deployment Questions for a Collector Group	84
Capacity Planning for a Collector Group	84
Defining a Collector Group	84
Defining a Collector Group in the Classic Skylar One User Interface	90
Editing a Collector Group	94
Editing a Collector Group in the Classic Skylar One User Interface	94
Collector Groups and Load Balancing	95
Tuning Collector Groups in the silo.conf File	96
Load Balancing and Device State	98
Collector Affinity	98
Failover for Collector Groups for Component Devices	99
Collector Groups for Merged Devices	99
Creating a Collector Group for Data Storage Only	100
Deleting a Collector Group	101
Deleting a Collector Group in the Classic Skylar One User Interface	101
Assigning a Collector Group for a Single Device	102
Assigning a Collector Group for a Single Device in the Classic Skylar One User Interface	102
Aligning the Collector Group in a Device Template	102
Changing the Collector Group for One or More Devices	103
Managing the Host Files for a Collector Group	103
Processes for Collector Groups	103
Enabling and Disabling Concurrent PowerShell for Collector Groups	105
Enabling Concurrent PowerShell on All Collector Groups	105
Disabling Concurrent PowerShell on All Collector Groups	106
Enabling Concurrent PowerShell on a Specific Collector Group	106
Disabling Concurrent PowerShell on a Specific Collector Group	107
Enabling and Disabling Concurrent SNMP for Collector Groups	107
Enabling and Disabling Concurrent SNMP for All Collector Groups	108

Enabling and Disabling Concurrent SNMP for Collector Groups	108
Enabling Multi-tenancy for Collector Groups	109
Aligning Collector Groups to Organizations	110
Aligning Collector Groups to Organizations in the Classic Skylar One User Interface	110
Daily Health Tasks	111
What is a Healthy Skylar One System?	112
Skylar One Self-Healing	115
Monitoring System Events	115
Searching the System Logs	116
Deleting Entries from the System Logs	117
Monitoring System Processes	118
Viewing the List of System Processes	118
Recommended System Maintenance	119
Monitoring the Status of Each Appliance	119
Logging in Skylar One Version 11.3.0 and Later	123
Configuring TLS Certificates	124
Forwarding Local Syslog Messages to Remote Systems	125
Specifying Alternate Inbound TCP or UDP Ports	126
Adjusting the Priority Filter for Inbound Messages	128
Filtering or Discarding Inbound Messages	129
Sending Logs via Syslog to a Remote Server in Skylar One Version 11.2.x and Earlier	130
Monitoring User Actions and Events on the Audit Logs Page	130
Viewing the List of Audit Logs	131
Generating Reports on Audit Logs	132
Using auditd to Monitor Sensitive Files	132
Files Logged by Default	133
Modifying the List of Files to be Logged	134
Monitoring the Status of Data Collectors	134
Updating Skylar One	136
Skylar One Upgrade Planning and Checklist	139
Planning the Update	139
Skylar One Upgrade Checklist	140

Skylar One Recommended Upgrade Paths	143
The System Updates Page	143
Scheduling Maintenance Windows	145
Pre-Upgrade Best Practices for Skylar One	146
Verifying PowerPack Version Compatibility	146
Backing Up SSL Certificates	147
Setting the Timeout for PhoneHome Watchdog	147
Running the Pre-Upgrade Test for PhoneHome Database Server	148
Adjusting the Timeout for Slow Connections	149
Running a Pre-Import Script to Prevent 12.5.x Patching Issues	150
Running the System Status Script Before Upgrading	153
Running the System Status Script	153
Updating the Skylar One Distributed Architecture	153
Downloading the Update	154
Importing the Update	155
Staging the Update	156
Automatic Staging	157
Manually Staging an Update	157
Monitoring Staging	158
Running the Pre-Upgrade Check	159
Running the Pre-Upgrade Check	159
Potential Issues to Address	160
CentOS 5 Failure	160
Collector Group Membership	160
Eligibility Failure	160
Enabled Failure	160
Free Disk-Space Failure	160
Host File Failure	160
Patch-Hook Ownership Failure	161
RPM Database Failure	161
RPM Package Failure	161
Putting All Skylar One Appliances into Maintenance Mode	162

Deploying the Update	162
Troubleshooting System Update	164
Using the sysuptb Troubleshooting Tool	164
Available Commands	164
Using the phtb Troubleshooting Tool	168
Available Commands	169
Monitoring Deployment	169
Running the Post-Upgrade Recovery Script for 12.5.x Patching Issues	170
Installing Additional RPMs on a Skylar One Appliance	172
Remove Skylar One Appliances from Maintenance Mode	172
Updating Skylar One Extended Architecture	173
Automatically Upgrading MariaDB	173
Manually Upgrading MariaDB	174
Download RPMs to Skylar One Appliances	175
Manually Upgrade Two Database Servers Configured for High Availability or Disaster Recovery	176
Step 1: On the Secondary Database Server	176
Step 2: On the Primary Database Server	177
Step 3: On the Secondary Database Server	179
Manually Upgrade Three Database Servers Configured for High Availability and Disaster Recovery	180
Step 1: On the Secondary Database Server	180
Step 2: On the Primary Database Server	181
Step 3: On the Secondary Database Server	183
Step 4: On the Disaster Recovery Database Server	184
Manually Upgrading Standalone Database Servers, All-In-One Appliances, Data Collectors, and Message Collectors	186
Rebooting Appliances in the Skylar One Distributed Stack	187
Rebooting the Administration Portal	188
Rebooting Multiple Administration Portals	188
Rebooting a Single Administration Portal	188
Rebooting Data Collectors and Message Collectors	188
Rebooting Data Collectors and Message Collectors from the Appliance Manager page	188

Rebooting Data Collectors and Message Collectors from the Command Line	189
Rebooting Standalone All-In-One Appliance and Standalone Database Server	190
Rebooting Two Database Servers Configured for Disaster Recovery	190
Rebooting Two Database Servers in a High Availability Cluster	191
Rebooting Three Database Servers Configured for High Availability and Disaster Recovery	193
Upgrading to Aurora 3 RDS (MySQL 8.0)	194
Before You Upgrade to Aurora 3	194
Performing the Aurora 3 Upgrade	195
Restoring the SSL Certificates	195
Resetting the Timeout for PhoneHome Watchdog	196
Clearing System Cache	197
Updating Default PowerPacks	197
Configuring Subscription Billing	198
Upgrading Skylar One Extended Architecture	199
Workflow	200
Prerequisites	200
Resizing the Disks on the Compute Node	201
Installing ORAS	202
Obtaining Your Harbor Credentials	203
Upgrading to 12.3.x	203
Step 1: Preupgrade	203
Step 2: Disable the Scylla Cluster	204
Step 3: Upgrade the Skylar One Extended Architecture	205
Step 4: Upgrade the Skylar One Distributed Architecture	206
Upgrading to 12.2.x	206
Step 1: Preupgrade	206
Step 2: Disable the Scylla Cluster	207
Step 3: Upgrade the Skylar One Extended Architecture	208
Step 4: Upgrade the Skylar One Distributed Architecture	209
Skylar One Self-Monitoring	210
The Workflow for Skylar One Self-Monitoring	211
PowerPacks Required for Self-Monitoring	211

ScienceLogic Support Pack	211
Data Pull Support	212
Skylar One: Operational Insights PowerPacks	212
Credentials Required for Self-Monitoring	214
Creating an SNMP Credential for Self-Monitoring	214
Creating Database Credentials for Self-Monitoring	216
Enabling Connectivity on Port 7707	217
Discovering Your Skylar One Devices	218
Aligning Skylar One Self-Monitoring Dynamic Applications	219
Aligning the Correct Credentials to Your Devices	220
Configuring Run Book Automations to Populate Dashboards	221
Verifying Your Devices	222
Configuring the Run Book Automations	223
Additional Self-Monitoring Resources	223
Monitoring and Maintaining Skylar One	224
Monitoring and Managing User Access	225
Viewing Information about Each Access Session	225
Deleting a User's Session	226
Limiting the Number of Simultaneous User Sessions	226
Viewing Lockouts and Unlocking Lockouts	226
Global Settings for Lockouts	227
Audit Logs	227
Improved User Session Control with tmux	227
Managing Scheduled Tasks	228
Recommended System Maintenance	229
Viewing the List of Schedules	230
Enabling or Disabling One or More Schedules	230
Deleting One or More Schedules	231
Putting the Database Server into Maintenance Mode	231
Monitoring Overall System Usage and Statistics	231
Viewing an Overview of All Events	232
Viewing Events by Appliance and Event Source	233

Admin Notifier	234
How Does the Admin Notifier Work?	235
Admin Notifier in the Classic Skylar One User Interface	235
What Issues Does the Admin Notifier Monitor?	235
Diagnostic Tools	239
Viewing Information About ScienceLogic Processes	240
Viewing the List of ScienceLogic Processes	240
Debugging a Process and Viewing Debug Logs	241
Viewing Information About Unhandled Exceptions	243
Viewing the List of Unhandled Exceptions	243
Viewing the Output of the System Status Script	243
Viewing the Database Tables on the Database Server	243
Accessing the Database Tool	244
Disabling Normalization, Re-Enabling Normalization, and Backfilling Raw Data	245
Enable Logging for Data Pull Storage Objects	248
Enable	248
Disable	249
Controlling Developer Log Settings	249
Generating Developer Logs for the Default User Interface (AP2)	250
Generating Developer Logs for the Classic User Interface	251
Changing Administrator Passwords	253
Disabling phpMyAdmin	254
Changing the Password for the Default User Interface Account	254
Changing the Password for the Default Console User	255
Changing the Password for the Web Configuration Utility	255
Changing Database Passwords	255
Configuring a New MySQL Password on Database Appliances	256
Configuring a New MySQL Password on Collector Appliances	258
Editing Silo.Conf	259
Updating the master.system_settings_licenses Table	259
Recovering the Root MySQL Password	260
Recovering the MySQL SNMP User Account on Data Collector	261

Changing the MariaDB Password on Skylar One Appliances	261
Changing the IP Address of a Skylar One Appliance	263
Changing the IP Address on an All-In-One Appliance	264
Step 1. Stop the EM7 Service	264
Step 2. Change the IP Address in the silo.conf File	264
Step 3. Change the IP Address in the /etc/hosts File	264
Step 4. Change the IP Address in the Network Interface Configuration File	265
Step 5. Update the IP Address in the MySQL Database	267
Step 6. Reboot the Appliance	267
Step 7. Confirm the Change in Skylar One	267
Changing the IP Address on a Database Server	267
Step 1. Stop the EM7 Service	268
Step 2. Change the IP Address in the silo.conf File	268
Step 3. Change the IP Address in the /etc/hosts File	268
Step 4. Change the IP Address in the Network Interface Configuration File	269
Step 5. Update the IP Address in the MySQL Database	271
Step 5a: For Database Servers Configured with PhoneHome	271
Step 5b For Clustered Database Appliances (using HA, DR, or HA+DR)	273
Step 6. Reboot the Appliance	275
Step 7. Change the Database Appliance IP Address in the Administration Portals, Data Collectors, and Message Collectors	276
Step 8. Confirm the Change in Skylar One	276
Changing the IP Address on a Data Collector or Message Collector	276
Using the Web Configuration Utility to Change the IP Address of a Data Collector or Message Collector	277
Using the Command Line to Change the IP Address of a Data Collector or Message Collector	278
Confirming the IP Address Change on the Appliance Manager Page	278
Changing Domain Name Servers (DNS) and Host Names on a Skylar One Appliance	279
Changing Name Servers on a Skylar One Appliance	280
Changing Name Servers on Administration Portals	280
Changing Hostnames on a Skylar One Appliance	281
Backup Management	283
Types of Skylar One Backups	285

Configuration Backups	285
What Does a Configuration Backup Include?	286
Full Backups	287
What Does a Full Backup Include?	288
Disaster Recovery Backups	288
What Does a Disaster Recovery Backup Include?	288
The Workflow for Backing Up and Restoring Skylar One	288
Planning for Skylar One Backups	289
Creating Backup Credentials	289
Creating an S3 Backup Credential	290
Creating a Basic/Snippet Backup Credential	291
Configuring Backups	291
Adding Files to Include in Configuration Backups	292
Defining a Configuration Backup	292
Defining a Full Backup	294
Defining a Disaster Recover Backup	296
Enabling tmux	298
Mounting Backup Files	299
Mounting NFS Shares	299
Mounting SMB Shares	300
Restoring Backups	302
Restoring a Configuration Backup from an S3 Bucket	302
Restoring a Configuration Backup from a Remote NFS or SMB Share	307
Restoring a Full Backup from an S3 Bucket	313
Restoring a Full Backup from a Remote NFS or SMB Share	317
Restoring a DR Backup from an S3 Bucket	322
Restoring a DR Backup from a Remote NFS or SMB Share	325
Unmounting Backup Files	326
Retaining Backups	327
Retaining Full Backups	327
Retaining DR Backups	328
Additional Configuration for Solaris NFS Remote Shares	328

Performing Configuration and Full Backups on the DR Database Server	329
Configuration Backup on a Disaster Recovery Database Server	330
Full Backup on a Disaster Recovery Database Server	330
Viewing License Data	332
Viewing License Information	332
Subscription Data	334
Ensuring Accurate Data	335
Viewing Subscription Usage	335
Current License Usage Totals	336
Historical License Usage Totals	336
Current License Totals	337
Viewing Delivery Status	339
Manually Uploading License Usage to ScienceLogic	340
Downloading the Daily License Usage File	340
Manually Uploading the Daily License Usage File to ScienceLogic	340
Uploading the ScienceLogic Receipt	341
Data Retention Settings for Licensing	341
CAC Authentication	343
Using CAC Authentication	345
Prerequisites	346
Importing SSL Certificates	347
Extracting the Common Name from a Certificate for Authentication	348
Defining the Client Certificate Chain	349
Verifying SSL Certificate File Import and Resolving Issues	351
Clearing the Skylar One Cache and Restarting NGINX	352
Testing the Configuration	352
Troubleshooting CAC Authentication	353
Failed to Identify Personal Identity Verification (PIV) Card	353
Failed CAC Authentication After Disaster Recovery (DR) Failover	354
Failed CAC Authentication After Setting Up High Availability (HA)	354
Accessing the Appliance without CAC Authentication	355
Special Circumstance: Multiple Levels of Intermediate Certificates	355

Authentication Profiles and Resources	357
Authentication Profiles	358
Viewing the List of Authentication Profiles	358
The "default" Authentication Profile	359
Creating an Authentication Profile	360
Editing an Authentication Profile	362
Deleting One or More Authentication Profiles	362
Authentication Resources	362
Viewing the List of Authentication Resources	363
The "Internal" Resource	364
Creating an LDAP/AD Authentication Resource	364
Creating an SSO Authentication Resource	370
Editing an Authentication Resource	376
Deleting an Authentication Resource	377
Installing a TLS (SSL) Certificate	378
Using TLS Certificates	379
Certificates for ScienceLogic Servers	379
Requesting a Commercial TLS Certificate	379
Creating Your Own Certificate	381
Installing the Certificate on a Skylar One Appliance	383
Managing Host Files	386
Viewing the List of Host Entries	387
Creating a New Host Entry	387
Editing a Host Entry	388
Using an Existing Host File Entry to Create a New Host File Entry (Save As)	388
Deleting One or More Host Entries	389

Chapter

1

Introduction

Overview

This manual describes the tasks that System Administrators who monitor and maintain the health of Skylar One (formerly SL1) must perform, and the tools they can use to perform those tasks.

This chapter covers the following topics:

<i>Who Should Read This Manual?</i>	15
<i>What's In This Manual?</i>	15
<i>Requirements</i>	15

Who Should Read This Manual?

This manual is intended for System Administrators who must monitor and maintain the health of Skylar One.

This manual describes tasks on the **System** menu (the **[System]** tab in the classic user interface) that are related to the maintenance and monitoring of Skylar One. This manual also includes advanced tasks that are performed at the console or in an SSH session.

What's In This Manual?

This manual includes information on global settings, collector groups, upgrading Skylar One, health tasks, maintenance tasks, licensing, and tools for troubleshooting and debugging.

Requirements

To follow some of the steps listed in this manual, you must have administrator-level access to the console of your Skylar One appliances.

Chapter

2

Global Settings

Overview

In Skylar One (formerly SL1), global settings allow you to define default behavior that applies to all elements in the platform. For settings that affect devices, you can override global settings with device-level settings.

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:

Global Settings for API	17
Global Settings for Appliances	18
Global Settings for Asset Automation	22
Global Settings for System Behavior	23
Global Settings for Data Retention	40
Global Settings for Inbound Email and Outbound Email	48
Global Settings for Login Alert Messages	51
Global Settings for Password Reset Emails	51
Global Settings for Security	53
Global Settings for System Thresholds	54

<i>Global Settings for Interface Thresholds</i>	59
<i>Settings in Silo.conf</i>	66
<i>Disabling the User Interface on a Database Server</i>	76

Global Settings for API

The **REST API Settings** page (System > Settings > API) allows you to define global parameters that affect the behavior of the REST API. When defined, these parameters affect all interaction with the API.

NOTE: This page is available only to administrator users.

To edit the settings in the REST API Settings page:

1. Go to the **REST API Settings** page (System > Settings > API).
2. In the **REST API Settings** page, edit the values in one or more of the following fields:
 - **Internal Request Account.** Specify the user account that allows Skylar One to make API requests without a password.
 - **X-EM7-run-as Header Support.** Specifies whether administrator users can make API requests using the permissions of another user without that user's password. Choices are:
 - *Disabled.* Administrator users cannot make API requests using the permissions of another user.
 - *Enabled (Admin only).* Administrator users can include the X-EM7-run-as Header to make API requests using the permissions of another user.
 - **Logging.** Specifies which logs Skylar One will write to when tickets are created or updated using the API. Choices are:
 - *Transaction Logging Only (System Logs).* If a ticket is created or updated using the API, Skylar One will write the standard entry to the audit log that indicates a user performed a write-operation using the API. However, Skylar One will not write to the ticket log for the ticket that was created or updated.
 - *Normal (Ticket and System Logs).* If a ticket is created or updated using the API, Skylar One will write to the audit log and to the ticket log for the ticket that was created or updated.
 - **X-EM7-suppress-logging Header Support.** If *Normal (Ticket and System Logs)* is selected in the **Logging** field, this field specifies whether the X-EM7-suppress-logging header can be used when an administrator creates or updates a ticket using the API. If the X-EM7-suppress-logging header is used when creating or updating a ticket, Skylar One will not write to the ticket log for the ticket that was created or updated. Choices are:
 - *Disabled.* The X-EM7-suppress-logging header cannot be used.
 - *Enabled (Admin only).* The X-EM7-suppress-logging header can be used to stop Skylar One from writing to the ticket log for the ticket that was created or updated.

- **Send Notification.** When a ticket is created or updated, Skylar One can automatically send notification emails to the ticket assignee and ticket watchers. This option specifies the conditions under which Skylar One will send notification emails when tickets are created or updated using the API. Choices are:
 - *Only if X-EM7-send-notification:1 is sent.* Skylar One will send notification emails for a ticket only when the X-EM7-send-notification header is set to 1.
 - *Sent after every write operation.* Skylar One will send notification emails for every API request that creates or updates a ticket.

3. Click the **[Save]** button to save changes in this page.

Global Settings for Appliances

The **Appliance Manager** page (System > Settings > Appliances) provides for global appliance configuration and management for your entire system or stack. This includes collector group and load distribution, version information, license status, and other items that are important when you upgrade.

During upgrade, table cells will highlight known, pending action items that must be done to successfully complete an upgrade, such as highlighting a Skylar One appliance that is running a different version of Skylar One than the Database Server.

This page is useful for ensuring that every Data Collector is assigned to a Collector Group before you begin an upgrade. In some cases, the Data Collector might be assigned to an empty Collector Group, if the collector is new.

You can also use this page to ensure that Data Collector load is near or below the system requirements for each collector.

From the **Appliance Manager** page, you can also:

- Click the wrench icon (🔧) to edit the properties for each Skylar One appliance.
- Click the toolbox icon (🔧) to access the **Web Configuration Utility** for each Skylar One appliance.
- Click the agent endpoint icon (🌐) to access the **Agent Endpoint Configuration** modal for any Skylar One appliance that has a Skylar One Gen-1 agent installed.
- Click the magnifying-glass icon (🔍) to view the output of the system status script for each Skylar One appliance.
- Click the lock icon (🔒) to get a one-time password for each Skylar One appliance.
- Click the lightning bolt icon (⚡) to run the **Enterprise Database: Collector Config Push** process (config_push.py) on any SL1 Collector appliance.
- Click the delete icon (🗑) to delete a Skylar One appliance.

To edit and view information about a Skylar One appliance:

1. Go to the **Appliance Manager** page (System > Settings > Appliances).
2. Locate the Skylar One appliance you want to edit. Click its wrench icon (🔧). The fields in the top pane are populated with values from the selected Skylar One appliance.
3. You can edit one or more of the following fields:

- **Host Name.** Name of the appliance.
- **IP Address.** Primary IP address for the appliance.
- **Model Type/Module Type.** Type of appliance. If an appliance is added with the wrong appliance type, Skylar One generates a critical error to notify the user. The types include:
 - *All In One Server*
 - *Database*
 - *Administration Portal*
 - *Data Collection Unit*
 - *Message Collection Unit*

NOTE: The combination appliance with a Database Server and an Administration Portal on a single appliance will appear with **Module Type** of *Database*. The combination appliance with a Message Collection Unit and a Data Collection Unit will appear with **Module Type** of *Data Collection Unit*.

- *PowerFlow (Skylar OnePowerFlow)*
- **Description.** Description of the appliance.
- **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.

4. You can also edit two optional fields for Data Collectors or Message Collectors:

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

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

If you are using AWS RDS with your Skylar One System, you must define the **DB User** and **DB Password** for each Data Collector or Message Collector.

NOTE: ScienceLogic recommends that you vary the Data Collector and Message Collector database credentials for enhanced per-appliance security. This greatly enhances the security of your central database by disallowing a successful attack to go unnoticed on your Data Collector and then succeed without failure on the central database.

5. You can view the following information about each appliance that appears on the **Appliance Manager** page:

- **Name.** Name of the appliance.
- **IP Address.** Primary IP address for the appliance.
- **Module Type.** Type of appliance.
- **Collector Group.** For Data Collectors and All-In-One Appliances, specifies the Collector Group associated with the appliance.
- **Description.** Description of the appliance.
- **Build.** Specifies the latest build installed on the appliance.

NOTE: If a Skylar One appliance is running a different version of Skylar One than the Database Server, the corresponding cell in the **Build** column will be highlighted.

- **MariaDB.** Specifies the version of MariaDB running on the All-In-One Appliance, Database Server, Data Collector, or Message Collector.
- **Platform.** The current operating system platform version for the appliance. Potential values include *e17* for appliances running on Oracle Linux 7 (OL7), *e18* for appliances running on OL8, *NULL*, or *Unknown*. The platform value is highlighted when the primary Database Server is not running on the same platform version as the other appliances in the system.
- **Capacity.** For Database Servers, specifies the licensed capacity of the appliance.
- **Allocation.** For Data Collectors, specifies the number of devices aligned with the appliance.
- **ID.** Unique numeric ID, automatically assigned by the platform to each appliance in the **Appliance Manager** page.
- **Validated.** Specifies whether the license is valid.
- **Endpoint.** Skylar One Agent endpoint for the Gen 1 Agent.

- **Needs Reboot?**. Specifies whether the appliance requires reboot to add latest kernel or security updates. This column is updated every 30 minutes. Hover your mouse to determine why the reboot is required and information about kernel version, packages, and last reboot.
- **Task Manager Paused?**. Specifies whether the task manager service (em7) is paused. This value is updated every two minutes.
- **Access**. Indicates whether the appliance is shared with all organizations or is private.

NOTE: The **Access** column displays only for Administrator user accounts and user accounts assigned to the System organization.

- **Edit Date**. Date the appliance's information was discovered or last edited.
- **Edit User**. User who last edited the appliance's information.
- **Create Date**. Date and time the appliance was registered and licensed.

6. To view the Web Configuration Utility for an appliance, where you can track license data, interfaces, and other device settings, click the Appliance Manager icon (🔗). Use the same login credentials that you used to log into Skylar One, and close the pop-up window for the Utility when you are done.
7. If a Skylar One appliance is running a different version of Skylar One than the Database Server, that appliance is highlighted in the **Appliance Manager** page. The version number, if known, is listed in the **Build** column.
8. For all Skylar One appliances, Skylar One runs the system status script every 15 minutes. You can click the logs icon (🔗) to view the results of the latest system status script.
9. If you are logging in to the "sl1admin" account on an appliance, you can click the padlock (🔒) icon for that appliance to get a one-time password. For more information, see "Using the sl1admin Account" in the *Role-Based User Accounts* chapter of the **Organizations and Users** manual.
10. For Data Collectors and Message Collectors, you can click the lightning bolt icon (⚡) to manually force the Database Server to send the latest configuration information.

NOTE: The delete icon (🚫) does not appear for Database Servers that are not configured for High Availability or Disaster Recovery. The bomb icon does not appear for Database Servers that are configured as the primary database in a High Availability or Disaster Recovery configuration.

11. Click the **[Save]** button to save any changes. Click the **[Save As]** button to save your changes to a new appliance name.

The Web Configuration Utility


The Web Configuration Utility allows you to configure system-level settings for your appliances. Each appliance includes access to the Web Configuration Utility.

The Web Configuration Utility adds an additional layer of security to Skylar One by segregating administrative functions from the rest of the user interface and by exposing system-level settings and diagnostic tools that might otherwise require command-line access to the appliance. The Web Configuration Utility can be accessed only through an HTTPS connection and requires its own administrator-level password.

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`
2. Type the address of the Web Configuration Utility into the address bar of your browser, replacing "ip-address-of-appliance" with the IP address or public-facing fully qualified hostname of the appliance.
3. You will be prompted to type your username and password. Log in as **em7admin** with the appropriate password. After logging in, the main **Configuration Utility** page appears:

NOTE: For better security, change your Web Configuration Utility credentials. If you leave your Web Configuration Utility credentials as the stock credentials, you greatly increase the risk that your central database could be attacked, or that you could be locked out of the Web Configuration Utility.



4. In the **Configuration Utility**, you can license a Skylar One appliance, configure interfaces, and edit settings for the Skylar One appliance and the Database Server if applicable.
 - For details on using the **Configuration Utility** to license a Skylar One appliance, see the manual *Installation and Initial Configuration*.
 - For details on using the **Configuration Utility** to inform Data Collectors, Message Collectors, and Administration Portals when you change the IP address of a Database Server, see the section on *Changing IP Addresses*.

Global Settings for Asset Automation

The **Asset Automation** page (System > Settings > Assets) allows you to define the default behavior for all asset records.

For each standard asset field, you can specify:

- Whether the field can be automatically populated by Skylar One.
- Whether the field's value should be automatically updated by Skylar One.
- Whether or not Skylar One should generate an event if the field's value changes.

You can define the default behavior for each standard field in the following asset pages:

- **Asset Properties**
- **Asset Maintenance & Service**
- **Asset Configuration**
- **Asset Licenses**
- **Asset IP Networks**
- **Asset Components**

The defined behavior will be applied to every asset record in Skylar One.

For more details on asset records and enabling automation for asset records, see the manual **Asset Management and Vendors**.

Global Settings for System Behavior

To define or edit the settings in the **Behavior Settings** page:

1. Go to the **Behavior Settings** page (System > Settings > Behavior).
2. On the **Behavior Settings** page, edit the values in one or more of the following fields:
 - **Interface URL.** URL for accessing the user interface. This value should be in URL format and can be up to 64 characters in length. Do not include a trailing forward slash ("") at the end of the Interface URL. When Skylar One generates URLs for tickets or events (for example, in email messages), the trailing forward slash will be automatically included.
 - **Require TLS Validation for Gen 1 Agent.** When this checkbox is selected, Gen 1 agents that operate outside of the Extended Architecture will require TLS validation to upload data.

NOTE: To enable this validation, all Data Collectors and Message Collectors that will ingest agent data must have a valid and signed TLS certificate.

- **Password Expiration.** Specifies whether or not the password for a user account will expire and if so, when the password will expire. Choices are:
 - **Disabled.** Passwords do not expire or are managed externally to Skylar One.
 - **30 Days.** Passwords will expire after 30 days.
 - **60 Days.** Passwords will expire after 60 days.
 - **90 Days.** Passwords will expire after 90 days.

- *180 Days.* Passwords will expire after 180 days.
 - *365 Days.* Passwords will expire after 365 days.
- **Password Reset Interval.** The minimum amount of time that must pass before a user can change a password. For example, if the value in this field is *2 Hours*, a user can change a password every two hours. This applies to users changing their own passwords and administrators changing other users' passwords. Values range from 1 hour to 24 hours, in increments of one hour.
- **Password Hash Method.** Specifies how user passwords will be encrypted for storage in the ScienceLogic database. You can choose the hashing algorithm that works best for your enterprise. Choices are:
 - *SHA-512.* As of 10.2.0, this is the default value. Previous passwords will use their previous hash method until the password is changed.
 - *Automatic (PHP Password API)*
- **Password Minimum Length.** Specifies the minimum number of alphanumeric characters allowed for the password. You can specify any value from 1 to 99. The default value is "8" characters.
- **User Login Session Timeout.** Select the amount of idle time that can pass without any user activity before that user's Skylar One session expires. For better security, use a shorter time frame. By default, user sessions expire after 10 minutes of inactivity.
- **Page Auto-Refresh Keeps User Session Active.** Specifies whether or not a user's Skylar One session stays active when a page in the system auto refreshes. This setting is disabled by default, and ScienceLogic recommends keeping it disabled for security purposes. If this setting is enabled, a user's Skylar One session will remain active if they are on a page that auto-refreshes, even if the amount of time specified in the **User Login Session Timeout** field elapses.
- **Account Lockout Type.** If a user enters incorrect login information multiple times in a row, that user will be locked out of the user interface. In this field, you can select how the lockout will be applied. Choices are:
 - *Lockout by IP Address.* All subsequent login attempts from the IP address will be denied once a lockout for that IP address has been identified. Use this option to isolate all access from the remote IP address for any account until a system administrator has validated the lockout and cleared it.
 - *Lockout by Username and IP Address.* All subsequent login attempts by this username from the IP address will be denied. This will permit other users from the same IP address to continue to access the system as long as they are not locked out.
 - *Lockout by Username (default).* All subsequent login attempts by this username from any IP address will be denied.
 - *Disabled.* Lockouts are disabled. This setting can leave your system vulnerable to attacks and as such, it is not recommended.

- **Account Lockout Attempts.** Number of times a user can enter incorrect login information before a lockout occurs. Choices are 1 time through 10 times.
- **Login Delay.** To prevent unauthorized users from using brute-force login attempts, you can set a login delay in this field. After each failed login, Skylar One will not allow another attempt for the number of seconds specified in this field. Choices are:
 - *Disabled.* Skylar One does not enforce a delay between failed logins.
 - *1 Second.* After a failed login, Skylar One will not allow another attempt for one second.
 - *2 seconds.* After a failed login, Skylar One will not allow another attempt for two seconds.
 - *4 seconds.* After a failed login, Skylar One will not allow another attempt for four seconds.
 - *8 seconds.* After a failed login, Skylar One will not allow another attempt for eight seconds.
- **Single Instance Login (Admins).** Specifies whether more than one instance of a single username can be logged in to the user interface at the same time. Defines the default behavior for users of account type "Administrator". You can specify the following types of behavior:
 - *Disabled.* Multiple instances of the same account name can be logged in to the user interface. There are no requirements or limitations on any of the instances. None of the instances will be automatically logged out.
 - *Session can be transferred after.* If you select one of these options, the second instance of a user account can log in only after the first instance of the account is inactive. In Skylar One, an account is considered "inactive" if the user has not performed any tasks or navigated within the user interface. You can specify how long the first instance must be inactive before the second instance can log in. When the second instance successfully logs in to the user interface, the browser where the first instance is logged in will display the following message: "User id 'account name' logged in from a different browser and transferred this session."

NOTE: If this field is set to any value other than *disabled*, you can still override an earlier instance. If you try to log in to the user interface and there is another instance of the account already logged in to the user interface, the login page will display the following message: "User id 'account name' is already logged in to the system. To transfer the session, check 'Transfer Session' and log in."

- If you select the **Transfer Session** checkbox, this logs the first instance out of the user interface and allows the second instance to log in to the user interface. The browser where the first instance was logged in will display the following message: "User id 'account name' logged in from a different browser and transferred this session."

- *Other (manual entry)*. Allows you to enter a custom value, in seconds. When the first instance of a user account is inactive in the user interface for the specified number of seconds, the first instance is logged out and the second instance is allowed.

NOTE: To support single instance login in the current Skylar One user interface ("AP2"), you must make the appropriate settings in this field and then perform additional steps. For more information, see the section on [Configuring Single Instance Login in AP2](#).

- **Single Instance Login (Users)**. Specifies whether more than one instance of a single username can be logged in to the user interface at the same time. Defines the default behavior for users of account type "User". You can specify the following types of behavior:

- *Disabled*. Multiple instances of the same account name can be logged in to the user interface. There are no requirements or limitations on any of the instances. None of the instances will be automatically logged out.
- *Session can be transferred after*. If you select one of these options, the second instance of a user account can log in only after the first instance of the account is inactive. In Skylar One, an account is considered "inactive" if the user has not performed any tasks or navigated within the user interface. You can specify how long the first instance must be inactive before the second instance can log in. When the second instance successfully logs in to the user interface, the browser where the first instance is logged in will display the following message: "User id 'account name' logged in from a different browser and transferred this session."

NOTE: If this field is set to any value other than *disabled*, you can still override an earlier instance. If you try to log in to the user interface and there is another instance of the account already logged in to the user interface, the login page will display the following message: "User id 'account name' is already logged in to the system. To transfer the session, check 'Transfer Session' and log in."

- If you select the **Transfer Session** checkbox, this logs the first instance out of the user interface and allows the second instance to log in to the user interface. The browser where the first instance was logged in will display the following message: "User id 'account name' logged in from a different browser and transferred this session."

- *Other (manual entry)*. Allows you to enter a custom value, in seconds. When the first instance of a user account is inactive in the user interface for the specified number of seconds, the first instance is logged out and the second instance is allowed.

NOTE: To support single instance login in the current Skylar One user interface ("AP2"), you must make the appropriate settings in this field and then perform additional steps. For more information, see the section on [Configuring Single Instance Login in AP2](#).

- **Account Lockout Duration.** Specifies how long a user will be locked out of the user interface. Choices are 1 hour - 24 hours, in one hour increments. The shorter the duration, the sooner the system is again vulnerable to a potential attack; however, a shorter duration will also allow a user to attempt to log in again without the help of a system administrator.
- **Lockout Contact Information.** This contact information will be displayed when a user is locked out of the user interface. Can be any combination of alphanumeric characters, up to 255 characters in length. This information should allow the user to contact his/her administrator to unlock the account.
- **Login Header Title.** HTML title of the login page. This text will appear at the very top of the browser on the login page.
- **System Identifier.** Unique name for the current Skylar One system. Can be up to 128 characters in length. This field is useful for companies or organizations with multiple Skylar One systems. If a value is provided in this field, Skylar One will include a "system identifier" value in each event generated by the current Skylar One system. This allows users to easily determine the source Skylar One system associated with the event.
- **Ping & Poll Timeout (Msec.).** This field specifies the number of milliseconds the discovery tool or availability polling will wait for a response after pinging a device. After the specified number of milliseconds have elapsed, the poll will timeout. The choices are between 100 and 5000 milliseconds.
- **SNMP Poll Timeout (Msec.).** This field specifies the number of milliseconds the discovery tool will wait for a response after sending an SNMP query to a device. After the specified number of milliseconds have elapsed, the SNMP poll will timeout. The choices are between 100 and 5000 milliseconds.
- **SNMP Failure Retries.** This field specifies the number of times the discovery tool will try to communicate with a device after a timeout or failure. After that number of times has been met, the discovery tool will not retry unless the user manually restarts the discovery process. The choices are 0-6.
- **Initially Discovered Interface Poll Rate.** This field specifies the frequency with which Skylar One will poll newly discovered interfaces. This setting does not affect interfaces that have been previously discovered with a different value in this field or interfaces for which the **Frequency** field has been manually edited in the Interface Properties page.

- **DHCP Community Strings (Comma separated)**. SNMP "read only" community string to use during discovery. This is required only if DHCP servers and devices use a different SNMP community string than other devices in the network. If the community string specified in the **Discovery Control Panel** page (System > Manage > Classic Discovery or System > Manage > Discovery in the classic user interface) does not work for DHCP devices, Skylar One will automatically use the community string specified in this field.
- **Strip FQDN From Inbound Email Device Name**. In Events from Email policies, specifies how Skylar One will match the regular expression for device name. Choices are:
 - *Enabled*. Skylar One will search the text string in the incoming email and match all characters up to the first period that appears in the text string. If multiple devices match the characters up to the first period (for example, my_device.1 and my_device.2), Skylar One will align the event with the matching device with the highest Device ID.
 - *Disabled*. Skylar One will search the text string in the incoming email for a match for the device name. The text string must include an exact match to the regular expression (defined in the Events from Email policy), including any text following a period in the device name. If Skylar One does not find an exact match in the incoming email, Skylar One creates an entry in the system log.
- **Inbound Email Alert Message**. In each event policy, the **First Match String** and **Second Match String** fields specify the string or regular expression used to correlate the event with a log message. To trigger an event, the text of a log message must match the value in the **First Match String** and **Second Match String** fields in that event's policy. For Events from Email policies, this field specifies whether only the email message body will be written to the device log or whether both the email message subject and email message body will be written to the device log. Choices are:
 - *Email Message Body Only*. Only the email message body is written to the device log. The **First Match String** and **Second Match String** fields can examine and match only the email message body.
 - *Email Message Subject and Body*. Both the email message body and the email message subject are written to the device log. The **First Match String** and **Second Match String** fields can examine and match against both the email message body.

NOTE: The global setting **Inbound Email Alert Message** affects how events are triggered. This field does not affect the **Regex Pattern** field in the Event from Email policy. The **Regex Pattern** field in an Event from Email policy specifies which device log to write to.

- **Event Console Ticket Life Ring Button Behavior**. Specifies how the life-ring icon (⌚) in the Event Console will behave. Choices are:

- *Create/View EM7 Ticket.* When you click the life-ring icon (⌚) for an event in the **Event Console**, Skylar One will display the **Ticket Editor** page, where you can define a ticket and automatically associate it with the selected event. This is the default behavior.
- *Create/View External Ticket.* If an external ticket is aligned with an event, when you click the life-ring icon (⌚) for that event (from the **Event Console**), Skylar One spawns a new window and displays the external ticket (as specified in the **force_ticket_uri** field). If an external ticket is not yet aligned with an event, when you click the life-ring icon (⌚) for that event, Skylar One sets a "request" flag for the ticket and displays an acknowledgment that a new ticket has been requested. You can then use the "request" in run book logic, to create the ticket on the external system.

CAUTION: If you select *Create/View External Ticket* in the **Event Console Ticket Life Ring Button Behavior** field, you can no longer create tickets from the **Event Console**.

- **Automatic Ticketing Emails.** Specifies whether ticket watchers will automatically receive email notification when a ticket is created or changes status. Choices are:
 - *Enabled.* This is the default value. When you select this option, Skylar One automatically sends email notifications to all watchers when a ticket is created, assigned, or updated.
 - *Disabled.* When you select this option, Skylar One does not automatically send email notifications to all watchers when a ticket is created, assigned, or updated.

- **Force Child Ticket State and Status Inheritance.** This checkbox specifies whether or not the state and status of a ticket is applied to child tickets.
- **Skylar Options (JSON).** This field enables you to override the hard-coded default options for exporting metadata to Skylar. To do so, enter values in the same JSON structure as those that appear in the `config.py` file. For example:

```
DEFAULT_OPTIONS = { "metadata": { "intervals": { "snapshot": 60, "cleared_events": 5 }, "snapshot": { "batch_sizes": { "asset_basic": 500, "asset_ip_config": 10000, "perf_index_label": 5000, "device_config": 20000, }, }, "cleared_events": { "query_range": 60, }, } }
```

- **Prevent Browser Saved Credentials.** This checkbox specifies whether Skylar One will allow the browser to cache login credentials and perform auto-complete in the login page. By default, the user interface will allow browsers to cache login credentials. Choices are:
 - *Selected.* The user interface will not allow browsers to cache credentials and use auto-complete in the login page. Use this setting to comply with PCI DSS and other security protocols.
 - *Not Selected.* This is the default setting. The user interface will allow browsers to cache credentials and use auto-complete in the login page. The implementation of this functionality varies between browsers.
- **Prevent Loading Interface in External Frames.** If you select this checkbox, other pages cannot be loaded in external frames in the same browser session that includes Skylar One. This is a security measure, to prevent click-jacking attacks.
- **Hide Perpetual License Count.** Specifies whether to display the device count graph in the **System Usage** page (System > Monitor > System Usage). The default behavior is to hide the graph in the **System Usage** page. Users might find this graph useful to troubleshoot licensing issues. For a description of the **System Usage** page, see the [Monitoring Overall System Usage and Statistics](#) section.
- **Hide "New" button on the Ticket Editor.** If you select this checkbox, the **Ticket Editor** page will not display the **[New]** button. This field is unselected by default.
- **Enable Unique Asset Tag to Organization Constraint.** Select this option if you want to prevent an asset from being moved to another organization if the new organization is already aligned to an asset with the same Asset Tag. Do not select this option if you want to allow assets with the same Asset Tags in the same organization. This field is selected by default.
- **Display Previous Login In Footer.** If you select this checkbox, the user interface will display information about the last successful login to the user interface and the last failed login (if applicable) in the footer. The user interface will display the following in the lower left of the footer:
 - **Last Login:** *mm-dd-yyyy hh-mm.*
 - **Failed Login:** *mm-dd-yyyy hh-mm.*

- ***Ignore trap agent-addr varbind***. If you select this checkbox, Skylar One will align incoming SNMP trap messages with the forwarding device (last hop) instead of searching for the IP address of the originator of the trap.
- ***Enable Selective PowerPack Field Protection***. If you select this checkbox, the following fields will ***not*** be updated when you update a PowerPack:
 - Event Policy > ***Operational State***
 - Event Policy > ***Event Severity***
 - Event Policy > ***Event Message***
 - Event Policy > ***Occurrence Count***
 - Event Policy > ***Occurrence Time***
 - Event Policy > ***Expiry Delay***
 - Event Policy > ***Detection Weight***
 - Event Policy > ***External Event ID***
 - Event Policy > ***External Category***
 - Event Policy > ***Use multi-match***
 - Event Policy > ***Use message-match***
 - Event Policy > ***Topology Suppression***
 - Dynamic Application > Properties > ***Operational State***
 - Dynamic Application > Properties > ***Poll Frequency***
 - Dynamic Application > Properties > ***Disable Data Rollup***
 - Dynamic Application > Collection > ***Custom Attribute***
 - Dynamic Application > Collection > ***Asset / Formlink***
 - Dynamic Application > Collection > ***Change Alerting***
 - Dynamic Application > Collection > ***Hide Object***
 - Dynamic Application > Collection > ***Component Identifiers***
 - Dynamic Application > Presentation > ***Active State***
 - Dynamic Application > Threshold > ***Override Threshold Value***
 - Dynamic Application > Threshold > ***Numeric Range: High***
 - Dynamic Application > Threshold > ***Numeric Range: Low***
 - Dynamic Application > Threshold > ***Threshold Value***
 - Device Class > ***Device Dashboard***

- **Hide "Create a Ticket" in Toolbox menu.** If you select this checkbox, the **Toolbox** menu (three stacked horizontal lines in the upper-left corner in the classic user interface) will not display the *Create a Ticket* option. This field is unselected by default.
- **Hide "other" filesystem types.** If you select this checkbox, file systems of type "other" (which includes XFS file systems) will not be discovered and monitored. This checkbox is selected by default.
- **Default Country.** Specifies the country that will be selected by default in each page where the user specifies a country. The user can override this default value in each page.
- **System Timezone.** Specifies the default timezone for Skylar One. In each page where the user can select a timezone, this value will be selected by default. The user can override this default value in each page. Skylar One also uses this default value to perform timezone conversions when no user timezone setting is available. For example, if Skylar One sends an email to an address not associated with a user, any timestamps contained in the email will use the value from the **System Timezone** field. You can select from a list of all timezones. The default value is "UTC".
- **NFS Detection Disable.** If selected, this checkbox prevents Skylar One from monitoring and reporting on NFS "shared" file systems. Skylar One will monitor and report only on local file systems.
- **Port Polling Type.** Specifies how Skylar One should poll devices to discover open ports. The choices are:
 - *Half Open.* Uses a faster TCP/IP connection method and does not appear on the device's logs.
 - *Full Connect.* Uses the standard TCP/IP connection to detect open ports.
- **Initial Discovery Scan Level.** Specifies the data to be gathered during the initial discovery session. You can override this setting for a single discovery session in the **Discovery Session Editor** modal page. The options are:
 - *0. Model Device Only.* Discovery tool will discover if device is up and running and if so, collect the make and model of the device. Skylar One will then generate a device ID for the device, so it can be managed by Skylar One.
 - *1. Initial Population of Apps.* Discovery tool will search for Dynamic Applications to associate with the device. Discovery tool will attempt to collect data for the aligned Dynamic Applications. Discovery will also perform *0. Model Device Only* discovery.
 - *2. Discover SSL Certificates.* Discovery tool will search for SSL certificates and retrieve SSL data. Discovery tool will also perform *1. Initial Population of Apps* and *0. Model Device Only*.
 - *3. Discover Open Ports.* Discovery tool will search for open ports. Discovery tool will also perform *2. Discover SSL Certificates*, *1. Initial Population of Apps*, and *0. Model Device Only*. If your system includes a firewall and you select *3. Discover Open Ports*, discovery might be blocked and/or might be taxing to your network.

- 4. *Advanced Port Discovery*. Discovery tool will search for open ports, using a faster TCP/IP connection method. Discovery tool will also perform 2. *Discover SSL Certificates*, 1. *Initial Population of Apps*, and 0. *Model Device Only*. If your system includes a firewall and you select 4. *Advanced Port Discovery*, some auto-discovered devices might remain in a pending state (purple icon) for some time after discovery. These devices will achieve a healthy status, but this might take several hours.
- 5. *Deep Discovery*. Discovery tool will use nmap to retrieve operating system name and version. Discovery will also scan for services running on each open port and can use this information to match devices to device classes. Discovery tool will search for open ports, using a faster TCP/IP connection method. Discovery tool will also perform 2. *Discover SSL Certificates*, 1. *Initial Population of Apps*, and 0. *Model Device Only*. For devices that don't support SNMP, option 5. *Deep Discovery* allows you to discover devices that don't support SNMP and then align those devices with a device class other than "pingable".

CAUTION: Option 5. *Deep Discovery* is compute-intensive and might significantly tax your network if used as the default setting. ScienceLogic recommends that you use this option on a per-discovery basis by selecting it in the **Discovery Session Editor** page.

- **Rediscovery Scan Level (Nightly)**. Specifies the data to be gathered/updated each night during auto-discovery. The auto-discovery process will find any changes to previously discovered devices and will also find any new devices added to the network. The options are the same as for **Initial Discovery Scan Level**.

TIP: ScienceLogic recommends that you delete all unused PowerPacks from your Skylar One system to improve the performance of the nightly auto-discovery process.

- **Discovery Scan Throttle**. Specifies the amount of time a discovery process should pause between each IP address or hostname in a discovery session. (You specify the list of IP addresses or hostnames for a discovery session in the **IP Address/Hostname Discovery List** field in the **Discovery Session Editor** page.) Pausing discovery processes between IP addresses or hostnames spreads the amount of network traffic generated by discovery over a longer period of time. The choices are:
 - *Disabled*. Discovery processes will not pause.
 - *1000 Msec to 10000 Msec*. A discovery process will pause for a random amount of time between half the selected value and the selected value.
- **Port Scan All IPs**. Specifies whether Skylar One should scan all IP addresses on a device for open ports. You can override this setting for a single discovery session in the **Discovery Session Editor** modal page. The choices are:
 - *0. Disabled*. Skylar One will scan only the primary IP address (the one used to communicate with Skylar One) for open ports.

- 1. *Enabled*. Skylar One will scan all discovered IP addresses for open ports.
- **Port Scan Timeout**. Length of time, in milliseconds, after which Skylar One should stop trying to scan an IP address for open ports and begin scanning the next IP address (if applicable). You can override this setting for a single discovery session in the **Discovery Session Editor** modal page. Choices are between 60,000 and 1,800,000 milliseconds.
- **Restart Windows Services (Agent required)**. Specifies whether Skylar One should automatically restart failed Windows services that have been defined on the device with a startup type of "automatic". To use this feature, the managed device must be running the agent SNMP Informant, WMI Edition. For assistance or information on purchasing and installing this agent, please contact ScienceLogic. Users must also supply a value in the **SNMP Write** field in the **Device Properties** page for the device. The choices are:
 - 0. *Disabled*. Skylar One will not automatically restart failed Windows services that have been defined on the device with a startup type of "automatic".
 - 1. *Enabled*. Skylar One will automatically restart failed Windows services that have been defined on the device with a startup type of "automatic".
- **Hostname Precedence**. Specifies which name Skylar One will use for each discovered device. Choices are:
 - *SNMP System Name*. Use the device name specified in the device's SNMP System MIB. If *SNMP System Name* is selected and Skylar One cannot find an SNMP name for the device, Skylar One will assign the name returned by the DNS Reverse Lookup. If Skylar One cannot find a DNS Reverse Lookup name for the device, Skylar One will use the device's Admin Primary IP address as the device name in Skylar One.
 - *DNS Reverse Lookup*. Use the device name specified in the device's reverse-lookup record.
- **Event Interface Name Format**. Specifies the format of the network interface name that you want to appear in events. If you selected *Interface Alias* for the deprecated *Interface Name Precedence* field in a previous release of Skylar One, the format for existing interfaces is set to {alias}. If you selected *Interface Name* for the deprecated *Interface Name Precedence* field in a previous release of Skylar One, the format for existing interfaces is set to {name}. The default format is {name}. You can use a combination of string text and the following tokens to define the interface name format for events, such as string_{name}, string_{alias}, {name}{alias}, or {ifdesc}:
 - {alias}
 - {name}
 - {state}
 - {ifdesc}
 - {if_id}
 - {did}

- {ifindex}
- {ifphysaddress}
- {iftype}
- {ifspeed}
- {ifhighspeed}
- {ifoperstatus}
- {ifadminstatus}

- **DNS Hostnames.** If Skylar One will use the DNS Reverse Lookup name as the device name (see the description of the field **Hostname Precedence**), this field specifies whether Skylar One will use the fully-qualified domain name or only the hostname for each discovered device. Choices are:
 - *Strip Device Name (Hostname).* Skylar One will use only the device name as the DNS hostname for each device.
 - *Use Full Domain Name (FQDN).* Skylar One will use the fully-qualified domain name as the device name for each device.
- **Event Clearing Mode.** Describes how clearing an event will affect correlated events. Choices are:
 - *Clear Selected Only.* Clear only the selected events. If a parent event is cleared, the previously suppressed child events will appear in the **Event Console**.
 - *Clear All in Group.* When parent event is cleared, all child events correlated with parent event will be cleared. This is the default behavior.
- **Maintenance Minimum Severity.** Specifies the minimum severity required for an event to be suppressed during device maintenance and user maintenance for devices. The default value is *Healthy*, which causes all events to be suppressed. Choices are *Healthy*, *Notice*, *Minor*, *Major*, or *Critical*.
- **Patch Maintenance Minimum Severity.** If you schedule Device Maintenance and have defined a **Patch Window** within the larger maintenance interval, this field allows you to specify the event severity that will trigger the beginning of the **Patch Window**. The first event that both matches the severity in this field and occurs within the larger maintenance window triggers the start of the **Patch Window**. Choices are *Healthy*, *Notice*, *Minor*, *Major*, or *Critical*.
- **SSL Certificate Expiry Soon.** Specifies, in number of days, when Skylar One should generate an event for an SSL Certificate that is about to expire. The choices range from 1 day to 9 months.
- **SSL Certificate Expiry Imminent.** Specifies, in number of days, when Skylar One should generate a more urgent event for an SSL Certificate that is about to expire. The choices range from 1 day to 9 months.
- **Asset Warranty Expiry.** Specifies, in number of days, when Skylar One should generate an event for an asset warranty that is about to expire. The choices range from 1 day to 9 months.

- **Domain Name Expiry.** Specifies, in number of days, when Skylar One should generate an event for a domain's registration that is about to expire. The choices range from 1 day to 9 months.
- **Validate Phone Number.** Specifies whether or not phone numbers entered into the user interface must be in US format. Choices are:
 - *Disabled.* Phone numbers are not required to be in US format.
 - *Enabled.* Phone numbers must be in US format.
- **Dashboard Maximum Series Count Per Widget.** This field allows you to select the maximum number of time-series lines that can appear in a single **Multi-series Performance** widget. Choices are 8-25. Increasing this setting might cause longer load times in the **[Dashboards tab]** page.
- **Responder API Base URL.** This field lets you update the Responder API Base URL, which is required if you want to align PowerShell Dynamic Applications to the agent. In Skylar One version 11.2.0 and later, Skylar One completes this field for you.
- **Component Device Map Update Mode.** This field specifies how Skylar One rebuilds relationships between component devices that are created by Dynamic Applications (DCM-R). Choices are:
 - *[Periodic].* (Default). A rebuild of the component device map occurs at a set interval, which is defined in the **DCM+R Rebuild** process (System > Settings > Processes). By default, the process runs every five minutes. Note that the Operating State for the **DCM+R Rebuild** process must be set to "Enabled" before periodic component device map updates will work. For more information, see [Viewing Information About ScienceLogic Processes](#).
 - *DCM-R Triggers.* A rebuild of the component device map occurs immediately after topology changes are registered.

CAUTION: Setting the **Component Device Map Update Mode** to "DCM-R Triggers" might impact your system if your environment has large topology trees or if the environment experiences frequent simultaneous topology changes.

- **Enable CBQoS Collection.** If selected, Skylar One will collect configuration data about Class-Based Quality-of-Service (CBQoS) from interfaces that are configured for CBQoS. If selected, you can enable collection of CBQoS metrics per-interface. The collected CBQoS metrics are displayed in Device Performance reports associated with the device that contains those interfaces. This setting is disabled by default. (For more information about Device Performance reports, see the manual [Monitoring Device Infrastructure Health](#).)

- **Enable Variable Rate Interface Counters.** If selected, enables more accurate collection of data from interfaces. If enabled, when Skylar One retrieves data from an interface, that data is stored in the ScienceLogic database along with the timestamp associated with the exact collection time. Before normalization occurs, Skylar One applies an interpolation function that spaces the data at regular time intervals. For example, suppose you have specified that Skylar One should collect interface data every five minutes. However, due to network traffic across the Data Collectors, Skylar One might collect data from an interface at 13:01 and then 13:05. Because the ScienceLogic normalization process expects data that has been collected every five minutes, Skylar One first applies an interpolation to the data to prepare the data for normalization. With **Enable Variable Rate Interface Counters** enabled, graphing interpolates between two collected data points without a limit of the distance between those data points. However, performance graphs will not display interpolation between two points where there is no supporting collected data, or "data gap", for a collection time when this feature is disabled.
- **Enable Concurrent SNMP Collection.** If selected, enables Concurrent SNMP Collection for all SNMP collection. Concurrent SNMP Collection allows multiple collection tasks to run at the same time with a reduced load on Data Collectors. Concurrent SNMP Collection also prevents missed polls and data gaps because collection will execute more quickly. For details see the manual *SNMP Dynamic Application Development*.

NOTE: If the "Data Collection: SNMP Collector" process is disabled on the **Process Manager** page (System > Settings > Admin Processes), this concurrent SNMP collection option is disabled.

NOTE: Concurrent SNMP Collection is not available in military unique deployments (MUD) or STIG-compliant deployments.

- **Enable Concurrent Network Interface Collection.** If selected, enables asynchronous concurrent SNMP collection for all network interfaces. This provides better scalability for large networks by allowing multiple collection tasks to run at the same time with a reduced load on Data Collectors.

NOTE: If the "Data Collection: SNMP Collector" process is disabled on the **Process Manager** page (System > Settings > Admin Processes), this concurrent network interface collection option is disabled.

NOTE: Concurrent network interface collection is not available in military unique deployments (MUD) or STIG-compliant deployments.

- **New UI Default.** Starting with Skylar One 11.1.0, the new Skylar One user interface ("AP2") is the default user interface. If you want to make the classic user interface the default interface, de-select this option. If your 11.1.0 Skylar One system was installing using an ISO, the Skylar One user interface is set as the default, but if your Skylar One system was upgraded with a patch, the classic user interface will still be set as the default user interface.
- **Include PowerPack Sensitive Fields.** If selected, lets you include sensitive fields when sharing a PowerPack. These sensitive fields include passwords and SSH keys.
- **Prefer Global Device Summary Dashboard Over Category/Class.** If you select this checkbox, the global default device dashboard will be displayed as the default in the **Device Summary** page instead of the device dashboard assigned to the device category or device class of the device. For more information about device dashboards, see the *Classic Dashboards* manual.
- **Enhanced OID Translation.** If selected, ensures that varbind OIDs that use multi-dimensional indexes are translated correctly. The symbolic translation of the known portion of the OID is included in the log message associated with the trap. Enabling the **Enhanced OID Translation** option might affect performance on large environments with a large number of traps.
- **Enable Concurrent PowerShell Collection.** If selected, enables concurrent PowerShell collection for all PowerShell collection, which allows multiple collection tasks to run at the same time with a reduced load on Data Collectors.

NOTE: If the "Data Collection: PowerShell Collector" process is disabled on the **Process Manager** page (System > Settings > Admin Processes), this concurrent PowerShell collection option is disabled.

NOTE: Concurrent PowerShell Collection is not available in military unique deployments (MUD) or STIG-compliant deployments prior to Skylar One 12.5.1.

- **Report Size Estimation.** If selected, enables the **Row Count Estimate** field for custom reports on the Run Report page (Reports > Run Report). This field provides an estimate of the number of rows that will appear in the report before Skylar One generates the report. The estimate changes based on the selections you make for the report. You can use this field to manage the size of the generated report by adding or removing items from the report as needed.
- **Enable Snippet Framework Collection.** If selected, enables Skylar One to process Snippet Framework Dynamic Applications for collection. You can disable Snippet Framework Dynamic Application processing in your Skylar One system by clearing this checkbox.
- **Enable L2 Topology Collection and Processing.** If selected, Skylar One will perform Layer-2 (L2) topology collection and processing for each device that supports it, and then generate topology maps from the discovered L2 relationships.
- **Enable L3 Topology Collection and Processing.** If selected, Skylar One will perform Layer-3 (L3) topology collection and processing for each device that supports it, and then generate topology maps from the discovered L3 relationships.

- **Enable CDP Topology Collection and Processing.** If selected, Skylar One will perform Cisco Discovery Protocol (CDP) topology collection and processing for each device that supports it, and then generate topology maps from the discovered CDP relationships.

NOTE: CDP is a proprietary protocol developed by Cisco and is not supported by all network hardware. If your network includes both CDP enabled and non-CDP network switches and routers, the topology data reported by the CDP enabled devices might not be accurate. In Skylar One, if a conflict exists between the collected CDP topology data and the collected Layer-2 topology data, the CDP topology data takes precedence. In some cases, the ScienceLogic Layer-2 data might be more accurate. Therefore, if your network includes both CDP enabled and non-CDP network switches and routers, you might want to disable CDP topology collection. For details, see the [Views](#) manual.

- **Enable LLDP Topology Collection and Processing.** If selected, Skylar One will perform Link Layer Discovery Protocol (LLDP) topology collection and processing for each device that supports it, and then generate topology maps from the discovered LLDP relationships.

IMPORTANT: Enabling a discovered device configured with CDP or LLDP topology in Skylar One will cause the device to provide information on its neighbor. This information only identifies that there is a neighbor device, not which is the parent or the child. This may cause the parent-child relationship to switch, which requires you to manually reverse the issue within the Skylar One user interface. Skylar One allows you to manually build parent-child relationships between specific device categories. For more information, see [Defining Parent and Child Devices](#) in the [Events](#) manual.

- **L2: Collect from all VLANs.** If selected, Skylar One will perform Layer-2 (L2) topology collection for all virtual local area networks (VLANs) discovered across your system. If not selected, Skylar One performs L2 topology collection for only VLAN 1, which is the default VLAN on most network switches. By default, Skylar One compiles a list of VLANs for which data should be collected. A VLAN is added to the list for collection only if its state is 1 (operational) and its type is 1 (ethernet).
- **CDP & LLDP Enhanced Custom Attribute.** When enhanced topology is enabled for CDP or LLDP, you can select from this drop-down field a device custom attribute that will be used to try to form the relationship.

3. Click the **[Save]** button to save changes in this page.

Configuring Single Instance Login for AP2

The **Single Instance Login** fields on the **Behavior Settings** page enable you to specify whether more than one instance of a single username can be logged in to the user interface at the same time.

In the classic Skylar One user interface, you can configure single instance login using just those fields. However, for the current Skylar One user interface ("AP2"), you must complete several additional steps.

To configure single instance login in AP2:

1. Go to the **Behavior Settings** page (System > Settings > Behavior).
2. Make the appropriate selections in the **Single Instance Login (Admins)** and **Single Instance Login (Users)** fields, and then click **[Save]**.
3. Either go to the console of the Skylar One Database Server or use SSH to access the Skylar One All-In-One Appliance.
4. Log in as user **em7admin**.
5. At the command line, open the `nextui.env` file in the `vi` editor:

```
sudo vi /opt/em7/nextui/nextui.en
```

6. Un-set the environment variable `AUTH_CACHE=300000` by adding `#` as a prefix to that line.
7. Save and exit the `nextui.env` file.
8. Restart the `nextui` server:

```
sudo systemctl restart nextui
```

Global Settings for Data Retention

The **Data Retention Settings** page (System > Settings > Data Retention) allows you to define parameters for log and data retention.

These settings apply to all logs and all collected data. However, you can override these system settings on a case-by-case basis. For example, you can define data-retention thresholds for a device in the **Device Thresholds** page. The settings you define for the specific device override the settings in the **Data Retention Settings** page.

NOTE: For details on data roll-up and data normalization, see [Normalization and Roll-Up of Performance Data](#).

From the **Data Retention Settings** page, you can edit how long the platform stores log entries and collected data. To edit the settings for data retention:

1. Go to the **Data Retention Settings** page (System > Settings > Data Retention).
2. On the **Data Retention Settings** page, you can drag sliders to change the value of each field or manually enter values in the fields to the right of the sliders. You can edit the value for one or more of the following fields:
 - **Audit Logs.** Number of months to retain log entries in the **Audit Logs** page (System > Monitor > Audit Logs). Log entries that are older than the specified number of months are automatically deleted. The default value is 3 months.

- **Event Logs.** Number of days to retain event logs. Event history data is used to generate the **Event Overview** page (System > Monitor > Event Overview). Log entries that are older than the specified number of months are automatically deleted. The default value is 3 months.
- **Access Logs.** Number of months to retain log entries in the **Access Sessions** page (System > Monitor > Access Logs). Log entries that are older than the specified number of months are automatically deleted. The default value is 12 months.
- **System Logs.** Number of days to retain log entries in the **System Logs** page (System > Monitor > System Logs). Log entries that are older than the specified number of days are automatically deleted. The default value is 31 days.
- **Collection Unit Data Buffer.** Number of days each Data Collector and Message Collector should store collected data. Choices are 1-10 days. Data that has been retrieved by the Database Server will be stored on the Data Collector(s) and optional Message Collector(s) for the specified number of days and then automatically deleted from the server(s). This setting does not apply to All-In-One Appliances. The default value is 2 days.
- **Ad-hoc and Scheduled Reports.** Number of days Skylar One will retain Quick Reports and Scheduled Reports in the **Scheduled Report Archive** page (Scheduled Job > Report Archive > Archived Job button). Possible values are 0 - 365, in days. If you use the default value of 0, Skylar One will remove files older than 30 days from the populated directory:
`/opt/em7/gui/ap/www/em7/libs/od_templates/populated`.
- **Raw Performance Data.** Number of days to retain performance data collected from devices. This setting applies to all performance data types, except for bandwidth data. Performance data that is older than the specified number of days is automatically deleted. This is the default system-wide value. The value in the **Device Thresholds** page for each device can override this value. The default value is 7 days.
- **Hourly Rollup Performance Data.** Number of days to retain hourly normalized performance data for devices. This setting applies to all performance data types, except for bandwidth data. Hourly normalized performance data that is older than the specified number of days is automatically deleted. This is the default system-wide value. The value in the **Device Thresholds** page for each device can override this value. The default value is 90 days.
- **Daily Rollup Performance Data.** Number of months to retain daily normalized performance data for devices. This setting applies to all performance data types, except for bandwidth data. Daily normalized performance data that is older than the specified number of months is automatically deleted. This is the default system-wide value. The value in the **Device Thresholds** page for each device can override this value. The default value is 24 months.
- **Configuration Data.** Number of days to retain data from Dynamic Applications of type "configuration". The value in the **Device Thresholds** page for each device can override this value. The default value is 7 days.
- **Journal Data.** Number of days to retain collected data from Dynamic Applications of type "journal". The value in the **Device Thresholds** page for each device can override this value. The default value is 60 days.
- **Bandwidth Data.** Number of days to retain bandwidth data and CBQoS data collected from each interface on a device. Bandwidth data that is older than the specified number of days is automatically deleted. The value in the **Device Thresholds** page for each device can override this value. The default value is 31 days.

- **Hourly Rollup Bandwidth Data.** Number of days to retain hourly normalized data and hourly normalized CBQoS data for each interface on a device. Hourly normalized data that is older than the specified number of days is automatically deleted. The value in the **Device Thresholds** page for each device can override this value. The default value is 90 days.
- **Daily Rollup Bandwidth Data.** Number of months to retain daily normalized data and daily normalized CBQoS data for each interface on a device. Daily normalized data that is older than the specified number of months is automatically deleted. The value in the **Device Thresholds** page for each device can override this value. The default value is 24 months.
- **Bandwidth Billing Data.** Number of months to retain data collected by each bandwidth billing policy. Bandwidth billing data that is older than the specified number of months is automatically deleted. The default value is 24 months.
- **Device Logs Age.** Number of days to retain each device log. Log records that are older than the specified number of days are automatically deleted. The value in the **Device Thresholds** page for each device can override this value. The default value is 90 days.
- **Device Logs Max.** Maximum number of records to store in each device log. When this number is exceeded, the oldest entries will be deleted. The value in the **Device Thresholds** page for each device can override this value. The default value is 10,000 records.
- **Raw ITSM Data.** Before the value for a metric in an IT Service policy is calculated, a copy of all the device data that will be aggregated is saved. This setting is the number of days to retain the un-aggregated copies of device data associated with each IT Service. The default value is 31 days.
- **ITSM Service Metrics Data.** Number of days to retain values for metrics in IT Service policies. The default value is 30 days, with a maximum of 30 days.
- **Hourly Rollup ITSM Service Metrics Data.** Number of days to retain hourly normalized values for metrics in IT Service policies. The default value is 90 days, with a maximum of 90 days.
- **Daily Rollup ITSM Service Metrics Data.** Number of months to retain daily normalized values for metrics in IT Service policies. The default value is 12 months.
- **ITSM Key Metrics Data.** Number of days to retain values for key metrics in IT Service policies (Health, Availability, and Risk). The default value is 30 days, with a maximum of 30 days.
- **Hourly Rollup ITSM Key Metrics Data.** Number of days to retain hourly normalized values for key metrics in IT Service policies (Health, Availability, and Risk). The default value is 90 days, with a maximum of 180 days.
- **Daily Rollup ITSM Key Metrics Data.** Number of months to retain daily normalized values for key metrics in IT Service policies (Health, Availability, and Risk). The default value is 24 months.
- **SSL Certificate Purge Timeout.** Specifies the number of days after which SSL certificate data will be deleted. The default value is 730 days.
- **Ports Data Retention.** Specifies the number of days after which expired port data will be marked for deletion during the hourly maintenance process. The default value is 730 days.
- **Services Data Retention.** Specifies the number of days after which expired services data will be marked for deletion during the hourly maintenance process. The default value is 24 hours.

- **Filesystems Data Retention.** Specifies the number of hours after which expired filesystems data will be marked for deletion during the hourly maintenance process. The default value is 24 hours.
- **Schedules Purge Timeout.** Specifies the number of days after which expired schedules will be deleted. The default value is 730 days.
- **Subscriber Device Configuration Data.** For users with a subscriber license. Number of months to retain the files and database tables that contain configuration information for a device. Default value is 2 months.
- **Subscriber Device Usage Data.** For users with a subscriber license. Number of months to retain the files and database tables that contain usage information for a device. Default value is 2 months.
- **Subscriber System Configuration Data.** For users with a subscriber license. Number of months to retain the files and database tables that contain configuration information for the Skylar One system. Default value is 3 months.
- **Subscriber System Usage Data.** For users with a subscriber license. Number of months to retain the files and database tables that contain usage information for the Skylar One system. Default value is 3 months.
- **Subscriber Device Type Data.** For users with a subscriber license. Number of months to retain the files and database tables that map each device to a device category, as per your subscriber license. Default value is 3 months.
- **Subscriber Daily Delivery Data.** For users with a subscriber license. Number of months to retain the "crunched" license usage data that is calculated each day using the Subscriber Device Configuration Data, Subscriber System Configuration Data, Subscriber System Usage Data, and Subscriber Device Type Data. Skylar One will not prune data that has not yet been delivered to the ScienceLogic Licensing and Billing server. Default value is 3 months.

3. Click the **[Save]** button to save any changes to the data-retention settings.

NOTE: In Skylar One, normalized data does not include polling sessions that were missed or skipped. So for normalized data, null values are not included when calculating maximum values, minimum values, or average values.

TIP: You might want to retain normalized data for longer periods of time and non-normalized data for shorter periods of time. This allows you to save space and still create historical reports.

Normalization and Roll-Up of Performance Data

Normalization and *roll-up* are the ways in which Skylar One processes collected performance data for display and storage. Note the following important distinctions:

- **Raw data** is the data exactly as it was collected from a device or application.
- **Normalized and rolled up data** is data for which Skylar One has calculated summary statistics (sample size, count, maximum value, minimum value, mean value, average value, sum, and standard deviation) over a period of time.

Collection of Raw Data

Collector	Collected Data and Intervals
Dynamic Applications	<p>Collects raw performance data from a device at the following intervals:</p> <ul style="list-style-type: none"> • 1 minute • 2 minutes • 3 minutes • 5 minutes • 10 minutes • 15 minutes • 30 minutes • 1 hour • 2 hours • 6 hours • 12 hours • 24 hours <p>For performance Dynamic Applications, you specify this interval in the <i>Poll Frequency</i> field, in the Properties Editor page (System > Manage > Applications > Create or click the  icon).</p>
IT Services	<p><i>IT Service policies</i> can generate raw performance data for an IT service by aggregating raw performance data from devices in the policy at the following intervals:</p> <ul style="list-style-type: none"> • 1 minute • 2 minutes • 3 minutes • 5 minutes • 10 minutes • 15 minutes • 30 minutes • 1 hour • 2 hours • 6 hours • 12 hours • 24 hours <p>You can specify the interval at which the IT Service policy collects and aggregates data in the <i>Aggregation Frequency</i> field, in the</p>

Collector	Collected Data and Intervals
	<p>IT Service Editor page (Registry > IT Services > IT Service Manager > Create or click the  icon).</p>
Bandwidth	<p>Collects raw bandwidth data from a network interface at the following intervals:</p> <ul style="list-style-type: none"> • 1 minute • 5 minutes • 10 minutes • 15 minutes • 30 minutes • 60 minutes • 120 minutes <p>You can specify the frequency at which Skylar One collects raw data for a specific interface by selecting the interval in the Frequency field, in the Interface Properties page (Registry > Networks > Interfaces > interface wrench icon and click the  icon for the given interface).</p>
Additional Performance Data	<p>Skylar One collects additional raw performance data about availability, latency, file systems, and statistics generated by monitoring policies for DNS availability, Email round-trip time, system processes, system services, port availability, web-content availability, and SOAP/XML transactions. By default, Skylar One collects this data every 5 minutes.</p>

Data Normalization and Rollup

Skylar One rolls up performance data so that reports with a larger timespan do not become difficult to view and to save storage space in the database. When Skylar One rolls up data, Skylar One groups data into larger sets and calculates the average value for the larger set.

Skylar One supports two types of rollup:

- **Hourly.** Groups and averages data that is collected at intervals of 60 minutes or less. Skylar One rolls up data and calculates an average hourly value for each metric. Hourly samples include samples from the top of the hour to the end of the hour. For example, for an hourly rollup of data collected at 1 minute intervals between 1:00 and 2:00, the first data point would be the one collected at 01:00:00 and the last would end at 01:59:00.
- **Daily.** Daily rollup groups and averages all collected data. Skylar One rolls up data and calculates an average daily value for each metric. Daily samples include samples from the beginning of the day until the end of the day. For example, for a daily rollup of data collected at 1 minute intervals, the first data point would be the one collected at 00:00:00 and the last data point would be the one collected at 23:59:00.

Skylar One rolls up raw performance data as follows:

Frequency of Raw Collection	Rollup
Every 1 minute	60 minutes, 24 hours
Every 2 minutes	60 minutes, 24 hours
Every 3 minutes	60 minutes, 24 hours
Every 5 minutes	60 minutes, 24 hours
Every 10 minutes	60 minutes, 24 hours
Every 15 minutes	60 minutes, 24 hours
Every 30 minutes	60 minutes, 24 hours
Every 60 minutes	60 minutes, 24 hours
Every 120 or longer	24 hours

Before Skylar One normalizes date, Skylar One transforms the data. To transform the data, Skylar One does the following:

- For bandwidth data and data from Dynamic Applications of type "Performance", Skylar One derives rates from counter metrics. The rate from counter metrics are expressed in units-per-polling_interval. For example, rates for 5 minute collections are expressed as units-per-5-minutes.
- For data from Dynamic Applications of type "Performance", Skylar One evaluates presentation formulas. Counter metrics are first transformed into rates before evaluation.

NOTE: During the data transform steps, Skylar One does not directly rollup the raw data in the database tables.

When Skylar One rolls up data, Skylar One must normalize that data, as follows:

NOTE: As a new piece of data is collected by Skylar One, the hourly normalization and daily normalization is calculated. Skylar One does not wait for the end of an hour or the end of a day to calculate the hourly and daily normalization.

- Groups and orders the data
- Determines the sample size
- Calculates the count
- Determines the maximum value
- Determines the minimum value
- Calculates the mean value
- Calculates the average value
- Calculates the sum
- Determines the standard deviation

NOTE: In Skylar One, normalized data does not include polling sessions that were missed or skipped. For normalized data, null values are not included when calculating sample size, maximum values, minimum values, or average values.

Example

Suppose that every five minutes, Skylar One collects data about file system usage on the device named *my_device*. As each raw data point is collected, Skylar One normalizes and rolls up the collected data for file system usage for *my_device*. Skylar One does the following:

1. Apply any necessary data transforms (as discussed in the previous section).
2. Repeat the following for both hourly normalization and daily normalization:
 - a. If this is the first data point for an hourly normalization or a daily normalization, insert summary statistics for that one data point
 - Sample size = 1
 - Average = value of new data point
 - Max = value of new data point
 - Min = value of new data point
 - Sum = value of new data point
 - Standard Deviation = 0
 - b. For all subsequent data points for an hourly normalization or a daily normalization, update the summary statistics of the existing rollup bucket
3. If there are no gaps in collection, the summary statistics for hourly normalization will represent 12 data points and the summary statistics for daily normalization will represent 288 data points.

Storage of Raw and Rolled Up Data

There are two ways you can define how long Skylar One should store raw data and rolled up and normalized data:

- You can define system-wide, default settings in the **Data Retention Settings** page (System > Settings > Data Retention). These settings apply to all collected data. However, you can override these system settings in the **Device Thresholds** page (Devices > Classic Devices > wrench icon > Thresholds, or Registry > Devices > Device Manager > wrench icon > Thresholds in the classic user interface).
- For IT Service policies, aggregated device data is saved to a new database table specifically for the IT service policy. For each IT Service policy, data is normalized and rolled up. You define the data retention settings for an individual IT Service policy in the **IT Service Editor** page (Registry > IT Services > IT Service Manager > Create or click the  icon). These settings override the data retention settings in the **Data Retention Settings** page (System > Settings > Data Retention).

Global Settings for Inbound Email and Outbound Email

The **Email Settings** page (System > Settings > Email) allows you to define how Skylar One will send and receive email. Skylar One automatically sends email when tickets are updated, when automation actions are triggered, and to monitor email round-trip time. Email can be sent to the platform to create tickets and/or events.

From the **Email Settings** page, you can edit the global email parameters. To do so:

1. Go to the **Email Settings** page (System > Settings > Email).
2. In the **Email Settings** page, you can edit the value for one or more of the following fields:
 - **Authorized Email Domains.** One or more SMTP domains that will be used by Skylar One. Skylar One will use these domains to receive incoming email. This list of domains should include:
 - All domains used for loopback addresses in email round-trip monitoring policies.
 - All domains used to generate tickets from emails.
 - All domains used to receive event messages from third-party monitoring systems.
 - Each entry in this field must be a fully-qualified email domain and cannot exceed 64 characters. If you include a list of domains, separate the list with commas.
 - Each domain in this field must be managed by the Database Server. This means that a DNS MX record must already exist or be created for each domain specified in this field. Each DNS MX record must map the domain to the Database Server. When creating the DNS MX record, use the fully-qualified name of the Database Server as the name of the email server.
 - **System From Email Address.** The email address from which Skylar One will send all outbound email.

NOTE: Some outbound email servers, such as Gmail, might overwrite the **System From Email Address** value and instead use the email address of the authenticated user.

- **Email Formal Name.** Name that will appear in the "from" field in email messages sent from Skylar One. This value can be any alphanumeric value, up to 64 characters in length.
- **Email Gateway.** IP address or fully-qualified name of Skylar One's SMTP Relay server. If Skylar One is to send outgoing messages, this field must be defined. Examples of when Skylar One sends outgoing email messages are:
 - Automatically in response to Tickets from Email policies.
 - Automatically in response to changes in a ticket (ticket is assigned, edited, or resolved).

- Automatically based on Ticket Escalation policies.
- Automatically when executing Email Round-Trip Monitoring policies.
- Automatically when executing Run Book policies that include email actions.
- Automatically based on Report Jobs policies.
- Manually, when a user selects the **Send Message** page from the ticket panel pages.

Each Database Server and All-In-One Appliance includes a built-in SMTP Relay server. The fully-qualified name of Skylar One SMTP Relay server is the same as the fully-qualified name of the Database Server or All-In-One Appliance.

If Skylar One cannot use its built-in SMTP relay server to route email messages directly to their destination server (for example, due to firewall rules or DNS limitations), Skylar One can use another relay server. You can specify the IP address or fully-qualified name of the relay server in this field. Make sure you have configured your network to allow the Database Server or All-In-One Appliance to access this SMTP Relay server.

NOTE: The **Email Gateway** field must be configured to use the appropriate port number to use, which is designated by a preceding colon. When no port number is specified, Skylar One uses the default SMTP port (25).

- **Email Gateway Alt.** IP address or fully-qualified name of the secondary SMTP Relay server. If the SMTP Relay server specified in the previous field fails or is unavailable, Skylar One will use the secondary SMTP Relay server. Make sure you have configured your network to allow the Database Server or All-In-One Appliance to access this SMTP Relay server.
- **Escalation Notify Subject.** Default "Subject" text in emails generated by Ticket Escalation policies. This field can include any combination of variables and text. The field can include up to 64 characters, including one or more variables:

The **Escalation Notify Subject** field can include one or more of the following variables:

Variable	Source	Description
%1 (one)	Event	Entity type.
%2	Event	Sub-entity type.
%3	Event Policy	Event policy ID.
%4	Event	Text string of the user name that cleared the event.
%5	Event	Timestamp of when event was deleted.
%6	Event	Timestamp for event becoming active.
%7	Event	Event severity (1-5), for compatibility with previous versions of the platform. 1=critical, 2=major, 3=minor, 4=notify, 5=healthy.
%A	Account	Username.
%a	Entity	IP address.

Variable	Source	Description
%B	Organization	Organization billing ID.
%b	Organization	Impacted organization.
%C	Organization	Organization CRM ID.
%c	Event	Event counter.
%D	Event	Timestamp of first event occurrence.
%d	Event	Timestamp of last event occurrence.
%E	Event Policy	External ID from event policy.
%e	Event	Event ID.
%F	Dynamic Alert	Dynamic Application alert id.
%f	Event Policy	Specifies whether event is stateful, that is, has an associated event that will clear the current event. 1 (one) = stateful; 0 (zero) = not stateful.
%G	Event Policy	Event Category.
%g	Asset	Asset serial.
%H	Event	URL link to event.
%h	Asset	Device ID associated with the asset.
%I (uppercase "eye")	Dynamic Alert	Table index for a Dynamic Application.
%i (lowercase "eye")	Asset	Asset Location.
%J	Ticket	Ticket subject.
%K	Asset	Asset Floor.
%k	Asset	Asset Room.
%M	Event	Event message.
%m	Automation	Automation policy note.
%N	Action	Automation action name.
%n	Automation	Automation policy name.
%O (uppercase "oh")	Organization	Organization name.
%o (lowercase "oh")	Organization	Organization ID.
%P	Asset	Asset plate.
%p	Asset	Asset panel.
%Q	Asset	Asset punch.
%q	Asset	Asset zone.
%R	Event Policy	Event policy cause/action text.
%r	System	Unique ID / name for the current Skylar One system.
%S	Event	Severity (Healthy - Critical).
%s	Event	Severity (0 - 4). 0=healthy, 1=notify, 2=minor, 3=major, 4=critical.
%T	Dynamic Alert	Dynamic Application alert threshold value.

Variable	Source	Description
%t	Ticket	Ticket ID.
%U	Asset	Asset rack.
%u	Asset	Asset shelf.
%V	Dynamic Alert	Dynamic Application alert result value.
%v	Asset	Asset tag.
%W	Asset	Asset make.
%w	Asset	Asset model.
%X	Event	Entity name.
%x	Event	Entity ID.
%Y	Event	Sub-entity name.
%y	Event	Sub-entity ID.
%Z	Event	Event source (1 - 8).
%z	Event	Event source (Syslog - Group).

Global Settings for Login Alert Messages

In Skylar One, administrators can add a customizable click-through alert message as a security measure at logon. Users will not be able to access the system until the user clicks the **[OK]** button to agree to the terms and conditions of use for that system.

To add a custom login alert message to Skylar One:

1. Go to the **Login Alert Editor** page (System > Settings > Login Alert Message).
2. In the **Alert Message** field, type the text of your login alert message.
3. After entering the login alert text, click the **[Save]** button.
4. When a user logs in, the alert message will display.

NOTE: On a STIG system, you can use the **Command Line Interface Login Message (STIG only)** field on the **Login Alert Message** page to change the banner text that appears after you sign into an account that has access to the Skylar One command-line user interface. You cannot enter or input HTML code in the text banner.

Global Settings for Password Reset Emails

The **Password Reset Email Editor** page (Password Reset Email Editor) allows ScienceLogic administrators to define the email message that is sent to ScienceLogic users who select the "I forgot my password" option from the **Login** page.

If the user enters a valid ScienceLogic username in the **Login** page and then selects the *I forgot my password* option, Skylar One will check the account information for that user. If the user's account information includes an email address, Skylar One will send the user an email message. The email message will include a link that allows the user to redefine their ScienceLogic password. The new password must meet the requirements defined in the **Password Strength** field and the **Password Shadowing** field for the user account. Skylar One will prompt the user to meet these requirements and display a description of those requirements.

The user can select the *I forgot my password* option up to ten times without responding to the sent email (using the link in the email to reset the password). After ten times, Skylar One will no longer send another email message to the user's email address. The user can continue to select the *I forgot my password* option, but Skylar One will not resend an email.

If the user's account information does not include an email address, Skylar One displays the message "Password recovery is not available for your account, please contact your system administrator".

If the user does not enter a valid ScienceLogic username in the **Login** page, the *I forgot my password* option is still displayed, but Skylar One does not send an email. This prevents intruders from guessing ScienceLogic account names.

If the user exceeds the number of login tries (defined in the **Behavior Settings** page), the "I forgot my password" option is not displayed in the **Login** page.

Defining the Email Message for "I forgot my password"

In the **Password Reset Email Editor** page (System > Settings > Password Reset Email), you can define the email that is sent from Skylar One when an end user selects the *I forgot my password* option from the **Login** page.

To define the email message sent by Skylar One:

1. Go to the **Password Reset Email Editor** page (System > Settings > Password Reset Email).
2. Supply a value in each of the following fields:
 - **Priority.** This will be the priority of the email message. Choices are:
 - *High.* Emails will be marked as high priority.
 - *Normal.* Emails will be marked as normal priority.
 - *Low.* Emails will be marked as low priority.
 - **Subject.** This will be the subject of the email message.
 - **Message.** This will be the body of the email message. ***The body must include the variable %L.*** This variable inserts the link to the page that allows the user to reset their ScienceLogic password.
3. You can include the following variables in the **Subject** field and the **Message** field:
 - **%L (uppercase "el").** The link to the page that allows the user to reset their password.
 - **%O (uppercase "oh").** The user's primary organization, as defined in the **Account Permissions** page for the user.

- **%fn (lowercase "eff" "en")**. The user's first name, as defined in the **Account Permissions** page for the user.
- **%ln (lowercase "el "en")**. The user's last name, as defined in the **Account Permissions** page for the user.

4. Click the **[Save]** button to save the email template.
5. When a user follows the link in the email, Skylar One displays the **Login** page, with the message "Your account has been reset. Please create a new password." The user must then enter their new password twice. The new password is recorded in Skylar One and replaces the previous (forgotten) password.

For example, you could define the following:

Subject: ScienceLogic | %O (automated message)

Message: Hello %fn %ln,

Your password for account %A has been reset.

Please use the following link to log in and choose a new password:

%L.

For the user "Keyser Soze", who is a member of the System organization, the following email would be sent:

Subject: ScienceLogic | System (automated message).

Hello Keyser Soze,

Your password for account ksoze has been reset.

Please use the following link to login and choose a new password:

https://name_or_IP_of_EM7_Administration_Portal/login.em7?prs=hash

Global Settings for Security

The **Security Settings** page (System > Settings > Security) allows you to define and view the status of global security settings for your Skylar One system.

You can define the following global security setting on this page:

- **Verify TLS Certificates**. When this setting is enabled, Skylar One services will validate TLS certificates that are presented to them and reject self-signed certificates. To enable this setting, select the checkbox and click **[Save]**.

On this page, you can also view the current status of the Enterprise Key Management Service (EKMS) for your Skylar One system. This service provides strong encryption for Skylar One credentials, and is enabled by default in versions 12.2.0 and later with no additional configuration required.

Global Settings for System Thresholds

The **System Threshold Defaults** page (System > Settings > Thresholds > System) allows you to define global thresholds for system latency, file system usage, counter rollovers, ICMP availability, number of component devices, interface inventory, and inbound messages.

These settings apply to all devices. However, you can override these system settings on a case-by-case basis. For example, you can define thresholds for a device's file systems in the **[Thresholds]** tab of the **Device Investigator** (or the **Device Thresholds** page in the classic Skylar One user interface). The settings you define for the specific device override the settings in the **System Threshold Defaults** page.

To edit the global settings for system thresholds:

1. Go to the **System Threshold Defaults** page (System > Settings > Thresholds > System).
2. In the **System Threshold Defaults** page, you can drag sliders to change to value of each field or edit a field manually. You can edit the value for one or more of the following fields:
 - **System Latency.** During polling, the platform initially pings monitored devices. The value in this field is the maximum number of milliseconds for the device to respond to Skylar One's ping (round-trip time divided by 2). The default value is 100 ms. When the latency threshold is exceeded, Skylar One generates an event for that device.
 - **System Availability.** During polling, Skylar One monitors devices for availability. Availability means the device's ability to accept connections and data from the network. The value in this field is the percent availability required of each device. The default value is 99%. When a device falls below this level of availability, Skylar One generates an event for that device.

During polling, a device has two possible availability values:

- 100%. Device is up and running.
- 0%. Device is not accepting connections and data from the network.

NOTE: Component devices use a Dynamic Application collection object to measure availability. Skylar One polls component devices for availability at the frequency defined in the Dynamic Application. For details, see the chapter on *Monitoring Device Availability and Latency* in the **Monitoring Device Infrastructure Health** manual.

NOTE: The **Ping & Poll Timeout (Msec)** setting in the **Behavior Settings** page (System > Settings > Behavior) affects how Skylar One monitors device availability. This field specifies the number of milliseconds the discovery tool and availability polls will wait for a response after pinging a device. After the specified number of milliseconds have elapsed, the poll will timeout.

- **File System Major.** Threshold that will trigger a "low disk space" event. The default threshold is 85%. When a device has used more disk space than the specified percentage, Skylar One will generate a "file system usage exceeded threshold" event with a status of "major".

- **File System Critical.** Threshold that will trigger a "low disk space" event. The default threshold is 95%. When a device has used more disk-space than the specified percentage, Skylar One will generate a "file system usage exceeded threshold" event with a status of "critical".

NOTE: If you hide a file system in the **Device Hardware** page (Devices > Hardware), Skylar One does not generate events for that file system.

- **Rollover Percent.** For any collected data that uses a 32-bit counter, you can specify how Skylar One determines that the counter has "rolled over", that is, has reached its maximum value, is reset to zero, and restarts counting. When this happens, the collected values go from the maximum value to a lower value. However, there are multiple circumstances under which a counter value can go from a higher value to a lower value:
 - Maximum value has been exceeded and counter was reset to zero.
 - Retrieved value was manually reset to zero on the external device.
 - Data was collected out-of-order, that is, due to a slowdown somewhere in the network, two counter values were stored out of sequence.

NOTE: For 64-bit counters, when the counter values go from a higher value to a lower value, Skylar One assumes that the counter has been manually reset or that the two values were collected out of order. Skylar One does not assume that the counter has rolled over.

The **Rollover Percent** field allows you to specify a threshold that indicates that a 32-bit counter has reached its maximum value and restarted counting. The default value is 20%. When Skylar One records a counter value that is lower than the previously collected value, the platform:

- Calculates the difference between the two counter values (the delta):

$$2^{32} - \text{Last Collected Value} + \text{Current Collected Value}$$
- Examines the value of the **Rollover Percent** threshold. If the delta is less than the specified percentage of the maximum possible value (2^{32}), Skylar One concludes that the 32-bit counter rolled over.
- For example, if you specified "25" in this field, Skylar One would determine if the delta is less than 25% of the maximum possible value. If the delta is less than 25% of the maximum possible value, Skylar One concludes that the 32-bit counter rolled over.
- When Skylar One determines a counter has rolled over, Skylar One uses the delta value when displaying the data point for this poll period.

NOTE: The *Rollover Percent* field applies only to 32-bit counters. If a 64-bit counter value goes from a higher value to a lower value, the change is treated as either a manual reset or an out-of-order collection.

- ***Out-of-order Percent***. For any collected data that uses a counter, you can specify how Skylar One determines that data has been collected out of order. When this data is collected out of order, the collected values go from a higher value to a lower value. However, there are multiple circumstances under which a counter value can go from a higher value to a lower value:
 - Maximum value has been exceeded and counter was reset to zero (for 32-bit counters only).
 - Data was collected out-of-order, that is, due to a slowdown somewhere in the network, two counter values were stored out of sequence.
 - Retrieved value was manually reset to zero on the external device.

The *Out-of-order Percent* field allows you to specify a threshold that indicates that data has been collected out of order. The default value is 50%. When Skylar One records a counter value that is lower than the previously collected value and the platform has determined that the value is not a rollover, Skylar One:

- Compares the current value to the last collected value:
$$\text{current value} / \text{last collected value}$$
- If the ratio of current value / last collected value is greater than the percent specified in the *Out-of-order Percent* field, Skylar One concludes that the data was collected out of order.
- When Skylar One determines a data point has been collected out of order, Skylar One uses the following value as the current value of the data point:
$$\text{last collected value} - \text{current collected value}$$

NOTE: If a 32-bit counter value goes from the maximum value to a lower value, and the current collected value does not meet the criteria for a rollover AND the current collected value does not meet the criteria for out-of-order, Skylar One concludes that the 32-bit counter was manually reset to zero (0). Skylar One uses the current collected value for this data point.

NOTE: If a 64-bit counter value goes from a higher value to a lower value, and the current collected value does not meet the criteria for out-of-order, Skylar One concludes that the 64-bit counter was manually reset to zero (0). Skylar One uses the current collected value for this data point.

- **Availability Ping Count.** If you select *ICMP* in the **Availability Port** field in the **Device Properties** page (Devices > Classic Devices > wrench icon) for a device, this field specifies the number of packets that should be sent during each availability check. The default value is "1".
- **Avail Required Ping Percentage.** If you select *ICMP* in the **Availability Port** field in the **Device Properties** page (Devices > Classic Devices > wrench icon) for a device, this field specifies the percentage of packets that must be returned during an availability check for Skylar One to consider the device available. The default value is "100%".
- **Process Runtime Threshold Low.** Threshold that will trigger a "process time exceeded" event. The default threshold is 80%. When a process has used more than 80% of its allowed **Run Length**, Skylar One will generate a "process time exceeded threshold" event with a status of "minor".
- **Process Runtime Threshold High.** Threshold that will trigger a "process time exceeded" event. The default threshold is 100%. When a process has used 100% of its allowed **Run Length**, Skylar One will generate a "process time exceeded threshold" event with a status of "major".

NOTE: *Run Length* is defined in the **Process Manager** page (System > Settings > Admin Processes).

- **Component Purge Timeout.** This field specifies the number of hours a device can be set to "vanished" before Skylar One purges the component device. When a device is purged, Skylar One stops trying to collect data about the component device. The purged device will not appear in reports or views on any pages in the user interface. When a device is purged, all of its configuration data and collected data is deleted from the Database Server. If you set this value to "0", component devices are never purged. You can override this threshold for a specific device in the **[Thresholds]** tab of the **Device Investigator** (or the **Device Thresholds** page in the classic Skylar One user interface) for the device.

NOTE: When a device is set to "vanished", all children of that device are also set to "vanished". When a device is purged, all children of that device are also purged.

- **Component Vanish Timeout Mins.** If Skylar One cannot retrieve information from a root device about a component device, this field specifies how many minutes to wait until putting the component device into "vanish" mode. When a device is set to "vanished", Skylar One stops trying to collect data about the component device. The vanished device will not appear in reports or views. The vanished device will appear in the **Vanished Device Manager** page. If you set this value to "0", component devices are never set to "vanished". You can override this threshold for a specific device in the **[Thresholds]** tab of the **Device Investigator** (or the **Device Thresholds** page in the classic Skylar One user interface) for the device.
- **Interface Inventory Timeout.** Specifies the maximum amount of time that the discovery processes will spend polling a device for the list of interfaces. After the specified time, Skylar One will stop polling the device, will not model the device, and will continue with discovery. The default value is 600,000 ms (10 minutes).

- During *initial discovery*, initiated from the Discovery Session Editor page (System > Manage > Classic Discovery > Create), Skylar One uses the value in this field if there is no differing value specified in the **Discovery Session Editor** page.
- During *re-discovery* (clicking the magnifying glass icon (🔍) in the Device Properties page), Skylar One will use the value in this field if there is no value is specified in the **[Thresholds]** tab of the **Device Investigator** (or the **Device Thresholds** page in the classic Skylar One user interface) for the device.
- During *nightly auto-discovery* (run automatically by Skylar One every night, to update device information), Skylar One uses the value in this field if no differing value is specified in the **[Thresholds]** tab of the **Device Investigator** (or the **Device Thresholds** page in the classic Skylar One user interface) for a device.
- **Maximum Allowed Interfaces.** Specifies the maximum number of interfaces per device. If a device exceeds this number of interfaces, Skylar One will stop scanning the device, will not model the device, and will continue with discovery. The default value is 10,000.
 - During *initial discovery*, initiated from the **Discovery Session Editor** page (System > Manage > Classic Discovery > Create), Skylar One uses the value in this field if there is no differing value specified in the **Discovery Session Editor** page.
 - During *re-discovery* (clicking the magnifying glass icon (🔍) in the Device Properties page), Skylar One will use the value in this field if there is no differing value is specified in the **[Thresholds]** tab of the **Device Investigator** (or the **Device Thresholds** page in the classic Skylar One user interface) for the device.
 - During *nightly auto-discovery* (run automatically by Skylar One every night, to update device information), Skylar One uses the value in this field if no differing value is specified in the **[Thresholds]** tab of the **Device Investigator** (or the **Device Thresholds** page in the classic Skylar One user interface) for a device.
- **Inbound Message Throttle Thresholds.** Specifies the maximum number of messages that can be received before Skylar One will notify the system administrator and discard the current batch of messages. The default message threshold is 25.
 - *Syslog per-IP.* Specifies the threshold for incoming syslog messages from a given IP address.
 - *Dynamic Alert per-device.* Specifies the threshold for incoming alerts for a Dynamic Application on a given device.
 - *SNMP Trap per-IP.* Specifies the threshold for incoming SNMP traps from a given IP address.

3. Click the **[Save]** button to save changes in this page.
4. All changes to this page are logged in the audit logs.

Global Settings for Interface Thresholds

The **Interface Thresholds Defaults** page (System > Settings > Thresholds > Interface) allows you to define global thresholds for interfaces.

The settings in the **Interface Thresholds Defaults** page apply to all interfaces. However, you can override these system settings on a case-by-case basis for each interface in the **Thresholds** tab on the **Interface Properties** page (Registry > Networks > Interfaces > interface wrench icon).

If you have specified that Skylar One should monitor an interface, Skylar One will collect data about the interface and also monitor performance thresholds for the interface. Skylar One will use either the default thresholds defined in the **Interface Thresholds Defaults** page (System > Settings > Thresholds > Interface) or the custom threshold you define in the **Thresholds** tab on the **Interface Properties** page (Registry > Networks > Interfaces > interface wrench icon). When the values for an interface exceed one or more thresholds, Skylar One will generate an event.

To define global thresholds for interfaces:

1. Go to **Interface Thresholds Defaults** page (System > Settings > Thresholds > Interface).
2. The following global thresholds are defined by default in the **Interface Thresholds Defaults** page:

NOTE: You can specify the unit of measure for all the metrics in **Bandwidth In** and **Bandwidth Out**. You can select **bps**, **kbps**, **Mbps** (the default), or **Gbps**.

Threshold	Default Value	Default Status
Utilization % In > Inbound Percent	65.000	Enabled
Utilization % Out > Outbound Percent	65.000	Enabled
Bandwidth In > Inbound Bandwidth	0.000	Disabled
Bandwidth Out > Outbound Bandwidth	0.000	Disabled
Errors % In > Inbound Error Percent	1.000	Enabled
Errors % Out > Outbound Error Percent	1.000	Enabled
Errors In > Inbound Errors	1000.000	Enabled
Errors Out > Outbound Errors	1000.000	Enabled
Discard % In > Inbound Discard Percent	1.000	Enabled
Discards % Out > Outbound Discard Percent	1.000	Enabled
Discards In > Inbound Discards	1000.000	Enabled
Discards Out > Outbound Discards	1000.000	Enabled

Threshold	Default Value	Default Status
Multicast % In > Rising Medium	30.000	Disabled
Multicast % In > Rising Low	20.000	Disabled
Broadcast % Out > Rising Medium	30.000	Disabled
Broadcast % Out > Rising Low	20.000	Disabled

3. Selecting the **Show Hidden Thresholds** checkbox displays the following default thresholds:

NOTE: You can specify the unit of measure for all the metrics in **Bandwidth In** and **Bandwidth Out**. You can select **bps**, **kbps**, **Mbps** (the default), or **Gbps**.

Threshold	Default Value	Default Status
Utilization % In > Rising High	0.000	Hidden
Utilization % In > Rising Medium	0.000	Hidden
Utilization % In > Rising Low	0.000	Hidden
Utilization % In > Falling Low	0.000	Hidden
Utilization % In > Falling Medium	0.000	Hidden
Utilization % In > Falling High	0.000	Hidden
Utilization % In > Inbound Percent	65.000	Enabled
Utilization % Out > Rising High	0.000	Hidden
Utilization % Out > Rising Medium	0.000	Hidden
Utilization % Out > Rising Low	0.000	Hidden
Utilization % Out > Falling Low	0.000	Hidden
Utilization % Out > Falling Medium	0.000	Hidden
Utilization % Out > Falling High	0.000	Hidden
Utilization % Out > Outbound Percent	65.000	Enabled
Bandwidth In > Rising High	0.000	Hidden
Bandwidth In > Rising Medium	0.000	Hidden
Bandwidth In > Rising Low	0.000	Hidden
Bandwidth In > Falling Low	0.000	Hidden

Threshold	Default Value	Default Status
Bandwidth In > Falling Medium	0.000	Hidden
Bandwidth In > Falling High	0.000	Hidden
Bandwidth In > Inbound Bandwidth	0.000	Disabled
Bandwidth Out > Rising High	0.000	Hidden
Bandwidth Out > Rising Medium	0.000	Hidden
Bandwidth Out > Rising Low	0.000	Hidden
Bandwidth Out > Falling Low	0.000	Hidden
Bandwidth Out > Falling Medium	0.000	Hidden
Bandwidth Out > Falling High	0.000	Hidden
Bandwidth Out > Outbound Bandwidth	0.000	Disabled
Errors % In > Rising High	0.000	Hidden
Errors % In > Rising Medium	0.000	Hidden
Errors % In > Rising Low	0.000	Hidden
Errors % In > Falling Low	0.000	Hidden
Errors % In > Falling Medium	0.000	Hidden
Errors % In > Falling High	0.000	Hidden
Errors % In > Inbound Error Percent	1.000	Enabled
Errors % Out > Rising High	0.000	Hidden
Errors % Out > Rising Medium	0.000	Hidden
Errors % Out > Rising Low	0.000	Hidden
Errors % Out > Falling Low	0.000	Hidden
Errors % Out > Falling Medium	0.000	Hidden
Errors % Out > Falling High	0.000	Hidden
Errors % Out > Outbound Error Percent	1.000	Enabled
Errors In > Rising High	0.000	Hidden
Errors In > Rising Medium	0.000	Hidden
Errors In > Rising Low	0.000	Hidden
Errors In > Falling Low	0.000	Hidden
Errors In > Falling Medium	0.000	Hidden

Threshold	Default Value	Default Status
Errors In > Falling High	0.000	Hidden
Errors In > Inbound Errors	1000.000	Enabled
Errors Out > Rising High	0.000	Hidden
Errors Out > Rising Medium	0.000	Hidden
Errors Out > Rising Low	0.000	Hidden
Errors Out > Falling Low	0.000	Hidden
Errors Out > Falling Medium	0.000	Hidden
Errors Out > Falling High	0.000	Hidden
Errors Out > Outbound Errors	1000.000	Enabled
Discards % In > Rising High	0.000	Hidden
Discards % In > Rising Medium	0.000	Hidden
Discards % In > Rising Low	0.000	Hidden
Discards % In > Falling Low	0.000	Hidden
Discards % In > Falling Medium	0.000	Hidden
Discards % In > Falling High	0.000	Hidden
Discards % In > Inbound Discard Percent	1.000	Enabled
Discards % Out > Rising High	0.000	Hidden
Discards % Out > Rising Medium	0.000	Hidden
Discards % Out > Rising Low	0.000	Hidden
Discards % Out > Falling Low	0.000	Hidden
Discards % Out > Falling Medium	0.000	Hidden
Discards % Out > Falling High	0.000	Hidden
Discards % Out > Outbound Discard Percent	1.000	Enabled
Discards In > Rising High	0.000	Hidden
Discards In > Rising Medium	0.000	Hidden
Discards In > Rising Low	0.000	Hidden
Discards In > Falling Low	0.000	Hidden
Discards In > Falling Medium	0.000	Hidden
Discards In > Falling High	0.000	Hidden

Threshold	Default Value	Default Status
Discards In > Inbound Discards	1000.000	Enabled
Discards Out > Rising High	0.000	Hidden
Discards Out > Rising Medium	0.000	Hidden
Discards Out > Rising Low	0.000	Hidden
Discards Out > Falling Low	0.000	Hidden
Discards Out > Falling Medium	0.000	Hidden
Discards Out > Falling High	0.000	Hidden
Discards Out > Outbound Discards	1000.000	Enabled
Broadcast % In > Rising High	0.000	Hidden
Broadcast % In > Rising Medium	30.000	Disabled
Broadcast % In > Rising Low	20.000	Disabled
Broadcast % In > Falling Low	0.000	Hidden
Broadcast % In > Falling Medium	0.000	Hidden
Broadcast % In > Falling High	0.000	Hidden
Broadcast % Out > Rising High	0.000	Hidden
Broadcast % Out > Rising Medium	30.000	Disabled
Broadcast % Out > Rising Low	20.000	Disabled
Broadcast % Out > Falling Low	0.000	Hidden
Broadcast % Out > Falling Medium	0.000	Hidden
Broadcast % Out > Falling High	0.000	Hidden
Broadcast In > Rising High	0.000	Hidden
Broadcast In > Rising Medium	0.000	Hidden
Broadcast In > Rising Low	0.000	Hidden
Broadcast In > Falling Low	0.000	Hidden
Broadcast In > Falling Medium	0.000	Hidden
Broadcast In > Falling High	0.000	Hidden
Broadcast Out > Rising High	0.000	Hidden
Broadcast Out > Rising Medium	0.000	Hidden
Broadcast Out > Rising Low	0.000	Hidden

Threshold	Default Value	Default Status
Broadcast Out > Falling Low	0.000	Hidden
Broadcast Out > Falling Medium	0.000	Hidden
Broadcast Out > Falling High	0.000	Hidden
Multicast % In > Rising High	0.000	Hidden
Multicast % In > Rising Medium	00.000	Hidden
Multicast % In > Rising Low	00.000	Hidden
Multicast % In > Falling Low	0.000	Hidden
Multicast % In > Falling Medium	0.000	Hidden
Multicast % In > Falling High	0.000	Hidden
Multicast % Out > Rising High	0.000	Hidden
Multicast % Out > Rising Medium	00.000	Hidden
Multicast % Out > Rising Low	00.000	Hidden
Multicast % Out > Falling Low	0.000	Hidden
Multicast % Out > Falling Medium	0.000	Hidden
Multicast % Out > Falling High	0.000	Hidden
Multicast In > Rising High	0.000	Hidden
Multicast In > Rising Medium	0.000	Hidden
Multicast In > Rising Low	0.000	Hidden
Multicast In > Falling Low	0.000	Hidden
Multicast In > Falling Medium	0.000	Hidden
Multicast In > Falling High	0.000	Hidden
Multicast Out > Rising High	0.000	Hidden
Multicast Out > Rising Medium	0.000	Hidden
Multicast Out > Rising Low	0.000	Hidden
Multicast Out > Falling Low	0.000	Hidden
Multicast Out > Falling Medium	0.000	Hidden
Multicast Out > Falling High	0.000	Hidden
Unicast % In > Rising High	0.000	Hidden
Unicast % In > Rising Medium	00.000	Hidden

Threshold	Default Value	Default Status
Unicast % In > Rising Low	00.000	Hidden
Unicast % In > Falling Low	0.000	Hidden
Unicast % In > Falling Medium	0.000	Hidden
Unicast % In > Falling High	0.000	Hidden
Unicast % Out > Rising High	0.000	Hidden
Unicast % Out > Rising Medium	00.000	Hidden
Unicast % Out > Rising Low	00.000	Hidden
Unicast % Out > Falling Low	0.000	Hidden
Unicast % Out > Falling Medium	0.000	Hidden
Unicast % Out > Falling High	0.000	Hidden
Unicast In > Rising High	0.000	Hidden
Unicast In > Rising Medium	0.000	Hidden
Unicast In > Rising Low	0.000	Hidden
Unicast In > Falling Low	0.000	Hidden
Unicast In > Falling Medium	0.000	Hidden
Unicast In > Falling High	0.000	Hidden
Unicast Out > Rising High	0.000	Hidden
Unicast Out > Rising Medium	0.000	Hidden
Unicast Out > Rising Low	0.000	Hidden
Unicast Out > Falling Low	0.000	Hidden
Unicast Out > Falling Medium	0.000	Hidden
Unicast Out > Falling High	0.000	Hidden

4. For each threshold, you can edit the following:

- **Value.** The value at which the threshold will trigger an event.
 - For thresholds that include the word *Rising*, when a value exceeds the specified value, Skylar One triggers an event.
 - For thresholds that include the word *Falling*, when a value falls below the specified value, Skylar One triggers an event.
 - For thresholds that do not include the word *Rising* or *Falling*, when a value exceeds the specified value, Skylar One triggers an event.

- **Status.** Specifies whether the threshold is active and whether the threshold will appear in the **Thresholds** tab on the **Interface Properties** page (Registry > Networks > Interfaces > interface wrench icon). Choices are:
 - *Enabled.* The threshold is applied to all interfaces and is monitored by Skylar One. The threshold appears in the **Thresholds** tab on the **Interface Properties** page (Registry > Networks > Interfaces > interface wrench icon). Users can edit the **Value** and **Status** of the threshold.
 - *Disabled.* The threshold is applied to all interfaces but is not monitored by Skylar One. The threshold appears in the **Thresholds** tab on the **Interface Properties** page (Registry > Networks > Interfaces > interface wrench icon) with a status of *Disabled*. In the **Thresholds** tab on the **Interface Properties** page, users can edit the **Value** and **Status** of the threshold.
 - *Hidden.* The threshold is not applied to all interfaces, and is not monitored by Skylar One. The threshold does not appear in the **Thresholds** tab on the **Interface Properties** page (Registry > Networks > Interfaces > interface wrench icon).
- **Unit of Measure.** For all the metrics under **Bandwidth In** and **Bandwidth Out**, you can select the unit of measure. Choices are:
 - kbps
 - Mbps
 - Gbps

Settings in Silo.conf

Every Skylar One appliance has a configuration file called **silo.conf**, which contains configuration information about the appliance itself, such as the IP address, licensing information, and directory locations. The default settings in **silo.conf** are configured automatically when the appliance is installed. The following section describes how you can add additional, non-default settings to **silo.conf**.

CAUTION: ScienceLogic recommends that you do not edit the values in these files without first consulting ScienceLogic. Incorrect values can severely disrupt platform operations.

NOTE: All settings in these **.conf** files are case-sensitive.

To edit the **silo.conf** file:

1. Either go to the console of the Skylar One appliance or use SSH to access the Skylar One appliance.
2. Open a shell session on the server.

3. Type the following at the command line:

```
sudo visilo
```

IMPORTANT: For ISO installs of Skylar One version 11.3.0 and later, password information in the **silo.conf** file is automatically encrypted when the file is modified using **visilo**. Users can no longer decrypt passwords in the **silo.conf** file.

NOTE: You can use the **ap_user** and **ap_pass** fields to define usernames and passwords for the Administration Portal that differ from the usernames and passwords used for the Database Server.

4. You can add or edit one or more of the following settings:

- **store_timeout**. You can edit this setting in the **silo.conf** file on each Database Server. When the Database Server pulls collected data back from Data Collectors and Message Collectors, each piece of data (called a storage object) must be stored within a set amount of time. The default timeout for a storage object is ten seconds. To change the timeout for all storage objects, add the following line to the **silo.conf** file on the Database Server:

```
store_timeout=xx
```

where **xx** is the timeout in seconds.

If you change this setting (for example, change the value to 30 seconds), you must stop and restart the high frequency, medium frequency, and low frequency data pull processes for the change to be applied.

NOTE: The **store_timeout** setting does not apply to All-In-One Appliances.

Data Pull settings are frequency-specific with Skylar One version 12.3.0. To apply a setting to a specific frequency of Data Pull prepend **lf_**, **mf_**, or **hf_** to the original setting name for low, medium, and high frequency Data Pull respectively.

NOTE: If a frequency-specific setting exists, it will take precedence over any generic value for the same setting.

The following two examples show this precedence:

Example 1

An **/etc/silo.conf** file setting **mf_num_storage_procs** to 10 while the low frequency and high frequency settings are set to 8 through the generic setting.

```
[DATA PULL]
```

```
num_storage_procs = 8
mf_num_storages_procs = 10
```

Example 2

An /etc/silo.conf file setting individual frequency memory limits along with Example 1 storage process settings.

```
[DATA_PULL]
num_storage_procs = 8
mf_num_storages_procs = 10
lf_memory_limit = 1073741824
mf_memory_limit = 4294967296
hf_memory_limit = 2147483648
```

- **eventmanager**. You can edit this setting in the **silo.conf** file on each Skylar One appliance. You can modify this default setting to allow API events to be processed on a Data Collector. The default configuration is:

```
eventmanager = internal,dynamic,syslog,trap
```

To allow a Data Collector to process API events, change this line to

```
eventmanager = internal,dynamic,syslog,trap,api
```

WARNING: Do not make any other changes to this setting or modify this setting on a Database Server or Data Collector.

- **report_memory_limit**. You can edit this setting in the **silo.conf** file on each Skylar One appliance that provides the user interface (an Administration Portal, Database Server, or All-In-One Appliance). If **report_memory_limit** is not defined in **silo.conf**, the default value is three gigabytes (3G). If reports are failing to be generated due to a lack of memory, you can increase this value.

To increase report memory, add the following line to the **[LOCAL]** section of **silo.conf** on each Skylar One appliance that provides the user interface for your system. In most cases, this will be the Administration Portal (for distributed system) or the All-In-One Appliance:

```
report_memory_limit=XY
```

where:

- **X** is a positive integer
- **Y** represents units. Value can be **K**(kilobytes), **M**(megabytes), or **G**(gigabytes),

For example, if reports are failing to be generated due to a lack of memory, you could add the following line to **silo.conf**:

```
report_memory_limit=4G
```

NOTE: You should add the *report_memory_limit* option to the **silo.conf** file on a Database Server only if there are no Administration Portals configured in your system.

NOTE: You must add the same *report_memory_limit* setting to every Administration Portal configured in your system.

- ***use_v1trap_envelope_addr***. You can edit this setting in the silo.conf file on each Data Collectors, Data Collectors that perform message collection, and All-In-One Appliances. In environments where Network Address Translation is performed on SNMP v1 trap messages sent to Skylar One, you can configure the platform to read the envelope address (the address of the host sending the trap) instead of the agent address (the IP address variable sent as part of the trap). If ***use_v1trap_envelope_addr*** is not defined in silo.conf, Skylar One will use the agent address for SNMP v1 trap messages.
 - To use the envelope address instead of the agent address for SNMP v1 trap messages, add the following line to the [LOCAL] section of silo.conf on Data Collectors, Data Collectors that perform message collection, and All-In-One Appliances

```
use_v1trap_envelope_addr=1
```

- To use the agent address for SNMP v1 trap messages, you can either omit the ***use_v1trap_envelope_addr*** setting or add the following line to the [LOCAL] section of silo.conf on Data Collectors, Data Collectors that perform message collection, and All-In-One Appliances

```
use_v1trap_envelope_addr=0
```

- ***disable_itil_compliance***. You can edit this setting in the silo.conf file on each If you enable this setting on an appliance that provides the user interface (an Administration Portal, Database Server, or All-In-One Appliance), the **Ticket Console** page on that appliance will include an option to delete tickets. The option to delete tickets will appear only to users that have been granted the Ticket: Delete access hook and users of type "administrator".

To enable this setting, add the following line to the [LOCAL] section of silo.conf on the appliance that provides the user interface (Administration Portal, Database Server, or All-In-One Appliance):

```
disable_itil_compliance=1
```

- ***suppress_ticket_link***. You can edit this setting in the **silo.conf** file on each Skylar One appliance that provides the user interface (an Administration Portal, Database Server, or All-In-One Appliance). If you enable this setting, automatic notifications that are generated when a ticket is created or updated will not include a hyperlink to the ticket.

To enable this setting, add the following line to the [LOCAL] section of **silo.conf** on Skylar One appliance that provides the user interface (Administration Portal, Database Server, or All-In-One Appliance):

```
suppress_ticket_link=1
```

- ***mailparse_interval***. You can edit this setting in the **silo.conf** file on each Database Server or All-In-One Appliance. The ***mailparse_interval*** setting defines how frequently the mail parsing process reads email messages from the mailbox. If the **mailparse_interval** setting is not defined in **silo.conf**, the default value is 60 seconds. When an email is received by Skylar One, the mail parsing process on the primary Database Server or All-In-One Appliance reads the email message from the mailbox file and sends it to one of the three processes responsible for acting on that email: the event engine (for events from email), the tickets from email process, or the round-trip email collection process.

To enable this setting, add the following line to the [LOCAL] section of **silo.conf** on each Database Server or All-In-One Appliance:

```
mailparse_interval=X
```

where **X** is the frequency at which the mailbox will be read, in seconds. Valid values are 15 seconds to 60 seconds.

- ***dynamic_collect_num_chunk_workers***. You can edit this setting in the **silo.conf** file on each Database Server or All-In-One Appliance. This setting represents the number of workers that handle collection requests. Skylar One first sorts collection requests into groups by execution environment and sends each group of collection requests (called a chunk) to a worker process. This worker process is called a chunk worker. For each chunk, a chunk worker creates the execution environment and creates a pool of request workers to process the collection requests. The number of chunk workers generally represents the number of PowerPacks that can be processed in parallel. The default value for this parameter is "2".

To change this setting, add the following line to the [LOCAL] section of **silo.conf** on each Database Server or All-In-One Appliance:

```
dynamic_collect_num_chunk_workers = [X]
```

where **X** is the number of chunk workers

NOTE: For more information about using ***dynamic_collect_num_chunk_workers***, see the section on [Tuning the Collector Load Balancing Process in the Silo.Config File](#).

- ***dynamic_collect_num_request_workers***. You can edit this setting in the silo.conf file on each Database Server or All-In-One Appliance. This setting represents the maximum number of request workers in each worker pool and generally represents the number of collections within a PowerPack that can be processed in parallel. The default value for this parameter is "2" or the number of cores on the Data Collector, whichever is greater.

To change this setting, add the following line to the [LOCAL] section of silo.conf on each Database Server or All-In-One Appliance:

```
dynamic_collect_num_request_workers = [X]
```

where:

- *X* is the maximum number of request workers in each worker pool.

NOTE: For more information about using ***dynamic_collect_num_request_workers***, see the section on [Tuning the Collector Load Balancing Process in the Silo.Config File](#).

- ***dynamic_collect_request_chunk_size***. You can edit this setting in the silo.conf file on each Database Server or All-In-One Appliance. This setting represents the maximum number of collection requests in a chunk and controls how many collections are processed by a single pool or request workers. The default value for this parameter is "200".

To change this setting, add the following line to the [LOCAL] section of silo.conf on each Database Server or All-In-One Appliance:

```
dynamic_collect_request_chunk_size = [X]
```

where:

- *X* is the maximum number of collection requests in a chunk.

NOTE: For more information about using ***dynamic_collect_request_chunk_size***, see the section on [Tuning the Collector Load Balancing Process in the Silo.Config File](#).

- ***read_timeout***. You can edit this setting in the silo.conf file on each Database Server. This setting controls the client read timeout for database connections to the collectors. This setting applies only to the ***Enterprise Database: Collector Config Push*** process (config_push.py) that runs on the primary Database Server.

WARNING: Change this value only if you are instructed to do so by ScienceLogic.

To change this setting, add the following line to the [CONFIG_PUSH] section of silo.conf on all Database Servers in your system.

```
read_timeout=X
```

where:

- ***X*** is the read timeout, in seconds.
- ***wait_timeout***. You can edit this setting in the silo.conf file on each Database Server. This setting controls the server wait timeout for database connections to the collectors. This setting applies only to the ***Enterprise Database: Collector Config Push*** process (config_push.py) that runs on the primary Database Server.

WARNING: Change this value only if you are instructed to do so by ScienceLogic.

To change this setting, add the following line to the [CONFIG_PUSH] section of silo.conf on all Database Servers in your system.

```
wait_timeout=X
```

where:

- ***X*** is the write timeout, in seconds.

- ***write_timeout***. You can edit this setting in the silo.conf file on each Database Server. This setting controls the client write timeout for database connections to the collectors. This setting applies only to the ***Enterprise Database: Collector Config Push*** process (config_push.py) that runs on the primary Database Server.

WARNING: Change this value only if you are instructed to do so by ScienceLogic.

```
write_timeout=X
```

where:

- *X* is the write timeout, in seconds.
- ***memory_limit***. You can edit this setting in the silo.conf file on each Database Server. This setting controls the memory limit for the ***Enterprise Database: Collector Config Push*** process. This setting applies only to the ***Enterprise Database: Collector Config Push*** process (config_push.py) that runs on the primary Database Server.

WARNING: Change this value only if you are instructed to do so by ScienceLogic.

To change this setting, add the following line to the [CONFIG_PUSH] section of silo.conf on all Database Servers in your system.

```
memory_limit=XY
```

where:

- *X* is a positive integer.
- *Y* represents units. Value can be ***KB*** (kilobytes), ***MB*** (megabytes), or ***GB*** (gigabytes).

- ***result_wait_timeout***. You can edit this setting in the silo.conf file on each Database Server. This setting controls the amount of time the parent ***Enterprise Database: Collector Config Push*** process will wait for a message from a child process before abandoning that process. This setting applies only to the ***Enterprise Database: Collector Config Push*** process (config_push.py) that runs on the primary Database Server

WARNING: Change this value only if you are instructed to do so by ScienceLogic.

To change this setting, add the following line to the [CONFIG_PUSH] section of silo.conf on all Database Servers in your system.

```
result_wait_timeout=X
```

where:

- *X* is the write timeout, in seconds.
- ***shutdown_timeout***. You can edit this setting in the silo.conf file on each Database Server. If the ***Enterprise Database: Collector Config Push*** process is terminated, this setting controls the amount of time the parent configuration process will wait for its child processes to stop before terminating itself and allowing the child processes to be inherited by init. This setting applies only to the ***Enterprise Database: Collector Config Push*** process (config_push.py) that runs on the primary Database Server.

WARNING: Change this value only if you are instructed to do so by ScienceLogic.

To change this setting, add the following line to the [CONFIG_PUSH] section of silo.conf on all Database Servers in your system.

```
shutdown_timeout=X
```

where:

- *X* is the write timeout, in seconds.

- **[PROC_VIRTUAL_MEM_LIMIT]**. By default, processes in Skylar One have a virtual memory limit of 1 GB. You can edit this section in the silo.conf file to overwrite the existing virtual memory limit for a given process in Skylar One to ensure that it does not fail by crossing its virtual memory limit.

To change this setting, add the [PROC_VIRTUAL_MEM_LIMIT] section to the silo.conf file. Below that section heading, specify the process you want to update and the new virtual memory limit for that process. Use the following format for each setting:

```
[process ID]=X
```

where:

- *[process ID]* is the ID of the process you want to update, as found in master.system_settings_procs.aid
- *X* is the new virtual memory limit, in bytes

For example, if you wanted to update a process with an ID of "12" with a new 2 GB memory limit, you would write the following under [PROC_VIRTUAL_MEM_LIMIT]:

```
12=2147483648
```

- **[ADHOC_REPORT_IN_BATCH]**. ad hoc reports are processed in a batch process. You can edit this section in the silo.conf file to overwrite the default timing values for certain ad hoc reporting settings.

To change these settings, under the [ADHOC_REPORT_IN_BATCH] section heading in the silo.conf file, specify the time value (in seconds) for each setting. The following settings are included in the [ADHOC_REPORT_IN_BATCH] section:

- *report_execution_delay*. This setting controls the amount of time between when a report is scheduled to start running and when it actually begins running. Its default value is 10.
- *ajax_start_delay*. This setting controls the amount of time elapsed before jQuery triggers the ajaxStart event. Its default value is 20.
- *ajax_stop_time*. This setting controls the amount of time elapsed before jQuery triggers the ajaxStop event after all AJAX requests have completed. Its default value is 1800.
- *ajax_frequency*. This setting controls the frequency with which jQuery fires AJAX requests. Its default value is 10.
- *ajax_frequency_decreased_after*. This setting controls the amount of time elapsed after which jQuery will fire AJAX requests less frequently than in the *ajax_frequency* setting. Its default value is 300.

- *ajax_decreased_frequency*. This setting controls the decreased frequency with which jQuery fires AJAX requests after the amount of time listed in the *ajax_frequency_decreased_after* setting has elapsed. Its default value is 60.
- *report_fail_check_time*. This setting controls the amount of time elapsed after which a running report will be considered to have failed. Its default value is 10800.
- *auto_page_refresh*. This setting controls the amount of time elapsed after which the **Scheduled Report Jobs** page (Report > Create Report > Scheduled Job / Report Archive) automatically refreshes. Its default value is 10.
- *about_to_start_time_check*. This setting controls the amount of time before a report job is scheduled to start that it will be labeled as "About to start" on the **Scheduled Report Jobs** page (Report > Create Report > Scheduled Job / Report Archive). Its default value is 30.
- *time_unit*. This setting controls the unit of time measurement for the ad hoc report settings. Its default value is "second".
- *ui_php_timeout*. This setting controls the amount of time elapsed after which an inactive Skylar One reports session will time out. Its default value is 1800.

5. To save your changes, click **Save** and then close the modal window.

NOTE: All changes to the **silo.conf** file are logged in the Skylar OneDatabase Server.

Disabling the User Interface on a Database Server

Database Servers are automatically configured to provide the user interface. If your Skylar One system includes an Administration Portal, you might want to disable the user interface capability on your Database Server(s). Perform the following steps to disable the user interface capability on a Database Server:

NOTE: To complete these steps, you must be familiar with how to edit a file using the vi text editor. If you need assistance with these steps, please contact ScienceLogic Support.

1. Log in to the console of the Database Server or use SSH to access the server as the **em7admin** user with the appropriate password.
2. Execute the following command to open the firewall rules file:

```
sudo vifirewalld
```

3. Add the following lines:

```
rule port port="443" protocol="tcp" reject
```

```
rule port port="80" protocol="tcp" reject
```

4. Save the file and exit the vi editor.
5. Execute the following commands to update and restart the firewall:

```
sudo /opt/em7/share/scripts/update-firewalld-conf.py
```

Chapter

3

Collector Groups

Overview

This chapter provides an overview of collector groups in Skylar One (formerly SL1). A **collector group**—sometimes referred to as a CUG—is a group of Skylar One Data Collectors that retrieve data from managed devices and applications so you can use that data in 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:

What is a Collector Group?	80
Installing, Configuring, and Licensing Data Collectors	80
Technical Information About Data Collectors	81
Viewing the List of Collector Groups	82
Creating a Collector Group	84
Editing a Collector Group	94
Collector Groups and Load Balancing	95
Collector Affinity	98
Creating a Collector Group for Data Storage Only	100
Deleting a Collector Group	101

<i>Assigning a Collector Group for a Single Device</i>	102
<i>Aligning the Collector Group in a Device Template</i>	102
<i>Changing the Collector Group for One or More Devices</i>	103
<i>Managing the Host Files for a Collector Group</i>	103
<i>Processes for Collector Groups</i>	103
<i>Enabling and Disabling Concurrent PowerShell for Collector Groups</i>	105
<i>Enabling and Disabling Concurrent SNMP for Collector Groups</i>	107
<i>Enabling Multi-tenancy for Collector Groups</i>	109

What is a Collector Group?

A **collector group**—sometimes referred to as a CUG—is a group of Skylar One Data Collectors. Data Collectors retrieve data from managed devices and applications. This collection occurs during initial discovery, during nightly updates, and in response to policies defined for each managed device. The collected data is used to trigger events, display data in the user interface, and generate graphs and reports.

You can group multiple Data Collectors into a collector group. Depending on the number of Data Collectors in your Skylar One system, you can define one or more collector groups. Each collector group must include at least one Data Collector.

On the **Collector Groups** page (Manage > Collector Groups)—or the **Collector Group Management** page (System > Settings > Collector Groups) in the classic Skylar One user interface—you can view a list of existing collector groups, add a collector group, and edit a collector group.

NOTE: System upgrades will only consider Data Collectors and Message Collectors that are members of a collector group.

Grouping multiple Data Collectors allows you to:

- Create a load-balanced collection system, where you can manage more devices without loss of performance. At any given time, the Data Collector with the lightest load manages the next discovered device.
- Optionally, create a redundant, high-availability system that minimizes downtime should a failure occur. If a Data Collector fails, one or more Collection servers in the collector group will handle collection until the problem is solved.

NOTE: If you are using a Skylar One All-In-One Appliance, most of the sections in this chapter do not apply to your system. For an All-In-One Appliance, a single, default collector group is included with the appliance; you cannot create any additional collector groups. However, you can [view information about the default collector group](#). You can also [create a virtual collector group](#), for data storage only. However, the other tasks described in this section do not apply to an All-In-One Appliance.

Installing, Configuring, and Licensing Data Collectors

Before you can create a collector group, you must install and license at least one Data Collector. For details on installation and licensing of a Data Collector, see the **Installation** manual.

After you have successfully installed, configured, and licensed a Data Collector, the platform automatically adds information about the Data Collector to the Database Server.

NOTE: For more information on using external credential services that store, collect, and retrieve secret data, see the chapter on "Using External Credential Services" in the **Discovery and Credentials** manual.

Technical Information About Data Collectors

You might find the following technical information about Data Collectors helpful when creating collector groups.

Duplicate IP Addresses

A single collector group **cannot** include multiple devices that use the same Admin Primary IP Address (this is the IP address the platform uses to communicate with a device). If a single collector group includes multiple devices that use the same Primary IP Address or use the same Secondary IP Address, the platform will generate an event. Best practice is to ensure that within a single collector group, all IP addresses on all devices are unique.

- During initial discovery, if a device is discovered with the same Admin Primary IP Address as a previously discovered device in the collector group, the later discovered device will appear in the discovery log, but will not be modeled in the platform. That is, the device will not be assigned a device ID and will not be created in the platform. The platform will generate an event specifying that a duplicate Admin Primary IP was discovered within the collector group.
- If you try to assign a device to a collector group, and the device's Admin Primary IP Address already exists in the collector group, the platform will display an error message, and the device will not be aligned with the collector group.

Open Ports

By default, Data Collectors accept connections only to the following ports:

- TCP 22 (SSH)
- TCP 53 (DNS)
- TCP 123 (NTP)
- UDP 161 (SNMP)
- UDP 162 (Inbound SNMP Trap)
- UDP 514 (Inbound Syslog)
- TCP 7700 (Web Configuration Utility)
- TCP 7707 (one-way communication from the Database Server)

For increased security, all other ports are closed.

Viewing the List of Collector Groups

The **Collector Groups** page displays a list of all collector groups in your Skylar One system.

For each collector group, the page displays the following:

- **ID**. Unique numeric identifier automatically assigned by Skylar One to each collector group.
- **Name**. Name of the collector group.
- **Devices Count**. Number of devices currently using the collector group for data collection.
- **Message Collectors**. The name(s) of the Message Collector(s) (if any) associated with the collector group.
- **Data Collectors**. The name(s) of the Data Collectors in the collector group.
- **Edit User**. User who created or last edited the collector group.
- **Edit Date**. Date and time the collector group was created or last edited.
- **Collector Failover**. Indicates if Data Collector failover is enabled or disabled for the collector group.
- **Enable Concurrent SNMP Collection**. Indicates if the collector group has concurrent SNMP collection enabled, disabled, or set to the systemwide default setting.
- **Enable Concurrent PowerShell Collection**. Indicates if the collector group has concurrent PowerShell collection enabled, disabled, or set to the systemwide default setting.
- **Enable Concurrent Network Interface Collection**. Indicates if the collector group has concurrent network interface collection enabled, disabled, or set to the systemwide default setting.
- **Collectors Available for Failover**. The number of Data Collectors that must be available before a Data Collector failover can occur, if Data Collector failover is enabled for the collector group.
- **Fallback Mode**. Indicates if fallback is automatic or manual, if Data Collector failover is enabled for the collector group.
- **Failover Delay (in minutes)**. The number of minutes Skylar One should wait after a Data Collector outage before redistributing the data collection tasks among the other Data Collectors in the collector group, if Data Collector failover is enabled for the collector group.
- **Failback Delay (in minutes)**. The number of minutes Skylar One should wait after the failed Data Collector is restored before redistributing data-collection tasks among the collector group, including the previously failed Data Collector, if Data Collector failover is enabled for the collector group.
- **Status**. Indicates the Oracle Linux 8 (OL8) conversion status for the collector group.
- **Organization(s)**. The organization(s) to which the collector group is assigned.

NOTE: The **Organization(s)** column displays only if *multi-tenancy is enabled for collector groups*. For more information about editing a collector group's organizations, see the section on [Aligning Collector Groups to Organizations](#).

TIP: If you do not see one of these columns on the **Collector Groups** page, click the **Select Columns** icon (⚙) to add or remove columns. You can also drag columns to different locations on the page or click on a column heading to sort the list of collector groups by that column's values. Skylar One retains any changes you make to the columns that appear on the **Collector Groups** page and will automatically recall those changes the next time you visit the 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.

Viewing the List of Collector Groups in the Classic Skylar One User Interface

To view the list of collector groups in the classic Skylar One user interface:

1. Go to the **Collector Group Management** page (System > Settings > Collector Groups).
2. The **Collector Group Registry** pane displays a list of all collector groups in your Skylar One system. For each collector group, the **Collector Group Management** page displays the following:
 - **Name.** Name of the collector group.
 - **ID.** Unique numeric identifier automatically assigned by Skylar One to each collector group.
 - **Organization.** The organization(s) to which the collector group is assigned. Click the organization icon (👤) to edit the collector group's organizations.

NOTE: The **Organization** column displays only if *multi-tenancy is enabled for collector groups*. For more information about editing a collector group's organizations, see the section on *Aligning Collector Groups to Organizations in the Classic Skylar One User Interface*.

- **# Collectors.** Number of Data Collectors in the collector group.
- **Msg Collector.** Name of the Message Collector(s) (if any) associated with the collector group.
- **# Devices.** Number of devices currently using the collector group for data collection.
- **Edit User.** User who created or last edited the collector group.
- **Edit Date.** Date and time the collector group was created or last edited.

Creating a Collector Group

You can group multiple Data Collectors into a collector group. Depending on the number of Data Collectors in your Skylar One system, you can define one or more collector groups. Each collector group must include at least one Data Collector.

Pre-Deployment Questions for a Collector Group

Consider the following questions before creating a new collector group. Your responses to these questions will help you determine how to create and name your new collector group:

- Will your collector group span regionally close data centers and be configured for maximum resilience?
- Will your users be required to know your collector group naming scheme, or will you provide a general collector group for them to use as a default (and use specialized collector groups for distinct use cases only)?
- Will your collector group be structured for minimum latency to the monitored endpoints?
- Consider the following questions about the resilience of your deployment:
 - What happens to the ability to monitor if a data center hosting an entire collector group goes offline?
 - Is the deployment resilient and will it perform well?
 - What is your failure mode? $100\% > 0\%$ or $!00\% > 50\% > 0\%$?

Capacity Planning for a Collector Group

In addition to deciding on your resiliency strategy, look at your failure mode and determine if you are allocating sufficient capacity to achieve a $100\% > 50\%$ capacity degradation on a data center failure before failing completely at 0%.

Consider the number of devices in your collector group and the number of Data Collectors in your collector group to determine if you have overloaded Data Collectors or underpowered Data Collectors.

Defining a Collector Group

To define a new collector group:

1. Go to the **Collector Groups** page (Manage > Collector Groups).
2. Click the **[Add Collector Group]** button. The **Add Collector Group** modal appears.
3. On the **Add Collector Group** modal, complete the following fields:

- **Collector Group Name.** Type a name for the collector group.
- **Virtual Collector Group (vCUG).** Toggle this option on to make the collector group a virtual collector group. Virtual collector groups do not contain any Data Collectors or Message Collectors and **Skylar One does not collect any data from devices aligned with a virtual collector group.** Instead of collecting data, virtual collector groups serve only as storage areas for historical data from decommissioned devices.
- **Generate Alert on Collector Outage.** Toggle this option on to specify that the platform should generate an event if a Data Collector has an outage, or toggle it off if the platform should not generate an event if a Data Collector has an outage.
- **All current and future organizations.** Toggle this option on to align the collector group to all of your Skylar One organizations, or toggle it off to specify the organizations to which you want to assign the collector group.
- **Limit access to specific organizations.** If you toggled off the **All current and future organizations** option, select the organization(s) to which you want to assign the collector group.

NOTE: The **All current and future organizations** and **Limit access to specific organizations** fields display only if [multi-tenancy is enabled for collector groups](#).

- **Message Collector Selection.** Select one or more available Message Collectors from the drop-down list to add it to the collector group.

NOTE: A single Message Collector can be used by multiple collector groups. When you align a single Message Collector with multiple collector groups, the single Message Collector might then be aligned with two devices (each in a separate collector group) that use the same primary IP address or the same secondary IP address. If this happens, Skylar One will generate an event.

- **Data Collector Selection.** Select one or more available Data Collectors from the drop-down list to add it to the collector group.

- **Concurrent SNMP Collection.** Specifies whether you want to enable concurrent SNMP collection. Concurrent SNMP collection uses asynchronous input/output for massive concurrency with lower system resource requirements. This means that Data Collectors can collect more data using fewer system resources. Concurrent SNMP collection also prevents missed polls and data gaps because collection will execute more quickly. For the selected collector group, this field overrides the value in the **Behavior Settings** page (System > Settings > Behavior). Your choices are:
 - *Use Systemwide Default.* The collector group will use the global settings for concurrent SNMP collection that has been configured on the **Behavior Settings** page (System > Settings > Behavior).
 - *Enabled.* Concurrent SNMP collection is enabled on this collector group regardless of the global setting on the **Behavior Settings** page.
 - *Disabled.* Concurrent SNMP collection is disabled on this collector group regardless of the global setting on the **Behavior Settings** page.

NOTE: If the "Data Collection: SNMP Collector" process is disabled on the **Process Manager** page (System > Settings > Admin Processes), this concurrent SNMP collection option is disabled.

NOTE: Concurrent SNMP Collection is not available in military unique deployments (MUD) or STIG-compliant deployments.

- **Concurrent PowerShell Collection.** Specifies whether you want to enable concurrent PowerShell collection for this collector group. Concurrent PowerShell collection allows multiple collection tasks to run at the same time with lower system resource requirements. This means that Data Collectors can collect more data using fewer system resources. Concurrent PowerShell collection also prevents missed polls and data gaps because collection will execute more quickly. The PowerShell Collector is an independent service running as a container on a Data Collector. For the selected collector group, this field overrides the value in the **Behavior Settings** page (System > Settings > Behavior). Your choices are:
 - *Use Systemwide Default.* The collector group will use the global settings for concurrent PowerShell collection that has been configured on the **Behavior Settings** page (System > Settings > Behavior).
 - *Enabled.* Concurrent PowerShell collection is enabled on this collector group regardless of the global setting on the **Behavior Settings** page.
 - *Disabled.* Concurrent PowerShell collection is disabled on this collector group regardless of the global setting on the **Behavior Settings** page.

NOTE: If the "Data Collection: PowerShell Collector" process is disabled on the **Process Manager** page (System > Settings > Admin Processes), this concurrent PowerShell collection option is disabled.

NOTE: Concurrent PowerShell Collection is not available in military unique deployments (MUD) or STIG-compliant deployments prior to Skylar One 12.5.1.

- **Concurrent Network Interface Collection.** Specifies whether you want to enable or disable concurrent network interface collection for this collector group. Concurrent network interface collection uses asynchronous SNMP collection for all network interfaces. This provides better scalability for large networks by allowing multiple collection tasks to run at the same time with a reduced load on Data Collectors. For the selected collector group, this field overrides the value in the **Behavior Settings** page (System > Settings > Behavior). Your choices are:
 - *Use Systemwide Default.* The collector group will use the global settings for concurrent network interface collection that has been configured on the **Behavior Settings** page (System > Settings > Behavior).
 - *Enabled.* Concurrent network interface collection is enabled on this collector group regardless of the global setting on the **Behavior Settings** page.
 - *Disabled.* Concurrent network interface collection is disabled on this collector group regardless of the global setting on the **Behavior Settings** page.

NOTE: If the "Data Collection: SNMP Collector" process is disabled on the **Process Manager** page (System > Settings > Admin Processes), this concurrent network interface collection option is disabled.

NOTE: Concurrent network interface collection is not available in military unique deployments (MUD) or STIG-compliant deployments.

- **Collector Failover.** This option is available only if you have at least two Data Collectors in the collector group. Specifies whether you want to maximize the number of devices to be managed or whether you want to maximize reliability. Your choices are:
 - *Off (Maximize Manageable Devices).* The collector group will be load-balanced only. At any given time, the Data Collector with the lightest load handles the next discovered device. If a Data Collector fails, no data will be collected from the devices aligned with the failed Data Collector until the failure is fixed.
 - *On (Maximize Reliability).* The collector group will be load-balanced and configured as a high-availability system that minimizes downtime. If one or more Data Collectors should fail, the tasks from the failed Data Collector will be distributed among the other Data Collectors in the collector group. ScienceLogic recommends that you use this setting.

- **Collectors Available for Failover.** This option is available only if you selected *On (Maximize Reliability)* in the **Collector Failover** field. Specifies the minimum number of Data Collectors that must be available (i.e., with a status of "Available [0]") before a Data Collector failover may occur.
 - For collector groups with only two Data Collectors, this field will contain the value "1 collector".
 - For collector groups with more than two Data Collectors, the field will contain values from a minimum of one half of the total number of Data Collectors up to a maximum of one less than the total number of Data Collectors. For example, for a collector group with eight Data Collectors, the possible values in this field would be 4, 5, 6, and 7.
 - Skylar One will never automatically increase the maximum number of Data Collectors that can fail in a collector group. For example, suppose you have a collector group with three Data Collectors. Suppose **Collectors Available For Failover** field is set to "2". If you add a fourth Data Collector to the collector group, Skylar One will automatically set the **Collectors Available For Failover** field to "3" to maintain the maximum number of Data Collectors that can fail as "one". However, you can override this automatic setting by manually changing the value in the **Collectors Available For Failover** field.

NOTE: If you set this to half of your available Data Collectors and a 50% Data Collector outage occurs and the remaining Data Collectors are down by one, no rebalance will occur. If you specify one-third of the total number of Data Collectors, then a rebalance will be attempted until your overall capacity falls below one-third of your Data Collectors, thereby maximizing your resiliency but minimizing the opportunity for your system to enter an unproductive rebalancing loop.

CAUTION: If the number of available Data Collectors is less than the value in the **Collectors Available For Failover** field, Skylar One will not fail over within the collector group. **Skylar One will not collect any data from the devices aligned with the failed Data Collector(s) until the failure is fixed on enough Data Collector(s) to equal the value in the Collectors Available For Failover field.** Skylar One will generate a critical event.

- **Fallback Mode.** This option is available only if you selected *On (Maximize Reliability)* in the **Collector Failover** field. Specifies how you want collection to behave when the outage is fixed. Your choices are:
 - *Automatic.* After the failed Data Collector is restored, Skylar One will automatically redistribute data-collection tasks among the collector group, including the previously failed Data Collector. **ScienceLogic recommends that you use this setting.**
 - *Manual.* After the failed Data Collector is restored, you will manually prompt Data Collector to redistribute data-collection tasks.

- **Failover Delay (minutes)**. This option is available only if you selected *On (Maximize Reliability)* in the **Collector Failover** field. Specifies the number of minutes Skylar One should wait after the outage of a Data Collector before redistributing the data-collection tasks among the other Data Collectors in the group. During this time, data will not be collected from the devices aligned with the failed Data Collector(s). The default minimum value for this field is 5 minutes. **ScienceLogic recommends that you set this field to 15 minutes.**
- **Fallback Delay (minutes)**. This option is available only if you selected *On (Maximize Reliability)* in the **Collector Failover** field and *Automatic* in the **Fallback Mode** field. Specifies the number of minutes Skylar One should wait after the failed Data Collector is restored before redistributing data-collection tasks among the collector group, including the previously failed Data Collector. The default minimum value for this field is 5 minutes. **ScienceLogic recommends that you set this field to 15 minutes.**

4. Click **[Save]**.

Defining a Collector Group in the Classic Skylar One User Interface

To define a new collector group in the classic Skylar One user interface:

1. Go to the **Collector Group Management** page (System > Settings > Collector Groups).
2. In the **Collector Group Management** page, click the **[Reset]** button to clear the values from the fields in the top pane.
3. Go to the top pane and enter values in the following fields:
 - **Collector Group Name**. Name of the collector group.
 - **Collector Failover**. Specifies whether you want to maximize the number of devices to be managed or whether you want to maximize reliability. Your choices are:
 - *Off (Maximize Manageable Devices)*. The collector group will be load-balanced only. At any given time, the Data Collector with the lightest load handles the next discovered device. If a Data Collector fails, no data will be collected from the devices aligned with the failed Data Collector until the failure is fixed.
 - *On (Maximize Reliability)*. The collector group will be load-balanced and configured as a high-availability system that minimizes downtime. If one or more Data Collectors should fail, the tasks from the failed Data Collector will be distributed among the other Data Collectors in the collector group. **ScienceLogic recommends that you use this setting.**
 - **Generate Alert on Collector Outage**. Specifies whether or not the platform should generate an event if a Data Collector has an outage. **ScienceLogic recommends that you select Yes for this setting.**
 - **Enable Concurrent SNMP Collection**. Specifies whether you want to enable Concurrent SNMP Collection. Concurrent SNMP Collection uses asynchronous I/O for massive concurrency with lower system resource requirements. This means that Data Collectors can collect more data using fewer system resources. Concurrent SNMP Collection also prevents missed polls and data gaps because collection will execute more quickly. For the selected collector group, this field overrides the value in the **Behavior Settings** page (System > Settings > Behavior). Your choices are:

- *Use systemwide default.* The collector group will use the global settings for Concurrent SNMP Collection configured in the **Behavior Settings** page (System > Settings > Behavior).
- *No.* Concurrent SNMP Collection is disabled on this collector group regardless of the global setting on the **Behavior Settings** page.
- *Yes.* Concurrent SNMP Collection is enabled on this collector group regardless of the global setting on the **Behavior Settings** page.

NOTE: If the "Data Collection: SNMP Collector" process is disabled on the **Process Manager** page (System > Settings > Admin Processes), this concurrent SNMP collection option is disabled.

NOTE: Concurrent SNMP Collection is not available in military unique deployments (MUD) or STIG-compliant deployments.

- **Enable Concurrent PowerShell Collection.** Specifies whether you want to enable Concurrent PowerShell Collection for this collector group. If you make no selection, the default behavior is to "Use systemwide default", which uses the global setting specified on the **Behavior Settings** page (System > Settings > Behavior). Your choices are:

- *Use systemwide default.* The collector group will use the global setting for Concurrent PowerShell Collection as it is configured on the **Behavior Settings** page (System > Settings > Behavior).
- *No.* Concurrent PowerShell Collection is disabled on this collector group regardless of the global setting on the **Behavior Settings** page.
- *Yes.* Concurrent PowerShell Collection is enabled on this collector group regardless of the global setting on the **Behavior Settings** page.

NOTE: If the "Data Collection: PowerShell Collector" process is disabled on the **Process Manager** page (System > Settings > Admin Processes), this concurrent PowerShell collection option is disabled.

NOTE: Concurrent PowerShell Collection is not available in military unique deployments (MUD) or STIG-compliant deployments prior to Skylar One 12.5.1.

- **Enable Concurrent Network Interface Collection.** Specifies whether you want to enable or disable Concurrent Network Interface Collection for this collector group. If you make no selection, the default behavior is to "Use systemwide default", which uses the global setting specified on the **Behavior Settings** page (System > Settings > Behavior). Your choices are:

- *Use systemwide default.* The collector group will use the global setting for Concurrent Network Interface Collection as it is configured in the **Behavior Settings** page (System > Settings > Behavior).
- *No.* Concurrent Network Interface Collection is disabled on this collector group regardless of the global setting on the **Behavior Settings** page.
- *Yes.* Concurrent Network Interface Collection is enabled on this collector group regardless of the global setting on the **Behavior Settings** page.

NOTE: If the "Data Collection: SNMP Collector" process is disabled on the **Process Manager** page (System > Settings > Admin Processes), this concurrent network interface collection option is disabled.

NOTE: Concurrent network interface collection is not available in military unique deployments (MUD) or STIG-compliant deployments.

- **Collector Selection.** Displays a list of available Data Collectors.
 - To assign an available Data Collector server to the collector group, simply highlight it. You can assign one or more Data Collectors to a collector group.
 - To assign multiple Data Collectors to the collector group, hold down the **<Ctrl>** key and click multiple Data Collectors.
- **Message Collector.** Displays a list of available Message Collectors.
 - To assign an available Message Collector to the collector group, simply highlight it. You can assign one or more Message Collectors to a collector group.
 - To assign multiple Message Collectors to the collector group, hold down the **<Ctrl>** key and click multiple Message Collectors.

NOTE: A single Message Collector can be used by multiple collector groups. When you align a single Message Collector with multiple collector groups, the single Message Collector might then be aligned with two devices (each in a separate collector group) that use the same primary IP address or the same secondary IP address. If this happens, Skylar One will generate an event.

- **Collectors Available for Failover.** Applies only if you selected "On (Maximize Reliability)" in the **Collector Failover** field. Specifies the minimum number of Data Collectors that must be available (i.e. with a status of "Available [0]") before a Data Collector failover may occur.
 - For collector groups with only two Data Collectors, this field will contain the value "1 collector".

- For collector groups with more than two Data Collectors, the field will contain values from a minimum of one half of the total number of Data Collectors up to a maximum of one less than the total number of Data Collectors.
- For example, for a collector group with eight Data Collectors, the possible values in this field would be 4, 5, 6, and 7.
- Skylar One will never automatically increase the maximum number of Data Collectors that can fail in a collector group. For example, suppose you have a collector group with three Data Collectors. Suppose **Collectors Available For Failover** field is set to "2". If you add a fourth Data Collector to the collector group, Skylar One will automatically set the **Collectors Available For Failover** field to "3" to maintain the maximum number of Data Collectors that can fail as "one". However, you can override this automatic setting by manually changing the value in the **Collectors Available For Failover** field.

NOTE: If you set this to half of your available Data Collectors and a 50% Data Collector outage occurs and the remaining Data Collectors are down by one, no rebalance will occur. If you specify one-third of the total number of Data Collectors, then a rebalance will be attempted until your overall capacity falls below one-third of your Data Collectors, thereby maximizing your resiliency but minimizing the opportunity for your system to enter an unproductive rebalancing loop.

CAUTION: If the number of available Data Collectors is less than the value in the **Collectors Available For Failover** field, Skylar One will not failover within the collector group. *Skylar One will not collect any data from the devices aligned with the failed Data Collector(s) until the failure is fixed on enough Data Collector(s) to equal the value in the Collectors Available For Failover field.* Skylar One will generate a critical event.

- **Fallback Mode.** Applies only if you selected *On (Maximize Reliability)* in the **Collector Failover** field. Specifies how you want collection to behave when the outage is fixed. You can specify one of the following:
 - *Automatic.* After the failed Data Collector is restored, Skylar One will automatically redistribute data-collection tasks among the collector group, including the previously failed Data Collector. ScienceLogic recommends that you use this setting.
 - *Manual.* After the failed Data Collector is restored, you will manually prompt Data Collector to redistribute data-collection tasks by clicking the lightning bolt icon (⚡) for the collector group.
- **Failover Delay (minutes).** Applies only if you selected *On (Maximize Reliability)* in the **Collector Failover** field. Specifies the number of minutes Skylar One should wait after the outage of a Data Collector before redistributing the data-collection tasks among the other Data Collectors in the group. During this time, data will not be collected from the devices aligned with the failed Data Collector(s). The default minimum value for this field is 5 minutes. ScienceLogic recommends that you set this field to 15 minutes.

- ***Fallback Delay (minutes)***. Applies only if you selected *On (Maximize Reliability)* in the **Collector Failover** field and *Automatic* in the **Fallback Mode** field. Specifies the number of minutes Skylar One should wait after the failed Data Collector is restored before redistributing data-collection tasks among the collector group, including the previously failed Data Collector. The default minimum value for this field is 5 minutes. ScienceLogic recommends that you set this field to 15 minutes.

4. Click the **[Save]** button to save the new collector group.
5. To assign devices to the collector group, see the section on [Aligning Single Devices with a Collector Group in the Classic Skylar One User Interface](#) and the section on [Aligning a Device Group with a Collector Group](#).

Editing a Collector Group

To edit a collector group:

1. Go to the **Collector Groups** page (Manage > Collector Groups).
2. Click the **Actions** icon (⋮) of the collector group you want to edit and then select *Edit*. The **Edit Collector Group** modal appears.
3. The fields in the **Edit Collector Group** modal are populated with values from the selected collector group. You can edit one or more of the fields. For a description of each field, see the section on [Defining a Collector Group](#).
4. Click **[Save]** to save any changes to the collector group.

Editing a Collector Group in the Classic Skylar One User Interface

From the **Collector Group Management** page, you can edit an existing collector group. You can add or remove Data Collectors and change the configuration from load-balanced to failover (high availability).

To edit a collector group in the classic Skylar One user interface:

1. Go to the **Collector Group Management** page (System > Settings > Collector Groups).
2. In the **Collector Group Management** page, go to the **Collector Group Registry** pane at the bottom of the page.
3. Find the collector group you want to edit. Click its wrench icon (🔧).
4. The fields in the top pane are populated with values from the selected collector group. You can edit one or more of the fields. For a description of each field, see the section on [Defining a Collector Group in the Classic Skylar One User Interface](#).
5. Click the **[Save]** button to save any changes to the collector group.

Collector Groups and Load Balancing

To perform initial discovery, Skylar One uses a single, selected Data Collector from the collector group. This allows you to troubleshoot discovery if there are any problems.

After each discovered device is modeled (that is, after Skylar One assigns a device ID and creates the device in the database), Skylar One distributes devices among the Data Collectors in the collector group. The newest device is assigned to the Data Collector currently managing the lightest load.

This process is known as ***Collector load balancing***, and it ensures that the work performed by the Dynamic Applications aligned to the devices is evenly distributed across the Data Collectors in the collector group.

Skylar One performs Collector load balancing in the following circumstances:

- A new Data Collector is added to a collector group
- New devices are discovered
- Failover or failback occurs within a collector group (if failover is enabled)
- A user clicks the lightning bolt icon (⚡) for a collector group to manually force redistribution
- Devices in DCM or DCM-R trees will be loaded on the Data Collector currently assigned to the DCM or DCM-R tree rather than being distributed across the collector group. DCM or DCM-R trees will be rebalanced as an aggregate when rebalancing occurs to an available Data Collector with sufficient capacity to sustain the load.

NOTE: Whenever a device is load-balanced from one Data Collector to another, whether due to failover or regular load balancing, the device state information is not transferred to the new Data Collector.

NOTE: The lightning bolt icon (⚡) appears only for collector groups that contain more than one Data Collector. For collector groups with only one Data Collector, this icon is grayed out. This icon does not appear for All-In-One Appliances.

When all of the devices in a collector group are redistributed, Skylar One will assign the devices to Data Collectors so that all Data Collectors in the collector group will spend approximately the same amount of time collecting data from devices.

Collector load balancing uses two metrics:

- ***Device Rating***. A device's rating is the total elapsed time consumed by either 1) all of the Dynamic Applications aligned to the device, or 2) collecting metrics from the device's interfaces, whichever is greater. A Collector's load is the sum of the ratings of the devices assigned to the Collector. The balancer tries to evenly divide the work performed by Collectors by assigning devices to Collectors using the device ratings and Collector loads.
- ***Collector Load***. The sum of the device ratings for all of the devices assigned to a collector.

Skylar One performs the following steps during Collector load balancing:

1. Searches for all devices that are not yet assigned to a collector group.
2. Determines the load on each Data Collector by calculating the device rating for each device on a Data Collector and then summing the device ratings.
3. Determines the number of new devices (less than one day old) and old devices on each Data Collector.
4. On each Data Collector, calculates the average device rating for old devices (sum of the device ratings for all old devices divided by the number of old devices). If there are no old devices, sets the average device rating to "1" (one).
5. On each Data Collector, assigns the average device rating to all new devices (devices less than one day old).
6. Assigns each unassigned device (either devices that are not yet assigned or devices on a failed Data Collector) to the Data Collector with the lightest load. Add each newly assigned device rating to the total load for the Data Collector.

Tuning Collector Groups in the silo.conf File

With the addition of execution environments to Skylar One, Skylar One sorts data collections into a two-process-pool model.

Skylar One sorts collection requests into groups by execution environment. These groups of collection requests are called "chunks". Each chunk contains a maximum of 200 collection requests, all of which use the same execution environment. Skylar One sends each chunk to a chunk worker.

The chunk worker determines the appropriate execution environment for the chunk, deploys the execution environment, and starts a pool of request workers in the execution environment.

The request workers then process the actual collection requests contained in the chunks and perform the actual data collection.

NOTE: For more information about ScienceLogic Libraries and execution environments, see the manual *ScienceLogic Libraries and Execution Environments*.

The following settings are available in the master.system_settings_core database table for tuning globally in a stack, or [in the Silo.Config file](#) for tuning locally on a single Data Collector:

Parameter Name	Description	Runtime Default
dynamic_collect_num_chunk_workers	The number of chunk workers. In general, this value controls the number of execution environments that can be processed in parallel.	2

Parameter Name	Description	Runtime Default
dynamic_collect_num_request_workers	The maximum number of request workers in each worker pool. In general, this value controls the number of collections using an execution environment that can be processed in parallel.	"2" or the number of cores on the Data Collector, whichever is greater
dynamic_collect_request_chunk_size	The maximum number of collection requests in a chunk. This value controls how many collections are processed by each pool of request workers.	200

NOTE: The database values for these parameters are "Null" by default, which specifies that Skylar One should use the runtime defaults.

The maximum total number of worker processes used during a scheduled collection is generally `dynamic_collect_num_chunk_workers` x `dynamic_collect_num_request_workers`.

There might be circumstances where adjustment is necessary to improve the performance of collection.

Example 1: Additional Environments Required

You might need to adjust the values of the collection processes when scheduled collection requires more than two environments.

Because the default number of chunk workers is "2", Skylar One can simultaneously process chunks of collection requests for a maximum of two virtual environments. If the collection requests require more than two virtual environments, you can increase parallelism by setting `dynamic_collect_num_chunk_workers` to match the number of environments.

If you increase `dynamic_collect_num_chunk_workers`, you might want to decrease `dynamic_collect_num_request_workers` to avoid performance problems caused by too many request workers.

If you cannot increase `dynamic_collect_num_chunk_workers` because doing so would result in too many request workers, you can decrease `dynamic_collect_request_chunk_size` to give collection requests for each environment a "fairer share" of the chunk workers.

NOTE: Smaller chunk sizes require more resources to establish the virtual environments and establish more pools of request workers to process the chunks. Conversely, if you want to use fewer resources for establishing virtual environments and creating pools of request worker pools, and you want to use more resources for collection itself, increasing `dynamic_collect_request_chunk_size` allows more collection requests to be processed by each pool of request workers.

Example 2: Input/Output Bound Collections

You might need to adjust the values of the collection processes when collection requests are input/output (I/O) bound with relatively large latencies.

In this scenario, you can increase `dynamic_collect_num_request_workers` to improve parallelism. If you increase `dynamic_collect_num_request_workers`, you might want to decrease `dynamic_collect_num_chunk_workers` to avoid performance problems caused by too many request workers.

CAUTION: Increasing the number of collection processes will increase CPU and memory utilization on the Data Collector, so be careful when increasing the values dramatically.

Before adjusting `dynamic_collect_num_request_workers`, you need to know the following information:

- The number of CPU cores in the Data Collector
- The current CPU utilization of Data Collector
- The current memory utilization of Data Collector

Start by setting `dynamic_collect_num_request_workers` to equal the number of CPUs plus 50%. For example: with 8 cores, start by setting `dynamic_collect_num_request_workers` to 12. If that is insufficient, you can then try 16, 20, 24, and so forth.

If data collections are terminating early, it means that collections are not completed within the 15-minute limit. If this is the case, wait 30 minutes to see results after adjusting the collection values.

Load Balancing and Device State

It is important to note that, whenever a device is load-balanced from one Data Collector to another, whether due to failover or regular load balancing, the device state information is not transferred to the new Data Collector.

In the time immediately after load balancing, if an interface or another aspect of the device changes state at the same time as the device is load-balanced, then the new Data Collector might not register an event that triggers based on a device state change.

Also, if a Dynamic Application that depends on cached data being present attempts to collect data, but the cached data is not yet present on the new Data Collector, the Dynamic Application might initially fail to collect data.

Collector Affinity

Collector Affinity specifies the Data Collectors that are allowed to run collection for Dynamic Applications aligned to component devices. You can define Collector Affinity for each Dynamic Application. Choices are:

- **Default.** If the Dynamic Application is auto-aligned to a component device during discovery, then the Data Collector assigned to the root device will collect data for this Dynamic Application as well. For devices that are not component devices, the Data Collector assigned to the device running the Dynamic Application will collect data for the Dynamic Application.

NOTE: *Default* might appear as the **Collector Affinity** setting for older Dynamic Applications. However, it is no longer an option when creating or editing a Dynamic Application.

- **Root Device Collector.** The Data Collector assigned to the root device will collect data for the Dynamic Application. This guarantees that Dynamic Applications for an entire DCM tree will be collected by a single Data Collector. You might select this option if:
 - The Dynamic Application has a cache dependency with one or more other Dynamic Applications.
 - You are unable to collect data for devices and Dynamic Applications within the same Device Component Map on multiple Data Collectors in a collector group.
 - The Dynamic Application will consume cache produced by a Dynamic Application aligned to a non-root device (for instance, a cluster device).
 - The Dynamic Application includes snippet code with a **root_device** tuple.
- **Assigned Collector.** The Dynamic Application will use the Data Collector assigned to the device running the Dynamic Application. This allows Dynamic Applications that are auto-aligned to component devices during discovery to run on multiple Data Collectors. This is the default setting. You might select this option if:
 - The Dynamic Application has no cache dependencies with any other Dynamic Applications.
 - You want the Dynamic Application to be able to make parallel data requests across multiple Data Collectors in a collector group.
 - The Dynamic Application can be aligned using mechanisms other than auto-alignment during discovery (for instance, manual alignment or alignment via Device Class Templates or Run Book Actions).

Failover for Collector Groups for Component Devices

If you specified **Default** or **Root Device Collector** for Dynamic Applications, and the single Data Collector in the collector group for component devices fails, users must create a new collector group with a single Data Collector and manually move the devices from the failed collector group to the new collector group. For details on manually moving devices to a new collector group, see the section on [Changing the Collector Group for One or More Devices](#).

Collector Groups for Merged Devices

You can merge a physical device and a component device. There are two ways to do this:

- From the Actions menu in the **Device Properties** page (Devices > Classic Devices > wrench icon) for either the physical device or the component device.
- From the **Actions** menu in the **Device Manager** page (Devices > Classic Devices, or Registry > Devices > Device Manager in the classic user interface), select *Merge Devices* to merge devices in bulk.

You can unmerge a component device from a physical device. You can do this in two ways:

- From the **Actions** menu in the **Device Properties** page (Devices > Classic Devices > wrench icon) for either the physical device or the component device, select *Unmerge Devices* to unmerge devices.
- From the **Actions** menu in the **Device Manager** page Devices > Classic Devices, or Registry > Devices > Device Manager in the classic user interface), select *Unmerge Devices* to unmerge devices in bulk.

When you merge a physical device and a component device, the device record for the component device is no longer displayed in the user interface; the device record for the physical device is displayed in user interface pages that previously displayed the component device. For example, the physical device is displayed instead of the component device in the **Device Components** page (Devices > Device Components) and the **Component Map** page (Device Component Map). All existing and future data for both devices will be associated with the physical device.

If you manually merge a component device with a physical device, Skylar One allows data for the merged component device and data from the physical device to be collected on different Data Collectors. Data that was aligned with the component device can be collected by the collector group for its root device. Data aligned with the physical device can be collected by a different collector group.

NOTE: You can merge a component device with only one physical device.

Creating a Collector Group for Data Storage Only

You can create a **virtual collector group** (vCUG) that serves as a storage area for all historical data from decommissioned devices.

The virtual collector group will store all existing historical data from all aligned devices, but will not perform collection on those devices. The virtual collector group will not contain any Data Collectors or any Message Collectors. **Skylar One will stop collecting data from devices aligned with a virtual collector group.**

To define a virtual collector group in the default Skylar One user interface:

1. Go to the **Collector Groups** page (Manage > Collector Groups).
2. Click the **[Add Collector Group]** button. The **Add Collector Group** modal appears.
3. On the **Add Collector Group** modal, complete the following fields:

- **Collector Group Name.** Type a name for the collector group.
- **Virtual Collector Group (vCUG).** Toggle this option on to make the collector group a virtual collector group. Virtual collector groups do not contain any Data Collectors or Message Collectors and **Skylar One does not collect any data from devices aligned with a virtual collector group.** Instead of collecting data, virtual collector groups serve only as storage areas for historical data from decommissioned devices.

4. Leave all other fields set to the default values. Do not include any Data Collectors or Message Collectors in the collector group.
5. Click **[Save].**
6. To assign devices to the virtual collector group, see the section on [aligning single devices with a collector group](#) and the section on [aligning a device group with a collector group](#).

To define a virtual collector group in the classic Skylar One user interface:

1. Go to the **Collector Group Management** page (System > Settings > Collector Groups).
2. In the **Collector Group Management** page, click the **[Reset]** button to clear values from the fields in the top pane.
3. Go to the top pane and enter a name for the virtual collector group in the **Collector Group Name** field.
4. Leave all other fields set to the default values. Do not include any Data Collectors or Message Collectors in the collector group.
5. Click the **[Save]** button to save the new collector group.
6. To assign devices to the virtual collector group, see the section on [aligning single devices with a collector group](#) and the section on [aligning a device group with a collector group](#).

Deleting a Collector Group

To delete a collector group:

1. Go to the **Collector Groups** page (Manage > Collector Groups).
2. Click the **Actions** icon () of the collector group you want to delete and then select *Delete*.

Deleting a Collector Group in the Classic Skylar One User Interface

From the **Collector Group Management** page, you can delete a Collector Group. When you delete a collector group, those Data Collectors become available for use in other collector groups.

NOTE: Before you can delete a collector group, you must move all aligned devices to another collector group. For details on how to do this, see the section [Changing the Collector Group for One or More Devices](#).

To delete a collector group in the classic Skylar One user interface:

1. Go to the **Collector Group Management** page (System > Settings > Collector Groups).
2. In the **Collector Group Management** page, go to the **Collector Group Registry** pane at the bottom of the page.
3. Find the collector group you want to delete. Click its delete icon (Delete icon).

Assigning a Collector Group for a Single Device

After you have defined a collector group, you can align devices with that collector group.

To assign a collector group to a device:

1. From the **Devices** page, click the name of the device that you want to assign to a collector group. The **Device Investigator** page opens for that device.
2. On the **Device Investigator** page, click the **[Settings]** tab.
3. Click the **[Edit]** button. This enables you to change your device settings.
4. In the **Collection Poller** field, select the name of the collector group that you want to use for collection on the device.
5. Click **[Save]**.

Assigning a Collector Group for a Single Device in the Classic Skylar One User Interface

After you have defined a collector group, you can align devices with that collector group.

To assign a collector group to a device in the classic Skylar One user interface:

1. Go to the **Device Manager** page (Devices > Classic Devices, or Registry > Devices > Device Manager in the classic user interface).
2. In the **Device Manager** page, find the device you want to edit. Click its wrench icon (Wrench icon). The **Device Properties** page appears.
3. On the **Device Properties** page, select a collector group from the **Collection** fields.
4. Click the **[Save]** button to save the change to the device.

Aligning the Collector Group in a Device Template

You can specify a collector group in a device template. Then, when you apply the device template to a device, either through discovery or when you apply the device template to a device group or selection of devices, the specified collector group is automatically associated with the device(s). Optionally, you can later edit the collector group for each device.

For more details on device templates and device groups, see the manual **Device Groups and Device Templates**.

Changing the Collector Group for One or More Devices

You can change the collector group for multiple devices simultaneously. This is helpful if you want to reorganize devices or collector groups. If you want to delete a collector group, you first must first move each aligned device to another collector group. In this situation, you might want to change the collector group for multiple devices simultaneously.

To change the collector group for multiple device simultaneously:

1. Go to the **Device Manager** page (Devices > Classic Devices, or Registry > Devices > Device Manager in the classic user interface).
2. In the **Device Manager** page, click on the heading for the **Collection Group** column to sort the list of devices by collector group.
3. Select the checkbox for each device that you want to move to a different collector group.
4. In the **Select Action** field (in the lower right), go to **Change Collector Group** and select a collector group.
5. Click the **[Go]** button. The selected devices will now be aligned with the selected collector group.

Managing the Host Files for a Collector Group

The **Host File Entry Manager** page allows you to edit and manage host files for all of your Data Collectors from a single page in the Skylar One system. When you create or edit an entry in the **Host File Entry Manager** page, Skylar One automatically sends an update to every Data Collector in the specified collector group.

The **Host File Entry Manager** page is helpful when:

- The Skylar One system does not reside in the end-customer's domain
- The Skylar One system does not have line-of-sight to an end-customer's DNS service
- A customer's DNS service cannot resolve a host name for a device that the Skylar One system monitors

For details, see the section on [Managing Host Files](#).

Processes for Collector Groups

For troubleshooting and debugging purposes, you might find it helpful to understand the ScienceLogic processes that affect a collector group.

NOTE: You can view the list of all processes and details for each process in the **Process Manager** page (System > Settings > Admin Processes).

- The **Enterprise Database: Collector Task Manager process (em7_ctaskman)** process distributes devices between Data Collectors in a collector group, to load-balance the collection tasks. The process runs every 60 seconds and also checks the license on each Data Collector. The "Enterprise Database: Collector Task Manager" process (*em7_ctaskman.py*) redistributes devices between collectors when:
 - A collector group is created.
 - A new Data Collector is added to a collector group.
 - Failover or failback occurs within a collector group.
 - A user clicks on the lightning bolt icon (⚡) for a collector group, to manually force redistribution.
- The **Enterprise Database: Collector Data Pull processes** retrieves information from each Data Collector in a collector group. The process pulls data from the *in_storage* tables on each Data Collector. The retrieved information is stored in the Database Server.
 - *Enterprise Database: Collector Data Pull, High F (em7_hfpulld)*. Retrieves data from each Data Collector every 15 seconds (configurable).
 - *Enterprise Database: Collector Data Pull, Low F (em7_lfpulld)*. Retrieves data from each Data Collector every five minutes.
 - *Enterprise Database: Collector Data Pull, Medium (em7_mfpulld)*. Retrieves data from each Data Collector every 60 seconds.
- The **Enterprise Database: Collector Config Push process (config_push.py)** updates each Data Collector with information on system configuration, configuration of Dynamic Applications, and any new or changed policies. This process runs once every 60 seconds and checks for differences between the configuration tables on the Database Server and the configuration tables on each Data Collector. The list of tables to be synchronized is stored in *master.definitions_collector_config_tables* on the Database Server.
- **Asynchronous Processes** (for example, discovery or programs run from the **Device Toolbox** page). Asynchronous processes need to be run immediately and cannot wait until the "Enterprise Database: Collector Config Push" process (*config_push.py*) runs and tells the Data Collector to run the asynchronous process. Therefore, Skylar One uses a stored procedure and the "EM7 Core: Task Manager" process (*em7*) to trigger asynchronous processes on both the Database Server and Data Collector.
 - If a user requests an asynchronous process, a stored procedure on the Database Server inserts a new row in the table *master_logs.spool_process* on the Database Server.
 - Every three seconds, the "EM7 Core: Task Manager" process (*proc_mgr.py*) checks the table *master_logs.spool_process* on the Database Server for new rows.

- If the asynchronous process needs to be started on a Data Collector, a stored procedure on the Database Server inserts the same row into the table master_logs.spool_process on the Data Collector.
- Every three seconds, the "EM7 Core: Task Manager" process (*em7*) checks the table master_logs.spool_process on the Data Collector for new rows.
- If the "EM7 Core: Task Manager" process (*em7*) on the Data Collector finds a new row, the specified asynchronous process is executed on the Data Collector.
- If a Database Server becomes the passive node in a High Availability configuration, the "EM7 Core: Task Manager" process (*em7*) stops core services.

Enabling and Disabling Concurrent PowerShell for Collector Groups

To improve the process of collecting data via PowerShell, you can enable Concurrent PowerShell Collection. Concurrent PowerShell Collection allows multiple collection tasks to run at the same time with a reduced load on Data Collectors. Concurrent PowerShell Collection also prevents missed polls and data gaps because collection will execute more quickly. As a result, Data Collectors can collect more data using fewer system resources.

When you use the PowerShell Collector for Concurrent PowerShell Collection, the collection process can bypass failed or paused collections, reduce collection time, and reduce the number of early terminations (sigterms) that occur with data collection. The PowerShell Collector is an independent service running as a container on a Data Collector.

You can enable one or more collector groups to use concurrent PowerShell collection, and you can collect metrics for concurrent PowerShell collection.

NOTE: Concurrent PowerShell Collection is for PowerShell Performance and Performance Configuration Dynamic Application types and does not include Snippet Dynamic Applications which happen to run PowerShell commands.

NOTE: Concurrent PowerShell Collection is not available in military unique deployments (MUD) or STIG-compliant deployments prior to Skylar One 12.5.1.

For more details on concurrent PowerShell collection, see the manual *Monitoring Windows Systems with PowerShell*, the chapter on *Concurrent PowerShell*.

Enabling Concurrent PowerShell on All Collector Groups

To enable concurrent PowerShell collection service for all collector groups:

1. Go to the **Database Tool** page (System > Tools > DB Tool).

NOTE: The **Database Tool** page is available only in versions of Skylar One prior to 12.2.1 and displays only for users that have sufficient permissions to access the page.

2. Enter the following in the **SQL Query** field:

```
INSERT INTO master.system_custom_config (`field`, `field_value`)
VALUES ('enable_powershell_service', '1');
```

Disabling Concurrent PowerShell on All Collector Groups

To disable concurrent PowerShell collection service for all collector groups:

1. Go to the **Database Tool** page (System > Tools > DB Tool).

NOTE: The **Database Tool** page is available only in versions of Skylar One prior to 12.2.1 and displays only for users that have sufficient permissions to access the page.

2. Enter the following in the **SQL Query** field:

```
UPDATE master.system_custom_config SET field_value=0 where
field='enable_powershell_service';
```

Enabling Concurrent PowerShell on a Specific Collector Group

To enable concurrent PowerShell collection for a specific collector group:

1. Go to the **Database Tool** page (System > Tools > DB Tool).

NOTE: The **Database Tool** page is available only in versions of Skylar One prior to 12.2.1 and displays only for users that have sufficient permissions to access the page.

2. Enter the following in the **SQL Query** field:

```
INSERT INTO master.system_custom_config (`field`, `field_value`,
`cug_filter`) VALUES ('enable_powershell_service_CUGx', '1',
'collector_group_ID');
```

where:

collector_group_ID is the collector group ID. You can find this value in the **Collector Group Management** page (System > Settings > Collector Groups).

Disabling Concurrent PowerShell on a Specific Collector Group

To disable concurrent PowerShell collection for a specific collector group:

1. Go to the **Database Tool** page (System > Tools > DB Tool).

NOTE: The **Database Tool** page is available only in versions of Skylar One prior to 12.2.1 and displays only for users that have sufficient permissions to access the page.

2. Enter the following in the **SQL Query** field:

```
UPDATE master.system_custom_config SET field_value=0 where  
field='enable_powershell_service_CUGx';
```

where:

collector_group_ID is the collector group ID. You can find this value in the **Collector Group Management** page (System > Settings > Collector Groups).

Enabling and Disabling Concurrent SNMP for Collector Groups

To increase the scale for SNMP collection, you can enable **Concurrent SNMP Collection**.

Concurrent SNMP Collection uses the standalone container called the Skylar One SNMP Collector.

The SNMP Collector is an independent service that runs as a container on a Data Collector. When you enable Concurrent SNMP Collection, each Data Collector will contain four (4) SNMP Collector containers.

NOTE: On each Data Collector, Skylar One will restart each of the SNMP Collector containers periodically to ensure that each container remains healthy. When one SNMP Collector container is restarted, the other three SNMP Collector containers continue to handle the workload.

With Concurrent SNMP Collection, SNMP collection tasks can run in parallel. A single failed task will not prevent other tasks from completing.

Concurrent SNMP Collection provides:

- Improved throughput for SNMP Dynamic Applications
- Reduced use of resources on each Data Collector
- More dependable collection from high-latency Devices

NOTE: Concurrent SNMP Collection is not available in military unique deployments (MUD) or STIG-compliant deployments.

Enabling and Disabling Concurrent SNMP for All Collector Groups

NOTE: This feature is disabled by default.

To enable Concurrent SNMP Collection in Skylar One:

1. Go to the **Behavior Settings** page (System > Settings > Behavior).
2. Check the **Enable Concurrent SNMP Collection** field.
3. Click **[Save]**.

NOTE: If the "Data Collection: SNMP Collector" process is disabled on the **Process Manager** page (System > Settings > Admin Processes), this concurrent SNMP collection option is disabled.

TIP: If you do not want all of your Skylar One Collectors to use Concurrent SNMP Collection, you can specify which Collector Units should use it in [Enabling a Collector Group to Use Concurrent SNMP Collection](#).

Enabling and Disabling Concurrent SNMP for Collector Groups

Depending on the needs of your Skylar One environment, you can enable or prevent a collector group from using concurrent SNMP collection.

To enable Concurrent SNMP Collection with a Skylar One collector group:

1. Go to the **Collector Group Management** Page (System > Settings > Collector Groups).
2. Click the wrench icon (🔧) for the collector group you want to edit. The fields at the top of the page are updated with the data for that collector group.
3. Select an option in the **Enable Concurrent SNMP Collection** drop-down field:
 - **Use system-wide default.** Select this option if you want this collector group to use or not use Concurrent SNMP Collection based on the **Enable Concurrent SNMP Collection** field on the **Behavior Settings** page. This is the default.
 - **Yes.** Select this option to enable Concurrent SNMP Collection for this collector group, even if you did not enable it on the **Behavior Settings** page.
 - **No.** Select this option to prevent this collector group from using Concurrent SNMP Collection, even if you did enable it on the **Behavior Settings** page.

NOTE: If the "Data Collection: SNMP Collector" process is disabled on the **Process Manager** page (System > Settings > Admin Processes), this concurrent SNMP collection option is disabled.

4. Update the remaining fields as needed, and then click **[Save]**.

Enabling Multi-tenancy for Collector Groups

To support multi-tenancy, Skylar One allows you to align collector groups with one, multiple, or all organizations in Skylar One. When you align an organization to a collector group, you control who can view details about that collector group and who can apply the collector group in Skylar One.

By default, newly created collector groups are aligned to all organizations. However, you can update the organization setting for a collector group if you have multi-tenancy enabled for collector groups.

When multi-tenancy is enabled:

- An administrative user can update the organization alignment for all collector groups.
- Non-administrative users can update all collector groups that are aligned to all organizations or the organizations to which the user belongs.

NOTE: If you enable multi-tenancy for collector groups, you might encounter a situation where a device is not aligned to a collector group if they do not belong to the same organization.

CAUTION: If you enable multi-tenancy for collector groups, ScienceLogic strongly recommends that you do not disable it at a later date.

To enable multi-tenancy for collector groups:

1. Go to the **Database Tool** page (System > Tools > DB Tool).

NOTE: The **Database Tool** page is available only in versions of Skylar One prior to 12.2.1 and displays only for users that have sufficient permissions to access the page.

2. Select "master" as the database.
3. Type the following in the **SQL Query** field:

```
UPDATE master.system_settings_core SET enable_cug_orgs=1
```

4. Click **[Go]**.

Aligning Collector Groups to Organizations

To align existing collector groups to organizations:

1. Go to the **Collector Groups** page (Manage > Collector Groups).
2. Click the **Actions** icon () of the collector group you want to edit and then select *Edit*. The **Edit Collector Group** modal appears.
3. On the **Edit Collector Group** modal, complete the following fields:
 - **All current and future organizations.** Toggle this option on to align the collector group to all of your Skylar One organizations, or toggle it off to specify the organizations to which you want to assign the collector group.
 - **Limit access to specific organizations.** If you toggled off the All current and future organizations option, select the organization(s) to which you want to assign the collector group.
4. Click **[Save]** to save any changes to the collector group.

Aligning Collector Groups to Organizations in the Classic Skylar One User Interface

To align existing collector groups to organizations in the classic Skylar One user interface:

1. Go to the **Collector Group Management** page (System > Settings > Collector Groups).
2. In the **Collector Group Registry** pane, click the Organization icon () of the collector group you want to align to an organization. The Align Organizations modal appears.
3. In the Align Organizations modal, complete the following fields:
 - **Collector Group Availability.** Select *All Organizations* to align the collector group to all of your Skylar One organizations, or select *Aligned Organizations Only* to specify the organizations to which you want to assign the collector group.
 - **Aligned Organizations.** If you selected *Aligned Organizations Only* in the **Collector Group Availability** field, select the organization(s) to which you want to assign the collector group.
4. Click **[Save]**.

Chapter

4

Daily Health Tasks

Overview

The tasks in this chapter help you monitor the health of your Skylar One (formerly SL1) system. You should perform these tasks daily (or more frequently, if you require) to gather information about the overall status of your Skylar One system and to maintain operational stability.

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 a Healthy Skylar One System?</i>	112
<i>Skylar One Self-Healing</i>	115
<i>Monitoring System Events</i>	115
<i>Monitoring System Processes</i>	118
<i>Monitoring the Status of Each Appliance</i>	119
<i>Logging in Skylar One Version 11.3.0 and Later</i>	123
<i>Monitoring User Actions and Events on the Audit Logs Page</i>	130
<i>Using auditd to Monitor Sensitive Files</i>	132
<i>Monitoring the Status of Data Collectors</i>	134

What is a Healthy Skylar One System?

The following table presents a broad list of focus areas for Skylar One health that are important to track for all Skylar One systems. Where a specific automated check is available, it is included in the table.

Focus Area	Check	Background	Specific Operation	Result
Patch level	Version has been updated within last 12 months	Software updates are released at least quarterly and include security and stability improvements.	Quarterly manual review of available and planned software updates from ScienceLogic	Plan to keep all Skylar One platforms updated within 12 months of the latest release.
Response time	API response times for standard requests are within five seconds	API response times are highly dependent on the size of the response, however all Skylar One systems should respond to a simple request without delay.	/api/organization/0	Returns a response set for the “system” organization (id 0).
Central storage capacity	At least 20% of local database storage is free and available for new data	The InnoDB database file will auto-expand but never shrink. When data is removed from the database, space is made available for future use.	Support PowerPack - Support: InnoDB Size	Built-in alerts for the Support PowerPack default to 80% used for major and 90% for critical.
CPU Consumption	System CPU utilization and load average	Both Collectors and Databases can become CPU constrained, leading to unhealthy performance characteristics.	Operating system measures of CPU utilization and load average	Load average should be at or below the system’s available core count. CPU utilization of a 5 min collection should not be above 80%, 70% if hyperthreading is enabled.

Focus Area	Check	Background	Specific Operation	Result
Memory Consumption	Avoid swap usage	The MariaDB database will make use of available memory for caching purposes over time, but no Skylar One system should require regular swapping, which can lead to extremely poor performance.	Operating system measure of swap usage	Swap usage < 50%
Performance Data Processing	The central system is keeping up with all collection processing.	It is normal to have some backlog of "MF" data, a busy system may normally have 10,000 rows or more between each processing cycle, but they should be completely processed within each cycle (backlog should not build).	Built-in MF rows-behind compared with MF object processing rate	Backlog time < 1 processing cycle
Event Processing	The central system is keeping up with all event processing.	It is normal to have some backlog of "HF" data. A busy system may normally have 10,000 rows or more between each processing cycle, but these rows should be completely processed	Built-in HF rows-behind compared with HF object processing rate	Backlog time < 1 processing cycle

Focus Area	Check	Background	Specific Operation	Result
		within each cycle (backlog should not build)		
Run Book Automation (RBA)	The central system is keeping up with all RBA processing.	The built-in RBA engine supports parallel execution and queuing of operations. This can be critical for time-sensitive notification and integration with external systems.	Built-in alerting in the RBA scheduler will notify if the system is falling behind.	No critical events starting with the following phrase: "The automation engine is still processing..."
Performance Data Collection	Collection of data is completing as scheduled.	Collection that is unbalanced or overloaded, or target devices that are misconfigured or unresponsive can result in collection not completing successfully.	Skylar One Operational Insights PowerPack	Check for occurrences of "sigterming" collection. The Operational Insights PowerPack makes this easy to navigate using a dashboard.
Asynchronous Message Processing	Message collection is keeping up with asynchronous syslog and SNMP trap messages.	Data Collectors, Message Collectors, and All-in-One appliances receive	Built-in alert for suppressing of messages from "spamming" devices	By default Skylar One will suppress messages from devices generating at a rate of > 25/sec/device with a built-in alert.
System Maintenance	Daily maintenance tasks are completing normally	The primary daily maintenance task (scheduled nightly outside of core business hours) is to prune old data from the Skylar One database,	Regular check of the system log	Daily maintenance tasks not being terminated due to an incomplete status.

Focus Area	Check	Background	Specific Operation	Result
		which is an essential activity for long term health.		
System Backup	Backups completing per schedule.	Skylar One supports both configuration-only and full backups. Both should be used since they support different recovery models	Regular check of the system log	The system log will show reports of backup completion and duration.

Skylar One Self-Healing

For the Skylar One classic user interface, Skylar One provides three self-healing jobs that are performed on the Database Server or All-In-One Appliance. These jobs:

- automatically set "s-em7-core:s-em7-core" as the owner and group for the file "silo.log"
- automatically restart the process em7_patch_manager if it is stuck in deactivating mode
- automatically kill queries that have run for longer than 1 hour and logs each query

Monitoring System Events

To view the entries in the **System Logs**:

1. Go to the **System Logs** page (System > Monitor > System Logs).

2. On the **System Logs** page, pay special attention to any log entry tagged as *Critical* or *Major*. These entries might require additional diagnostics.
3. For each log entry, the **System Logs** page displays:
 - **Date**. Date and time the log entry was generated.
 - **Module**. Name of the appliance that generated the log entry.
 - **Severity**. Specifies the severity assigned to the log entry. The choices are:
 - Healthy
 - Notice
 - Minor
 - Major
 - Critical
 - **Message**. Descriptive text included in the log entry.

Searching the System Logs

When viewing the **System Logs**, you might want to sort the entries by date or by log message. This is helpful when you want to view information about a specific occurrence of a system event.

You can also filter the list of logs by appliance, severity, and date.

To search the system logs:

1. Go to the **System Logs** page (System > Monitor > System Logs).

2. The first two filters at the top of the **System Logs** page let you filter the list of logs by appliance, severity, and date. Select a sort type from the first filter (Appliance, Severity, or From Date) and then type or select filter criteria in the next filter to the right.
3. Click the **[Search]** button to see the results of the filters.
4. The next two search fields at the top of the **System Logs** page allow you to search for log entries by message, date, or module.
 - **Search where.** Specifies the parameter you want to search by. You can select from the following:
 - *Search Message.* Searches all log entries for those that match the text that you enter in the regular expression field.
 - *Search Module ID.* Searches all log entries for those that have the same module ID text as that entered in the regular expression field.
 - *Search Date = (Y-m-d H:i:s).* Searches all log entries for those that have a date and time that is equal to the date entered in the regular expression field.
 - *Search Date > (Y-m-d H:i:s).* Searches all log entries for those that have a date and time that is later than the date entered in the regular expression field.
 - *Search Date Like (Y-m-d H:i:s).* Searches all log entries for those that have a date and time that is similar to the date entered in the regular expression field.
 - *Search Date Like != (Y-m-d H:i:s).* Searches all log entries for those that have a date and time that is **not** similar to the date entered in the regular expression field.
 - **regular expression.** In this field you manually enter the text to search for. You can use the following special characters in this field:
 - * Match zero or more characters preceding the asterisk. For example:
"dell*" would match "dell", "dell2650", "dell7250" and "dell1700N".
"**dell*" would match "mydell", "dell", "dell2650", "dell7250" and "dell1700N".
 - % Match zero or more characters preceding the asterisk. This special character behaves in the same way as the asterisk.
5. When you click the **[Search]** button, the **System Logs** page will be refreshed and will display only the log entries that match the search parameters.

Deleting Entries from the System Logs

To save space, you might want to remove some or all log entries from the system log.

There are two ways to delete entries from the **System Logs** page:

1. Go to the **System Logs** page (System > Monitor > System Logs).
2. In the **System Logs** page, click the **[Purge]** button to delete all entries from the System Logs.

Or:

1. Go to the **System Logs** page (System > Monitor > System Logs).
2. In the **System Logs** page, highlight each entry you want to delete. To select multiple entries, right-click while holding down the [**<Ctrl>**] key.
3. Click the **[Delete]** button to delete all the selected entries from the System Logs.

Monitoring System Processes

The **System Processes** page (System > Monitor > Admin System Processes or System > Monitor > System Processes in the classic user interface) allows you to view read-only information about the execution of Skylar One's system processes. System Processes gather, manipulate, and publish the data used in Skylar One. These system processes can be configured and debugged in the **Process Manager** page (System > Settings > Admin Processes).

NOTE: ScienceLogic recommends that you enable the debug option **only** while troubleshooting a problem while working with ScienceLogic Support or while following a troubleshooting guide, and that you then immediately turn off debugging when you have completed troubleshooting. Do not leave the debug option enabled during normal operation of Skylar One. When you turn on debugging, Skylar One will run significantly more slowly.

If you work frequently with ScienceLogic Support, ensure that you verify periodically which processes are running in debug mode, and disable debug mode for any processes that do not require it as soon as is reasonable. This will reduce the noise generated to your log files and will reduce load on the /var/log/em7 partition on the Database Server.

Viewing the List of System Processes

To view the list of system processes for all appliances:

1. Go to the **System Processes** page (System > Monitor > Admin System Processes).
2. The **System Processes** page displays the following for each process:
 - **Appliance.** The appliance where the process ran or is currently running. This field will contain the device name of the appliance.
 - **Process.** Name of the process.
 - **ID.** Unique numeric ID automatically assigned to the process by Skylar One.
 - **Start Time.** Date and time at which the process started running.
 - **End Time.** Date and time at which the process stopped running.
 - **Duration.** Amount of time, in hours, minutes, and seconds, for which the process ran.
 - **Frequency.** Frequency with which Skylar One launches the process. Possible values are:
 - **Asynchronous.** The process is launched in response to a system event or user request. Asynchronous events display a value of "-1" (negative one) in this column.

- *Always*. The process always runs while Skylar One is running. Always running processes display a value of "0" (zero) in this column.
- The process runs at intervals in minutes ranging from *1 Minute* to *1440 Minutes (Daily)*.
- **Percent**. Percent of *Run Length* (defined in the **Process Manager** page) currently in use by the process.
- **Instances**. This field is not currently in use.
- **Max Instances**. Maximum number of instances of the process that have run in parallel.
- **Processed**. Number of records processed by this run of the process.
- **Errors**. Number of errors encountered by this run of the process.

Recommended System Maintenance

ScienceLogic recommends that you take the following actions on a regular basis to reduce outages as much as possible.

Daily:

- Review "Skylar One Operational Insights: Database Performance" classic dashboard
- Review "Skylar One Operational Insights: Collector Performance" classic dashboard
- Review "Skylar One Operational Insights: System Log Summary" classic dashboard
- Review "Skylar One Operational Insights: Backup History" classic dashboard

Weekly:

- Run the System Status Script and review
 - Address every error item in the report
 - Read Knowledge Base articles
 - Open tickets for issues when help from Skylar One Support is needed

Monthly:

- Review capacity items
 - You must understand License Usage and how to project future capacity

Quarterly:

- Audit User Profile access to verify that it meets expected needs
- Audit DNS servers and Timeservers on all collectors

Monitoring the Status of Each Appliance

The **Appliance Manager** page (System > Settings > Appliances) provides for global appliance configuration and management for your entire system or stack. This includes collector group and load distribution, version information, license status, and other items that are important when you upgrade.

During upgrade, table cells will highlight known, pending action items that must be done to successfully complete an upgrade, such as highlighting a Skylar One appliance that is running a different version of Skylar One than the Database Server.

This page is useful for ensuring that every Data Collector is assigned to a Collector Group before you begin an upgrade. In some cases, the Data Collector might be assigned to an empty Collector Group, if the collector is new.

You can also use this page to ensure that Data Collector load is near or below the system requirements for each collector.

From the **Appliance Manager** page, you can also:

- Click the wrench icon (🔧) to edit the properties for each Skylar One appliance.
- Click the toolbox icon (🔧) to access the **Web Configuration Utility** for each Skylar One appliance.
- Click the agent endpoint icon (👤) to access the **Agent Endpoint Configuration** modal for any Skylar One appliance that has a Skylar One Gen-1 agent installed.
- Click the magnifying-glass icon (🔍) to view the output of the system status script for each Skylar One appliance.
- Click the lock icon (🔒) to get a one-time password for each Skylar One appliance.
- Click the lightning bolt icon (⚡) to run the **Enterprise Database: Collector Config Push** process (config_push.py) on any SL1 Collector appliance.
- Click the delete icon (🗑) to delete a Skylar One appliance.

To edit and view information about a Skylar One appliance:

1. Go to the **Appliance Manager** page (System > Settings > Appliances).
2. Locate the Skylar One appliance you want to edit. Click its wrench icon (🔧). The fields in the top pane are populated with values from the selected Skylar One appliance.
3. You can edit one or more of the following fields:
 - **Host Name**. Name of the appliance.
 - **IP Address**. Primary IP address for the appliance.
 - **Model Type/Module Type**. Type of appliance. If an appliance is added with the wrong appliance type, Skylar One generates a critical error to notify the user. The types include:
 - *All In One Server*
 - *Database*
 - *Administration Portal*
 - *Data Collection Unit*

- *Message Collection Unit*

NOTE: The combination appliance with a Database Server and an Administration Portal on a single appliance will appear with **Module Type** of *Database*. The combination appliance with a Message Collection Unit and a Data Collection Unit will appear with **Module Type** of *Data Collection Unit*.

- *PowerFlow (Skylar OnePowerFlow)*
- **Description.** Description of the appliance.
- **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.

4. You can also edit two optional fields for Data Collectors or Message Collectors:
 - **DB User.** User name that can access the MariaDB database on the Data Collector or Message Collector.
 - **DB Password.** Password that allows access the MariaDB database on the Data Collector or Message Collector.

If you are using AWS RDS with your Skylar One System, you must define the **DB User** and **DB Password** for each Data Collector or Message Collector.

NOTE: ScienceLogic recommends that you vary the Data Collector and Message Collector database credentials for enhanced per-appliance security. This greatly enhances the security of your central database by disallowing a successful attack to go unnoticed on your Data Collector and then succeed without failure on the central database.

5. You can view the following information about each appliance that appears on the **Appliance Manager** page:

- **Name.** Name of the appliance.
- **IP Address.** Primary IP address for the appliance.
- **Module Type.** Type of appliance.

- **Collector Group.** For Data Collectors and All-In-One Appliances, specifies the Collector Group associated with the appliance.
- **Description.** Description of the appliance.
- **Build.** Specifies the latest build installed on the appliance.

NOTE: If a Skylar One appliance is running a different version of Skylar One than the Database Server, the corresponding cell in the **Build** column will be highlighted.

- **MariaDB.** Specifies the version of MariaDB running on the All-In-One Appliance, Database Server, Data Collector, or Message Collector.
- **Platform.** The current operating system platform version for the appliance. Potential values include *e17* for appliances running on Oracle Linux 7 (OL7), *e18* for appliances running on OL8, *NULL*, or *Unknown*. The platform value is highlighted when the primary Database Server is not running on the same platform version as the other appliances in the system.
- **Capacity.** For Database Servers, specifies the licensed capacity of the appliance.
- **Allocation.** For Data Collectors, specifies the number of devices aligned with the appliance.
- **ID.** Unique numeric ID, automatically assigned by the platform to each appliance in the **Appliance Manager** page.
- **Validated.** Specifies whether the license is valid.
- **Endpoint.** Skylar One Agent endpoint for the Gen 1 Agent.
- **Needs Reboot?.** Specifies whether the appliance requires reboot to add latest kernel or security updates. This column is updated every 30 minutes. Hover your mouse to determine why the reboot is required and information about kernel version, packages, and last reboot.
- **Task Manager Paused?.** Specifies whether the task manager service (em7) is paused. This value is updated every two minutes.
- **Access.** Indicates whether the appliance is shared with all organizations or is private.

NOTE: The **Access** column displays only for Administrator user accounts and user accounts assigned to the System organization.

- **Edit Date.** Date the appliance's information was discovered or last edited.
- **Edit User.** User who last edited the appliance's information.
- **Create Date.** Date and time the appliance was registered and licensed.

6. To view the Web Configuration Utility for an appliance, where you can track license data, interfaces, and other device settings, click the Appliance Manager icon (🔗). Use the same login credentials that you used to log into Skylar One, and close the pop-up window for the Utility when you are done.
7. If a Skylar One appliance is running a different version of Skylar One than the Database Server, that appliance is highlighted in the **Appliance Manager** page. The version number, if known, is listed in the **Build** column.
8. For all Skylar One appliances, Skylar One runs the system status script every 15 minutes. You can click the logs icon (🔗) to view the results of the latest system status script.
9. If you are logging in to the "sl1admin" account on an appliance, you can click the padlock (🔒) icon for that appliance to get a one-time password. For more information, see "Using the sl1admin Account" in the *Role-Based User Accounts* chapter of the **Organizations and Users** manual.
10. For Data Collectors and Message Collectors, you can click the lightning bolt icon (⚡) to manually force the Database Server to send the latest configuration information.

NOTE: The delete icon (🚫) does not appear for Database Servers that are not configured for High Availability or Disaster Recovery. The bomb icon does not appear for Database Servers that are configured as the primary database in a High Availability or Disaster Recovery configuration.

11. Click the **[Save]** button to save any changes. Click the **[Save As]** button to save your changes to a new appliance name.

Logging in Skylar One Version 11.3.0 and Later

In Skylar One version 11.3.0 and later, configuration files for rsyslog were completely updated. This new configuration gives you the option of configuring TLS to send or receive syslog message, forwarding logs to a security information and event management (SIEM) tool, filtering inbound logs, and other features. These options are described in detail in the following sub-topics.

WARNING: Any existing modifications you made to your rsyslog configurations to support log forwarding, filtering, or TLS reception before Skylar One version 11.3.0 will be removed. To re-configure any custom rules using the appropriate syntax, see the sub-topics below.

The following options are available:

- [Configure TLS to send or receive syslog messages](#)
- [Forward locally generated syslogs to one or more external systems for auditing or processing](#)
- [Specify alternate inbound TCP or UDP ports for listening for syslog messages](#)
- [Adjust the priority filter for inbound messages](#)
- [Filter or discard inbound messages](#)

Configuring TLS Certificates

You will need to configure the TLS certificates before you can send or receive syslogs using SSL/TLS.

Requirements:

- A PEM-encoded CA Certificate
- A PEM-encoded Certificate
- A PEM-encoded Private Key
- All three files uploaded to the `/etc/pki/rsyslog/` directory

NOTE: When the ***Require TLS Validation for Gen 1 Agent checkbox*** is selected on the **Behavior Settings** page (System > Settings > Behavior), Gen 1 agents that operate outside of the Extended Architecture will require TLS validation to upload data. To enable the ***Require TLS Validation for Gen 1 Agent checkbox***, all Data Collectors and Message Collectors that have agents uploading data to them must have a valid and signed certificate so that the TLS validation can complete the upload request.

To configure the TLS certificates:

1. Go to the console of the Skylar One server or use SSH to access the server and log in as user **em7admin** with the password you configured during setup.
2. Locate the `/etc/rsyslog.d/siteconfig.d/global_tls.conf` file and edit the file to reference the CA Certificate, Certificate, and Private Key. Uncomment the global section if needed. For example:

```
# TLS certificates must be defined here when using inbound TLS or TLS
forwarding.
```

```
global (
    DefaultNetstreamDriver="gtls"
    DefaultNetstreamDriverCAFfile="/etc/pki/rsyslog/exampleCA.crt"
    DefaultNetstreamDriverCertFile="/etc/pki/rsyslog/mycert.crt"
    DefaultNetstreamDriverKeyFile="/etc/pki/rsyslog/mycert.key"
)
```

3. Save the file and check the configuration syntax by running the following command:

```
sudo rsyslogd -N1
```

4. You can optionally restart the **rsyslog** service:

```
sudo service rsyslog restart
```

Forwarding Local Syslog Messages to Remote Systems

Requirements:

- One or more destinations that can accept syslog messages with either UDP, TCP or TCP with TLS
- The IP address or FQDN of each destination system; the FQDN must be resolvable from the sending appliance
- The port number for each destination system
- The protocol for each destination system

To forward local syslog messages to remote systems:

1. Go to the console of the Skylar One server or use SSH to access the server and log in as user **em7admin** with the password you configured during setup.
2. Locate the **/etc/rsyslog.d/siteconfig.d/log_forwarding.conf** file and edit the file to add the destinations.

The file contains several examples in addition to the following examples:

Example 1: If you are forwarding to a server using UDP, add the following line to the bottom of the configuration file, substituting in your IP or FQDN for the **Target** and your port for the **Port**:

```
action(name="UDP_Forward" type="omfwd" Target="192.0.2.5" Port="514"
Protocol="udp")
```

Example 2: If you are forwarding a subset of high-priority messages using TCP with TLS, add the following line to the bottom of the configuration file, substituting in your IP/FQDN for the **Target** and your port for the **Port**.

```
if prifilt("*.err;*.emerg;*.alert;") then {

action(name="TCP_Forward" type="omfwd" Target="192.0.2.200"
Port="1514" Protocol="tcp" StreamDriver="gtls" StreamDriverMode="1"
StreamDriverAuthMode="anon")

}
```

3. Repeat step 2 if you want to add multiple destinations. When adding multiple destinations you will need to alter the name property so it is unique. It can be as simple as **TCP_Forward1**, **TCP_Forward2**, etc... ScienceLogic strongly recommends using the UDP protocol and having no more than 2 destinations. Each destination increases the bandwidth usage and using TCP or TCP w/TLS will increase the processing overhead.
4. Save the file and check the configuration syntax by running the following command:

```
sudo rsyslogd -N1
```

5. You can optionally restart the **rsyslog** service:

```
sudo service rsyslog restart
```

Specifying Alternate Inbound TCP or UDP Ports

1. Go to the console of the Skylar One server or use SSH to access the server and log in as user **em7admin** with the password you configured during setup.

2. Locate the `/etc/rsyslog.d/siteconfig.d/inbound_alternates.conf` file and edit the file to add the destinations. The file contains several examples in addition to the example here:

The file contains several examples in addition to the following examples:

Example 1: To accept TCP with TLS-encrypted inbound messages, add the following section (there should only be one module section in the configuration):

```
module(  
    load="imtcp"  
    StreamDriver.Name="gtls"  
    StreamDriver.Mode="1"  
    StreamDriver.Authmode="anon"  
)
```

Add the input section below the module section to specify the network port you would like to use (the default 514 is already in use, and ScienceLogic does not support disabling it in rsyslog):

```
input(  
    type="imtcp"  
    port="1514"  
    ruleset="ruleset-networksocket"  
)
```

Example 2: To accept TCP inbound messages on an alternate port, add the following section to specify the network port you would like to use (the default 514 is already in use, and ScienceLogic does not support disabling it in rsyslog):

```
input(  
    type="imptcp"  
    port="9514"  
    ruleset="ruleset-networksocket"  
)
```

NOTE: The type is correct as `imptcp`.

3. When setting up alternate input ports, do not change the ruleset line. It should always read as `ruleset="ruleset-networksocket"`. If it does not read this, inbound messages will not be processed properly.
4. When opening up additional inbound ports you will also be required to allow the ports through the local firewall as well.
5. Save the file and check the configuration syntax by running the following command:

```
sudo rsyslogd -N1
```

6. You can optionally restart rsyslog now or proceed to other sections to configure inbound or outbound messages:

```
sudo service rsyslog restart
```

Adjusting the Priority Filter for Inbound Messages

The following procedure lets you adjust the priority filter for Inbound messages from Message Collectors and Data Collectors:

1. Go to the console of the Skylar One appliance or use SSH to access the appliance and log in as user `em7admin` with the password you configured during setup.
2. Locate the `/etc/rsyslog.d/siteconfig.d/inbound_message_filter.conf` file and edit the file to include or uncomment the example line and adjust the priority filter to meet your needs. You can refer to the following examples, but there should only be one line to adjust the filter:

Example 1: To accept all messages, use the following line:

```
set $.priority_filter = prifilt("*.*");
```

CAUTION: ScienceLogic does not recommend using this setting, as it can result in message floods.

Example 2: Accept "auth" messages regardless of severity, in addition to the default priorities:

```
set $.priority_filter = prifilt
("auth.*,*.err,* emerg,* alert,local7.*,local6.*,local5.*,local4.*,lo
cal3.*,local2.*,local1.*,local0.*");
```

3. Save the file and check the configuration syntax by running the following command:

```
sudo rsyslogd -N1
```

4. You can optionally restart rsyslog now or proceed to other sections to configure inbound or outbound messages:

```
sudo service rsyslog restart
```

Filtering or Discarding Inbound Messages

The following procedure lets you filter or discard inbound messages from Message Collectors and Data Collectors:

1. Go to the console of the Skylar One appliance or use SSH to access the appliance and log in as user **em7admin** with the password you configured during setup.
2. Locate the **/etc/rsyslog.d/siteconfig.d/inbound_message_filter.conf** file and edit the file to include a block to match against the desired messages and a stop statement.

For example:

```
if ($programname startswith "noisy_program") then {  
    stop  
}
```

NOTE: This is not the recommended way to limit the inbound messages to the system. The preferred method is to alter the sending device to only send relevant syslog messages to the system. This functionality is provided for limited use, because sending large amounts of messages in to the system only to be discarded increases processing load and can impact performance. Always limit messages at the source if available.

NOTE: The matching syntax is fairly flexible. Please refer to the official Rsyslog documentation at <https://www.rsyslog.com/doc/master/index.html> for available variables and match logic.

3. Save the file and check the configuration syntax by running the following command:

```
sudo rsyslogd -N1
```

4. You can optionally restart rsyslog now or proceed to other sections to configure inbound or outbound messages:

```
sudo service rsyslog restart
```

Sending Logs via Syslog to a Remote Server in Skylar One Version 11.2.x and Earlier

WARNING: This method of sending logs to a remote server was deprecated in Skylar One version 11.3.0 and later. For more information on the new method of sending logs to a remote server, see [Logging in Skylar One Version 11.3.0 and Later](#).

To send logs to a syslog server or a security information and event management (SIEM) tool in Skylar One version 11.2.0 or earlier:

1. Edit the file `/etc/rsyslog.conf` and include the following text to send the `audit.log` to rsyslog:

```
#audit log  
  
$ModLoad imfile  
  
$InputFileName /var/log/audit/audit.log  
  
$InputFileTag tag_audit_log:  
  
$InputFileStateFile audit_log  
  
$InputFileSeverity info  
  
$InputFileFacility local6  
  
$InputRunFileMonitor
```

2. If rsyslog is not already configured to send to a remote server, configure the remote server at this time.
3. Restart rsyslog for the changes to take effect:

```
[root@server ~]# systemctl restart rsyslog
```

Monitoring User Actions and Events on the Audit Logs Page

The **Audit Logs** page (System > Monitor > Audit Logs) provides an audit trail for Skylar One. The **Audit Logs** page displays a record of actions in Skylar One that are generated by **users** or by **managed elements**. These actions are organized by organization.

Some of the actions that are logged in the **Audit Logs** page include:

- User logins to Skylar One
- The sl1admin user requests a one-time password
- Organization name changes
- Appliance IP address changes
- The addition, editing, or deletion of elements in Skylar One

NOTE: Entries for the addition, editing, and deletion of elements includes the affected device ID, when applicable.

- The installation, editing, or uninstallation of PowerPacks, including when a PowerPack is imported or installed from Global Manager to a Stack
- Manually triggered discovery sessions
- Events and cleared events
- The creation, editing, or deletion of event policies
- Session ID renewal events
- Devices being set to maintenance mode or devices no longer being in maintenance mode, including user-initiated changes in maintenance mode that are performed on one or more devices
- The unalignment of Dynamic Applications from devices and the deletion of that data
- The creation, editing, or deletion of Dynamic Applications
- The creation, editing, or deletion of Run Book Automation policies
- The addition or deletion of Reports
- The situation where a user gets an "Access Denied" page as a result of attempting to access a page for which he or she does not have permission
- Asset Record changes
- User-defined changes to settings on the **Data Retention Settings** page (System > Settings > Data Retention)
- Changes to settings on the **System Threshold Defaults** page (System > Settings > Thresholds > System)
- API requests that use a PUT, POST, or DELETE method
- All actions that renew the Skylar One session ID
- Updates to interface collection states

NOTE: By default, the **Audit Logs** page displays a list of actions associated with all organizations.

Viewing the List of Audit Logs

To view the list of log entries in the **Audit Logs** page:

1. Go to the **Audit Logs** page (System > Monitor > Audit Logs).
2. The **Audit Logs** page displays all actions that are performed by users or managed elements in Skylar One. For each action, the **Audit Logs** page displays:
 - **Date/Time**. Date and time the action occurred and the log entry was created.
 - **Organization**. Organization associated with the action.
 - **Source**. Source of the log entry. This usually describes where the action took place. For example, if you change the contact information for your account, an entry will be made in the audit log, and the source will be "Contact Information."
 - **Message**. Text of the log entry.

Generating Reports on Audit Logs

You can export the entries on the **Audit Logs** page as one of the following report types:

- Acrobat document (.pdf)
- Web page (.html)
- Excel spreadsheet (.xlsx)
- OpenDocument Spreadsheet (.ods)
- Comma-separated values (.csv)

When you create a report in the **Audit Logs** page, Skylar One includes only those logs that appear in the current view of the page. If you filter the entries on the **Audit Logs** page, only those logs that meet the filter criteria and currently appear on the page will appear in the report.

To generate an audit logs report:

1. From the **Audit Logs** page, click the **[Report]** button. The **Export current view as a report** window appears.
2. In the **Output Format** field, select the report format type.
3. Click **[Generate]**.

Using auditd to Monitor Sensitive Files

As an additional security measure, the ScienceLogic Security team monitors the Skylar One SaaS environment with the Linux Auditing System (**auditd**). The auditd system provides enhanced logging that identifies when changes are made to specific files and directories in Skylar One that contain sensitive data. Starting with Skylar One version 11.1.0, you can send audit logs externally to a security information and event management (SIEM) tool when needed.

The auditd logging at **/var/log/audit/audit** provides logging based on the configured rules that are shipped with Skylar One version 11.1.0 or later. You can send syslog messages to a syslog server or a SIEM tool.

NOTE: Starting with Skylar One version 12.1.0, audisdp configuration options are part of **auditd.conf**. In addition, the **plugins.d** directory was moved under **/etc/audit**. You can now check the status of auditd and its plug-ins can now be checked by running the `service auditd state` command.

Skylar One monitors specific files and directories for changes. The topics below describe what is logged by default, how to modify what is logged, and how to send logs externally.

Files Logged by Default

Starting in Skylar One version 11.1.0, a file called **rules.d/70_silo_watchlist.rules** is created and added to the **/etc/audit** directory.

The following files are monitored:

```
-w /etc/silo.conf -p war -k silowatch  
  
-w /etc/php.ini -p wa -k silowatch  
  
-w /etc/yum.conf -p wa -k silowatch  
  
-w /etc/postfix/main.cf -p wa -k silowatch  
  
-w /etc/postfix/master.cf -p wa -k silowatch  
  
-w /etc/nginx/nginx.conf -p wa -k silowatch  
  
-w /etc/yum.conf -p wa -k silowatch  
  
-w /etc/ntp.conf -p wa -k silowatch  
  
-w /etc/my.cnf -p wa -k silowatch  
  
-w /etc/my.cnf.d/ -p wa -k silowatch  
  
-w /etc/php.d/ -p wa -k silowatch  
  
-w /etc/skel/ -p wa -k silowatch  
  
-w /etc/yum/ -p wa -k silowatch  
  
-w /etc/yum.repos.d/ -p wa -k silowatch  
  
-w /etc/ntp/ -p wa -k silowatch  
  
-w /etc/chrony.d/ -p wa -k silowatch
```

```
-w /etc/chrony.keys -p wa -k silowatch
```

```
-w /etc/chrony.conf -p wa -k silowatch
```

Modifying the List of Files to be Logged

To add or change files that are monitored:

1. Log in to the Skylar One appliance with the admin account and sudo to the root user.
2. Edit the file called **/etc/audit/rules.d/70_silo_watchlist.rules**. Add the following line to the file.

```
-w /etc/.custom_alignment.conf -p wa -k silowatch
```

NOTE: This file is not included in Skylar One version 11.1.0 by default.

3. Run the following command to load the files into the **audit.rules** file:

```
[root@server ~]# augenrules --load
```

NOTE: You can add or remove specific files that you want to monitor from this list. Including the slash at the end of the monitored path (such as **/etc/ntp/**) will monitor the entire directory and everything stored within it.

Monitoring the Status of Data Collectors

The **Collector Status** page displays the status of each Data Collector and Message Collector in your system.

NOTE: This page does not appear in All-In-One Appliances.

Data Collectors retrieve data from managed devices and applications. This collection occurs during initial discovery, during nightly updates, and in response to policies defined for each managed device. The collected data is used to trigger events, display data in the user interface, and generate graphs and reports.

Message Collectors receive and process inbound, asynchronous syslog and trap messages from monitored devices. In most distributed systems, dedicated **Message Collector** appliances perform message collection. A single **Message Collector** can handle syslog and trap messages from devices that are monitored by multiple **Data Collectors**.

To perform collection, you must define a Collector Group and align it with at least one Data Collector. If your Collector Group includes multiple Data Collectors, you can configure the Collector Group for high-availability. For details, see the section on [Collector Groups](#).

To ensure the health of your system, you should periodically check on the status of the Data Collectors and Message Collectors. To access the **Collector Status** page:

1. Go to the **Collector Status** page (System > Monitor > Collector Status).
2. For each Data Collector in your system, the **Collector Status** page displays the following:
 - **Collector Name**. Name of the Data Collector or Message Collector.
 - **Collector ID**. Unique numeric ID automatically assigned to the Data Collector or Message Collector by Skylar One.
 - **Collector Address**. IP address of the Data Collector or Message Collector.
 - **Group ID**. Unique numeric ID of the [Collector Group](#) associated with the Data Collector or Message Collector.
 - **Group Name**. Name of the [Collector Group](#) associated with the Data Collector or Message Collector.
 - **Last State Change**. Date and time the platform last polled the status of the Data Collector or Message Collector.
 - **Collector State**. Operating state of the Data Collector or Message Collector.

Chapter

5

Updating Skylar One

Overview

This chapter provides an overview of the **System Updates** page, detailed steps for performing a Skylar One (formerly SL1) upgrade, and detailed steps on upgrading MariaDB, upgrading PowerPacks, and performing reboots.

WARNING: As of version 12.2.0, Skylar One can be deployed *only* on Oracle Linux 8 (OL8) operating systems. If you are upgrading from a version of Skylar One prior to 12.1.1 that does not already run on OL8, you *must* first upgrade to Skylar One 12.1.1 or 12.1.2 and then convert to OL8 before upgrading to Skylar One 12.2.0 or later.

Depending on the version of Skylar One that you are currently running, you might be required to upgrade to one or more earlier versions of Skylar One before you can upgrade to 12.2.0 or later.

For more information about supported upgrade paths, see the upgrade notes included in the [Skylar One release notes](#) for the release you want to upgrade to as well as the [OL8 Conversion Resource Center](#) on the ScienceLogic Support portal, which includes links to numerous resources such as the [Oracle Linux 8 Conversion Guide](#). The conversion guide includes prerequisites, full instructions for converting to OL8 for all deployment types, FAQs, and other helpful information to walk you through the OL8 conversion process.

All older Skylar One systems with OL7 are still operable, but ScienceLogic no longer supports them, and the systems might not be secure.

WARNING: See the section [*Running a Pre-Import Script to Prevent 12.5.x Patching Issues*](#) before importing a 12.5.x patch bundle. This section includes important warnings for users who still have Python 2 content in their Skylar One system or are upgrading from 12.1.x or 12.2.x versions of the platform.

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>Skylar One Upgrade Planning and Checklist</i>	139
<i>Skylar One Recommended Upgrade Paths</i>	143
<i>The System Updates Page</i>	143
<i>Scheduling Maintenance Windows</i>	145
<i>Pre-Upgrade Best Practices for Skylar One</i>	146
<i>Verifying PowerPack Version Compatibility</i>	146
<i>Backing Up SSL Certificates</i>	147
<i>Setting the Timeout for PhoneHome Watchdog</i>	147
<i>Running the Pre-Upgrade Test for PhoneHome Database Server</i>	148
<i>Adjusting the Timeout for Slow Connections</i>	149
<i>Running a Pre-Import Script to Prevent 12.5.x Patching Issues</i>	150
<i>Running the System Status Script Before Upgrading</i>	153
<i>Updating the Skylar One Distributed Architecture</i>	153
<i>Updating Skylar One Extended Architecture</i>	173
<i>Automatically Upgrading MariaDB</i>	173
<i>Manually Upgrading MariaDB</i>	174
<i>Rebooting Appliances in the Skylar One Distributed Stack</i>	187
<i>Upgrading to Aurora 3 RDS (MySQL 8.0)</i>	194
<i>Restoring the SSL Certificates</i>	195
<i>Resetting the Timeout for PhoneHome Watchdog</i>	196
<i>Clearing System Cache</i>	197
<i>Updating Default PowerPacks</i>	197
<i>Configuring Subscription Billing</i>	198

Skylar One Upgrade Planning and Checklist

Updating Skylar One is a multiple-step process that you should review and plan out carefully before starting. The checklists and tables below list the different aspects that you will need to consider before upgrading to any Skylar One release.

Planning the Update

Before upgrading Skylar One, perform the following steps that are specific to your organization:

1. Read the [release notes](#) to determine:
 - What is fixed?
 - What is new?
 - What has changed?
 - What has been deprecated?
 - What are the supported upgrade paths?
 - What are the important notes that you need to know before upgrading?
2. Read the **Known Issues** for the release. These can be found in the release notes and at <https://support.scienceologic.com/s/knowledge/known-issues>.
3. Identify all integrations and third-party applications that access the Skylar One database or manipulate data on Skylar One. Determine how to disable these integrations during the deployment and re-enable after deployment.
4. Identify activities and customers that will be affected by maintenance windows and schedule and inform appropriately.
5. Identify any custom work, such as changes to PowerPacks, run book automations, event policies, and dashboard widgets, and ensure that all custom work is backed up so you can restore it if necessary.

NOTE: Ensure that each Skylar One node or appliance has at least 4 GB of free space in the `/var` partition for pre-upgrade and deployment. Ensure that each Skylar One node or appliance has at least 1 GB of free space in `/` (the root partition) to allow you to deploy the upgrade; however, depending on the appliance type and version you are upgrading from, you might need up to 3 GB of free space in `/` (the root partition).

NOTE: The option to enable a Military Unique Deployment (MUD) configuration is not available for Skylar One 12.1.x or 12.2.0 installations or upgrades.

Skylar One Upgrade Checklist

The checklist below provides the required steps you should consider before updating your Skylar One platform to a major release version:

Step	Actions
□ 1.	Determine the recommended upgrade path availability: <ol style="list-style-type: none">Review the Skylar One Recommended Upgrade Paths.Identify your target Skylar One release version.Confirm that upgrade limitations or constraints do not apply.
□ 2.	Review your infrastructure system requirements: <ol style="list-style-type: none">Verify that your hardware specifications meet the minimum requirements for the target version at the ScienceLogic Support site:<ul style="list-style-type: none">Customer Premise-VirtualCustomer Premise-HardwareCloud-AWSCloud-AzureEnsure that sufficient disk space is available on all appliances.Check network connectivity between all Skylar One components.Verify DNS resolution for all Skylar One appliances.Confirm that the current Skylar One version is supported for a direct upgrade to the target version.Check for any required intermediate upgrades.Verify that all Skylar One appliances are running the same version.Ensure that all system packages are up to date.
□ 3.	Consider the Skylar One requirements: <ol style="list-style-type: none">Review the compatibility of currently installed PowerPacks with the target Skylar One version.Check for any custom scripts or integrations that might need updating.Verify that the database schema is compatible with the target version of Skylar One.Ensure that all required licenses are valid and up to date.
□ 4.	Review the Skylar One Upgrade Workflow steps: <ol style="list-style-type: none">Plan the update.Schedule maintenance windows.Review pre-upgrade best practices for Skylar One.Verify PowerPack version compatibility.

Step	Actions
	<p>5. <i>Back up SSL certificates.</i></p> <p>6. <i>Set the timeout for PhoneHome Watchdog.</i> (Required only if upgrading from Skylar One 11.1.x or earlier.)</p> <p>7. <i>Run the Pre-Upgrade Test for PhoneHome Database Servers.</i> (Required only if upgrading from Skylar One 11.1.x or earlier.)</p> <p>8. <i>Adjust the timeout for slow connections.</i></p> <p>9. <i>Run a pre-import script to prevent 12.5.x patching issues.</i> (Required only if upgrading from 12.1.x or 12.2.x to 12.5.x.)</p> <p>10. <i>Run the system status script</i> on the Database Server or All-In-One Appliance before upgrading.</p> <p>11. <i>Update the Skylar One Distributed Architecture</i> using the System Update tool.</p> <p>12. <i>Upgrade MariaDB</i>, if there is a version difference.</p> <p>13. <i>Reboot Skylar One appliances</i>, if needed.</p> <p>14. <i>Upgrade to Aurora 3 RDS (MySQL 8.0).</i> (Required only if deployed on AWS and currently on Aurora 2.)</p> <p>15. <i>Restore SSL certificates.</i></p> <p>16. <i>Reset the timeout for PhoneHome Watchdog.</i> (Required only if previously adjusted.)</p> <p>17. <i>Clear your system cache.</i></p> <p>18. <i>Update default PowerPacks.</i></p> <p>19. <i>Configure Subscription Billing</i> (required one time only).</p>
□ 5.	<p>Run the pre-upgrade Health Checks:</p> <ol style="list-style-type: none"> Run the built-in <i>pre-upgrade check tool</i> . Verify that all services are running and healthy. Check for any active alarms or critical events. Review system logs for any recurring errors or warnings. Perform a database integrity check. Verify that backup systems are functioning and up to date. Ensure that each Skylar One node or appliance has at least 4 GB of free space in the <code>/var</code> partition for pre-upgrade and deployment. Ensure that each Skylar One node or appliance has at least 1 GB of free space in <code>/</code> (the root partition) to allow you to deploy the upgrade. Depending on the appliance type and version you are upgrading from, you might need up to 3 GB of free space in <code>/</code> (the root partition).
□ 6.	<p>Upgrade preparation steps:</p> <ol style="list-style-type: none"> <i>Create a full and a configuration backup</i> of all Skylar One appliances. Document your current configuration settings. Notify users of the planned maintenance window.

Step	Actions
□ 7.	<p>Update the Skylar One Distributed Architecture:</p> <ol style="list-style-type: none"> a. <i>Download the update.</i> b. <i>Import the update.</i> c. <i>Stage the update.</i> d. <i>Run the pre-upgrade check.</i> e. <i>Place all Skylar One appliances in Maintenance Mode.</i> f. <i>Deploy the update.</i> g. <i>Run the post-upgrade recovery script.</i> (Required only if upgrading from 12.1.x or 12.2.x to 12.5.x and you did not run the <i>pre-import script to prevent 12.5.x patching issues.</i>) h. <i>Remove all Skylar One appliances from Maintenance Mode</i>

Skylar One Recommended Upgrade Paths

For full details on recommended upgrade paths, see the online documentation on [Updating Skylar One](#).

The System Updates Page

The **System Updates** page (System > Tools > Updates) allows you to update the software on your Skylar One appliances.

You must first download the update file to the local computer. You can then *import* the software update through the user interface.

After you import a software update to your Skylar One system, the Skylar One system can automatically *stage* the software update. Staging is when the software is copied to each ScienceLogic appliance. Staging allows Skylar One to simultaneously apply the software changes to each ScienceLogic appliance, regardless of the speed of the connection to each ScienceLogic appliance. You can allow the Skylar One system to automatically stage the software or you can manually stage the software.

After the software update is staged, you can *run preupgrade checks* and then *deploy* the software.

WARNING: To apply updates to an existing Data Collector, that Data Collector must be a member of a Collector Group. In some Skylar One systems, users might have to create a Collector Group for a single Data Collector before applying updates.

NOTE: To conserve disk space on Data Collectors and Message Collectors, after an update, Skylar One removes previous Docker images.

The **System Updates** page displays a list of deployed updates and the date and time each update was applied. The following information displays for each update in the list:

- **EM7 Version.** Displays the version number for each Skylar One patch that has been imported (and optionally staged and deployed) to the Skylar One system. Skylar One releases include new features and address issues with previous releases.
- **OS Version.** Displays the version number for each platform patch that has been imported (and optionally staged and deployed) to the Skylar One system. Platform patches include updates to the ScienceLogic operating system, including newer versions of industry-standard packages.
- **Update Signature.** Name of the entity that released the patch and type of patch. Usually "ScienceLogic Official Release".
- **Imported On.** Date and time the software update was loaded onto the Skylar One system.
- **Imported By.** Skylar One user who loaded the software update onto the Skylar One system.

- **Import Status.** Status of the import process. Clicking on the log icon displays the log file associated with importing the selected software. Possible values are:
 - *In Progress.* Software is currently being imported by the system.
 - *Complete.* Software has been imported successfully.
 - *Failed.* Software import has failed due to an unexpected condition. Contact ScienceLogic Support for assistance.
- **Staging Status.** Status of the staging process. Clicking on the log icon displays the log file associated with staging the selected software. Possible values are:
 - --. No staging request is active and software has not been staged on any Skylar One appliances.
 - *Scheduled.* Skylar One is aware of the staging request and is preparing for staging.
 - *In Progress.* Staging is in progress but has not completed.
 - *Complete.* Staging has completed, and all appliances are ready to deploy the software.
 - *Incomplete.* Staging has completed, and one or more appliances are ready to deploy the software.
 - *Cancelled.* User manually cancelled the staging process.
 - *Outdated.* The current update is not the latest or has already been installed.
 - *Failed.* An unexpected error occurred in the staging process. Contact ScienceLogic Support for assistance.

NOTE: If a software update has been set to manually stage, the Staging Status column will include an asterisk (*).

- **Preupgrade Status.** Status of the preupgrade checks. Possible values are:
 - --. No preupgrade checks are active.
 - *Complete.* Preupgrade checks have completed successfully.
 - *Incomplete.* Preupgrade checks have not yet completed.
 - *Failed.* An unexpected error occurred in the preupgrade checks. Contact ScienceLogic Support for assistance.

- **Deployment Status.** The current deployment state. Possible values are:
 - *--*. No deployment request is active and software has not been deployed on any Skylar One appliances.
 - *Scheduled*. Skylar One is aware of the deployment request and is preparing for deployment.
 - *In Progress*. Deployment is in progress but has not completed.
 - *Complete*. Deployment has completed, and all appliances are updated.
 - *Incomplete*. Deployment has completed, and one or more appliances are updated.
 - *Cancelled*. User manually canceled the deployment.
 - *Outdated*. The current update is not the latest or has already been installed.
 - *Failed*. An unexpected error occurred in the deployment process. Contact ScienceLogic Support for assistance.
- **Deployment Status Date.** Specifies the date and time the software update was last deployed.

From the **System Updates** page, you can also click the **[Appliance List]** button to view a list of appliances in your Skylar One stack. The following information displays for each appliance in the list:

- **Appliance Name.** The name assigned to the Skylar One appliance.
- **IP.** The appliance's IP address.
- **mid.** The appliance's module ID.
- **Module Type.** The Skylar One appliance type.
- **Build.** The version of Skylar One currently running on the appliance.
- **MariaDB.** The version of MariaDB currently running on the appliance.
- **Platform.** The current operating system platform version for the appliance.
- **Task Manager Paused?** Indicates whether the task manager service is paused.
- **Patch Eligible.** Indicates if the appliance is eligible for patching.
- **Patch Enabled.** Indicates if patching is enabled or disabled for the appliance.

NOTE: To enable or disable patching for an appliance, select the checkbox for that appliance on the **Appliance List** modal, select *Enable Selected Appliances* or *Disable Selected Appliances* from the **Select Action** menu, and then click **[Go]**.

Scheduling Maintenance Windows

Upgrading Skylar One includes a minimum of two and possibly four maintenance windows:

- **Import and stage update and run the pre-upgrade script.** These steps can take place prior to the day of upgrade and **do not affect Skylar One functionality**. ScienceLogic suggest you perform these steps at least three days before the planned upgrade and ideally a week before the planned upgrade.

- **Deploy update.** On the day of the upgrade, put all Skylar One appliances in maintenance mode. The Skylar One system will not be available during this procedure. Update the Skylar One Distributed systems.
- **Update MariaDB (if required).** The Skylar One system will not be available during this procedure. Refer to the release notes for your current release to determine if you must upgrade MariaDB.
- **Reboot Appliances (if required).** Individual Skylar One appliances will not be available during these procedures. Refer to the release notes for your current release to determine if you must reboot all Skylar One appliances after upgrading.

Identify activities and users that will be affected by these maintenance windows, and schedule the maintenance windows appropriately. Be sure to communicate all downtime with users.

Pre-Upgrade Best Practices for Skylar One

Before you upgrade, check the following:

- Review the hardware specifications of all the appliances in your system to ensure they meet the requirements for the current usage of your system. For more details about sizing and capacity for your specific environment, contact your Customer Success Manager and see <https://support.scienceologic.com/s/skylar-one/system-requirements>.
- Verify that recent backups are available for your system.
- Ensure that each Skylar One appliance has a valid license.
- Ensure that a Data Collector is a member of a Collector Group if you are applying updates to an existing Data Collector. In some Skylar One systems, users might have to create a Collector Group for a single Data Collector.
- Ensure that each Data Collector is listed as "Available" to the Database Server. To check, see the **Collector Status** page (System > Monitor > Collector Status).

Verifying PowerPack Version Compatibility

Before upgrading Skylar One, you should verify whether any PowerPacks currently running on your system are newer than the PowerPacks included in the Skylar One release to which you are upgrading.

If your Skylar One system currently has a PowerPack version that is newer than the one included with the release you are upgrading to, you might see spurious error messages post-upgrade. To avoid these error messages:

1. Before installing the Skylar One update, go to the **Device Components** page (Devices > Device Components).
2. Find each root device associated with the PowerPacks you do not want to update and select their checkboxes.
3. Click the **Select Action** field and choose *Change Collection State: Disabled (recursive)*, and then click the **[Go]** button.
4. Wait five minutes after disabling collection.

5. Install the Skylar One update.
6. After the Skylar One update is complete, go to the **Device Components** page (Devices > Device Components).
7. Select the checkboxes for all affected root devices.
8. Click the **Select Action** field and choose *Change Collection State: Enabled (recursive)*, and then click the **[Go]** button.

Backing Up SSL Certificates

To back up your SSL Certificates:

1. Log in to the console of the Database Server or SSH to the Database Server.
2. Open a shell session and type the following commands at the shell prompt:

```
cp /etc/nginx/silossl.key /etc/nginx/silossl.key.bak
```

```
cp /etc/nginx/silossl.pem /etc/nginx/silossl.pem.bak
```

3. Repeat these steps on each Database Server in your Skylar One system.

Setting the Timeout for PhoneHome Watchdog

NOTE: This section applies to users who are upgrading from Skylar One 11.1.x or earlier and have an existing PhoneHome configuration. If you are upgrading from Skylar One 11.2.0 or later or you do not have a pre-11.2.0 PhoneHome configuration, you can ignore this section.

You can manually adjust the settings for the PhoneHome Watchdog server to reduce CPU consumption during the upgrade process. To do this:

1. Log in to the console of the Data Collector as the root user or open an SSH session on the Data Collector.
2. At the command line, type the following:

```
phonehome watchdog view
```

3. You should see something like the following:

```
Current settings:  
autosync: yes  
interval: 20  
state: enabled  
autoreconnect: yes  
timeoutcount: 2  
check: default
```

4. Note the settings for *interval* and *timeoutcount*, so you can restore them after the upgrade.

5. To change the settings for Skylar One upgrade, type the following at the command line:

```
sudo phonehome watchdog set interval=120;  
sudo phonehome watchdog set timeoutcount=2;  
systemctl stop em7_ph_watchdog;  
systemctl start em7_ph_watchdog;
```

6. Repeat the steps in this section on each Data Collector.
7. Repeat the steps in this section on each Message Collector.
8. Repeat the steps in this section on each Database Server.

Running the Pre-Upgrade Test for PhoneHome Database Server

NOTE: This section applies to users who are upgrading from Skylar One 11.1.x or earlier and have an existing PhoneHome configuration. If you are upgrading from Skylar One 11.2.0 or later or you do not have a pre-11.2.0 PhoneHome configuration, you can ignore this section.

When upgrading to Skylar One version 11.2.0 or later, a pre-upgrade test in the system update procedure checks for existing PhoneHome Database Servers. Specifically, this test looks for PhoneHome token IDs inside the `/home/phonehome0/config.json` file, and fails if the value of the `id` field is less than or equal to "0".

In versions of Skylar One prior to version 11.2.0, the primary PhoneHome Database Server was not self-registered with a token, causing it to have an `id` value of "0".

Therefore, if you are currently using a version of Skylar One prior to version 11.2.0 and you are using PhoneHome communication with your Skylar One Collectors, you or your Skylar One administrator **must** complete the following one-time manual configuration steps prior to upgrading to version 11.2.0 or later.

CAUTION: Do not attempt to upgrade to Skylar One version 11.2.0 or above until all pre-upgrade tests are successful on all PhoneHome Database Servers.

To register a token for your primary PhoneHome Database Server:

1. Log in to the console of the Database Server or use SSH to access the server.
2. To determine if your PhoneHome Database Server is registered, type the following command and locate its `id` value:

```
cat /home/phonehome0/config.json
```

3. If a PhoneHome Database Server has an `id` value of "0", type the following command and locate the `id` of the current appliance:

```
phonehome status
```

4. Type the following command and locate the PhoneHome token:

```
phonehome token <id from step 3>
```

5. Type the following command to register the PhoneHome token:

```
phonehome register <token from step 4>
```

6. Repeat steps 2-5 for all PhoneHome Database Servers that have an `id` value of "0".

7. Type the following command to ensure that all of your PhoneHome Database Servers are synced:

```
phonehome sync
```

8. Repeat step 2 and confirm that all Database Servers have `id` values greater than "0".

TIP: After you have *downloaded* but not yet *installed* the ISO, you can also go to the **System Updates** page (System > Tools > Updates) and check the **Preupgrade Status** column to determine if the upgrade has successfully passed the pre-upgrade test. If the **Preupgrade Status** is *Incomplete*, click the magnifying glass icon (🔍) to determine which appliance is failing during the pre-upgrade test. *Do not* install the upgrade until the **Preupgrade Status** status displays as *Complete*.

Adjusting the Timeout for Slow Connections

If you have slow connections between Skylar One appliances, you can adjust the timeout values for staging and deploying upgrades.

To adjust the timeouts:

1. Log in to the console of the Database Server or SSH to the Database Server.
2. Open a shell session and type the following at the shell prompt:

```
sudo pcli set-patcher-param staging_wait_time <timeout_in_seconds>
```

where:

`<timeout_in_seconds>` is the timeout value, in seconds, for staging for each Skylar One appliance. The default value is 1800 seconds (30 minutes). You can increase this value for slow connections.

3. Type the following at the shell prompt:

```
sudo pcli set-patcher-param deploy_wait_time <timeout_in_seconds>
```

where:

`<timeout_in_seconds>` is the timeout value, in seconds, for deploying to each Skylar One appliance. The default value is 3600 seconds (1 hour). You can increase this value for slow connections.

NOTE: If you are upgrading from a version of Skylar One prior to version 8.14.0, see [Adjusting the Timeout for Skylar One 8.12 and Prior Releases](#).

Running a Pre-Import Script to Prevent 12.5.x Patching Issues

WARNING: *With the Skylar One 12.5.1 Juneau release, support for Python 2 is now deprecated in Skylar One.* Any custom Python 2 content you have developed will stop working if you upgrade to 12.5.1 or later. This includes all Dynamic Application snippets, Execution Environments, Run Book Actions, and ScienceLogic Libraries that use Python 2. These items must be recreated utilizing ScienceLogic-provided enablement tools or migrated to Python 3 **before** you upgrade to this release.

In addition, because of the removal of Python 2 support, if you are upgrading to 12.5.x from a 12.1.x or 12.2.x release, you must execute a script on your active or primary Database Server or Data Engine **before you import the 12.5.x patch bundle**. Failure to do so will prevent you from upgrading to 12.5.x. This issue prevents patch hook RPM from getting installed on the Skylar One appliances in your stack, which in turn prevents you from upgrading to the newer silouupdate RPM that is needed to complete the system update (deployment) process on your Skylar One appliances.

This section describes how to download and run this pre-import script.

For additional information, see the following knowledge base article: <https://support.scienceologic.com/s/article/19093>

IMPORTANT: If you attempt to upgrade from 12.1.x or 12.2.x to xxx without first running the script, your upgrade will ultimately fail. If this occurs, you can *run a post-upgrade script* that will update the silouupdate RPM file on all of your appliances.

NOTE: If you are not upgrading from version 12.1.x or 12.2.x to 12.5.x, you can skip this section. This issue does not impact users upgrading from 12.3.x to 12.5.x.

To download and run the pre-import script required for users upgrading from a 12.1.x or 12.2.x release to 12.5.x:

1. Before importing the 12.5.x upgrade file, go the **Release Files** page for 12.5.x on the ScienceLogic Support site and download the 12-5-1-pre-import-fix.pyc script: <https://support.scienceologic.com/s/release-file/aBtVL0000001Dsr0AE/preupgrade-script>.
2. On the active or primary Database Server or Data Engine, verify the checksum:

```
sha512sum 12-5-1-pre-import-fix.pyc
```

```
9fb796859a9aef4e8053f6e904bfca35ebdc5dfbffc97f3e76aa08395a727670 12-  
5-1-pre-import-fix.pyc
```

3. Set the correct permissions on the script:

```
chmod 755 12-5-1-pre-import-fix.pyc
```

4. Test that you are able to execute the script by running usage help:

```
sudo /opt/em7/bin/python3.6 12-5-1-pre-import-fix.pyc -h
```

5. Validate that you can see a usage help similar to this:

```
usage: 12-5-1-pre-import-fix.pyc [-h] [-v] [--timeout seconds]

[-m module_id [module_id ...]]

Fix patcher status script

optional arguments:

-h, --help show this help message and exit

-v, --verbose enable debug logs

--timeout seconds timeout in seconds

-m module_id [module_id ...], --mid module_id [module_id ...]

SL1 appliance ID to run the test on
```

6. Make sure that most of the appliances in the stack are available. The script fixes the problems only on the Data Collectors and Message Collectors that are available and skips any unavailable ones. Note that you might need to run the script again for appliances that were unavailable but became available later. You can also run it for a specific set of appliances by specifying their appliance/module IDs using the `-m` flag.
7. Execute the script. Module selection using the `-m` flag to specify one or more appliance IDs is optional. If you do not select any module, it will automatically select all available Skylar One appliances in the stack:

```
sudo /opt/em7/bin/python3.6 12-5-1-pre-import-fix.pyc [-m <mid_1>
<mid_2> ... <mid_n>]
```

NOTE: If you have a large stack, ScienceLogic recommends running the script as a background task as a root user, as shown below:

```
touch /tmp/12-5-1-pre-import-fix.log

chown em7admin:em7admin /tmp/12-5-1-pre-import-fix.log

chmod 644 /tmp/12-5-1-pre-import-fix.log

nohup sudo /opt/em7/bin/python3.6 12-5-1-pre-import-fix.pyc [-m
<mid_1> <mid_2> ... <mid_n>] 2>&1 >/tmp/fix-patcher-status.log &
```

You can then exit from the shell as the root user and can follow the progress of the script by tailing the log file `12-5-1-pre-import-fix.log`.

8. Validate if there are any errors in the log file or in the command output. You might need to look at any reported problem. Fix any issues and then re-run the script for the failing appliances.
9. After you have ensured that most of the appliances in your stack have been fixed by the script, you can import the 12.5.x patch bundle and proceed with your upgrade.

Running the System Status Script Before Upgrading

Skylar One includes a script, **system_status.sh**, that provides diagnostic data for each node or appliance in your Skylar One system.

NOTE: On Skylar One systems prior to 10.2.0, after running the system status script, you must ensure that the file `/var/lob/em7/silo.log` has the owner and group "s-em7-core".

Running the System Status Script

If you are running Skylar One version SL 8.14.0 or later, Skylar One automatically runs the system status script every 15 minutes on each node or appliance in your Skylar One system.

ScienceLogic recommends that you view the output from the system status script before upgrading:

1. In Skylar One, go to the **Appliance Manager** page (System > Settings > Appliances).
2. Locate the Skylar One appliance that you want to view diagnostic information about.
3. Click on its magnifying-glass icon (🔍) to view the output of the system status script for that appliance.
4. If the output includes errors and you need help fixing them, contact ScienceLogic Customer Support to fix the errors before upgrading.
5. Repeat for each node or appliance in your Skylar One system.

TIP: To get the very latest status before upgrading, *manually run the system status script* on each Database Server or All-In-One Appliance.

NOTE: If you are upgrading from a Skylar One 8.12.x system, see *Running the System Status Script on Skylar One 8.12.x Releases*. If you are upgrading from a Skylar One 8.10.x or prior release, see *Running the System Status Script on Skylar One 8.10 and Prior Releases*.

Updating the Skylar One Distributed Architecture

Any Skylar One Distributed Architecture system running 8.6.0 or later can be upgraded by importing, staging, and deploying a single update file.

NOTE: Ensure that each Skylar One node or appliance has at least 4 GB of free space in the `/var` partition for pre-upgrade and deployment. Ensure that each Skylar One node or appliance has at least 1 GB of free space in `/` (the root partition) to allow you to deploy the upgrade; however, depending on the appliance type and version you are upgrading from, you might need up to 3 GB of free space in `/` (the root partition).

Upgrading the Skylar One Distributed Stack includes the following steps:

- [*Download the update.*](#)
- [*Import the update.*](#)
- [*Stage the Update.*](#)
- [*Run the pre-upgrade check.*](#)
- [*Place all Skylar One appliances in Maintenance Mode.*](#)
- [*Deploy the update.*](#)
- [*Remove all Skylar One appliances from Maintenance Mode*](#)

NOTE: Database Servers and Data Engines in your Skylar One stack are deployed first; other appliance types cannot be deployed until the Database Servers and Data Engines are finished deploying.

NOTE: If you are upgrading from a Skylar One system between 8.1.1 and 8.5.0, see [*Upgrading the Distributed Architecture on Skylar One Versions 8.5.0 and Earlier.*](#)

If you are currently running a Skylar One version prior to 8.12.0, you must go to the **System Updates** page and disable automatic staging (System > Tools > Updates > Actions > Disable automatic staging).

If you have previously used manual staging, you must also perform these additional steps:

1. Select all updates in the EM7 Releases pane and select all updates in the ScienceLogic OS pane.
2. In the **Select Action** menu, select *Unstage Update (remove staging policy override)*. Click **[Go]**.
3. For software that was previously staged with automatic staging, *Unstage Update (remove staging policy override)* does not affect staging.

Downloading the Update

Before you can load a patch or update onto your instance of the Skylar One system, you must first download the patch or update to your local computer.

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

To download the patch or update:

1. Log in to <https://support.scienceologic.com>. Use your ScienceLogic customer account and password to access this site.
2. From the **Skylar One** menu, select *Support*. The **Skylar One** platform downloads page appears.
3. Find the release you want to download and click its name. The **Release Version** page appears.
4. In the **Release Files** pane, click the hyperlink for the release you want to download. If you need to, click **[View All]** to expand the list of available release files.
5. On the **Release File Details** page, click the **[Download File]** button for the release image or release patch you want to download. The file is then downloaded to your local computer.

Importing the Update

To import a product update on to your Skylar One system:

1. In the Skylar One system, go to the **System Updates** page (System > Tools > Updates).
2. In the **System Updates** page, click the **[Import]** button.
3. In the **Import a new update** modal page, browse to the product update file and select it.
 - If you select the **Auto Stage** button, the Skylar One system will begin staging as soon as the import is completed.
 - If you do not select the **Auto Stage** button, you must click the staging button() after import is completed. You can do so at any time after import has completed.
 - For more information on automatic staging and manual staging, see the section on "Staging" in the **System Administration** manual.
4. Click the **[Import]** button.
5. In the **System Updates** page, the *Import Status* column can have one of the following statuses:
 - *In Progress*. Software is currently being imported by the Skylar One system.
 - *Complete*. Software has been imported successfully.
 - *Failed*. Software import has failed due to an unexpected condition. Contact ScienceLogic Support for assistance.
 - *Missing Base*. The Skylar One system cannot import this software until another software package has been imported. The dependency is for compression purposes. Check the log for a message stating which software package needs to be imported.
6. The update file or patch file is imported to Skylar One system and appears in the **System Updates** page.

NOTE: For details on the import process, go to the **System Updates** page, find the entry for the software you are interested in, go to its **Import Status** column, and click the log icon (🔗).

Staging the Update

After you import a software update to your Skylar One system, you must **stage** the software update. During staging, the Skylar One system copies the software update to each Skylar One appliance. Staging allows Skylar One to simultaneously apply the software changes to each Skylar One appliance, regardless of the speed of the connection to each Skylar One appliance. The Skylar One system stages updates per import. You can choose to automatically stage imports or manually stage import.

For easiest troubleshooting, ScienceLogic recommends that you manually stage imports.

NOTE: A staging mode called "enhanced file upload" enables pre-staging operations and file upload operations to happen in separate worker processes, which does not count file upload time toward the "staging wait time" setting. This mode is enabled by default, and ScienceLogic recommends that you keep it enabled due to the large amount of data that needs to be pushed for Skylar One updates. However, if you prefer to update this setting, you can make the following changes in the "master.system_settings_patcher" table:

- To change the number of file upload workers (12 by default), insert or update the following parameter:
`param="num_file_upload_workers" and value="number of workers"`
- To change the number of pre-staging workers (25 by default), insert or update the following parameter:
`param="pool_size" and value="number of workers"`
- To disable enhanced file upload mode, insert or update the following parameter:
`param="use_enhanced_file_upload" and value="0"`

The **Staging Status** column on the **System Updates** page can have one of the following statuses:

- *--*. No staging request is active and software has not been staged on any Skylar One appliances.
- *Scheduled*. The Skylar One system is aware of the staging request and is preparing for staging.
- *In Progress*. Staging is in progress but has not completed. The page displays the percentage complete as staging progresses.
- *Complete*. Staging has completed, and all appliances are ready to deploy the software.
- *Incomplete*. Staging has completed, and one or more appliances are ready to deploy the software.
- *Canceled*. User manually canceled the staging process.
- *Outdated*. The current update is not the latest or has already been installed.
- *Failed*. An unexpected error occurred in the staging process. Contact ScienceLogic Support.

NOTE: For details on the staging process, go to the **System Updates** page, find the entry for the software you are interested in, go to its **Staging Status** column, and click the log icon (☞).

After the software update is imported and staged, you can deploy the software.

Automatic Staging

To enable automatic staging:

1. In Skylar One, go to the **System Updates** page (System > Tools > Updates).
2. In the **System Updates** page, click the **[Import]** button.
3. In the **Import a new update** modal page, browse to the product update file and select it. If you select the **Auto Stage** button, the Skylar One system will begin staging as soon as the import is completed.
4. After import, in the **System Updates** page, the **Staging Status** column will display the number of ScienceLogic appliances that have been successfully staged compared to the total number of ScienceLogic appliances

To disable automatic staging:

1. In Skylar One, go to the **System Updates** page (System > Tools > Updates).
2. In the **System Updates** page, click the **[Import]** button.
3. In the **Import a new update** modal page, browse to the product update file and select it.
4. If you *do not select* the **Auto Stage** button, you must click the staging button (☞) after import is completed. You can do so at any time after import has completed.

Manually Staging an Update

You can manually stage a software update:

- If you imported an update but do not want to stage it immediately.
- If you add another ScienceLogic appliance to your Skylar One system and need to apply software updates.
- If staging failed on one or more ScienceLogic appliances.
- If you want to ensure that a previous staging process was successful.

When you manually stage a software update, Skylar One checks the status of the software updated on each ScienceLogic appliance. Skylar One then stages the software update **only to those Skylar One appliances that have not yet been staged** for this software update.

To manually stage a software update:

1. In Skylar One, go to the **System Updates** page (System > Tools > Updates).
2. Locate the software update you want to stage and click its staging icon (☞). The software update will be copied to each ScienceLogic appliance that has not yet been staged.

3. The *Staging Status* column will display the number of ScienceLogic appliances that have been successfully staged compared to the total number of ScienceLogic appliances.

Monitoring Staging

For Skylar One versions 8.12.0 and later, you can monitor the staging process:

1. Either go to the console of the Database Server or use SSH to access the Database Server.
2. Type the following at the shell prompt:

monitor_stage

You should see something like the following image:

Every 1.0s: bash -c monitor_stage

```
#####
##          Monitor Staging Process for DB
#####

System Update Vitals
=====

| Service Name | Status | RunTime | HealthCheck |
| em7_patch_manager | active(running) | 3days | Good |
| siloudate-spool | active(running) | 3days | Good |
| siloudate-pkgserver | active(running) | 3days | Good |
| mariadb | active(running) | 3days | Good |
| Disk Availability | - | - | Good |
=====

Staging Process Stats
=====
Staging Schedule ID      : 2
Staging Status           : Completed
Number of Staging processes : 0
Patch Hook Completion Summary : [Total = 201, Installed = 201, Failed = 0]
Staging Started at       : 2019-04-22 19:49:52
Staging Completed at     : 2019-04-22 20:52:42
Staging Run Time         : 01:02:50
Staging Summary          : [Eligible=202, Complete=202, Failed=0, TimedOut=0]
  \ Completed mid(s)    : [978,1,523,23,6,51,1031,46,14,563,955,518,537,1024,548,940,2,34,22,553,513,528,
533,39,541,7,24,56,64,582,573,588,109,73,598,606,88,99,103,72,633,622,80,613,856,114,599,113,578,60,92,617,122,638,
643,647,139,140,672,662,154,158,134,671,148,166,177,123,173,683,657,693,165,891,653,667,684,185,189,703,716,709,217
,210,200,199,227,195,241,727,216,736,237,732,723,228,708,744,737,757,246,1014,763,252,257,262,771,782,267,273,813,8
99,293,778,804,788,792,803,301,279,278,796,294,277,904,818,307,311,824,322,965,825,823,844,343,861,568,848,356,838,
875,361,330,366,371,626,334,860,339,344,862,887,378,377,387,391,558,392,402,920,916,930,416,410,433,944,906,905,753
,428,422,438,403,948,995,447,443,961,453,463,697,452,464,972,982,477,979,981,592,484,980,473,934,483,1000,764,498,1
015,502,999,508,491]

Press CTRL-C to exit
```

3. In the `monitor_stage` results, look for the following information:

- ***System Update Vitals***. Displays the current status of the services that are required for System Update.
- ***Staging Process Stats***. Displays status of staging on all Skylar One appliances.

Running the Pre-Upgrade Check

After importing and staging an update, you can run a pre-upgrade check before deploying. The pre-upgrade check will ensure that all criteria are met before deploying.

NOTE: If you are upgrading from Skylar One 8.14 or earlier, see [Running the Pre-Upgrade Check for Skylar One 8.14 or Earlier](#).

The pre-upgrade check examines the following:

- Is each Skylar One Appliance eligible to be updated?
- Are updates enabled on each Skylar One Appliance?
- Are any of the Skylar One Appliances running CentOS 5?
- Is this hostfile on each Skylar One Appliance correctly configured?
- Is each Data Collector and Message Collector in a Collector Group?
- Is there enough free space on the disk to perform the upgrade?
- Is the RPM database corrupted?
- Are the RPM packages corrupted?
- Does the patch hook directory have the correct owner assigned?
- Are the CRM templates on High Availability and Disaster Recovery systems out of date?
- Does /etc/init.d/mysql exist? (If it does not, it creates the file.)

NOTE: The pre-upgrade check skips any Skylar One appliances that have been deleted since the last upgrade.

Running the Pre-Upgrade Check

To run a pre-upgrade check:

1. Go to the **System Updates** page (System > Tools > Updates).
2. Find the upgrade that you want to deploy.
3. Click the purple checkmark at the end of the row. The pre-upgrade check will run.
4. If a pre-upgrade criterion fails, the **[Deploy]** button will be disabled for the selected row.
5. To view the output from the pre-upgrade check, click on the magnifying-glass icon (🔍) in the selected row.
6. If the pre-upgrade check finds a failure, see the list below for possible causes.
7. Fix all failures before deploying the update.

Potential Issues to Address

CentOS 5 Failure

CentOS 5 is no longer supported by System Update. If one or more Data Collectors are running CentOS5, the pre-upgrade check will fail. Contact your Customer Success Manager to determine how to upgrade your Data Collectors.

Collector Group Membership

This test checks that each Data Collector and Message Collector is a member of a Collector Group.

If a Data Collector or Message Collector is not a member of a Collector Group, the pre-upgrade test will define the appliance as "not eligible for patching".

To fix this error, add the Data Collector or Message Collector to a Collector Group.

Eligibility Failure

The most common reasons for eligibility failure are:

- The Skylar One appliance is not licensed or the license has expired
- The Skylar One appliance cannot be reached over the network
- The Data Collector has failed over
- The Skylar One appliance is not configured
- The Data Collector is waiting to be returned to service
- The Data Collector is not assigned to a Collector Group

Enabled Failure

By default, all Skylar One appliances are enabled for patching.

However, if you have used a command-line tool to exclude a Skylar One appliance from updates, the pre-upgrade check will fail. To fix this error, include the Skylar One appliance for updates.

Free Disk-Space Failure

This test checks the root partition and requires 1GB of free disk space. If the root partition does not have 1GB of free disk space, the pre-upgrade check will fail.

If the root partition does not have 1GB of free disk space, you must archive or delete files that are no longer required or add a new empty disk and resize the filesystem.

Host File Failure

This test validates the /etc/hosts file for the presence of an IPv6 entry for localhost, which is required by System Update.

If /etc/hosts does not include an IPv6 entry for localhost, the pre-upgrade test automatically adds the required entry.

Check the following in case of failure:

- The /etc/hosts file exists
- The /etc/hosts can be edited by root

Patch-Hook Ownership Failure

If the owner of the patch hook directory ((/var/lib/em7/patch_hook) is incorrect, the pre-upgrade test automatically fixes the ownership. However, if this error occurs, check for the following:

- The patch hook directory (/var/lib/em7/patch_hook) does not exist
- The s-em7-core user or the s-em7-core group does not exist

RPM Database Failure

If the RPM database fails the pre-upgrade test, the RPM database is corrupted.

To recover the RPM database:

1. Either go to the console of the Database Server or use SSH to access the Database Server. Log in with the credentials you defined when you installed the Database Server.
2. At the shell prompt, enter the following:

```
mkdir -p /tmp/rpm.bak
cp /var/lib/rpm/* /tmp/rpm.bak
rm -f /var/lib/rpm/__db*
rpm --rebuilddb -vv
rpm -q kernel
```

3. If the last command returns a value, you can delete the backup directory using the following command.

```
rm -Rf /tmp/rpm.bak
```

RPM Package Failure

If one or more RPM packages failed the pre-upgrade test, possible causes are:

- Packages are not staged, and hence some files are missing. This can be caused due to a failed staging or a timeout during staging. You can try to stage again. You can also [adjust the timeout for staging](#).
- Duplicate packages
- Conflicting packages
- Unmet dependencies

Duplicate Packages:

1. Either go to the console of the Database Server or use SSH to access the Database Server. Log in with the credentials you defined when you installed the Database Server.
2. At the shell prompt, enter the following command:

```
sudo package-cleanup --dupes
```

3. If there are duplicate packages, use the following command to remove them:

```
sudo package-cleanup --cleandupes --removenewestdupes
```

Conflicting Packages

1. Look for conflicting packages in the staging log
2. Verify that the package is a part of Skylar One ISO or patch bundle
3. If the package is not part of the Skylar One ISO or patch bundle, uninstall the package.

Unmet dependencies

You will need to reset the staging status of the appliance and stage it again. Contact ScienceLogic Customer Success for help in resetting the staging status.

Putting All Skylar One Appliances into Maintenance Mode

NOTE: ScienceLogic recommends that you perform these steps during a maintenance window.

Immediately before deploying a software update, ScienceLogic recommends that you put all Skylar One appliances in maintenance mode. This will prevent spurious error messages and events during the deployment.

To enable user maintenance mode for all the Skylar One appliances in your Skylar One system:

1. Go to the **Appliance Manager** page (System > Settings > Appliances). Note the list of Skylar One appliances in your system.
2. Go to the **Device Manager** page (Devices > Classic Devices, or Registry > Devices > Device Manager in the classic user interface) and select the checkbox for each Skylar One appliance in your Skylar One system. This includes both primary and secondary Database Servers.
3. In the **Select Action** drop-down list, select *Change User Maintenance Mode: Enabled without Collection*. This option puts the selected devices into user maintenance mode with collection disabled. The devices will remain in this state until you or another user disables user maintenance mode.
4. Click the **[Go]** button.

Deploying the Update

During deployment, avoid the following tasks:

- Running integrations and third-party applications that access the Skylar One database or manipulate data on Skylar One
- Running discovery sessions
- Running nightly discovery
- Bringing HA/DR out of maintenance mode
- Adding new Skylar One appliances

- Importing a new patch
- Adding Data Collectors to a Collector Group
- Removing Data Collectors from a Collector Group
- Rebalancing a Collector Group
- Killing processes related to patching and upgrading
- Run reporting jobs
- Unpausing the proc_mgr process

When you deploy an update, the update is installed on all nodes or appliances that have already been staged.

NOTE: Database Servers and Data Engines in your Skylar One stack are deployed first; other appliance types cannot be deployed until the Database Servers and Data Engines are finished deploying.

When you deploy an update, Skylar One checks to ensure that you have already deployed all required updates. If you have not, Skylar One will generate an error message specifying the updates you must deploy before continuing with the current update.

During deployment, the **Deployment Status** column on the **System Updates** page can have one of the following statuses:

- *--*. No deployment request is active, and software has not been deployed on any Skylar One appliances.
- *Scheduled*. The Skylar One system is aware of the deployment request and is preparing for deployment.
- *In Progress*. Deployment is in progress but has not completed.
- *Complete*. Deployment has completed, and all appliances are updated.
- *Incomplete*. Deployment has completed, and one or more, but not all, appliances are updated.
- *Canceled*. User manually canceled the deployment.
- *Outdated*. The current update is not the latest or has already been installed.
- *Failed*. An unexpected error occurred in the deployment process. Contact ScienceLogic Support.

To deploy a software update on your nodes or appliances:

1. Make sure that you have imported and staged the update file.
2. Go to the **System Updates** page (System > Tools > Updates).
3. In the **System Updates** page, find the software update you want to deploy. Click the lightning bolt icon (⚡) to deploy the software. If Skylar One is still staging the patch when you click the lightning-bolt icon (⚡), Skylar One will wait until staging has completed before deploying the updates to each ScienceLogic appliance.
4. The software update will be deployed to all appliances in your Skylar One system that have already been staged. If one or more appliances in your Skylar One system have been successfully staged, Skylar One will deploy the update to those appliances.

NOTE: For details on the deployment process, go to the **System Updates** page, find the entry for the software you are interested in, go to its **Deployment Status** column, and click the log icon ().

Troubleshooting System Update

You can use the **sysuptb** troubleshooting tool to determine issues with System Update and to generate diagnostic information about the update. You can also use the **phtb** tool to troubleshoot issue with the PhoneHome configuration.

These tools can be useful when System Update does not work as expected, or if you have issues with the PhoneHome configuration or with communication between appliances and the Database Server. These tools are available on all Skylar One appliances starting with Skylar One version 10.2.0, and the tools are backwards-compatible to Skylar One version 8.12.0.

Using the **sysuptb** Troubleshooting Tool

To use the **sysuptb** troubleshooting tool:

1. Either go to the console of any Skylar One appliance or use SSH to access the appliance
2. Enter the following at the shell prompt:

```
sudo sysuptb -h
```

3. For more information about each argument, enter the following at the shell prompt:

```
sudo sysuptb <argument> -h
```

Available Commands

- The following command executes all troubleshooting tests for System Update:

```
sudo sysuptb all <optionally -x name_of_test_to_exclude>
```

TIP: To learn more about a test run, use this command: `sudo sysuptb help <test-name>`

- Example:

```
sudo sysuptb all

Executing filestore tests
912 / 912 [-----]
-] 100.00% 14 p/s
Filestore test summary: [Total: 912, Intact: 912, Incomplete: 0,
Corrupt: 0]
Executing test for deleted appliances in patch history
No deleted appliances were found in the patch history
Executing test for invalid file id in patch schedules
No patch schedules were found to have invalid file id
Executing test for RPM database corruption
RPM database is intact
Executing test to check if filestore is empty
Filestore has 1026 files
Executing test for deactivating services
Service test summary: [Total: 1, Active: 1, Inactive: 0, Healed:
0, Skipped: 0, Failed (to heal): 0]
Executing test for free disk space
Free disk space test summary: [Total: 2, Pass: 2, Failed: 0]
Executing test for service errors
Service error test summary: [Total: 2, Without Errors: 2,
Restarted: 0, Failed: 0]
Executing hosts file check for IPV6 entry (::1) for localhost
An entry for ::1 is already present in the hosts file
Proxy is not configured for yum.
Executing test for hung yum process
No yum processes found
Yum process summary: [Total: 0, Hung: 0]
```

- The following command searches the logs for errors that match a service name and restarts services if any errors are found.

```
sudo sysuptb check-service-error <optionally -s name_of_service>
```

If you do not provide the name of a service, the command searches the logs for errors for silouupdate-pkgserver.service and silouupdate-spool.service.

- Example:

```
sudo sysuptb check-service-error

Executing test for service errors
Service error test summary: [Total: 2, Without Errors: 2,
Restarted: 0, Failed: 0]
```

- The following command removes deleted Skylar One appliances from the history of system updates so that they Skylar One does not search for them during update.

```
sudo sysuptb clear-mids
```

- Example:

```
sudo sysuptb clear-mids

Executing test for deleted appliances in patch history
No deleted appliances were found in the patch history
```

- The following command cancels all schedule updates that include an invalid ID for the patch file.

```
sudo sysuptb clear-schedule
```

- Example:

```
sudo sysuptb clear-schedule

Executing test for invalid file id in patch schedules
No patch schedules were found to have invalid file id
```

- The following command checks the filestore of downloaded packages for corrupt files and marks the corrupt files as incomplete.

```
sudo sysuptb filestore
```

- Example:

```
sudo sysuptb filestore

Executing filestore tests
912 / 912 [-----]
-----] 100.00% 14 p/s
Filestore test summary: [Total: 912, Intact: 912, Incomplete: 0,
Corrupt: 0]
```

- The following command checks the file system for available free space.

```
sudo sysuptb free-space <optionally, -d path_for_drive = minimum_size>
```

If you do not provide the path and minimum size of the directory, the command examines /var to make sure it has 300MB of free space and / to make sure it has 1GB of free space.

- Example:

```
sudo sysuptb free-space --disk /var=300MB
Executing test for free disk space
Free disk space test summary: [Total: 1, Pass: 1, Failed: 0]
```

- The command checks for update services that are stuck in a deactivating state and then heals them.

```
sudo sysuptb heal-service <optionally -s service_name>
```

If you do not specify a service, the command examines the service silouupdate-manager.service. Starting with Skylar One 12.1.0, silouupdate-manager.service replaced em7_patch_manager.service.

- Example:

```
sudo sysuptb heal-service
Executing test for deactivating services
Service test summary: [Total: 1, Active: 1, Inactive: 0, Healed: 0, Skipped: 0, Failed (to heal): 0]
```

- The following command checks the /etc/hosts file for an entry for IPv6 for the current server (like a loopback address). If no entry exists, the command adds ::1 to the /etc/hosts file.

```
sudo sysuptb hosts
```

- Example:

```
sudo sysuptb hosts
Executing hosts file check for IPV6 entry (>::1) for localhost
An entry for ::1 is already present in the hosts file
```

- The following command check is the filestore that holds the upgrade packages is empty.

```
sudo sysuptb is-filestore-empty
```

- Example:

```
sudo sysuptb is-filestore-empty
Executing test to check if filestore is empty
Filestore has 1026 files
```

- The following command checks the RPM database on /var/lib/rpm for corruption. If the command detects corruption, the output includes steps for remediation.

```
sudo sysuptb rpmdb
```

- Example:

```
sudo sysuptb rpmdb  
Executing test for RPM database corruption  
RPM database is intact
```

- The following command checks for a yum process which is hung.

```
sudo sysuptb yum-proc <optionally, -t timeout_in_minutes>
```

If you do not specify a running time, in minutes, the command searches for yum processes that have been running for more than 120 minutes.

- Example:

```
sudo sysuptb yum-proc  
Executing test for hung yum process  
No yum processes found  
Yum process summary: [Total: 0, Hung: 0]
```

- The following command checks if yum is configured with proxy. If so, the command removes the proxy configuration.

```
sudo sysuptb yum-proxy
```

- Example:

```
sudo sysuptb yum-proxy  
Proxy is not configured for yum.
```

Using the **phtb** Troubleshooting Tool

To use the **phtb** troubleshooting tool:

1. Either go to the console of a Skylar One appliance using PhoneHome communication or use SSH to access the appliance.
2. Enter the following at the shell prompt:

```
sudo phtb -h
```

3. For more information about each argument, enter the following at the shell prompt:

```
sudo phtb <argument> -h
```

TIP: To learn more about a test run, use this command: `sudo phtb help <test-name>`

Available Commands

- The following command checks destinations for SSH connectivity issues:

```
sudo phtb destination
```

- The following command checks the target host for SSH connectivity issues:

```
sudo phtb probe-host
```

- The following command checks connectivity to the proxy host, if configured:

```
sudo phtb proxy
```

Monitoring Deployment

To monitor the deployment process:

1. Either go to the console of the Database Server or use SSH to access the Database Server.
2. Enter the following command at the shell prompt:

```
monitor_deploy
```

3. You should see something like the following figure:

```

X root@dw-sm-db1-patch-10-2-12-113:/tmp
Every 1.0s: bash -c monitor_deploy

#####
##      Monitor Deploy Process for DB
#####

System Update Vitals
=====
| Service Name | Status | RunTime | HealthCheck |
| em7_patch_manager | active(running) | 2h18min | Good |
| silouupdate-spool | active(running) | 2h28min | Good |
| silouupdate-pkgserver | active(running) | 2h27min | Good |
| mariadb | active(running) | 21h | Good |
=====

Deployment Process Stats
=====
Deployment Schedule ID      : 4
Deployment Status           : Completed
No. Of system_patcher processes : 0
No. Of deploy-module processes : 0
Deployment Started at      : 2019-04-25 16:06:30
Deployment Completed at     : 2019-04-25 16:18:54
Deployment Run Time         : 00:12:24
Version Updated Count       : 4 [8.13.0.smailappan_EM_26803_stop_building_e15_throwawayr428]
Deployment Summary          : [Eligible=4, Complete=2, Failed=1, TimedOut=1]
  \_ Completed mid(s)      : [2,1]
  \_ Failed mid(s)         : [6]
  \_ TimedOut mid(s)       : [5]

Module Level Status
=====
proc_mngr Status           : Unpaused
deploy-module Status        : Completed
deploy-module Completed at  : 2019-04-25 16:16:35

Press CTRL-C to exit

```

- **System Update Vitals.** Displays the current status of the services that are required for System Update.
- **Deployment Process Stats.** Displays status of deployment on all Skylar One appliances.
- **Module Level Status.** Displays the status of the three deployment steps.

Running the Post-Upgrade Recovery Script for 12.5.x Patching Issues

If you attempt to upgrade from 12.1.x or 12.2.x to 12.5.x without first [running the required pre-import script](#), your upgrade will ultimately fail. If this occurs, you can run a post-upgrade script to update the silouupdate RPM file on all of your Skylar One appliances.

NOTE: If you are not upgrading from 12.1.x or 12.2.x to 12.5.x, or you ran the pre-import script and your upgrade was successful, you can skip this section.

To download and run the post-upgrade script:

1. Go to the **Release Files** page for xxx on the ScienceLogic Support site and download the `12-5-1-recovery.pyc` script: <https://support.scienceologic.com/s/contentdocument/069VL00000Xa57tYAB>.
2. On the active or primary Database Server or Data Engine, verify the checksum:

```
sha512sum 12-5-1-recovery.pyc
```

```
3cf3343e96d5d8eade8bf9a0cbf8ba7a1641c3fbaee39234cf7c6b72e239c822 12-5-1-recovery.pyc
```

3. Set the correct permissions on the script:

```
chmod 755 12-5-1-recovery.pyc
```

4. Test that you are able to execute the script by running usage help:

```
/opt/em7/bin/python3 12-5-1-recovery.pyc -h
```

5. Execute the script. Module selection using the `-m` flag to specify one or more appliance IDs is optional. If you do not select any module, it will automatically select all available Skylar One appliances in the stack:

```
/opt/em7/bin/python3 12-5-1-recovery.pyc [-m <mid_1> <mid_2> ... <mid_n>] [-v]
```

NOTE: If you have a large stack, ScienceLogic recommends running the script as a background task as a root user, as shown below:

```
touch /tmp/12-5-1-recovery.log
```

```
chmod 644 /tmp/12-5-1-recovery.log
```

```
nohup /opt/em7/bin/python3 12-5-1-recovery.pyc [-m <mid_1> <mid_2> ... <mid_n>] [-v] 2>&1 >/tmp/fix-patch-hook.log &
```

You can then exit from the shell as the root user and can follow the progress of the script by tailing the log file: `tail -f /tmp/12-5-1-recovery.log`

Installing Additional RPMs on a Skylar One Appliance

For certain patch releases, ScienceLogic might require additional RPMs to be installed on specific appliance types. If an RPM install is required, the release notes will indicate the additional RPMs to install on each specific appliance type.

To install additional RPMs on an appliance, perform the following steps:

1. Download the RPM files provided by ScienceLogic to your local machine.
2. Log in as root at the appliance console.
3. Copy each of the downloaded RPM files to the appliance. To copy the downloaded files, perform the following command as root at the console of the appliance:

```
scp <username-on-local-machine>@<ip-address-of-your-local-machine>:<full-path-to-rpm-on-your-local-machine> <full-path-on-appliance-to-copy-to>
```

4. Use the following command to run the RPM installer for each of the RPM files:

```
rpm -U <name-of-rpm-file>.rpm
```

5. If you have not yet done so, apply the latest patch to your Skylar One system.

Remove Skylar One Appliances from Maintenance Mode

To disable user maintenance mode for all the Skylar One appliances in your Skylar One system:

1. Go to the **Appliance Manager** page. Note the list of Skylar One appliances in your system.
2. Go to the **Device Manager** page (Devices > Classic Devices, or Registry > Devices > Device Manager in the classic user interface) and select the checkbox for each Skylar One appliance in your Skylar One system.
3. In the **Select Action** drop-down list, select *Change User Maintenance Mode: Disabled*. This option disables user maintenance mode for the selected devices.
4. Click the **[Go]** button.

CAUTION: Refer to the release notes for your current release to determine if you must *upgrade MariaDB* after upgrading.

CAUTION: Also refer to the release notes for your current release to determine if you must *reboot all Skylar One appliances* after upgrading.

Updating Skylar One Extended Architecture

New installations of Skylar One Extended Architecture are available only on SaaS deployments.

For existing on-premises deployments of Skylar One Extended Architecture, see the section on [Upgrading Skylar One Extended Architecture](#).

Automatically Upgrading MariaDB

Most Skylar One updates require you to upgrade MariaDB after you update Skylar One. Refer to the release notes for your patch release to determine if you must upgrade MariaDB.

Skylar One version 12.2.1.1 and later use the command `silouupdate upgrade-mariadb` to upgrade the MariaDB server. This command also updates MariaDB-client, MariaDB-common, and MariaDB-shared RPMs in addition to the MariaDB Server RPM.

NOTE: Versions of Skylar One before 12.2.1.1 use the now-deprecated `module_upgrade_mariadb` script to upgrade MariaDB rather than the command `silouupdate upgrade-mariadb` that is described in this section.

WARNING: If you are upgrading from a version of Skylar One prior to 12.2.3, then after upgrading Skylar One, you must also upgrade MariaDB 10.4.x to version 10.6.x. Failure to perform this MariaDB upgrade can cause major functionality issues in Skylar One. For the specific MariaDB version that is required, see the release notes for the Skylar One version to which you are upgrading.

CAUTION: You should store all custom configuration settings for each MariaDB database in the file `/etc/siteconfig/mysql.siteconfig`. If you have added custom settings to the file `/etc/my.cnf.d/silo_mysql.cnf`, those changes will be overwritten each time you upgrade MariaDB. Before upgrading, copy any custom settings to the file `/etc/siteconfig/mysql.siteconfig`. Skylar One will save these custom settings and apply them after you upgrade MariaDB.

TIP: To reduce spurious events, you can put the Database Server in maintenance mode while you upgrade MariaDB. For details, see the chapter on [Putting the Database Server into Maintenance Mode](#)

The `silouupdate upgrade-mariadb` command:

- Upgrades the following Skylar One appliances:

- All Database Servers
- All-In-One Appliances
- Data Collectors
- Message Collectors
- Upgrades High Availability (HA) and Disaster Recovery (DR) systems
- Includes a "test only" option before executing upgrade
- Enforces upgrading the primary Database Server before upgrading secondary Database Server and the Data Collectors.
- Will skip Skylar One appliances that have already been updated
- Logs entire sequence of commands and output for later analysis
- Stores log files in /data/logs/module_upgrade_mariadb.log and /data/logs/.upgrade_mariadb.log
- Checks for differences between current configuration and version you are about to install and spawns an alert. To skip this check, use the -s -s option
- Starting with Skylar One version 12.3.7, lets you enable or disable GPG check for RPM packages and run the yum transactions accordingly

To upgrade MariaDB in newer versions of Skylar One, perform the following steps:

1. Either go to the console of the Database Server or use SSH to access the Database Server.
2. At the shell prompt, enter the following command:

```
silouupdate upgrade-mariadb
```

TIP: To upgrade all modules, you can use the command `silouupdate upgrade-mariadb -m all`.

3. To see all the options for the `silouupdate upgrade-mariadb` command, enter the following command at the shell prompt:

```
silouupdate upgrade-mariadb -h
```

Manually Upgrading MariaDB

NOTE: Refer to the release notes for your current release to determine if you must upgrade MariaDB. Not every Skylar One update requires an upgrade of MariaDB.

CAUTION: ScienceLogic strongly recommends that you upgrade MariaDB using the script described in [Automatically Upgrading MariaDB with a Script](#).

TIP: To reduce spurious events, you can put the Database Server in maintenance mode while you upgrade MariaDB. For details, see the chapter on [Putting the Database Server into Maintenance Mode](#)

If you prefer to upgrade MariaDB manually, the following sections describe how to upgrade the MariaDB server for different Skylar One appliance types and architectures.

When you update MariaDB, you must update the following Skylar One appliances:

- All Database Servers
- All-In-One Appliances
- Data Collectors
- Message Collectors

CAUTION: You should store all custom configuration settings for each MariaDB database in the file `/etc/siteconfig/mysql.siteconfig`. If you have added custom settings to the file `/etc/my.cnf.d/silo_mysql.cnf`, those changes will be overwritten each time you upgrade MariaDB. Before upgrading, copy any custom settings to the file `/etc/siteconfig/mysql.siteconfig`. Skylar One will save these custom settings and apply them after you upgrade MariaDB.

Download RPMs to Skylar One Appliances

Before upgrading MariaDB, you must copy the RPMs from the primary Database Server to the Database Servers, All-In-One Appliances, Data Collectors, and Message Collectors in your Skylar One system. To do this.

1. Open an SSH session to the Database Server.
2. To download the latest RPMs from the Database Server, enter the following at the shell prompt:

For Skylar One version 10.1.6 and higher, download all of the packages listed when you enter the command:

```
cat /opt/em7/share/db_packages

wget --output-document /tmp/MariaDB-server.rpm
http://localhost:10080/<mariadb-server-pkg-from-db_packages>
wget --output-document /tmp/galera-4.rpm
http://localhost:10080/<galera-4-pkg-from-db_packages>
wget --output-document /tmp/socat.rpm http://localhost:10080/<socat-
pkg-from-db_packages>
```

3. Verify if the downloaded packages are valid (not corrupt or incomplete downloads) by entering the following commands:

```
rpm -qip /tmp/MariaDB-server.rpm  
rpm -qip /tmp/galera-4.rpm  
rpm -qip /tmp/socat.rpm
```

If any errors are reported, try restarting silouupdate-pkgserver, using the following command, and retry downloading and verifying the RPM files again.

```
systemctl restart silouupdate-pkgserver
```

4. Use SCP or another secure copy program to copy these files to the /tmp directory on each Database Server, All-In-One Appliance, Data Collector, and Message Collector:
 - MariaDB-server.rpm
 - galera-4.rpm

CAUTION: To conserve disk space, ScienceLogic recommends you delete the RPMs from the /tmp directory the on the Database Servers, All-In-One Appliances, Data Collectors, and Message Collectors in your Skylar One system after you successfully upgrade MariaDB.

Manually Upgrade Two Database Servers Configured for High Availability or Disaster Recovery

To upgrade a High Availability or Disaster Recovery cluster, perform the following steps:

WARNING: The system will be unavailable when performing these steps.

Step 1: On the Secondary Database Server

You must put the secondary Database Server in maintenance mode. To do this:

1. Open an SSH session to the Database Server.
2. At the shell prompt, assume root privileges:

```
sudo -s
```

3. When prompted, enter the administrator password.

4. At the shell prompt, enter the following command:

```
coro_config
```

The following menu appears:

- 1) Enable Maintenance
- 2) Option Disabled
- 3) Promote DRBD
- 4) Stop Pacemaker
- 5) Resource Status
- 6) Quit

5. Enter "1".

Step 2: On the Primary Database Server

1. To determine the current installed version of the RPMs, enter the following command:

```
sudo rpm -qa ^MariaDB-*
```

2. To stop Skylar One and MariaDB, enter the following commands at the shell prompt:

```
sudo systemctl stop em7
sudo systemctl stop mariadb.service
```

3. To stop the MySQL resource, enter the following command:

```
sudo crm resource stop mysql
```

4. To save the current enabled state for mariadb.service, enter the following command:

```
export MSRV='sudo systemctl is-enabled mariadb.service'
```

5. Check the version of MariaDB-server that you are running.

```
rpm -q MariaDB-server
```

WARNING: You **must** follow the steps below that correspond to your version of MariaDB. Step 6 is specific to MariaDB-server version 10.1.x, while Step 7 is specific to MariaDB-server versions 10.4.12 and higher.

6. **MariaDB-server version 10.1.x.** If you are running **MariaDB-server version 10.1.x:**

WARNING: Do these steps in order. Doing the steps in any other order will result in unintended consequences.

- a. Remove MariaDB-server by using the following commands:

```
sudo rpm --nodeps -ev MariaDB-server
```

- b. Replace the Galera package and install the new MariaDB-server package by using the following commands:

```
sudo yum --disablerepo=* swap -- remove galera -- install  
/tmp/galera-4.rpm  
sudo yum --disablerepo=* install /tmp/MariaDB-server.rpm
```

7. **MariaDB-server version 10.4.12 and higher.** If you are running **MariaDB-server version 10.4.12 or higher**, upgrade the MariaDB-server package and dependent packages (galera-4 and socat) by using the following commands:

```
sudo yum --disablerepo=* install /tmp/galera-4.rpm /tmp/socat.rpm  
sudo yum --disablerepo=* upgrade /tmp/MariaDB-server.rpm
```

8. To remove incompatible backup packages, enter the following command:

```
sudo yum remove percona-xtrabackup
```

NOTE: If the "yum remove" command fails, it means that the package does not exist on the Skylar One appliance. You can ignore the error message.

9. To regenerate the configuration file for MariaDB, enter the following command:

```
sudo /opt/em7/share/scripts/generate-my-conf.py -f -o  
/etc/my.cnf.d/silo_mysql.cnf
```

10. To re-start MariaDB, enter the following command:

```
sudo systemctl daemon-reload  
sudo systemctl start mariadb
```

11. To restart the MySQL resource, enter the following command:

```
sudo crm resource start mysql
```

12. To restore the mariadb.service enabled state, enter the following command:

```
sudo systemctl ${MSRV}:-1 mariadb.service
```

13. To upgrade the internal configuration for the database, enter the following:

```
sudo mysql_upgrade -u root -p
```

14. To restart the em7 service, enter the following commands:

```
sudo systemctl start em7  
sudo rpm -qa ^MariaDB-*
```

Step 3: On the Secondary Database Server

1. Determine the current installed version of the RPMs using the following command:

```
sudo rpm -qa ^MariaDB-*
```

2. To save the current enabled state for mariadb.service, enter the following:

```
export MSRV=`sudo systemctl is-enabled mariadb.service`
```

3. Check the version of MariaDB-server that you are running.

```
rpm -q MariaDB-server
```

WARNING: You **must** follow the steps that correspond to your version of MariaDB. Step 4 is specific to MariaDB-server version 10.1.x, while Step 5 is specific to MariaDB-server versions 10.4.12 and higher.

4. **MariaDB-server 10.1.x:** If you are running *MariaDB-server version 10.1.x*:

WARNING: Do these steps in order. Doing the steps in any other order will result in unintended consequences.

- a. Remove MariaDB-server by using the following command:

```
sudo rpm --nodeps -ev MariaDB-server
```

- b. Replace the Galera package and install the new MariaDB-server package by using the following commands:

```
sudo yum --disablerepo=* swap -- remove galera -- install
/tmp/galera-4.rpm
sudo yum --disablerepo=* install /tmp/MariaDB-server.rpm
```

5. **MariaDB-server 10.4.12 or higher:** If you are running *MariaDB-server version 10.4.12 or higher*, upgrade the MariaDB-server package and dependent packages (galera-4 and socat) by using the following commands:

```
sudo yum --disablerepo=* install /tmp/galera-4.rpm /tmp/socat.rpm
sudo yum --disablerepo=* upgrade /tmp/MariaDB-server.rpm
```

6. To remove incompatible backup packages, enter the following command:

```
sudo yum remove percona-xtrabackup
```

NOTE: If the "yum remove" command fails, it means that the package does not exist on the Skylar One appliance. You can ignore the error message.

7. To regenerate the configuration file for MariaDB, enter the following command:

```
sudo /opt/em7/share/scripts/generate-my-conf.py -f -o  
/etc/my.cnf.d/silo_mysql.cnf
```

8. To restore the mariadb.service enabled state, enter the following command:

```
sudo systemctl ${MSRV}:-1 mariadb.service
```

9. To take the secondary Database Server out of maintenance mode, enter the following command at the shell prompt:

```
sudo -s
```

10. When prompted, enter the administrator password.

11. At the shell prompt, enter the following command:

```
coro_config
```

The following prompt appears:

- 1) Disable Maintenance
- 2) Option Disabled
- 3) Promote DRBD
- 4) Stop Pacemaker
- 5) Resource Status
- 6) Quit

12. Enter "1".

Manually Upgrade Three Database Servers Configured for High Availability and Disaster Recovery

To upgrade a High Availability/Disaster Recovery cluster, perform the following steps:

WARNING: The system will be unavailable when performing these steps.

Step 1: On the Secondary Database Server

You must put the secondary Database Server in maintenance mode. To do this:

1. Open an SSH session to the Database Server.
2. At the shell prompt, assume root privileges:

```
sudo -s
```

3. When prompted, enter the administrator password.
4. At the shell prompt, enter the following command:

```
coro_config
```

5. The following prompt appears:

```
1) Disable Maintenance  
2) Option Disabled  
3) Promote DRBD  
4) Stop Pacemaker  
5) Resource Status  
6) Quit
```

6. Enter "1".

Step 2: On the Primary Database Server

1. Determine the current installed version of the RPMs:

```
sudo rpm -qa ^MariaDB-*
```

2. Stop Skylar One and MariaDB using the following commands:

```
sudo systemctl stop em7  
sudo systemctl stop mariadb.service
```

3. Stop the MySQL resource:

```
sudo crm resource stop mysql
```

4. Save the current enabled state for the mariadb.service:

```
export MSRV='sudo systemctl is-enabled mariadb.service'
```

5. Check the version of MariaDB-server that you are running:

```
rpm -q MariaDB-server
```

WARNING: You **must** follow the steps that correspond to your version of MariaDB.

6. **MariaDB-server version 10.1.x.** If you are running MariaDB-server version 10.1.x:

WARNING: Do these steps in order. Doing the steps in any other order will result in unintended consequences.

- a. Remove MariaDB-server by using the following command:

```
sudo rpm --nodeps -ev MariaDB-server
```

- b. Replace the Galera package and install the new MariaDB-server package by using the following commands:

```
sudo yum --disablerepo=* swap -- remove galera -- install  
/tmp/galera-4.rpm  
sudo yum --disablerepo=* install /tmp/MariaDB-server.rpm
```

7. **MariaDB-server version 10.4.12 and higher.** If you are running MariaDB-version 10.4.12 or higher, upgrade the MariaDB-server package and dependent packages (galera-4 and socat) by using the following commands:

```
sudo yum --disablerepo=* install /tmp/galera-4.rpm /tmp/socat.rpm  
sudo yum --disablerepo=* upgrade /tmp/MariaDB-server.rpm
```

8. Remove incompatible backup packages:

```
sudo yum remove percona-xtrabackup
```

NOTE: If the "yum remove" command fails, it means that the package does not exist on the Skylar One appliance. You can ignore the error message.

9. Regenerate the configuration file for MariaDB:

```
sudo /opt/em7/share/scripts/generate-my-conf.py -o -f  
/etc/my.cnf.d/silo_mysql.cnf
```

10. Restart MariaDB:

```
sudo systemctl daemon-reload  
sudo systemctl start mariadb
```

11. Restart the MySQL resource:

```
sudo crm resource start mysql
```

12. Restore the mariadb.service enabled state:

```
sudo systemctl ${MSRV}:-1 mariadb.service
```

13. Upgrade the internal configuration for the database:

```
sudo mysql_upgrade -u root -p
```

14. Restart the em7 service:

```
sudo systemctl start em7
```

Step 3: On the Secondary Database Server

1. Save the current enabled state for the mariadb.service:

```
export MSRV=`sudo systemctl is-enabled mariadb.service`
```

2. Check the version of MariaDB-server that you are running:

```
rpm -q MariaDB-server
```

WARNING: You **must** follow the steps that correspond to your version of MariaDB.

3. **MariaDB-server version 10.1.x.** If you are running **MariaDB-server version 10.1.x**:

WARNING: Do these steps in order. Doing the steps in any other order will result in unintended consequences.

- a. Remove MariaDB-server:

```
sudo rpm --nodeps -ev MariaDB-server
```

- b. Replace the Galera package and install the new MariaDB-server package by using the following commands:

```
sudo yum --disablerepo=* swap -- remove galera -- install  
/tmp/galera-4.rpm  
sudo yum --disablerepo=* install /tmp/MariaDB-server.rpm
```

4. **MariaDB-server version 10.4.12 and higher.** If you are running MariaDB-version 10.4.12 or higher, upgrade the MariaDB-server package and dependent packages (galera-4 and socat) by using the following commands:

```
sudo yum --disablerepo=* install /tmp/galera-4.rpm /tmp/socat.rpm  
sudo yum --disablerepo=* upgrade /tmp/MariaDB-server.rpm
```

5. Remove incompatible backup packages:

```
sudo yum remove percona-xtrabackup
```

NOTE: If the "yum remove" command fails, it means that the package does not exist on the Skylar One appliance. You can ignore the error message.

6. Regenerate the configuration file for MariaDB:

```
sudo /opt/em7/share/scripts/generate-my-conf.py -o -f  
/etc/my.cnf.d/silo_mysql.cnf
```

7. Restore the mariadb.service enabled state:

```
sudo systemctl ${MSRV}:-1 mariadb.service
```

8. Assume root privileges:

```
sudo -s
```

9. When prompted, enter the administrator password.

10. At the shell prompt, enter the following command:

```
coro_config
```

11. The following prompt appears:

- 1) Disable Maintenance
- 2) Option Disabled
- 3) Promote DRBD
- 4) Stop Pacemaker
- 5) Resource Status
- 6) Quit

12. Enter "1".

Step 4: On the Disaster Recovery Database Server

1. Open an SSH session to the Disaster Recovery Database Server.

2. Assume root privileges:

```
sudo -s
```

3. When prompted, enter the administrator password.

4. Save the current enabled state for the mariadb.service:

```
export MSRV=`sudo systemctl is-enabled mariadb.service`
```

5. Check the version of MariaDB-server that you are running:

```
rpm -q MariaDB-server
```

WARNING: You **must** follow the steps that correspond to your version of MariaDB.

6. **MariaDB-server version 10.1.x.** If you are running **MariaDB-server version 10.1.x**:

WARNING: Do these steps in order. Doing the steps in any other order will result in unintended consequences.

a. Remove MariaDB-server:

```
sudo rpm --nodeps -ev MariaDB-server
```

b. Replace the Galera package and install the new MariaDB-server package by using the following commands:

```
sudo yum --disablerepo=* swap -- remove galera -- install  
/tmp/galera-4.rpm  
sudo yum --disablerepo=* install /tmp/MariaDB-server.rpm
```

7. **MariaDB-server version 10.4.12 and higher.** If you are running MariaDB-version 10.4.12 or higher, upgrade the MariaDB-server package and dependent packages (galera-4 and socat) by using the following commands:

```
sudo yum --disablerepo=* install /tmp/galera-4.rpm /tmp/socat.rpm  
sudo yum --disablerepo=* upgrade /tmp/MariaDB-server.rpm
```

8. Remove incompatible backup packages:

```
sudo yum remove percona-xtrabackup
```

NOTE: If the "yum remove" command fails, it means that the package does not exist on the Skylar One appliance. You can ignore the error message.

9. Regenerate the configuration file for MariaDB:

```
sudo /opt/em7/share/scripts/generate-my-conf.py -o -f  
/etc/my.cnf.d/silo_mysql.cnf
```

10. Restore the mariadb.service enabled state:

```
sudo systemctl ${MSRV}:-1 mariadb.service
```

Manually Upgrading Standalone Database Servers, All-In-One Appliances, Data Collectors, and Message Collectors

To upgrade MariaDB on one or more Database Servers that are not configured for high availability or disaster recovery, a single All-In-One Appliance, one or more Data Collectors, or one or more Message Collectors, perform the following steps:

WARNING: The Database Server, All-In-One Appliance, Data Collector, or Message Collector will be unavailable when performing these steps.

1. Go to the console or open an SSH session to the Skylar One appliance.
2. Stop Skylar One and mariadb:

```
sudo systemctl stop em7
sudo systemctl stop mariadb.service
```

3. Save the current enabled state for the mariadb.service:

```
export MSRV='sudo systemctl is-enabled mariadb.service'
```

4. Check the version of MariaDB-server that you are running:

```
rpm -q MariaDB-server
```

WARNING: You **must** follow the steps that correspond to your version of MariaDB.

5. **MariaDB-server version 10.1.x.** If you are running **MariaDB-server version 10.1.x**:

WARNING: Do these steps in order. Doing the steps in any other order will result in unintended consequences.

- a. Remove MariaDB-server:

```
sudo rpm --nodeps -ev MariaDB-server
```

- b. Replace the Galera package and install the new MariaDB-server package by using the following commands:

```
sudo yum --disablerepo=* swap -- remove galera -- install
/tmp/galera-4.rpm
sudo yum --disablerepo=* install /tmp/MariaDB-server.rpm
```

6. **MariaDB-server version 10.4.12 and higher.** If you are running MariaDB-version 10.4.12 or higher, upgrade the MariaDB-server package and dependent packages (galera-4 and socat) by using the following commands:

```
sudo yum --disablerepo=* install /tmp/galera-4.rpm /tmp/socat.rpm  
sudo yum --disablerepo=* upgrade /tmp/MariaDB-server.rpm
```

7. Remove incompatible backup packages:

```
sudo yum remove percona-xtrabackup
```

NOTE: If the "yum remove" command fails, it means that the package does not exist on the Skylar One appliance. You can ignore the error message.

8. Regenerate the configuration file for MariaDB:

```
sudo /opt/em7/share/scripts/generate-my-conf.py -o -f  
/etc/my.cnf.d/silo_mysql.cnf
```

9. Restart MariaDB:

```
sudo systemctl daemon-reload  
sudo systemctl start mariadb
```

10. Restore the mariadb.service enabled state:

```
sudo systemctl ${MSRV}:-1 mariadb.service
```

11. Upgrade the internal configuration for the database:

```
sudo mysql_upgrade -u root -p
```

12. Restart the em7 service:

```
sudo systemctl start em7
```

13. Repeat all the steps in this section on each non-HA/DR Database Server, All-In-One Appliance, Data Collector, and Message Collector.

Rebooting Appliances in the Skylar One Distributed Stack

NOTE: Refer to the release notes for your current release to determine if you must reboot all Skylar One appliances. Not every Skylar One update requires rebooting.

When an upgrade requires a reboot, use the steps listed in this section to reboot all Skylar One appliances in the Distributed stack.

Rebooting the Administration Portal

You can reboot Administration Portals either from the user interface or from the command line.

Rebooting Multiple Administration Portals

If your Skylar One system includes multiple Administration Portals, you can remotely reboot Administration Portals from another Administration Portal. To do so:

1. Go to the **Appliance Manager** page (System > Settings > Appliances).
2. Select the checkboxes for the Skylar One appliances you want to reboot.
3. In the **[Select Action]** menu, select **Reboot** and click the **[Go]** button.
4. Click the **[OK]** button when the "Are you sure you want to reboot the selected appliances?" message is displayed.
5. During the reboot, the user interface for the affected Administration Portal unavailable.
6. When the reboot has completed, the **Audit Logs** page (System > Monitor > Audit Logs) will include an entry for each appliance that was rebooted.

Rebooting a Single Administration Portal

If your Skylar One system include only a single Administration Portal, perform the following steps to reboot that Administration Portal:

1. Either go to the console of the Database Server or use SSH to access the Database Server.
2. Log in as **em7admin** with the appropriate password.
3. At the shell prompt, execute the following:

```
python -m silo_common.admin_toolbox <appliance_ID> "/usr/bin/sudo  
/usr/sbin/shutdown -r +1"
```

where:

- *appliance_ID* is the appliance ID for the Data Collector, Message Collector, or Administration Portal.

Rebooting Data Collectors and Message Collectors

You can reboot Data Collectors and Message Collectors either from the user interface or from the command line.

Rebooting Data Collectors and Message Collectors from the Appliance Manager page

From the Skylar One user interface, perform the following steps to reboot a Data Collector or a Message Collector:

1. Go to the **Appliance Manager** page (Appliance Manager).
2. Select the checkbox for each Skylar One appliance you want to reboot.
3. In the **[Select Action]** menu, select **Reboot** and click the **[Go]** button.
4. Click the **[OK]** button when the "Are you sure you want to reboot the selected appliances?" message is displayed.
5. During the reboot, go to the **System Logs** page (System > Monitor > System Logs). You should see this message:

```
Major: Could not connect to module (5) database USING SSL=TRUE:  
Error attempting to connect to database with SSL enabled True: (2003,  
'Can't connect to MySQL server on '10.2.12.77' (113 "No route to  
host")')
```

6. When the reboot has completed, the **Audit Logs** page (System > Monitor > Audit Logs) will include an entry for each appliance that was rebooted.

TIP: After upgrading, to ensure proper data collection, go to the **Appliance Manager** page (Appliance Manager, locate one of the Data Collector or Message Collector appliances, and click the lightning bolt icon (⚡) to force configuration push for that appliance.

Rebooting Data Collectors and Message Collectors from the Command Line

From the console of the Database Server or SSH to the Database Server, perform the following steps to reboot Data Collector or Message Collector:

1. Either go to the console of a Database Server or SSH to access the Database Server.
2. Log in as **em7admin** with the appropriate password.
3. At the shell prompt, execute the following:

```
python -m silo_common.admin_toolbox <appliance_ID> "/usr/bin/sudo  
/usr/sbin/shutdown -r +1"
```

where:

- *appliance_ID* is the appliance ID for the Data Collector, Message Collector, or Administration Portal.

TIP: After upgrading, to ensure proper data collection, go to the **Appliance Manager** page (Appliance Manager, locate one of the Data Collector or Message Collector appliances, and click the lightning bolt icon (⚡) to force configuration push for that appliance.

Rebooting Standalone All-In-One Appliance and Standalone Database Server

Perform the following steps to reboot a standalone All-In-One Appliance or a standalone Database Server:

1. Either go to the console or use SSH to access the Skylar One appliance.
2. Log in as ***em7admin*** with the appropriate password.
3. On the Skylar One appliance, pause the system and shutdown MariaDB.

```
sudo touch /tmp/.proc_mgr_pause
sudo systemctl stop mariadb
```

4. Reboot the Skylar One appliance:

```
sudo reboot
```

5. After the Skylar One appliance has rebooted, either go to the console or use SSH to access the Skylar One appliance.
6. Log in as ***em7admin*** with the appropriate password.
7. Un-pause the Skylar One appliance:

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

Rebooting Two Database Servers Configured for Disaster Recovery

Perform the following steps to reboot two Database Servers configured for Disaster Recovery:

1. Either go to the console of the **primary** Database Server or use SSH to access the primary Database Server.
2. Log in as ***em7admin*** with the appropriate password.
3. Check the status of both Database Servers by typing the following command at the shell prompt:

```
drbadm status
```

4. Pause the system and shut down MariaDB on the **primary** Database Server. Enter the password for the **em7admin** user when prompted:

```
sudo touch /tmp/.proc_mgr_pause
```

```
sudo systemctl stop pacemaker
```

```
sudo systemctl stop mariadb
```

NOTE: You can skip this step if your Skylar One system is on AWS with RDS.

5. Reboot the **primary** Database Server:

```
sudo reboot
```

6. After the **primary** appliance has rebooted, log in to the console of the **primary** Database Server again.
7. Run the following command on the **primary** Database Server:

```
coro_config
```

8. Select the **[Promote DRBD]** option.

NOTE: If you are doing a simple reboot of the primary, you do not need to promote the cluster again, as Skylar One automatically does it for you.

9. Execute the following commands on the **primary** Database Server:

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

10. Enter the password for the **em7admin** user and confirm the command when prompted.
11. Log in to the **secondary** Database Server as the **em7admin** user using the console or SSH.
12. Execute the following command on the **secondary** Database Server to reboot the appliance:

```
sudo reboot
```

13. Enter the password for the **em7admin** user when prompted.

Rebooting Two Database Servers in a High Availability Cluster

Perform the following steps to reboot two Database Servers in a high availability cluster:

1. Either go to the console of the **secondary** Database Server or use SSH to access the **secondary** Database Server.
2. Log in as **em7admin** with the appropriate password.
3. Check the status of both Database Servers. To do this, enter the following at the shell prompt:

```
cat /proc/drbd
```

Your output will look like this:

```
1: cs:Connected ro:Secondary/Primary ds:UpToDate/UpToDate C r----
```

```
ns:17567744 al:0 bm:1072 lo:0 pe:0 ua:0 ap:0 ep:1 wo:b oos:12521012
```

NOTE: If your output includes "ro:Secondary/Primary", but does not include "UpToDate/UpToDate", data is being synchronized between the two appliances. You must wait until data synchronization has finished before rebooting.

4. Stop the cluster service on the **secondary** Database Server:

```
sudo systemctl stop pacemaker
```

5. Enter the password for the **em7admin** user when prompted.
6. Either go to the console of the **primary** Database Server or use SSH to access the **primary** Database Server.
7. Log in as **em7admin** with the appropriate password.
8. Pause the system and stop the cluster service on the **primary** Database Server. Enter the password for the **em7admin** user when prompted:

```
sudo touch /tmp/.proc_mgr_pause
sudo systemctl stop pacemaker
sudo systemctl stop mariadb
```

9. Reboot the **primary** Database Server:

```
sudo reboot
```

10. After the **primary** Database Server has rebooted, either go to the console of the **primary** Database Server or use SSH to access the **primary** Database Server.
11. Log in as **em7admin** with the appropriate password.
12. Execute the following command on the **primary** Database Server:

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

13. Enter the password for the **em7admin** user and confirm the command when prompted.
14. Either go to the console of the **secondary** Database Server or use SSH to access the **secondary** Database Server.
15. Log in as **em7admin** with the appropriate password.
16. Reboot the **secondary** Database Server:

```
sudo reboot
```

17. Enter the password for the **em7admin** user when prompted.

Rebooting Three Database Servers Configured for High Availability and Disaster Recovery

Perform the following steps to reboot three Database Servers configured for high availability and disaster recovery. In this configuration, two Database Servers are configured as a High Availability cluster and one Database Server is configured for Disaster Recovery.

1. Either go to the console of the **secondary** Database Server in the HA cluster or use SSH to access the **secondary** Database Server in the HA cluster,
2. Log in as **em7admin** with the appropriate password.
3. Check the status of both Database Servers in the HA cluster. To do this, enter the following at the shell prompt:

```
cat /proc/drbd
```

Your output will look like this:

```
10: cs:Connected ro:Secondary/Primary ds:UpToDate/UpToDate C r----  
ns:17567744 al:0 bm:1072 lo:0 pe:0 ua:0 ap:0 ep:1 wo:b oos:12521012
```

NOTE: If your output includes "ro:Secondary/Primary", but does not include "UpToDate/UpToDate", data is being synchronized between the two appliances. You must wait until data synchronization has finished before rebooting.

4. Stop the cluster service with the following command on the **secondary** Database Server in the HA cluster:

```
sudo systemctl stop pacemaker
```

5. Enter the password for the **em7admin** user when prompted.
6. Either go to the console of the **primary** Database Server in the HA cluster or use SSH to access the **primary** Database Server in the HA cluster.
7. Log in as **em7admin** with the appropriate password.
8. Pause the system and stop the cluster service on the **primary** Database Server in the HA cluster :

```
sudo touch /tmp/.proc_mgr_pause  
sudo systemctl stop pacemaker
```

9. Enter the password for the **em7admin** user when prompted
10. Reboot the **primary** Database Server in the HA cluster:

```
sudo reboot
```

11. After the **primary** Database Server in the HA cluster has rebooted, either go to the console of the **primary** Database Server in the HA cluster or use SSH to access the **primary** Database Server in the HA cluster.

12. Execute the following command on the **primary** Database Server in the HA cluster:

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

13. Enter the password for the *em7admin* user and confirm the command when prompted.

14. Either go to the console of the **secondary** Database Server in the HA cluster or use SSH to access the **secondary** Database Server in the HA cluster.

15. Log in as *em7admin* with the appropriate password.

16. Reboot the **secondary** Database Server in the HA cluster:

```
sudo reboot
```

17. Enter the password for the *em7admin* user when prompted.

18. Either go to the console of the Database Server for Disaster Recovery or use SSH to access the Database Server for Disaster Recovery.

19. Log in as *em7admin* with the appropriate password.

20. Reboot the Database Server for Disaster Recovery:

```
sudo reboot
```

21. Enter the password for the *em7admin* user when prompted.

Upgrading to Aurora 3 RDS (MySQL 8.0)

NOTE: This section applies only to users who deploy Skylar One on AWS and are currently running on Aurora 2 RDS (MySQL 5.7). If you do not deploy Skylar One on AWS or you are already running on Aurora 3, you can skip this section.

If you deploy Skylar One using AWS, you are currently running on Aurora 2 RDS (MySQL 5.7), and are upgrading to one of the following Skylar One versions, you can upgrade to Aurora 3 RDS (MySQL 8.0):

- Skylar One 12.1.2
- Skylar One 12.3.0 or later

NOTE: Aurora 3 upgrades are not supported for Skylar One 12.2.x.

Before You Upgrade to Aurora 3

Before you upgrade to Aurora 3, you must have already done the following:

- Ensured that you have administrator access to the AWS data engines, RDS, and IAM from the AWS console.
- Already upgraded Skylar One to a version that supports Aurora 3.
- Created an RDS snapshot.
- Created an Ansible AWX instance using the "ScienceLogic Skylar One Deployment - AWX V5" AMI.

NOTE: To obtain this AMI, contact ScienceLogic Support.

- Logged in to the AWX instance using credentials provided by ScienceLogic.
- Created a jump host using AWX, if you do not already have one.
- Updated the deploy, em7admin, and jump host credentials with the correct information.
- Updated the DE01 and jump host IP addresses.

Performing the Aurora 3 Upgrade

To upgrade to Aurora 3:

- Ensure that your Skylar One appliances are in *maintenance mode*.
- Log in to your AWX instance using the credentials provided by ScienceLogic.
- Run the "AWS - RDS Aurora3 Upgrade" job with the following input:

aws_region:

client:

vpc_id:

- After the playbook has executed successfully, verify that your database and application perform the post-application validation check.

Restoring the SSL Certificates

To restore your SSL Certificates:

- Login to the console of the Database Server or SSH to the Database Server.
- Open a shell session.

3. Enter the following at the shell prompt:

```
cp /etc/nginx/silossl.key.bak /etc/nginx/silossl.key
cp /etc/nginx/silossl.pem.bak /etc/nginx/silossl.pem
```

4. Repeat these steps on each Database Server in your Skylar One system.

Resetting the Timeout for PhoneHome Watchdog

NOTE: This section applies to users who are upgrading from Skylar One 11.1.x or earlier and have an existing PhoneHome configuration. If you are upgrading from Skylar One 11.2.0 or later or you do not have a pre-11.2.0 PhoneHome configuration, you can ignore this section.

You can manually reset the settings for the PhoneHome Watchdog server back to the settings you used before the upgrade.

To edit the settings for the watchdog service:

1. Log in to the console of the Data Collector as the root user or open an SSH session on the Data Collector.
2. View your PhoneHome Watchdog settings:

```
phonehome watchdog view
```

Your output will look like the following:

```
Current settings:
autosync: yes
interval: 120
state: enabled
autoreconnect: yes
timeoutcount: 1
check: default
```

Note the settings for *interval* and *timeoutcount*, so you can restore them after the upgrade.

3. To change the settings for Skylar One upgrade, type the following at the command line:

```
sudo phonehome watchdog set interval=<previous setting>;
sudo phonehome watchdog set timeoutcount=<previous setting>;
systemctl stop em7_ph_watchdog;
systemctl start em7_ph_watchdog;
```

4. Repeat these steps on each Data Collector.
5. Repeat these steps on each Message Collector.
6. Repeat these steps on each Database Server.

Clearing System Cache

After upgrading Skylar One, you should clear your system cache to remove cached items from Skylar One and prevent several potential issues that can occur post-upgrade due to caching. To do so, go to Misc > Clear Skylar One Cache.

Updating Default PowerPacks

Every time you install a software update on your appliances, ScienceLogic recommends that you also install the updates for all the PowerPacks that were included in the software update.

ScienceLogic includes multiple PowerPacks in the default installation of Skylar One. When you apply an update to your system, new versions of the default PowerPacks will be automatically imported in to your system. If a PowerPack is included in an update and is not currently installed on your system, Skylar One will automatically install the PowerPack. If a PowerPack is included in an update and is currently installed on your system, Skylar One will automatically import (but not install) the PowerPack.

If PowerPacks have been imported into your system but have not been installed, the **Update** column appears in the **PowerPack Manager** page (System > Manage > PowerPacks). For each PowerPack that has been imported to your system but has not been installed, the lightning bolt icon (⚡) appears in the **Update** field on the **PowerPack Manager** page.

To install the updates for multiple PowerPacks:

1. Go to the **PowerPack Manager** page (System > Manage > PowerPacks) and click the checkbox for each PowerPack you want to install.
2. In the **Select Action** drop-down field (in the lower right), choose *Update PowerPack(s)*. Skylar One displays a warning message before updating the PowerPack(s).
3. Click the **[OK]** button to continue the installation.
4. Click the **[Go]** button. If you completed the update, updated information about the PowerPack will appear in the **PowerPack Manager** page. All the items in the PowerPack will be installed in your Skylar One system.

NOTE: You can install multiple PowerPacks with the **Select Action** drop-down list only if each selected PowerPack includes an embedded Installation Key. PowerPacks that do not include embedded Installation Keys will fail to install.

NOTE: If the **Enable Selective PowerPack Field Protection** checkbox on the **Behavior Settings** page (System > Settings > Behavior) is selected, certain fields in Event Policies, Dynamic Applications, and Device Classes will **not** be updated.

Configuring Subscription Billing

If your Skylar One system is configured to communicate with the ScienceLogic billing server, usage data will be sent automatically from your Skylar One system to the ScienceLogic billing server once a day. After the ScienceLogic billing server receives the usage data, Skylar One will automatically mark the license usage file as delivered.

Sending usage data to the ScienceLogic billing server ensures that your bill is accurate and that ScienceLogic can continue making improvements to the Skylar One products.

To determine if you have correctly configured Subscription Billing:

- Go to the **System Usage** page (System > Monitor > System Usage). Click the **[Subscription]** button and choose ***License Data Delivery Status***.
- For air-gapped Skylar One systems, the value of ***Summary Date*** should be within the past 48 hours.
- For Skylar One systems that connect to ScienceLogic, the value of ***Summary Date*** should be within the past 48 hours and the value of ***Delivery Status*** is 1.

For details on configuring subscription billing, see the ***Subscription Billing*** manual.

Chapter

6

Upgrading Skylar One Extended Architecture

Overview

This chapter provides detailed steps for performing an upgrade on the Skylar One (formerly SL1) Extended Architecture.

NOTE: New installations of Skylar One Extended Architecture are available only on SaaS deployments.

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 (***).

<i>Workflow</i>	200
<i>Prerequisites</i>	200
<i>Resizing the Disks on the Compute Node</i>	201
<i>Installing ORAS</i>	202
<i>Obtaining Your Harbor Credentials</i>	203
<i>Upgrading to 12.3.x</i>	203
<i>Upgrading to 12.2.x</i>	206

Workflow

The following sections describe the steps to plan and deploy a Skylar One update.

If you would like assistance planning an upgrade path that minimizes downtime, contact your Customer Success Manager.

The workflow for upgrading Skylar One is:

1. Plan the update.
2. Schedule maintenance windows.
3. Review pre-upgrade best practices for Skylar One.
4. Back up SSL certificates.
5. Set the timeout for PhoneHome Watchdog.
6. Adjust the timeout for slow connections.
7. Run the system status script on the Database Server or All-In-One before upgrading.
8. Upgrade the Skylar One Distributed Architecture using the System Update tool (System > Tools > Updates).
9. Remove Skylar One appliances from maintenance mode.
10. Upgrade the Extended Architecture.
11. Upgrade MariaDB, if needed.
12. Reboot Skylar One appliances, if needed.
13. Restore SSL certificates.
14. Reset the timeout for PhoneHome Watchdog.
15. Update the default PowerPacks.
16. Configure Subscription Billing (one time only). For details, see the ***Subscription Billing*** manual.

NOTE: For details on all steps in this list except step 10, see the section on [Upgrading Skylar One](#).

Prerequisites

- ScienceLogic recommends that for production systems, each Compute Cluster contains six (6) Compute Nodes. Lab systems can continue to use Compute Clusters that include only three (3) Compute Nodes.
- The Storage Cluster requires a (possibly additional) node to act as the Storage Manager.
- Perform the installation steps in the Installation manual to install these additional nodes (for the Computer Cluster and the Storage Cluster) before upgrading your existing nodes.
- Ensure that all nodes in the Skylar One Extended Architecture can access the internet.

- You must use the same password for the em7admin account during ISO installation of the Database Server and ISO installation of the appliances in the Skylar One Extended Architecture.

NOTE: To perform the upgrade, you must have a ScienceLogic customer account that allows you access to the Harbor repository page on the [ScienceLogic Support Site](#). To verify your access, go to <https://registry.scilo.tools/harbor/>. For more information about obtaining Harbor login credentials, contact your Customer Success Manager.

Resizing the Disks on the Compute Node

The Kafka Messaging service requires additional disk space on each Compute Node. Before upgrading, ensure that each disk on each existing Compute Node in the Compute Node cluster is at least 350 GB.

If each disk on each existing Compute Node is not at least 350 GB, perform the following steps on each Compute Node:

1. Resize the hard disk via your hypervisor to at least 350 GB.
2. Note the name of the disk that you expanded in your hypervisor.
3. Power on the virtual machine.
4. Either go to the console of the Compute Node or use SSH to access the Compute Node.
5. Open a shell session on the server.
6. Log in with the system password for the Compute Node.
7. At the shell prompt, enter:

```
sudo lsblk | grep <disk_size>
```

where:

disk_size is your hard disk size from step #1.

8. Note the name of the disk that you expanded in your hypervisor.
9. At the shell prompt, enter:

```
sudo fdisk /dev/<disk_name>
```

where:

disk_name is the name of the disk you want to expand.

10. Enter **p** to print the partition table.
11. Enter **n** to add a new partition.
12. Enter **p** to make the new partition the primary partition.
13. Select the default values for partition number, first sector, and last sector.
14. Enter **w** to save these changes
15. Restart the VM.

16. At the shell prompt, enter:

```
sudo fdisk -l
```

17. Notice that now another partition is present.

18. To initialize the new partition as a physical volume, enter the following at the shell prompt:

```
sudo pvcreate <partition_name>
```

19. To add the physical volume to the existing volume group, enter the following at the shell prompt:

```
sudo vgextend em7vg <partition_name>
```

20. To verify and confirm that the volume group has grown to the expected size, enter the following at the shell prompt:

```
sudo vgdisplay | grep "VG Size
```

Installing ORAS

If you have not already installed OCI Registry as Storage (ORAS), you will need to do so before you can upgrade the Skylar One Extended Architecture.

To do so:

1. Use SSH to access the Management Node. Open a shell session on the server. Log in with the System Password you defined in the ISO menu.
2. In the Management Node, navigate to the sl1x-deploy directory. To do this, enter the following at the shell prompt:

```
cd sl1x-deploy
```

3. Run the following commands:

```
sudo su
```

```
curl -LO https://github.com/oras-
project/oras/releases/download/v0.12.0/oras_0.12.0_linux_amd64.tar.gz
```

```
mkdir -p oras-install/
```

```
tar -zxf oras_0.12.0_*.tar.gz -C oras-install/
```

```
mv oras-install/oras /usr/bin/
```

```
rm -rf oras_0.12.0_*.tar.gz oras-install/
```

```
exit
```

Obtaining Your Harbor Credentials

You will need to know your Harbor username and CLI secret when you upgrade the Skylar One Extended Architecture. To obtain these credentials:

1. Log in to Harbor at: https://registry.scilo.tools/harbor/sign-in?redirect_url=%2Fharbor%2Fprojects
2. Click **[Login via OIDC Provider]**.
3. Click **[Customer Login]**.
4. Log in with the username and credentials that you use to access the ScienceLogic Support site (<https://support.scienceologic.com>).
5. Click the username in the upper right and select **User Profile**.
6. On the **User Profile** page:
 - Note the username.
 - Click the pages icon next to the **CLI secret** field to copy the CLI secret to cache.
7. Exit the browser session.

Upgrading to 12.3.x

IMPORTANT: Before upgrading to Skylar One 12.3.0 or later, you **must** already be running Skylar One on Oracle Linux 8 (OL8). If you are on a version of Skylar One prior to 12.2.0 and running on OL7, you must first upgrade to Skylar One 12.1.1 or 12.1.2 and then migrate to OL8 before you can upgrade to Skylar One 12.3.x. For an overview of potential upgrade paths and their required steps, see the appropriate 12.3.x [Skylar One release notes](#).

To upgrade the Skylar One Extended Architecture to 12.3.x from 12.1.x or 12.2.x instances running on Oracle Linux 8 (OL8), follow these steps:

1. [Complete preupgrade steps](#).
2. [Disable Scylla](#).
3. [Upgrade the Skylar One Extended Architecture](#).
4. [Upgrade the Skylar One Distributed Architecture](#).

Step 1: Preupgrade

1. If you have not already done so, you must [install ORAS](#) and [obtain your Harbor credentials](#), which you will need for step 4.
2. Use SSH to access the Management Node. Open a shell session on the server. Log in with the system password you defined in the ISO menu.

3. In the Management Node, navigate to the `sl1x-deploy` directory. To do this, enter the following at the shell prompt:

```
cd sl1x-deploy
```

4. Log in to Harbor repository:

```
oras login registry.scilo.tools/scienceLogic/
```

- Enter the username you used to [log in to the browser-based session of Harbor](#).
- Enter the password (CLI Secret) that you saved from the browser-based session of Harbor.

5. Exit out of the `sl1x-deploy` directory and download the deployment files:

```
cd /home/em7admin/
```

```
oras pull registry.scilo.tools/scienceLogic/sl1x-deploy:12.3.x
```

```
cd sl1x-deploy
```

6. Copy the inventory template file to the `sl1x-inv.yml` file:

```
cp sl1x-inv-template.yml sl1x-inv.yml
```

7. Edit the `sl1x-inv.yml` file to match your Skylar One Extended system:

```
vi sl1x-inv.yml
```

CAUTION: Do not remove colons when editing this file.

- Make sure that the `sl1_version` value is `sl1_version: 12.3.x`.
- Supply values in all the fields that are applicable. For details on the `sl1x-inv.yml` file, see the manual ***Installing Skylar One Extended Architecture***, which can be obtained by contacting ScienceLogic Support.
- Save your changes and exit the file (`:wq`).

8. Pull the Docker image that is referenced in the docker-compose file:

```
docker-compose -f docker-compose.external.yml pull
```

Step 2: Disable the Scylla Cluster

To disable the Scylla cluster:

1. Use SSH to access the Management Node. Open a shell session on the server. Log in with the System Password you defined in the ISO menu.
2. Open a text editor for the `sl1x-inv.yml` file:

```
vi /home/em7admin/sl1x-deploy/sl1x-inv.yml
```

3. Edit the file to add the following additional variables:

```
all:  
  
vars:  
  
install_aiml: false  
  
enableNonScyllaPipeline: true  
  
enableLegacyScyllaPipeline: false
```

4. Run the following command to remove services that used the previous configuration:

```
docker-compose -f docker-compose.external.yml run --rm deploy app-purge
```

5. In the `sl1x-inv.yml` file, remove the Storage Node and Storage Manager IP addresses from the list. For example, you would remove the following lines:

```
sn:  
  
hosts:  
  
vars:  
  
scylla_admin_username: em7admin  
  
scylla_admin_password: <Scylla password>  
  
sm:  
  
hosts:  
  
vars:  
  
scylla_manager_db_user: em7admin  
  
scylla_manager_db_password: <Scylla password>
```

6. Save your changes and exit the file (`:wq`).

Step 3: Upgrade the Skylar One Extended Architecture

To upgrade the Skylar One extended architecture:

1. Use SSH to access the Management Node. Open a shell session on the server. Log in with the System Password you defined in the ISO menu.

2. To upgrade the Management Node services, run the following script:

```
sudo bash package-update.sh
```

3. To upgrade RKE and Kubernetes on the Compute Nodes, run the following command:

```
docker-compose -f docker-compose.external.yml run --rm deploy cn
```

4. To update the Skylar One Extended Architecture system services, run the following command:

```
docker-compose -f docker-compose.external.yml run --rm deploy app
```

Step 4: Upgrade the Skylar One Distributed Architecture

Update your classic Skylar One appliances. For more information, see the section on [Updating Skylar One](#).

Upgrading to 12.2.x

IMPORTANT: Before upgrading to Skylar One 12.2.0 or later, you *must* already be running Skylar One on Oracle Linux 8 (OL8). If you are on a version of Skylar One prior to 12.2.0 and running on OL7, you must first upgrade to Skylar One 12.1.1 or 12.1.2 and then migrate to OL8 before you can upgrade to Skylar One 12.2.x. For an overview of potential upgrade paths and their required steps, see the appropriate 12.2.x [Skylar One release notes](#).

To upgrade the Skylar One Extended Architecture to 12.2.x from 12.1.x instances running on Oracle Linux 8 (OL8), follow these steps:

1. [Complete preupgrade steps](#).
2. [Disable Scylla](#).
3. [Upgrade the Skylar One Extended Architecture](#).
4. [Upgrade the Skylar One Distributed Architecture](#).

Step 1: Preupgrade

1. If you have not already done so, you must [install ORAS](#) and [obtain your Harbor credentials](#), which you will need for step 4.
2. Use SSH to access the Management Node. Open a shell session on the server. Log in with the system password you defined in the ISO menu.
3. In the Management Node, navigate to the sl1x-deploy directory. To do this, enter the following at the shell prompt:

```
cd sl1x-deploy
```

4. Log in to Harbor repository:

```
oras login registry.scilo.tools/scienceLogic/
```

- Enter the username you used to [log in to the browser-based session of Harbor](#).
- Enter the password (CLI Secret) that you saved from the browser-based session of Harbor.

5. Exit out of the `sl1x-deploy` directory and download the deployment files:

```
cd /home/em7admin/
```

```
oras pull registry.scilo.tools/scienceLogic/sl1x-deploy:12.2.x
```

```
cd sl1x-deploy
```

6. Copy the inventory template file to the `sl1x-inv.yml` file:

```
cp sl1x-inv-template.yml sl1x-inv.yml
```

7. Edit the `sl1x-inv.yml` file to match your Skylar One Extended system:

```
vi sl1x-inv.yml
```

CAUTION: Do not remove colons when editing this file.

- Make sure that the `sl1_version` value is `sl1_version: 12.2.x`.
- Supply values in all the fields that are applicable. For details on the `sl1x-inv.yml` file, see the manual ***Installing Skylar One Extended Architecture***, which can be obtained by contacting ScienceLogic Support.
- Save your changes and exit the file (`:wq`).

8. Pull the Docker image that is referenced in the docker-compose file:

```
docker-compose -f docker-compose.external.yml pull
```

Step 2: Disable the Scylla Cluster

To disable the Scylla cluster:

1. Use SSH to access the Management Node. Open a shell session on the server. Log in with the System Password you defined in the ISO menu.
2. Open a text editor for the `sl1x-inv.yml` file:

```
vi /home/em7admin/sl1x-deploy/sl1x-inv.yml
```

3. Edit the file to add the following additional variables:

```
all:  
  
vars:  
  
install_aiml: false  
  
enableNonScyllaPipeline: true  
  
enableLegacyScyllaPipeline: false
```

4. Run the following command to remove services that used the previous configuration:

```
docker-compose -f docker-compose.external.yml run --rm deploy app-purge
```

5. In the `sl1x-inv.yml` file, remove the Storage Node and Storage Manager IP addresses from the list. For example, you would remove the following lines:

```
sn:  
  
hosts:  
  
vars:  
  
scylla_admin_username: em7admin  
  
scylla_admin_password: <Scylla password>  
  
sm:  
  
hosts:  
  
vars:  
  
scylla_manager_db_user: em7admin  
  
scylla_manager_db_password: <Scylla password>
```

6. Save your changes and exit the file (`:wq`).

Step 3: Upgrade the Skylar One Extended Architecture

To upgrade the Skylar One extended architecture:

1. Use SSH to access the Management Node. Open a shell session on the server. Log in with the System Password you defined in the ISO menu.

2. To upgrade the Management Node services, run the following script:

```
sudo bash package-update.sh
```

3. To upgrade RKE and Kubernetes on the Compute Nodes, run the following command:

```
docker-compose -f docker-compose.external.yml run --rm deploy cn
```

4. To update the Skylar One Extended Architecture system services, run the following command:

```
docker-compose -f docker-compose.external.yml run --rm deploy app
```

Step 4: Upgrade the Skylar One Distributed Architecture

Update your classic Skylar One appliances. For more information, see the section on [Updating Skylar One](#).

Chapter

6

Skylar One Self-Monitoring

Overview

Skylar One (formerly SL1) includes various tools that enable your Skylar One system to monitor the appliances within your system as well as your other Skylar One stacks. This proactive system health monitoring helps prevent outages and other major issues.

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>The Workflow for Skylar One Self-Monitoring</i>	211
<i>PowerPacks Required for Self-Monitoring</i>	211
<i>Credentials Required for Self-Monitoring</i>	214
<i>Enabling Connectivity on Port 7707</i>	217
<i>Discovering Your Skylar One Devices</i>	218
<i>Aligning Skylar One Self-Monitoring Dynamic Applications</i>	219
<i>Aligning the Correct Credentials to Your Devices</i>	220
<i>Configuring Run Book Automations to Populate Dashboards</i>	221
<i>Additional Self-Monitoring Resources</i>	223

The Workflow for Skylar One Self-Monitoring

The workflow for Skylar One self-monitoring is:

1. *Ensure you have the latest version of the required PowerPacks.*
2. *Create the appropriate credentials.*
3. *Enable connectivity on port 7707 (not required in all scenarios)*
4. *Discover your Skylar One devices.*
5. *Align additional Dynamic Applications to your Skylar One devices using device templates.*
6. *Ensure the correct credentials are aligned to the Dynamic Applications.*
7. *Configure run book automations to populate dashboards with data about your Skylar One devices.*

PowerPacks Required for Self-Monitoring

To self-monitor Skylar One, you must first have the latest versions of the following PowerPacks:

- *ScienceLogic Support Pack*
- *Data Pull Support*
- *Skylar One: Operational Insights*

ScienceLogic Support Pack

Skylar One includes the "ScienceLogic Support Pack" PowerPack by default with all Skylar One releases.

This PowerPack was developed by ScienceLogic Support to help you monitor your Skylar One systems. It allows a Skylar One stack to monitor itself and other Skylar One stacks.

The "ScienceLogic Support Pack" PowerPack:

- ensures that all Skylar One appliances are monitored and running the same version of Skylar One
- ensures that all Skylar One appliances appear on the **Devices** page
- ensures that configuration files on Skylar One appliances and clusters are kept in sync
- monitors the health and configuration of MariaDB
- generates alerts when system resources like disk-space are approaching capacity
- monitors DNS entries for Database Servers and All-In-One Appliances
- monitors long-running processes and queries
- monitors out-of-memory conditions
- ensures that CRM and corosync configurations are up to date
- ensures that configuration files match among high-availability Database Servers
- monitors crucial Skylar One processes like data pull and config_push

- monitors "rows behind" situations and generates alerts
- monitors and fixes backlogs of email messages
- generates alerts if any of the following files differ:
 - /etc/my.cnf.d/silo_mysql.cnf
 - /etc/silo.conf
 - /etc/siteconfig/mysql.siteconfig
 - /etc/siteconfig/siloconf.siteconfig
- collects telemetry related to Device Groups

The Dynamic Applications included in the PowerPack should automatically align when you discover your Skylar One appliances for self-monitoring.

NOTE: For more information about the Dynamic Applications that are included in the "ScienceLogic Support Pack" PowerPack, see the following knowledge base article:
<https://support.science logic.com/s/article/2951>.

Data Pull Support

The "Data Pull Support" PowerPack is also included by default in all Skylar One releases.

It includes several Dynamic Applications that use SNMP to collect configuration and performance data related to data pull, which is the process by which Skylar One Database Servers retrieve collected data from Data Collectors and Message Collectors.

TIP: Because the PowerPack is included in Skylar One installations and upgrades by default, you cannot download it from the ScienceLogic Support site. Whenever you upgrade your Skylar One stack to the latest release, you should go to the **PowerPack Manager** page (System > Manage > PowerPacks) to check for a new version of the "Data Pull Support" PowerPack. For more information about updating default PowerPacks, see the section on "How Skylar One Updates Default PowerPacks" in the **PowerPacks** manual.

Skylar One: Operational Insights PowerPacks

To self-monitor Skylar One, you will also need either the "Skylar One: Operational Insights - On Premise" or "Skylar One: Operational Insights - SaaS" PowerPack, depending on your system configuration:

- "Skylar One: Operational Insights - On-Premise" is required for MariaDB databases.
- "Skylar One: Operational Insights - SaaS" is required for Aurora RDS databases or SaaS stacks.

These PowerPacks include a variety of Dynamic Applications, run book actions and automation policies, dashboards, and other tools that provide additional Skylar One platform health visibility using data collected by the "ScienceLogic Support Pack".

Unlike the "ScienceLogic Support Pack" and "Data Pull Support" PowerPacks, the "Skylar One: Operational Insights" PowerPacks **are not** included by default in Skylar One releases. To use them, you must first download them from the ScienceLogic Support site and then install them on your Skylar One system.

IMPORTANT: If you are downloading version 105 or later of the "Skylar One: Operational Insights - On-Premise" or "Skylar One: Operational Insights - SaaS" PowerPack, then you can follow the steps below to download and install the PowerPack directly.

However, if you are upgrading from version 104 or earlier, there are additional steps that you must follow. These steps are detailed in the user manual that can be downloaded along with the PowerPack file on the ScienceLogic Support Site at <https://support.scienceologic.com/s/skylar-one/powerpacks>.

NOTE: You can periodically check the ScienceLogic Support Site at <https://support.scienceologic.com/s/skylar-one/powerpacks> for any newer versions of the "Skylar One: Operational Insights - On-Premise" or "Skylar One: Operational Insights - SaaS" PowerPacks.

TIP: By default, installing a new version of a PowerPack overwrites all content from a previous version of that PowerPack that has already been installed on the target system. You can use the **Enable Selective PowerPack Field Protection** setting in the **Behavior Settings** page (System > Settings > Behavior) to prevent new PowerPacks from overwriting local changes for some commonly customized fields. For more information, see the section on [Global Settings](#).

NOTE: For details on upgrading Skylar One, see the relevant [Skylar One Platform Release Notes](#).

To download and install the PowerPack:

1. Search for and download the PowerPack from the **PowerPacks** page at the [ScienceLogic Support Center](#) (Skylar One > PowerPacks, login required).
2. In Skylar One, go to the **PowerPacks** page (System > Manage > PowerPacks).
3. Click the **[Actions]** button and choose *Import PowerPack*. The **Import PowerPack** dialog box appears.
4. Click **[Browse]** and navigate to the PowerPack file from step 1.
5. Select the PowerPack file and click **[Import]**. The **PowerPack Installer** modal displays a list of the PowerPack contents.
6. Click **[Install]**. The PowerPack is added to the **PowerPacks** page.

NOTE: If you exit the **PowerPack Installer** modal without installing the imported PowerPack, the imported PowerPack will not appear in the **PowerPacks** page. However, the imported PowerPack will appear in the **Imported PowerPacks** modal. This page appears when you click the **[Actions]** menu and select *Install PowerPack*.

Credentials Required for Self-Monitoring

NOTE: For more background on using credentials in Skylar One, see the *Discovery & Credentials* manual.

To self-monitor Skylar One, you must create a minimum of three credentials:

- *An SNMP credential*
- *A database credential for your Database Server and at least one more for your Data Collectors and Message Collectors*

Your Skylar One system includes default example credentials for each of these credential types. However, ScienceLogic does not recommend that you use or edit the example credentials. Instead, you should create new credentials specific to your system.

Creating an SNMP Credential for Self-Monitoring

Prior to creating the SNMP credential for Skylar One self-monitoring, you must configure the SNMP settings for your Skylar One appliances, including providing access to SNMPD for those appliances. To do this, you must be familiar with your company's policies on SNMP and the required SNMP settings.

Newly installed Skylar One systems are configured with default settings that allow read-only access for SNMPD and enable you to align the default example SNMP credentials that are included in Skylar One. ScienceLogic strongly recommends that you configure SNMPD access with community/user access for your Skylar One environment and do not use the default example SNMP credentials.

CAUTION: The following procedure is the only supported way of customizing the SNMPD configuration on a Skylar One system. Other SNMPD customizations might get overwritten without warning.

To configure SNMP for your Skylar One appliances:

1. Open an SSH session on the Skylar One appliance using either the **sl1user** account or the **em7admin** account.

NOTE: ScienceLogic recommends using the **sl1user** account, as this account has fewer privileges and is designed for basic system setup tasks. If you are unfamiliar with this account, see the "Role-based sl1user Account" section in the *Organizations & Users* manual.

If you opt to use the **em7admin** account, you must run the following command to bring up the menu:

```
sudo /usr/local/bin/slmenu
```

2. From the menu, select the **SNMP Configuration** option and then press Enter.
3. The **SNMP Configuration** menu displays your current SNMP settings, as well as options to change the SNMPv2 Community or SNMPv3 settings. After reviewing your current settings, do one of the following:
 - To change the SNMPv2 community string, use the arrow keys on your keyboard to navigate to that option on the menu and then press Enter. Proceed to step 4.
 - To change the SNMPv3 settings, use the arrow keys on your keyboard to navigate to that option on the menu and then press Enter. Proceed to step 5.

TIP: You can use the following special characters in the SNMPv2 community string and SNMPv3 settings:

```
+ _)(*&^%$#@!|}{":?><=-\][;/.,
```

4. In the **Change SNMPv2 Community** menu, enter a new community string that you want to use for SNMPv1/2 access in Skylar One. Alternatively, you can disable SNMPv1/2 access by deleting the string and leaving the field blank. When you are done, navigate to **OK** and press Enter. If you want to also change your SNMPv3 settings, proceed to step 5; otherwise, to leave your SNMPv3 settings as they are, proceed to step 6 to return to the **SNMP Configuration** menu.

5. In the **Change SNMPv3 Settings** menu, use the arrow keys on your keyboard to navigate to the following fields and update their values as needed for your configuration:
 - **SNMPv3 User**. Enter the username you want to use for SNMPv3 authentication.
 - **SNMPv3 Auth Password**. Enter the password you want to use to authenticate the SNMPv3 user. This password must contain at least 8 characters.
 - **SNMPv3 Auth Password (confirm)**. Re-enter the authentication password for the SNMPv3 user.
 - **SNMPv3 Auth Protocol**. Enter the authentication protocol for the SNMPv3 user, based on your company policies. The possible values for this field appear on the **Change SNMPv3 Settings** menu.
 - **SNMPv3 Privacy Password**. Enter the password you want to use to ensure data privacy for the SNMPv3 user.
 - **SNMPv3 Privacy Password (confirm)**. Re-enter the privacy password.
 - **SNMPv3 Privacy Protocol**. Enter the data privacy protocol used for data encryption and decryption, based on your company policies. The possible values for this field appear on the **Change SNMPv3 Settings** menu.

When you are done changing your SNMPv3 settings, navigate to **OK** and press Enter to return to the **SNMP Configuration** menu.

6. On the **SNMP Configuration** menu, review the updated SNMP settings. Repeat steps 3-5 as necessary to make any additional changes. To save the new settings, navigate to the **Save** option and press Enter.
7. On the **Save Settings Confirmation** menu, to confirm your new settings, navigate to the **Yes** option and press Enter.
8. Repeat these steps for all Skylar One appliances.

After you have configured your SNMP settings, you can create one or more SNMP credentials for Skylar One self-monitoring. To do so, follow the steps in the section on "Defining an SNMP Credential" in the **Discovery & Credentials** manual. Use the values you configured in the previous section when creating your SNMP credentials.

Creating Database Credentials for Self-Monitoring

For Skylar One self-monitoring, you will need **at least two** database credentials:

- One database credential for your Database Server or Aurora RDS
- One or more database credentials for your Data Collectors and Message Collectors

NOTE: If you have a SaaS managed Skylar One system, the RDS credential will be created for you. If it is not or you are unsure which one it is, contact ScienceLogic Support.

NOTE: If your Data Collectors and Message Collectors all use the same database user and password, then you can use a single credential for all of them. If they use unique database users and passwords, then you will need to create credentials for each Data Collector or Message Collector.

To create the database credentials, follow the steps in the section on "Defining a Database Credential" in the *Discovery & Credentials* manual. When doing so, use the following guidance for completing these fields:

- **Timeout (ms).** Enter "0".
- **Database Type.** Select MySQL.
- **Database Name.** Enter "master".
- **Database User.** Typically, this value will either be "clientdbuser" or "root". If you are unsure of which value to use, refer to your silo.conf file for your system.
- **Database Password.** Enter the password you have set for the database.
- **Hostname/IP.** Enter "%D".
- **Port.** For your Database Server or Aurora RDS credential, enter "7706". For your Data Collector and Message Collector credentials, enter "7707".

Enabling Connectivity on Port 7707

On-premises Data Collector and Message Collector appliances do not allow connections to port 7707 (MariaDB) from an external device. Therefore, if you want one on-premises Data Collector to monitor another Data Collector or Message Collector, then there are some additional steps you must take to enable connectivity on port 7707.

NOTE: Some Dynamic Applications collect data from MariaDB/ Aurora and should be able to connect to port 7706 (DB/AIO) or 7707 (Data/ Message Collector). Please ensure that any firewalls between the monitoring collector and the monitored appliances allow communication from the monitoring collector to the monitored appliances on port 7706 or 7707.

NOTE: If this scenario does not apply to your Skylar One configuration, then you can skip this section.

For the instructions below, consider the following scenario:

- "Collector A" is the Data Collector or Message Collector you want to monitor.
- "Collector B" is the Data Collector on which you want to monitor Collector A.

To enable connectivity on port 7707 in this scenario:

1. Log in to the Web Configuration Utility for Collector A using any web browser supported by Skylar One. The address of the Web Configuration Utility is in the following format:

```
https://<Collector-A-IP-address>:7700
```

NOTE: For AWS instances, **Collector-A-IP-address** is the public IP for the AWS instance for Collector A. 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 username and password, log in as the "em7admin" user with the appropriate password. After logging in, the main **Configuration Utility** page appears.
3. Click the **[Device Settings]** button. The **Settings** page appears.
4. On the **Settings** page, add the IP address for Collector B to the **Database IP Address** field, immediately after the IP address that is already listed for Collector A, separated by a comma and no space. For example, if the IP address for Collector A is 10.10.10.1 and the IP address for Collector B is 10.20.20.2, then you would enter "10.10.10.1,10.20.20.2" (without the quotation marks) in the **Database IP Address** field.
5. Click **[Save]** and log out of the Web Configuration Utility. With this change saved, the firewall rules on Collector A will be updated to allow connectivity to port 7706 (MariaDB) from the IP address for Collector B.

Discovering Your Skylar One Devices

After creating your credentials, you will then use those credentials to discover your Skylar One appliances.

To do so, follow the instructions found in the "Adding Devices Using Unguided Discovery" section of the **Discovery & Credentials** manual.

While configuring your discovery policy, make sure to do the following:

- Select the **SNMP credential** you created.
- Additionally, select one or more of the **database credentials** you created, depending on whether you are discovering the Database Server, one or more Data Collectors or Message Collectors, or both the Database Server and one or more collectors.
- When listing the IP addresses that you want to discover, include the IP addresses of all Skylar One appliances that you want to discover.
- Under **Advanced Options**, make sure the **Discover non-SNMP** toggle is turned off.

Upon discovering your Skylar One appliances, most SNMP-based Dynamic Applications should automatically be aligned to the appropriate devices.

TIP: After discovering your Skylar One appliances, you can go to the **Appliance Manager** page (System > Settings > Appliances) to confirm that they all appear on the page and that their settings are correct. For example, the **Device Name** on the **Devices** and **Device Manager** (Devices > Classic Devices, or Registry > Devices > Device Manager in the classic user interface) pages should match the **Device Name** on the **Appliance Manager** page.

Aligning Skylar One Self-Monitoring Dynamic Applications

After discovering your Skylar One appliances, you must ensure that they are properly set up for monitoring. The easiest way to do this is to apply the device templates that are included in the "ScienceLogic Support Pack," "Data Pull Support," and "Skylar One: Operational Insights" PowerPacks to ensure your Skylar One device classes have the necessary Dynamic Applications aligned to them.

To align the appropriate self-monitoring Dynamic Applications using device templates:

1. Go to the **Device Manager** page (Devices > Classic Devices, or Registry > Devices > Device Manager in the classic user interface).
2. In the **Device Class** column, use the filter field to search for the appropriate device class based on the table at the end of this section.
3. Select the checkbox for each row that contains the device class type that you want to modify.

NOTE: Because each of these device classes requires a different configuration, you should modify devices of only one device class at a time.

4. From the **Select Action** field at the bottom of the page, select *MODIFY by Template* and then click **[Go]**. The **Bulk Device Configuration** modal appears.
5. In the **Template** drop-down field at the top of the **Bulk Device Configuration** modal, select the appropriate device template based on the table at the bottom of this section.

NOTE: You can select only one device template at a time. For device classes that have multiple device templates, complete steps 5-7 for one template, then repeat the steps for the other templates.

6. Click the **[Dyn Apps]** tab and verify that a list of Dynamic Applications appear in the **Subtemplate Selection** pane.
7. Click **[Apply]**, and then click **[Confirm]**.
8. Repeat steps 1-7 as needed until all of the device classes needed for self-monitoring have all of the appropriate device templates applied to them, based on the following table:

Appliance Type	Device Class	Device Templates
Primary Database Server or cluster IP device	Skylar One Database	Support: Skylar One Active DB/VIP Skylar One: Database - On Premise Skylar One System Processes
Administration Portal	EM7 Admin Portal	Support: Skylar One Application Portal
Data Collector	Skylar One Data Collector	Support: Skylar One Collector Skylar One: Collectors Skylar One System Processes
Message Collector	Skylar One Message Collector	Support: Skylar One Message Collector Skylar One System Processes

NOTE: It can take Dynamic Applications up to 15 minutes to start collecting data after they are aligned.

Aligning the Correct Credentials to Your Devices

When you align the self-monitoring Dynamic Applications to your Skylar One appliances using the device templates that are included in the self-monitoring PowerPacks, those Dynamic Applications use the unedited versions of the "Skylar One CDB" and "Skylar One Collector Database" credentials that are included in the PowerPacks.

Now, you must edit those Dynamic Applications to use the *copies you made for those credentials* instead of the original versions that are included in the self-monitoring PowerPacks. This will prevent Skylar One from overwriting your self-monitoring configuration whenever you install updates for those PowerPacks in the future.

To ensure the correct credentials are aligned to your Skylar One devices:

1. Go to the **Device Manager** page (Devices > Classic Devices, or Registry > Devices > Device Manager in the classic user interface).
2. Locate one of your Skylar One devices and click its wrench icon (🔧). The **Device Administration** panel appears.
3. On the **Device Administration** panel, click the **[Collections]** tab. The **Dynamic Application Collections** page appears.
4. On the **Dynamic Application Collections** page, locate and select the checkboxes for any Dynamic Applications that have "Skylar One CDB" listed in the **Credential** column.

5. From the **Select Action** field at the bottom of the page, select the database credential you created for the Database Server or Aurora RDS from the list and then click **[Go]**. The **Dynamic Application Collections** page will refresh and the database credential you created should now be listed in the **Credential** column for the Dynamic Applications you selected.
6. Repeat steps 4 and 5 to replace the credential for any Dynamic Applications that currently have the "Skylar One Collector Database" credential in the **Credential** column with the database credential you created for Data Collectors and Message Collectors.

Configuring Run Book Automations to Populate Dashboards

The "Skylar One: Operational Insights" PowerPacks include two run book actions and automation policies: "Skylar One: Collector Data Collection" and "Skylar One: System Log Data Collection."

These automations collect data that is then used to populate the "Collector Performance" and "System Logs Summary" dashboards, respectively, with data relating to your monitored Skylar One appliances.

TIP: If you are using the *ScienceLogic Skylar One Operational Insights - On-Premise* Community PowerPack, please review the manual that is shipped with the PowerPack for more information on configuring run book automations. Be advised, the manual shipped with the PowerPack is not official ScienceLogic documentation.

Skylar One: Collector Data Collection

The "Skylar One: Collector Data Collection" automation is responsible for collecting Skylar One Collector-specific data that is used in the "Collector Performance" dashboard. The automation policy is configured to trigger the "Skylar One: Collector Data Collection" run book action four times once per hour. The automation helps identify Skylar One Collectors that have been discovered and match the following criteria required to appear in the "Collector Performance" dashboard:

- Data Collectors and Message Collectors should be *discovered as managed devices*.
- The discovered Skylar One Collector's **Device Name** on the **Devices** and **Device Manager** (Devices > Classic Devices, or Registry > Devices > Device Manager in the classic user interface) pages should match the name of the same collector from the **Appliance Manager** page (System > Settings > Appliances).
- The Skylar One Collectors should have the following two Dynamic Applications aligned and successfully collecting data:
 - Support: File System
 - Host Resource: Configuration
- The Data Collectors should be in a collector group (CUG). This does not apply to Message Collectors.

When the automation runs, it stores the data in a custom table, which is then read and displayed by the dashboard widget.

Skylar One: System Log Data Collection

The "Skylar One: System Log Data Collection" automation is responsible for collecting system log-specific data that is used in the "System Logs Summary" dashboard. The Automation is configured to trigger the "Skylar One: System Log Data Collection" run book action five times once per hour.

The automation parses the top problem logs, including SIGTERMs, PoolWorker logs, and unhandled exceptions.

When the automation runs, it stores the data in a custom table, which is then read and displayed by the dashboard widget.

NOTE: If the total system log count is greater than 6 million, the automation will not collect data; instead, the "System Logs Summary" dashboard will display a message stating that the log count is too high. If this occurs, delete older logs to bring the total count below 6 million.

Verifying Your Devices

Before you configure the run book automation policies required to populate the "Collector Performance" and "System Logs Summary" dashboards, you must verify that you have the necessary devices in your Skylar One system. The method for doing so will vary based on whether you have an on-premises Skylar One system or a SaaS Skylar One system. Both methods are described below.

On-Premises Skylar One Systems

To verify that you have the necessary devices:

1. Go to the **Device Manager** page (Devices > Classic Devices, or Registry > Devices > Device Manager in the classic user interface).
2. In the **Device Class** column, use the filter field to search for devices with the "Skylar One Database" device class.
3. Verify that all Database Servers are discovered.
4. If the Database Servers are not discovered, then you must *discover them*. Otherwise, if they are, you can skip to the next section.

SaaS Skylar One Systems

To verify that you have the necessary device:

1. Go to the **Device Manager** page (Devices > Classic Devices, or Registry > Devices > Device Manager in the classic user interface).
2. In the **Device Name** column, use the filter field to search for a device with the name "Skylar One Stats".
3. Verify that the "Skylar One Stats" device exists. If it does not yet exist, proceed to the next step. Otherwise, if it does, you can skip to the next section.
4. From the **[Actions]** menu, select *Create Virtual Device*. The **Create Virtual Device** modal appears.

5. Complete the following fields:
 - **Device Name**. Type "Skylar One Stats".
 - **Organization**. Select *System*.
 - **Device Class**. Select *Virtual Device / Dynamic App Emissary*.
 - **Collector**. Select *Self Monitoring Collector*.
6. Click **[Add]** and then exit the modal.

Configuring the Run Book Automations

To configure the run book automation policies required to populate the "Collector Performance" and "System Logs Summary" dashboards:

1. Go to the **Automation Policy Manager** page (Registry > Run Book > Automation).
2. In the **Automation Policy Name** column, use the filter field to search for a policy with the name "Skylar One: Collector Data Collection".
3. Click the wrench icon (🔧) for the "Skylar One: Collector Data Collection" automation policy. The **Automation Policy Editor** modal appears.
4. In the **Available Devices** field, use the filter field to search for the appropriate device:
 - For on-premises Skylar One deployments, search for the active Skylar One Database Server device.
 - For SaaS Skylar One deployments, search for the "Skylar One Stats" device.
5. Select the device from the **Available Devices** field and then click the right arrow icon to move it to the **Aligned Devices** field.
6. Click **[Save]** and then exit the **Automation Policy Editor** modal.
7. Repeat steps 1-6 to search for and configure the "Skylar One" System Log Data Collection" automation policy.

Additional Self-Monitoring Resources

For information about additional steps you can take in Skylar One to monitor your Skylar One system, see the following chapters in this manual:

- [Daily Health Tasks](#)
- [Monitoring and Maintaining Skylar One](#)
- [Diagnostic Tools](#)

Chapter

7

Monitoring and Maintaining Skylar One

Overview

This chapter describes how to manage user access, manage scheduled tasks, and more in 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>Monitoring and Managing User Access</i>	225
<i>Managing Scheduled Tasks</i>	228
<i>Putting the Database Server into Maintenance Mode</i>	231
<i>Monitoring Overall System Usage and Statistics</i>	231
<i>Viewing an Overview of All Events</i>	232
<i>Viewing Events by Appliance and Event Source</i>	233

Monitoring and Managing User Access

Skylar One lets multiple users log into the same Skylar One system at the same time. The time a user spends logged into Skylar One is known as a **user session**. You can end a user's session in Skylar One, and you can also limit the number of users that can be simultaneously logged into a Skylar One system.

The **Access Sessions** page allows administrators to monitor user logins and logouts to the user interface.

From this page, you can also:

- End a user's session.
- View a list of accounts that are locked out of the user interface due to invalid username and password.
- Unlock accounts that are locked out of the user interface.

Viewing Information about Each Access Session

The **Access Sessions** page displays a list of recent logins to the user interface. To view the **Access Sessions** page:

1. Go to the **Access Sessions** page (System > Monitor > Access Logs).
2. For each session, the **Access Sessions** page displays:
 - **User Account**. Username of person logging in to the user interface.
 - **User Display Name**. The username, email address, or preferred display name. This value is determined by the user's authentication resource settings.
 - **Last Address**. IP address from which the user accessed the user interface.
 - **State**. Current status of the user. The choices are:
 - *Active*. User is currently logged in to the user interface.
 - *Expired*. User's session in the user interface was killed.
 - *Logged Out*. User logged out of the user interface.
 - *Never Used*. User logged in to the user interface and did not perform any tasks before the session was killed.
 - **Login Time**. Date and time at which the user logged in.
 - **Last-Hit Time**. Date and time at which the user last loaded a page in the user interface.
 - **Logout Time**. Date and time at which the user logged out.
 - **Session Duration**. Length of time between login and logout.

Deleting a User's Session

From the **Access Sessions** page, you can end a user's session in the user interface. The user must log in again to access the user interface. The status of the session will be "expired".

To end a user's session:

1. Go to the **Access Sessions** page (System > Monitor > Access Logs).
2. In the **Access Sessions** page, find the session you want to end. Select the checkbox for that session.
3. Click the **Select Actions** field (in the lower right of the page) and then select *Kill user session*. Click the **[Go]** button
4. Each selected session is ended. The user associated with each selected session is logged out of the user interface. The status of the session changes to "expired".

NOTE: After ending a user's session, that user can immediately log in to the user interface again. To prevent a user from logging in to the user interface, you must disable the user's account. For information on user accounts, see the manual *Organizations and Users*.

Limiting the Number of Simultaneous User Sessions

If you get an "HTTP Response code was 429 (Too Many Requests)" error or a "User sessions at maximum" in Skylar One, you can adjust the **USER_MAX_SESSIONS** value in the **/opt/em7/nextui/nextui.conf** file:

1. SSH to the Skylar One appliance and log in as user **em7admin**.
2. At the command line, open the **nextui.conf** file in the **vi** editor:

```
sudo vi /opt/em7/nextui/nextui.conf
```

3. In the NextUI configuration file, set a new value for **USER_MAX_SESSIONS**, such as **USER_MAX_SESSIONS=1000** for 1,000 concurrent user sessions.
4. Save your changes and restart the NextUI service:

```
sudo systemctl restart nextui
```

Viewing Lockouts and Unlocking Lockouts

If a user enters incorrect login information multiple times in a row, that username, the user's IP address, or both will be locked out of the user interface.

To view lockouts or restore login privileges to locked out users:

1. Go to the **Access Sessions** page (System > Monitor > Access Logs).
2. In the **Access Sessions** page, click the **[Lockouts]** button.

3. The **Account Lockouts** modal page allows administrators to view a list of locked-out accounts and to restore login privileges to locked out users.
4. The **Account Lockouts** modal page displays the following about each lockout:
 - **Attempt Account**. Username that caused the lockout.
 - **From Address**. IP address from which the failed login attempts originated.
 - **Attempt Time**. Date and time at which lockout occurred.
 - **Tries**. Number of times user tried to log in to the user interface.
5. **To remove the lock for the user account** and allow logins from the username and/or IP address, click the delete icon (>Delete).

Global Settings for Lockouts

The platform includes global settings that define how lockouts behave. In the **Behavior Settings** page (System > Settings > Behavior), the following fields affect lock-outs:

- **Account Lockout Type**
- **Account Lockout Attempts**
- **Account Lockout Duration**
- **Lockout Contact Information**

Audit Logs

For additional information about users and their actions in the platform, you can view the **Audit Logs** page. The **Audit Logs** page provides a complete audit trail for the platform. The **Audit Logs** page displays a record of all actions in the platform that are generated by users or by managed elements. For details, see the section on [Audit Logs](#).

Improved User Session Control with tmux

The **tmux utility** is a terminal multiplexer that enables a number of terminals to be created, accessed, and controlled from a single screen, strengthens session-control mechanisms and aligns with industry-wide security practices.

Starting with Skylar One version 12.3.4, the tmux utility is disabled by default if you are on a non-STIG Skylar One deployment and access a Skylar One system using SSH.

NOTE: This is a change in behavior from Skylar One versions 12.2.1.1 through 12.3.3, where the tmux utility was enabled by default. For more information, see the [12.3.3 release notes](#).

If you are on a STIG-compliant Skylar One deployment, the tmux utility is enabled by default. ScienceLogic encourages non-STIG users enable the tmux utility as well.

If tmux is enabled, sessions are automatically locked after 15 minutes of idleness or if an unclean SSH disconnect or dropped SSH connection occurs. Upon login, Skylar One checks for and attaches any detached tmux session if it finds them; otherwise, it starts a new session.

The utility also facilitates advanced features like scroll-back buffering with search, built-in clipboarding, multiple sessions and panes, detaching or attaching sessions, and session supervision or sharing.

To enable the tmux utility in non-STIG deployments:

1. Either go to the console of the Skylar One appliance or use SSH to access the Skylar One appliance.
2. Open a shell session on the server.
3. Type the following at the command line to edit the **silo.conf** file:

```
sudo visilo
```

4. Change the following line in the **[OS_HARDENING]** section of the file to enable tmux:

```
TMUX = true
```

NOTE: If the **[OS_HARDENING]** heading does not already exist in the **silo.conf** file, you must add that immediately above the **TMUX = true** setting.

5. Save and quit the file (**:wq**).
6. Log out of Skylar One and then log back in. The tmux utility is now enabled.

For more information about tmux shortcuts and usage, see <https://tmuxcheatsheet.com/>.

NOTE: When using the command-line interface on a Skylar One appliance, the interface is meant for limited administrative purposes only.

TIP: When completing a task in a **tmux** session that will take longer than 15 minutes to complete, such as restoring from backup files, you should open a named **tmux** session to avoid the session timing out:

```
tmux new -t <session name>
```

Managing Scheduled Tasks

The **Schedule Manager** page (Registry > Schedules > Schedule Manager) allows you to view and manage all the scheduled processes you have defined in your system.

You can define scheduled processes in the following pages:

- Report Scheduler. (For more information, see the **Reports** manual.)
- My Work Schedule. (For more information, see the **Organizations and Users** manual.)
- Recurring Ticketing Scheduler. (For more information, see the **Ticketing** manual.)
- Discovery Control Panel. (For more information, see the **Discovery and Credentials** manual.)
- Dashboards. (For more information, see the **Dashboards** manual.)
- IT Service Editor. (For more information, see the **IT Services** manual.)
- Device Manager. (For more information, see the **Device Management** manual.)
- Backup Management. (For more information, see the section on **Configuration Backups**.)

Managed schedules are loaded in batches of 1,000 by default to prevent the Process Manager from becoming overloaded and stopping the scheduled jobs prematurely. Previously, all scheduled jobs were loaded simultaneously. You can update the default batch number of scheduled jobs by going to the Database Tool page (System > Tools > DB Tool) and entering the following in the **SQL Query** field, replacing `<batch_number>` with the number of scheduled jobs you want to load per batch:

```
INSERT INTO master.system_custom_config (field, field_value) VALUES
('load_schedules_in_batches', <batch_number>) ON DUPLICATE KEY UPDATE
field = 'load_schedules_in_batches', field_value = <batch_number>
```

NOTE: The **Database Tool** page is available only in versions prior to 12.2.1 and displays only for users that have sufficient permissions to access the page.

Recommended System Maintenance

ScienceLogic also recommends that you take the following actions on a regular basis to reduce outages as much as possible.

Daily:

- Review "Skylar One Operational Insights: Database Performance" classic dashboard
- Review "Skylar One Operational Insights: Collector Performance" classic dashboard
- Review "Skylar One Operational Insights: System Log Summary" classic dashboard
- Review "Skylar One Operational Insights: Backup History" classic dashboard

TIP: You can find the *Skylar One Operational Insights* PowerPack on the **PowerPacks** page at the ScienceLogic Support Site: <https://support.scienceologic.com/s/skylar-one/powerpacks>.

Weekly:

- Run the System Status Script and review:
 - Address every error item in the report
 - Read Knowledge Base articles

- Open tickets for issues when help from Skylar One Support is needed

Monthly:

- Review capacity items. You must understand License Usage and how to project future capacity

Quarterly:

- Audit User Profile access to verify that it meets expected needs
- Audit DNS servers and Timeservers on all collectors

Viewing the List of Schedules

The **Schedule Manager** page (Registry > Schedules > Schedule Manager) displays the following about each schedule:

- **Schedule Summary.** Displays the name assigned to the scheduled process.
- **Schedule Description.** Displays a description of the scheduled process.
- **Event ID.** Displays a unique, numeric ID for the scheduled process. Skylar One automatically creates this ID for each scheduled process.
- **sch id.** Displays a unique, numeric ID for the schedule. Skylar One automatically creates this ID for each schedule.
- **Context.** Displays the area of Skylar One upon which the schedule works.
- **Timezone.** Displays the time zone associated with the scheduled process.
- **Start Time.** Displays the date and time at which the scheduled process will begin.
- **Duration.** Displays the duration, in minutes, which the scheduled process occurs.
- **Recurrence Interval.** If applicable, displays the interval at which the scheduled process recurs.
- **End Date.** If applicable, displays the date and time on which the scheduled process will recur.
- **Last Run.** If applicable, displays the date and time the scheduled process most recently ran.
- **Owner.** Displays the username of the owner of the scheduled process.
- **Organization.** Displays the organization to which the scheduled process is assigned.
- **Visibility.** Displays the visibility level for the scheduled process. Possible values are "Private", "Organization", or "World".
- **Enabled.** Specifies if the scheduled process is enabled. Possible values are "Yes" or "No".

Enabling or Disabling One or More Schedules

You can enable or disable one or more scheduled process from the **Schedule Manager** page (Registry > Schedules > Schedule Manager). To do this:

1. Go to the **Schedule Manager** page (Registry > Schedules > Schedule Manager).
2. Select the checkbox icon for each scheduled process you want to enable or disable.
3. Click the **Select Action** menu and choose *Enable Schedules* or *Disable Schedules*.
4. Click the **[Go]** button.

Deleting One or More Schedules

You can delete one or more scheduled process from the **Schedule Manager** page (Registry > Schedules > Schedule Manager). To do this:

1. Go to the **Schedule Manager** page (Registry > Schedules > Schedule Manager).
2. Select the checkbox icon for each scheduled process you want to delete.
3. Click the **Select Action** menu and choose *Delete Schedules*.
4. Click the **[Go]** button.

Putting the Database Server into Maintenance Mode

You can now put the Database Server in maintenance mode and stop all pull processes from the Data Collectors. You can then perform database maintenance or network maintenance without generating events.

After maintenance is completed, you can put the Database Server out of maintenance mode. Pull processes from the the Data Collectors will resume from the point where they were paused.

The new commands are silostart and silostop.

To put a Database Server in maintenance mode:

1. Either go to the console of a Database Server or SSH to access the Database Server.
2. Log in as **em7admin** with the appropriate password.
3. At the shell prompt, execute the following:

```
silostop
```

To put a Database Server out of maintenance mode:

1. Either go to the console of a Database Server or SSH to access the Database Server.
2. Log in as **em7admin** with the appropriate password.
3. At the shell prompt, execute the following:

```
silostart
```

Monitoring Overall System Usage and Statistics

The **System Usage** page displays:

- Tables that show the type and number of each type of task performed by Skylar One
- An optional line graph that displays system usage. To enable the display of this graph, go to the **Behavior Settings** page (System > Settings > Behavior) and uncheck the **Hide Perpetual License Count** checkbox. The graph displays the following metrics over time:
 - **Capacity.** The total monitoring capacity of the system. This value is determined by the license (s) for the Database Server(s) or All-In-One Appliance(s) in the system.
 - **Number of Devices.** The number of devices currently discovered in the system.
 - **System Usage.** The amount of **Capacity** that the devices in the system are currently using. This value is the sum of the **Device Ratings** for all devices in the system. The **Device Rating** for each device is calculated daily and is based on the number of collections performed for that device.
- If you have a subscription license, you can also generate reports about subscription licensing.

To view the **System Usage** page, go to the **System Usage** page (System > Monitor > System Usage).

Viewing an Overview of All Events

The **Event Overview** page (System > Monitor > Event Overview) provides a graphical overview of all events in Skylar One.

The **Event Overview** page displays the following reports:

- **Number of Events by Severity.** This graph displays event distribution by severity for the last 24 hours and for the last 7 days.
 - The y-axis displays the number of events.
 - The x-axis displays severity.
 - The red line represents events in the last 24 hours.
 - The blue line represents events in the last 7 days.
 - Mousing over a data point in the red line displays the number of events of the specified severity in the last 24 hours.
 - Mousing over a data point on the blue line displays the number of events of the specified severity in the last 7 days.
- **Most Common Event Types.** This pie graph displays the ten most frequently occurring events for the last 7 days.
 - Each slice of the pie represents an event type. The legend on the left maps each slice color to an event and lists the actual number of events of that type.
 - The graph displays percent. Compared to the total number of occurrences for the top ten events, each slice displays the percent that belong to a specific event.
- **Mean Time-to-Resolution.** This bar graph displays the number of events generated in the last 24 hours, 7 days, 14 days, and 30 days, and their average resolution time.

- The y-axis displays the number of events.
- The x-axis displays the time span. There is a bar for 24 hours, 7 days, 14 days, and 30 days.
- The red bars represent the actual number of events associated with the time-to-resolution.
- The blue bars represent the average number of events associated with the time-to-resolution.
- Mousing over a bar displays the number of events associated with the time-to-resolution.

Viewing Events by Appliance and Event Source

The **Event Statistics** page displays a graph of the number of events processed by a selected Database Server, Data Collector, or Message Collector.

The **Event Statistics** page displays the following information:

- **Appliance.** In the field in the upper left, select from the list of all Database Servers, Data Collectors, and Message Collectors.
- **Event Type.** In the next field on the upper left, select from the list of event types. The choices are:
 - *API.* The event was generated by an external API.
 - *Dynamic.* Event was generated by a monitoring application running on the device.
 - *Email.* The event was generated by an incoming email.
 - *Internal.* Event was generated by Skylar One.
 - *Syslog.* Event was generated from standard system log generated by device.
 - *Trap.* Event was generated by an SNMP trap.

The graph displays the average number of events processed by the selected appliance, for the selected duration.

- The y-axis displays the average number of events.
- The x-axis displays time. The increments vary depending upon the selected data type (from the **[Options]** menu) and the date range (from the **Date Range Selection** pane).
- Mousing over any point in any line displays the value at that time-point in the **Mouse-over** column in the **Data Table** pane.
- You can use your mouse to scroll the report to the left and right.

Chapter

7

Admin Notifier

Overview

This chapter describes the **Admin Notifier**, which alerts administrator users upon login to any issues on monitored Skylar One (formerly SL1) appliances that could lead to an outage, such as the database running out of space.

After login, the Admin Notifier displays a banner if any of the filtered events are found so that you can take action on the issues immediately.

This chapter covers the following topics:

<i>How Does the Admin Notifier Work?</i>	235
<i>What Issues Does the Admin Notifier Monitor?</i>	235

How Does the Admin Notifier Work?

When certain events are present that could lead to a system outage, the Admin Notifier appears in Skylar One for administrator users at login and persists as a banner at the top of the screen until no qualifying events are found. The Admin Notifier monitors your Skylar One appliances for the events listed in the table in [What Issues Does the Admin Notifier Monitor?](#)

NOTE: If you are using the classic Skylar One user interface, the Admin Notifier has additional functions. See [Admin Notifier in the Classic Skylar One User Interface](#).

After logging in to Skylar One, administrators will see a modal if qualifying events are found, warning of events that might lead to system failure.

Click **[OK]** to continue or click **[View Events]** to go to the **Events** page to see a view filtered by admin events.

The System Warning banner will persist across the user interface until all of the qualifying events have been acknowledged.

To interact with the banner, you can click one of the buttons that appears on the banner:

- **Dismiss.** Hides the banner for the duration of your session or until another qualifying event is logged. If qualifying events are still present, the banner will reappear upon your next login.
- **Snooze.** Hides the banner for 15 minutes. If qualifying events are still present after 15 minutes, the banner will reappear.
- **View Events.** Opens the **Events** page with a filter applied that shows only the admin events.

Admin Notifier in the Classic Skylar One User Interface

The System Warning banner appears in the classic user interface.

In the classic user interface, you can select events to add to or remove from the System Warning banner. On the **Event Policy Manager** page (Registry > Events > Event Manager), the **Select Action** menu on the Event Policy Manager page includes the following options: *Enable Admin Banner Warning* and *Disable Admin Banner Warning* for selected events.

The **Event Policy Manager** page also includes a new **Admin Warning** column that specifies whether or not an event will be included in the System Warning Banner.

What Issues Does the Admin Notifier Monitor?

The following table includes the events that might appear in the Admin Notifier, the Dynamic Application that triggers the event, the severity of the event, and references that can help you resolve the problem before it causes an outage. Knowledge Base Articles (KBA) are found on the ScienceLogic Support website (<https://support.scienceologic.com/s/knowledge>).

Event Name	Dynamic Application	Severity	Reference
Support: Database Config - License expires in less than 14 days	Support: Database Configuration	Critical	Renew license. See the chapter on <i>Licensing and Configuring an Appliance</i> in the Installation manual or Knowledge Base Article 4978 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: DRBD Proxy License Expiration Eminent	Support: DRBD Proxy License Expiration	Critical	Renew license. See the chapter on <i>Licensing and Configuring an Appliance</i> in the Installation manual or Knowledge Base Article 4979 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
InnoDB Space: Critical Threshold	Support: InnoDB Size	Critical	See the chapter on <i>Device Thresholds and Data Retention</i> in the Device Management manual or Knowledge Base Article 49870 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: Skylar One Config - Appliance Not Licensed	Support: Skylar One Configuration	Critical	License your Skylar One appliance. See the chapter on <i>Licensing and Configuring an Appliance</i> in the Installation manual or Knowledge Base Article 4981 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: MariaDB Config - InnoDB Force Recovery Non-Zero Value	Support: MariaDB Configuration	Critical	See Knowledge Base Article 4982 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Poller: File system usage exceeded (critical) threshold	Internal	Critical	See Knowledge Base Article 4983 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: Appliance Validation Not Monitored	Support: Appliance Validation	Major	Discover all appliances as devices.
Support: Appliance Validation Minimum AIO Requirements Not Met	Support: Appliance Validation	Major	See Knowledge Base Article 4984 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: Appliance Validation Minimum AP Requirements Not Met	Support: Appliance Validation	Major	See Knowledge Base Article 4985 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).

Event Name	Dynamic Application	Severity	Reference
Support: Appliance Validation Minimum DB Requirements Not Met	Support: Appliance Validation	Major	See Knowledge Base Article 4986 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: Appliance Validation Minimum DC Requirements Not Met	Support: Appliance Validation	Major	See Knowledge Base Article 4987 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: Appliance Validation Minimum MC Requirements Not Met	Support: Appliance Validation	Major	See Knowledge Base Article 4988 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: API Message Backlog Has Exceeded Threshold	Support: Async Message Backlog Performance	Major	See Knowledge Base Article 4989 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: Syslog Message Backlog Has Exceeded Threshold	Support: Async Message Backlog Performance	Major	See Knowledge Base Article 4990 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: Trap Message Backlog Has Exceeded Threshold	Support: Async Message Backlog Performance	Major	See Knowledge Base Article 4991 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: Cluster Resource Template Outdated	Support: Cluster Configuration	Major	See Knowledge Base Article 4992 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: DNS Resolution - Invalid Response	Support: Database Configuration	Major	See Knowledge Base Article 4993 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: Skylar One will run out of space in less than 14 days	Support: DB Space Estimator	Critical	See Knowledge Base Article 4994 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).

Event Name	Dynamic Application	Severity	Reference
Mail Backlog: Critical	Support: Mail Backlog	Critical	See Knowledge Base Article 4995 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: MariaDB Config - File per table	Support: MariaDB Configuration	Critical	See Knowledge Base Article 4996 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: MariaDB Connections High - Critical	Support: MariaDB Performance	Critical	See Knowledge Base Article 4997 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: High Rows Behind High	Support: Rows Behind	Critical	See Knowledge Base Article 4998 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: Medium Rows Behind High	Support: Rows Behind	Major	See Knowledge Base Article 4999 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Siteconfig Check: Differing Variable	Support: Siteconfig Check	Critical	See Knowledge Base Article 5000 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Siteconfig Check: Missing Variable	Support: Siteconfig Check	Critical	See Knowledge Base Article 5001 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: Skylar One Config - Bad Timezone	Support: Skylar One Configuration	Critical	Reset timezone to UTC. See Knowledge Base Article 5002 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: Skylar One Config - Pause file detected	Support: Skylar One Configuration	Critical	The system is not currently collecting any information. See Knowledge Base Article 5003 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: Out of Memory - Process Has Been Killed	Syslog	Major	See Knowledge Base Article 5004 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).
Support: MariaDB Error	Syslog	Major	See Knowledge Base Article 5005 on the ScienceLogic Support website (https://support.scientelogic.com/s/knowledge).

Chapter

8

Diagnostic Tools

Overview

This chapter describes some diagnostic tools for troubleshooting and diagnosing problems in 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>Viewing Information About ScienceLogic Processes</i>	240
<i>Debugging a Process and Viewing Debug Logs</i>	241
<i>Viewing Information About Unhandled Exceptions</i>	243
<i>Viewing the Output of the System Status Script</i>	243
<i>Viewing the Database Tables on the Database Server</i>	243
<i>Disabling Normalization, Re-Enabling Normalization, and Backfilling Raw Data</i>	245
<i>Enable Logging for Data Pull Storage Objects</i>	248
<i>Controlling Developer Log Settings</i>	249

Viewing Information About ScienceLogic Processes

The **Process Manager** page allows you to view a list of ScienceLogic processes and optionally define parameters for those processes. These processes gather, manipulate, and publish the data used in Skylar One.

CAUTION: ScienceLogic recommends that you do not edit the values in this page without first consulting ScienceLogic. Incorrect values can severely disrupt ScienceLogic platform operations.

ScienceLogic processes fall into three scheduling categories or *Frequencies*:

- **Asynchronous.** The process is launched in response to a system event or user request.
- **Scheduled.** The process is launched on a regular schedule.
- **Always.** The process always runs while Skylar One is running.

Skylar One performs many tasks in parallel:

- Through a modular design, allowing functions to be distributed to multiple processing platforms.
- Through multi-processing, where multiple instances of a process run simultaneously.

Viewing the List of ScienceLogic Processes

To view the list of process in the **Process Manager** page:

1. Go to the **Process Manager** page (System > Settings > Admin Processes).
2. The **Process Manager** page displays information about each ScienceLogic process. The **Process Manager** page displays the following for each process:
 - **Process Name.** Name of the process.
 - **Program File.** Name of the executable file associated with the process.
 - **Frequency.** Frequency with which the platform launches the process. Possible values are:
 - **Asynchronous.** The process is launched in response to a system event or user request.
 - **Always.** The process always runs while Skylar One is running.
 - **Scheduled.** The process runs at intervals ranging from 1 Minute to Daily.
 - **Runtime Offset.** This field applies only to scheduled processes and allows the platform to stagger the launch of a process. The field specifies the number of minutes after the default scheduled time to execute a process. The default scheduled time at which processes are initially executed is midnight UTC. So if a process has a **Frequency** of 5 Minutes and the **Runtime Offset** is set to "2", the process will execute at two minutes past UTC midnight, seven minutes past UTC midnight, 12 minutes past UTC midnight, 17 minutes past UTC midnight, etc. Choices range from 0-1439.

- **Async Throttle.** This field applies only to asynchronous processes. This field indicates the number of jobs per process that can run simultaneously.
- **Batch Factor.** This field applies only to scheduled processes and determines how many multithreaded child processes are spawned on each execution of the process.

$$\text{number of tasks a process is responsible for completing} / \text{Batch Factor} = \text{number of child processes that will be spawned}$$
 - The number of tasks is typically determined by the number of devices the process is collecting data from.
 - The maximum number of child processes is limited by the number of CPUs installed in the Skylar One appliance that runs the process.

NOTE: **Batch Factor** defines the maximum number of worker processes or child processes. This value has precedence over the value specified in the section of this manual on *Tuning Collector Groups in the silo.conf File*.

- **Time Factor.** Determines how long the process can run before being stopped by the process manager. This setting only applies to asynchronous processes and scheduled processes. For asynchronous processes, this is the length of time an instance of the process can run. For scheduled processes, the value of **Time Factor** is used to calculate **Run Length**.

$$(\text{Frequency} * \text{Time Factor}) + \text{Frequency} = \text{Run Length}$$

For example, suppose a process runs every 15 minutes (as specified in the **Frequency** field). A **Time Factor** of 2 means the process is allowed to run for 45 minutes. A **Time Factor** of 0 means the process is allowed to run for 15 minutes.

- **Run Length.** Specifies how long the process can run before being stopped by the process manager. This number is based on the **Time Factor** for the process.
- **State.** Current operational state of the process. Possible values are:
 - *Enabled.* Process can run.
 - *Disabled.* Process cannot run.
- **Debug.** Specifies whether debugging information is enabled for the process. For more details on debugging a process, see the section [Debugging a Process](#).
- **ID.** Unique numeric ID assigned to each process by Skylar One.
- **Edited By.** Date and time the process settings were last edited.
- **Edit Date.** Date and time the process settings were last edited.

Debugging a Process and Viewing Debug Logs

When you debug a process, you tell Skylar One to use verbose logging for that process. You can then view Skylar One log file to view the logs.

There might be circumstances where you have narrowed down a problem to a specific ScienceLogic process (for example, based on an error message or event). When this happens, you might find it helpful to turn on debugging for that process and view the debug logs.

WARNING: ScienceLogic recommends that you enable the debug option *only* while troubleshooting a problem while working with ScienceLogic Support or while following a troubleshooting guide, and that you then immediately turn off debugging when you have completed troubleshooting. Do not leave the debug option enabled during normal operation of Skylar One. When you turn on debugging, Skylar One will run significantly more slowly.

NOTE: You cannot enable debug mode for the "Message Collection: SNMP Trap" or "Message Collection: Syslog" processes.

To enable the debug option for a process:

1. In the **Process Manager** page, find the process you want to edit. Select its wrench icon (🔧).
2. The **Process Editor** page appears and is populated with values for the selected process.
3. Edit the following field:
 - **Debug**. Enables or disables debugging information for a process. Select *Enabled*.
4. Click the **[Save]** button in the **Process Editor** page.
5. Log in to the console of the appliance where the process is running. Alternately, you can use SSH to open a shell session on the appliance. Log in as **em7admin** with the appropriate password.

TIP: To view a list of IP addresses for all appliances in your system, go to the **Appliance Manager** page (System > Settings > Appliances).

6. If the process you are debugging is a process that has a **Frequency** of *Always*, you must restart the process to make it pick up the new debug status (enabled). To restart the process, enter the following at the shell prompt:

```
sudo service process_name restart
```

For example, if you were debugging the process for the event engine, you would enter:

```
sudo service em7_event restart
```

7. Navigate to the directory **/var/log/em7**. View the file **silo.log**. The most recent entries will be posted at the end of the file.

8. After you have finished troubleshooting the process, remember to disable debugging. If the process has a **Frequency** of *Always*, you must restart the process to make it pick up the new debug status (disabled).

Viewing Information About Unhandled Exceptions

An **exception** specifies that something happened "out of the norm" that is preventing the software from executing the next step. Exceptions are a specific type of error, usually the result of invalid input, missing input, or a network error that prevents communication between software modules. For most exceptions, Skylar One will handle the exception by logging a specific error in the System Logs and will continue to run the process. However, *if the platform does not handle the exception*, the process will stop running, and Skylar One will generate an error message describing *the unhandled exception*.

Viewing the List of Unhandled Exceptions

To view the list of unhandled exceptions for all appliances:

1. Go to the **Unhandled Exceptions** page (System > Monitor > Unhandled Exceptions).
2. The **Unhandled Exceptions** page displays the following for each unhandled exception:
 - **Exception Filename**. Full path of the file where the exception occurred.
 - **Line**. Line number of the line in the file where the exception occurred.
 - **Exception Information**. Error message associated with the exception.
 - **First Occurrence**. Date and time of the first occurrence of the exception.
 - **Last Occurrence**. Date and time of the last occurrence of the exception.
 - **Count**. Number of times the exception has occurred.

Viewing the Output of the System Status Script

For each Database Server, Data Collector, and Message Collector, you can view the output of the system status script for that appliance. To do this:

1. Go to the **Appliance Manager** page (System > Settings > Appliances).
2. Locate the Database Server, Data Collector, or Message Collector that you want to view diagnostic information about.
3. Click on its magnifying-glass icon (🔍) to view the output of the system status script for that appliance.

Viewing the Database Tables on the Database Server

In some circumstances, you might need to view the contents of the database tables (the permanent tables are stored on the Database Server). There are two ways to do this:

- Using the built-in Database Tools in the **Database Tool** page (System > Tools > DB Tool).

NOTE: The **Database Tool** page is available only in versions of Skylar One prior to 12.2.1 and displays only for users that have sufficient permissions to access the page.

- Using the link to the phpMyAdmin interface in the **Appliance Manager** page (System > Settings > Appliances).

Accessing the Database Tool

The **Database Tool** page allows administrators to view information about the internal ScienceLogic databases and run SQL queries against those internal databases.

NOTE: The **Database Tool** page is available only in versions of Skylar One prior to 12.2.1 and displays only for users that have sufficient permissions to access the page.

CAUTION: Contact ScienceLogic for details on using the **Database Tool** page and troubleshooting databases. Do not make changes to the database or run the Optimizer Tool without guidance from ScienceLogic.

To access the database tool:

- Go to the **Database Tool** page (System > Tools > DB Tool).
- To run an SQL query** from the **Database Tool** page, enter values in the following fields:
 - Select Database.** Select a database to query.
 - SQL Query.** Enter an SQL query to execute against the selected database. For more information on each database and each table, use the options in the **[Actions]** menu.

NOTE: You must be familiar with SQL and know how to build a proper query before using the **Database Tool** page.

- Click the **[Go]** button to execute the query.
- The results from the query are displayed in the pane at the bottom of the page.
- To view the reports** about the a database(s), click the **[Actions]** menu. The following options are available:
 - Engines.** Displays status information about the server's storage engines. For each engine, the modal page displays a description of the engine, whether the engine is supported by Skylar One, and whether or not the engine supports transactions, XA, and save points.

- **Global Status.** Displays a list of global variables used in the database tables and the current value for each global variable.
- **InnoDB Variables.** Displays a list of InnoDB variables used in Skylar One and the value for each variable.
- **Open Tables.** Displays a list of currently open tables. For each table, the modal page displays the database name, table name, whether the table is currently in use, and whether the table is currently locked.
- **Optimizer Tool.** Leads to the **Database Optimizer Tool** page, where you can choose to optimize, repair, check, flush, or analyze all the tables in a database.

CAUTION: Contact ScienceLogic for details on using the **Database Optimizer Tool** page. Do not run the Optimizer Tool without guidance from ScienceLogic.

- **Processes.** Displays a list of running threads on the databases and tables. For each process, the modal page displays the connection ID, the database user who issued the statement, the host name of the client that issued the statement, the affected database, the command, the time in seconds that the thread has been in its current state, the state of the thread, and any available description of the process.
- **Table Status.** Displays the status of each database table in the platform. For each table, the modal page displays the table name, the database engine, database version, row format, number of rows, average row-length, length of the data file, maximum length of the data file, length of the index file, number of allocated but unused bytes, the next auto-increment value, the create time for the table, the update time for the table, the table's character set and collation, the live checksum value, options used with CREATE TABLE, and any comments.
- **Variables.** Displays a list of all database system variables used in Skylar One and the value of each variable.

Disabling Normalization, Re-Enabling Normalization, and Backfilling Raw Data

ScienceLogic does not recommend stopping normalization on Data Collectors. However, there are rare occasions where ScienceLogic Customer Support might ask you to disable normalization as part of troubleshooting.

To disable normalization:

1. Either go to the console of the Database Server or use SSH to access the Database Server.
2. Log in as user em7admin with the appropriate password.
3. Type the following at the command line:

```
sudo visilo
```

This is the file where users can customize the silo.conf file. In step #6, you will execute a command that sends these changes to the system silo.conf file.

4. In the LOCAL section, add the following line:

```
rollups_disabled=ON
```

5. Save your changes and exit the file (:wq).
6. Restart the data collection process to ensure they receive the change. Type the following at the command line:

```
sudo service em7_hfpulld restart
```

```
sudo service em7_lfpulld restart
```

```
sudo service em7_mfpulld restart
```

To re-enable normalization and normalize data that was collected while normalization was disabled:

1. Either go to the console of the Database Server or use SSH to access the Database Server.
2. Log in as user em7admin with the appropriate password.
3. Type the following at the command line:

```
sudo visilo
```

4. In the LOCAL section, add the following line:

```
rollups_disabled=OFF
```

5. Save your changes and exit the file (:wq).
6. Restart the data collection process to ensure collectors receive the change. Type the following at the command line:

```
sudo service em7_hfpulld restart
```

```
sudo service em7_lfpulld restart
```

```
sudo service em7_mfpulld restart
```

7. At the command line, type the following to normalize the data that was collected while normalization was disabled:

```
[/opt/em7/backend/data_normalizer_backfill.py --database, <database>
--dids <[device IDs]> --start <start date> --end <end date> --workers
<number of workers>
```

NOTE: To get help, at the shell prompt, type "/opt/em7/backend/data_normalizer_backfill.py -h".

where:

- **--database *database*.** Specifies the database that you want to backfill with normalized data. The choices are:
 - data_avail. Table that stores normalized data for availability.
 - data_cv. Table that stores normalized data for Web Content policies.
 - data_dns. Table that stores normalized data for DNS policies.
 - data_email. Table that stores normalized data for Email Round-Trip policies.
 - data_ports. Table that stores normalized data for TCP-IP Ports policies.
 - data_procs. Table that stores normalized data for System Processes policies.
 - data_services. Table that stores normalized data for Windows Services policies.
 - data_storage. Table that stores normalized data for file systems.
 - data_tv. Table that stores normalized data for SOAP/XML Transaction policies.
 - dynamic_app_data_appID. Table that stores normalized data for a Dynamic Application. Specify the application ID for the Dynamic Application.
- **--dids *device IDs*.** Specifies the device ID of the device or devices for which you want to normalize data.
 - You can specify a single device ID.
 - You can specify multiple device IDs, separated by commas and surrounded by square brackets.
 - If you do not specify any device IDs, Skylar One will normalize the specified data for all devices in your system.
- **--start *start date*.** The timestamp that specifies the data to normalize. Raw data with a timestamp at this time or later will be normalized. Skylar One will normalize data starting with this timestamp and ending with the end-date timestamp.
 - Specify the timestamp in the format yyyy-mm-dd hh:mm:ss, using a 24-hour clock. Surround the timestamp with single quotes.
- **--end *end date*.** The timestamp that specifies the data to normalize. Raw data with a timestamp at this time or earlier will be normalized. Skylar One will normalize data starting with the start-date timestamp and ending with this timestamp.
 - Specify the timestamp in the format yyyy-mm-dd hh:mm:ss, using a 24-hour clock. Surround the timestamp with single quotes.
- **--workers *workers*.** Number of worker processes to assign to this task. This field is optional. Please consult ScienceLogic Customer Support for suggestions on worker processes.

For example:

```
python /opt/em7/backend/data_normalizer_backfill.py --database dynamic_app_data_16 --start '2017-10-01 00:00:00' --end '2017-10-10 00:00:00' --workers 10
```

This command normalizes raw data collected by the Dynamic Application with an application ID of 16, associated with all subscriber devices (no device IDs specified, so defaults to "all devices"), and that was collected between midnight on October 1, 2017 and midnight on October 10, 2017. The `data_normalizer_backfill.py` code uses ten worker processes to perform the normalization.

Enable Logging for Data Pull Storage Objects

To investigate missed polls or slow database queries, you can temporarily enable logging for data pull storage objects. After you complete the diagnostics, you must disable logging for data pull storage objects, because the logging can affect the performance of data pull.

Enable

To enable logging for data pull storage objects:

1. Either go to the console of the Database Server or use SSH to access the Database Server.
2. Log in as an administrator.
3. At the shell prompt, enter the following:

```
sudo visilo
```

4. In the `silo.conf` file, add the following lines:

```
[DATAPULL]
```

```
log_storage_object_stats = 1
```

5. Save your changes to the file (`:wq`).
6. You must restart the data collection processes to ensure they receive the change. To do this, enter the following at the shell prompt:

```
sudo service em7_hfpulld restart
```

```
sudo service em7_lfpulld restart
```

```
sudo service em7_mfpulld restart
```

Disable

When you have completed your diagnostics, disable logging for data pull storage objects. To do this:

1. Either go to the console of the Database Server or use SSH to access the Database Server.
2. Log in as an administrator.
3. At the shell prompt, enter the following:

```
sudo visilo
```

4. In the silo.conf file, edit the following:

```
[DATAPULL]
```

```
log_storage_object_stats = 0
```

5. Save your changes to the file (:wq).
6. You must restart the data collection processes to ensure they receive the change. To do this, enter the following at the shell prompt:

```
sudo service em7_hfpulld restart
```

```
sudo service em7_lfpulld restart
```

```
sudo service em7_mfpulld restart
```

Controlling Developer Log Settings

In rare cases, you may need to modify log levels or suppression of certain logs in Skylar One, usually at the request of ScienceLogic Customer Support. To do so, you can use the following pages:

- **Skylar One Developer Logs** (System > Tools > Skylar One Developer Logs), which captures logs related to the default user interface (AP2), GraphQL, and associated services.
- **PHP Developer Logs** (System > Tools > PHP Developer Logs), which captures logs related to PHP execution, errors, and database queries for the classic Skylar One user interface.

Both of these pages are helpful for debugging and troubleshooting issues in Skylar One.

This section describes the options included on the **Skylar One Developer Logs** and **PHP Developer Logs** pages.

NOTE: These pages are available only for Administrator-level users in Skylar One.

Generating Developer Logs for the Default User Interface (AP2)

You can configure developer log settings for the default user interface (AP2) from the **Skylar One Developer Logs** page (System > Tools > Skylar One Developer Logs) and then generate logs filtered based on your configuration.

NOTE: This page is only available for Administrator-level users in Skylar One.

To generate developer logs for the default user interface (AP2):

1. Go to the **Skylar One Developer Logs** page (System > Tools > Skylar One Developer Logs).
2. In the **Log Levels** field, select the maximum log level you want to capture. In order of severity, your options are:
 - Error
 - Warning
 - Info
 - Debug
 - Trace
 - Log

NOTE: When you select a maximum log level, the log will capture messages of that severity level and higher. For example, if you select *Info*, the log will include messages with severities of Info, Warning, and Error.

3. In the **Capture log for the** field, select the log duration. This field determines the length of time during which the log info will be captured. Your options range from *next 1 minute* to *next 24 hours*.
4. Click **[Capture Log]**. The log will appear in the **Completed Logs** pane, along with its status and other information.
5. When the log file has finished compiling, you can click its hyperlink in the **Completed Logs** pane to download it. The completed log file includes a version of the `/var/log/s11/nextui.log` file, with the contents filtered based on the selections you made on the **Skylar One Developer Logs** page.

NOTE: Log files are automatically deleted 24 hours after they are created.

Generating Developer Logs for the Classic User Interface

You can configure developer log settings for the default user interface (AP2) from the **PHP Developer Logs** page (System > Tools > PHP Developer Logs) and then generate logs filtered based on your configuration.

NOTE: This page is only available for Administrator-level users in Skylar One.

To generate developer logs for the classic user interface:

1. Go to the **PHP Developer Logs** page (System > Tools > PHP Developer Logs).
2. In the **UI Developer Log Levels** section, select the types of messages that you want written to the user interface log file (`em7php.log`). Each type of message has an associated number; the log level is the sum of all enabled messages. The numbers and associated message types are:
 - **1. Critical**
 - **2. Error**
 - **4. Warning**
 - **8. Info**
 - **16. Debug**
 - **32. Trace**

NOTE: To determine the log level, sum the numbers associated with each type of message you want to enable. For example, if you want to enable Critical, Error, and Warning messages, then you would sum one, two, and four to get a log level value of seven.

3. In the **UI/REST MySQL Query Log Levels** section, select the types of messages you want written to the `mysqli.log` file. This log file collects every PHP-based call to MySQL and includes general information about the query. Determine the granularity of data you want by selecting one or more of the following checkboxes:
 - **Error**
 - **Warning**
 - **Info (non-error)**
 - **Request URI** (if you want the `mysqli.log` file to include the request uniform resource identifier)

4. In the **Advanced Settings** section, configure the suppressions and the date/time format you want to use by editing the following fields:
 - **Suppression List.** This list acts as a bitmask to log entries. For example, to suppress all entries for css-em7, you would enter "css.em7::127", where 127 is the sum of all possible log levels. You can specify multiple suppressions in the list, separated by commas.
 - **Datetime Format.** Specifies a user-defined date format that will be used for system logs. You can use any date variables supported by the PHP date function in this field.

NOTE: Seconds and milliseconds are always appended to the date/time stamp.

- **Include IP in log filenames.** Select this option to add the IP address from which the user is logged in to the name of each log file.

5. In the **Set Log Config** field, enter a length of time (in minutes) during which the log info will be captured.
6. Click **[Set Log Config]** to confirm your log configuration settings.
7. To download a log file, click the name of the log you want to download under **Download Logs**.

Chapter

9

Changing Administrator Passwords

Overview

This chapter describes how to change every administrator password used in Skylar One (formerly SL1).

NOTE: Appliances installed as an AWS EC2 instance have the "root" operating system account disabled by default. During the setup process, the user "ec2-user" is automatically added to the operating system configuration. The ec2-user account can be used to perform administrative tasks that require SSH command-line access. The ec2-user account is permitted to perform all operating system commands using the "sudo" command without a password.

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>Disabling phpMyAdmin</i>	254
<i>Changing the Password for the Default User Interface Account</i>	254
<i>Changing the Password for the Default Console User</i>	255
<i>Changing the Password for the Web Configuration Utility</i>	255
<i>Changing Database Passwords</i>	255

Disabling phpMyAdmin

The phpMyAdmin interface provides a web interface for viewing and managing MySQL databases. By default, you can log in to the Database Server using the phpMyAdmin interface to view and manage the MySQL databases on all Database Servers, Data Collectors, and Message Collectors in the system.

To disable phpMyAdmin, you must disable the service and then disable the ports on which the service runs. To do this:

1. Either go to the console of your Database Server or All-In-One Appliance, or use SSH to access the Database Server or All-In-One Appliance. Open a shell session on the server. Log in as an administrator.
2. Edit the file `/etc/siteconfig/firewalld-rich-rules.siteconfig`:

```
sudo vifirewalld
```

3. Add the following lines:

```
rule service name="phpmyadmin" reject\  
rule port port="8008" protocol="tcp" reject
```

4. Save your changes and exit the file (`:wq`).

Changing the Password for the Default User Interface Account

To change the password used by the **em7admin** user account to access the user interface:

1. Go to the **User Accounts** page (Registry > Accounts > User Accounts).
2. Click the wrench icon (🔧) for the em7admin user. The **Account Permissions** page appears.
3. Enter the new password in the **Change Password** field.
4. Re-type the new password in the **Confirm Password** field.

TIP: You can use the following special characters in the **em7admin** user account password:

```
+ _ ) ( * & ^ % $ # @ ! | } { " : ? > < = - \ ] [ ' ; / . ,
```

5. Click the **[Save]** button. A pop-up window appears, asking you to confirm the change.
6. Click **[OK]**. The message "Password Saved" is displayed.

Changing the Password for the Default Console User

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 the Password for the 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.

Changing Database Passwords

The following Skylar One appliances include a database instance:

- All-In-One Appliances
- Database Servers
- Data Collectors
- Message Collectors

By default, Skylar One appliances use the following user accounts to access appliance databases:

- ***clientdbuser***. This user is the default database user for MariaDB. This user has the same password as ***em7admin*** and ***root***, and the password is set during the initial installation. The ***clientdbuser*** does not have super privileges.
- ***ap_user***. This user is configured in the Database Server and is used by any appliance with a user interface (Database Server, All-In-One Appliance, Administration Portal) to access the database on a Database Server or All-In-One Appliance. By default, this user account has the user name ***apuser***.

Configuring a New MySQL Password on Database Appliances

WARNING: Exercise caution when manipulating MySQL user accounts. Do not use the following procedures unless you are confident in MySQL and know how to undo any changes, should something go wrong. Otherwise, contact ScienceLogic Support for assistance.

NOTE: If you are using a version of Skylar One prior to 11.3.0, follow the steps in this Support knowledge base article to change the password for ***dbuser***, which was used in older versions of Skylar One instead of ***clientdbuser***. <https://support.sciencelogic.com/s/article/1471>.

To change the password for the ***clientdbuser*** or ***ap_user*** account on the database appliance:

1. Either go to the console of the Database Server or All-In-One Appliance, or use SSH to access the server and log in as ***em7admin*** with the appropriate password.
2. If you have a high-availability (HA) cluster of database appliances, put the cluster into maintenance mode (`coro_config`). Otherwise, you can skip this step.
3. Stop Skylar One services by running following command:

```
sudo siloctl stop
```

4. Launch the MySQL prompt:

```
silo_mysql mysql
```

5. From the MySQL prompt, change the root password by running one of the following SQL queries, depending on which account password you are changing:

- For the **clientdbuser** account:

```
SET password FOR 'clientdbuser' = PASSWORD('<NEW_PASSWORD>');
```

- For the **ap_user** account:

```
SET password FOR 'ap_user' = PASSWORD('<NEW_PASSWORD>');
```

where **<NEW_PASSWORD>** is the password you want to configure for the account.

6. To effect the change immediately, you can run the following SQL query:

```
FLUSH PRIVILEGES;
```

7. Ensure you can access the database with the new password. Exit the MySQL interface (**exit**) and test the new password by running one of the following commands, entering the new password when prompted:

- For the **clientdbuser** account:

```
mysql -u clientdbuser -p
```

- For the **ap_user** account:

```
mysql -u ap_user -p
```

8. Restart Skylar One services:

```
sudo siloctl start
```

9. Configure the new password in the Database Server by updating the **/etc/silo.conf** file. To edit this file, run the following command:

```
visilo
```

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

- For the `clientdbuser` account:

```
[LOCAL]
```

```
dbpasswd = <NEW_PASSWORD>
```

```
[CENTRAL]
```

```
dbpasswd = <NEW_PASSWORD>
```

- For the `ap_user` account:

```
[CENTRAL]
```

```
ap_user = apuser
```

```
ap_pass = <NEW_PASSWORD>
```

11. Save the file (`:wq`) and enter `y` to move the changes to the `/etc/siteconfig/siloconf.siteconfig` file automatically.

12. Run the following command:

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

13. Repeat steps 9-12 on every Database Server or Administration Portal in your stack to update the passwords in the `/etc/silo.conf` file for those appliances as well.
14. If you have a high-availability (HA) cluster of database appliances and you put the cluster into maintenance mode in step 2, you can use (`coro_config`) again to remove it from maintenance mode.

Configuring a New MySQL Password on Collector Appliances

Perform the following steps to change the MySQL account password on an SL1 Collector:

1. Either go to the console of the Database Server, All-In-One Appliance, Data Collector, or Message Collector, or use SSH to access the server and log in as `em7admin` with the appropriate password.
2. Run the following command to launch the MySQL prompt:

```
silo_mysql mysql
```

3. From the MySQL prompt, change the root password by running the following SQL query:

```
SET PASSWORD FOR CURRENT_USER()= PASSWORD('new password');
```

4. To effect the change immediately, run the following SQL query:

```
FLUSH PRIVILEGES;
```

5. Ensure you can access the database with the new password. Exit the MySQL interface, and test by running the following command, entering the new password when prompted:

```
mysql -u clientdbuser -p
```

6. Configure the new password by updating the **/etc/silo.conf** file. To edit this file, run the following command:

```
visilo
```

7. In the **/etc/silo.conf** file, update the following section:

```
[LOCAL]
```

```
dbpasswd = <NEW_PASSWORD>
```

where **<NEW_PASSWORD>** is the password you want to configure for the account.

8. Save the file (**:wq**) and enter **y** to move the changes to the **/etc/siteconfig/siloconf.siteconfig** file automatically.
9. From the Skylar One user interface, go to the **Appliance Manager** page (System > Settings > Appliances), click the wrench icon (🔧) on the Collector, and then update the **DB User** and **DB Password** fields to reflect the new values. When you are done, click **[Save]**.
10. Go to the **Collector Status** page (System > Monitor > Collector Status) and confirm that the Collector has a **Collector State** of "Available".

Editing Silo.conf

To edit the **/etc/silo.conf** file:

1. Either go to the console of the Skylar One appliance, or use SSH to access the Skylar One appliance or and log in as **em7admin** with the appropriate password.
2. Type the following at the command line:

```
sudo visilo
```

3. Edit the file as needed.
4. Save and close the file (**:wq**).

Updating the master.system_settings_licenses Table

To update the **master.system_settings_licenses** table after you have changed the root password on a Data Collector or Message Collector:

1. Go to the **Appliance Manager** page (System > Settings > Appliances).
2. Locate the Data Collector or Message Collector in the list of appliances. Note the value in the **ID** column for the Data Collector or Message Collector.
3. Go to the **Database Tool** page (System > Tools > DB Tool).

NOTE: The **Database Tool** page is available only in versions of Skylar One prior to 12.2.1 and displays only for users that have sufficient permissions to access the page.

4. Enter the following in the **SQL Query** field, replacing `<new password>` with the new password and `<ID value of Collector>` with the value you noted in step 2:

```
UPDATE master.system_settings_licenses SET db_user='root', db_
pass=<new password> WHERE id=<ID value of Collector>;
```

If you want to update all Data Collectors and Message Collectors with the same password, enter the following in the SQL Query field, replacing `<new password>` with the new password:

```
UPDATE master.system_settings_licenses SET db_user='root', db_
pass='<new password>' WHERE function in (5,6);
```

5. Click the **[Go]** button.

Recovering the Root MySQL Password

To reset the root MySQL password if you become locked out:

1. Either go to the console of the Database Server or use SSH to access the server in CLI mode.
2. Log in as **em7admin** with the appropriate password.
3. Stop the em7 and mariadb services:

```
systemctl stop em7 mariadb
```

4. Start the mariadb service with the "`--skip-grant-tables`" option:

```
systemctl set-environment MYSQLD_OPTS="--skip-grant-tables" systemctl
start mariadb
```

5. Access the MySQL database:

```
mysql -u root mysql
```

6. Reset the root password from the MySQL prompt:

```
UPDATE user SET password=PASSWORD('<new password>') WHERE
User='root';
```

7. Stop the mariadb service again, unset the environment variable, and restart the service, using the following sequence of commands:

```
systemctl stop mariadb
```

```
systemctl unset-environment MYSQLD_OPTS
```

```
systemctl start mariadb
```

8. Ensure that you can access the MySQL database with the new password:

```
mysql -u root -p
```

9. Restart the em7 service:

```
systemctl start em7
```

10. Ensure that the password you set is also updated in the /etc/silo.conf `dbpasswd` variable. For more information, see [Configuring a New MySQL Password on Database Appliances](#)

Recovering the MySQL SNMP User Account on Data Collector

If you have removed the SNMP user account from the Data Collector's MySQL database in an attempt to harden your system, you must recover the account so that Skylar One can insert incoming SNMP traps into the database for processing.

To restore the SNMP user account:

1. Either go to the console of the Database Server or use SSH to access the server in CLI mode.
2. Log in as **em7admin** with the appropriate password.
3. Run the following command to restore the SNMP user account:

```
/opt/em7/share/scripts/em7_firstboot.d/30_trap_listener-db_init.sh
```

Changing the MariaDB Password on Skylar One Appliances

NOTE: This procedure should be used only when you do not know the current password and the Skylar One application cannot log in to the database.

Use the following instructions to change the MariaDB password in Database Servers, Data Collectors, Message Collectors, and All-In-One Appliances:

1. Either go to the console of the Database Server or use SSH to access the server in CLI mode.
2. Log in as **em7admin** with the appropriate password.

3. Determine the username that Skylar One uses for MariaDB access:

```
s11-config silo LOCAL dbuser
```

4. Stop all Skylar One services:

```
sudo siloctl stop
```

5. Access the MariaDB database with the super privileged account:

```
sudo /bin/mysql -u root mysql
```

6. Reset the password for the username that you identified in step 3 from the MariaDB prompt:

```
SET PASSWORD FOR '<username from step 3>@'%' = PASSWORD('<new password>');
```

where:

- *<username from step 3>* is the MariaDB username that you determined in step 3.
- *<new password>* is the updated password you want to establish for that username.

7. Exit the MariaDB database prompt:

```
\q
```

8. Edit the **silo.conf** file, as described in [Editing Silo.Config](#). Change the **dbpasswd** variable to the new password in both the [LOCAL] and [CENTRAL] sections.

NOTE: If you have clustered database appliances, be sure to update the **silo.conf** file for all cluster members.

NOTE: Upon saving, `visilo` will validate that the password works. If the password fails, ensure that you are typing it correctly, or that you set the password for the correct account from step 3.

9. Restart all Skylar One services:

```
siloctl start
```

Chapter

10

Changing the IP Address of a Skylar One Appliance

Overview

The IP address for an appliance is configured at installation. To change the IP address for an appliance after installation and preserve your Skylar One (formerly SL1) license, use the procedures in this section.

IMPORTANT: If you are running Skylar One on a cloud-based service like AWS or Azure, the following change procedures do not apply. The IP settings are provided automatically by DHCP, and configuration changes must be made in the DHCP provider's settings.

Moving a Skylar One appliance to a new network requires pre-planning. If your Skylar One configuration includes one or more Administration Portals, PhoneHome Collectors, or is configured for High Availability or Disaster Recovery, you must perform additional steps after changing IP addresses. The steps in this section allow you to change the IP address for a Skylar One appliance with minimal downtime.

NOTE: This procedure requires downtime, so plan to perform this procedure during a maintenance window.

CAUTION: Ensure console access to the appliance you are migrating in case of typographical or other errors that might prevent network access when changing IP addresses.

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>Changing the IP Address on an All-In-One Appliance</i>	264
<i>Changing the IP Address on a Database Server</i>	267
<i>Changing the IP Address on a Data Collector or Message Collector</i>	276
<i>Confirming the IP Address Change on the Appliance Manager Page</i>	278

Changing the IP Address on an All-In-One Appliance

To change the primary IP address of an All-In-One Appliance :

Step 1. Stop the EM7 Service

Before changing the IP address, you must stop the EM7 service. To stop the EM7 service:

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. Type the following at the command line:

```
sudo systemctl stop em7
```

Step 2. Change the IP Address in the silo.conf File

You must change the `ipaddress` value in the `silo.conf` configuration file.

1. Either go to the console of the Skylar One appliance or use SSH to access the server.
2. Type the following at the command line:

```
sudo visilo
```

3. Change the following line in the [LOCAL] section of the file to specify the new IP address:

```
ipaddress = new_IP_address
```

4. Save and quit the file (:wq).

Step 3. Change the IP Address in the /etc/hosts File

If the `/etc/hosts` file includes an entry for the appliance, update the entry with the new IP address.

1. Either go to the console of the Skylar One appliance or use SSH to access the server.
2. Type the following at the command line:

```
sudo vi /etc/hosts
```

3. If you see an IP address for the All-In-One Appliance, change the IP address to the new IP address.

NOTE: Updating the /etc/hosts/ file automatically restarts the dnsmasq service to ensure the service can read the updated file.

Step 4. Change the IP Address in the Network Interface Configuration File

NOTE: Be sure to set the IPADDR, PREFIX, GATEWAY and DNS# variables to the appropriate values for the new network. The PREFIX is the subnet mask in CIDR notation.

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. Type the following at the command line:

```
sudo ifconfig
```

4. Your output will look like this:

```
inet6 fe80::250:56ff:fe84:455f  prefixlen 64  scopeid 0x20<link>
  ether 00:50:56:84:45:5f  txqueuelen 1000  (Ethernet)
    RX packets 1774927  bytes 161985469 (154.4 MiB)
    RX errors 0  dropped 861  overruns 0  frame 0
    TX packets 1586042  bytes 158898786 (151.5 MiB)
    TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

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

5. Examine the output, find the interface that uses the old IP address, and note its name.
6. Use the vi editor to edit the settings for the interface. To do this, enter the following at the command line:

```
sudo vi /etc/sysconfig/network-scripts/ifcfg-interface name you noted  
in step #5
```

For example, from our output, we could enter:

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

7. The ifcfg file will look like this:

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

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

```
service network restart
```

Step 5. Update the IP Address in the MySQL Database

In this step, you must set the new IP address in the `master.system_settings_licenses` table so that when Skylar One is restarted, the new IP address is recognized as licensed.

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:

```
silo_mysql
```

4. At the mysql prompt, enter the following query:

```
UPDATE master.system_settings_licenses SET ip="[new IP address]"  
WHERE ip="[old IP address]" LIMIT 1;
```

For example:

```
[em7admin@hostname ~]$ silo_mysql  
  
MariaDB [(none)]> UPDATE master.system_settings_licenses SET  
ip="192.168.10.22" WHERE ip="10.1.1.240";  
Query OK, 1 row affected (0.01 sec)  
Rows matched: 1 Changed: 1 Warnings: 0  
  
MariaDB [(none)]>
```

5. Type "exit" to exit the MySQL session.

Step 6. Reboot the Appliance

Reboot the appliance to apply all of the changes you made.

The system will boot up and will start the interface with the new IP address. Skylar One will start up and will learn from the database that the new IP address matches its configuration file and the value in the database table. Therefore, Skylar One will keep the current license for the appliance.

Step 7. Confirm the Change in Skylar One

After changing the IP address for your Skylar One appliance, go to the **Appliance Manager** page (System > Settings > Appliances) and confirm that the correct IP address for that appliance appears in the **IP Address** column. If it does not, you must [update it](#).

Changing the IP Address on a Database Server

Changing the primary IP address of a Database Server requires additional steps if

- the Database Server resides in a High Availability configuration or a Disaster Recovery configuration
- the Database Server might connect to Data Collectors configured for PhoneHome

In addition, when you change the primary IP address of a Database Server, you must update the configurations for any Data Collectors, Message Collectors and Administration Portals that communicate with that Database Server.

WARNING: For Clustered Database Appliances (using HA, DR, or HA+DR), ScienceLogic recommends that you place the cluster into maintenance mode. Also, ScienceLogic recommends that you wait to wait to change the virtual IP address until all of the Database Servers have been moved to its new location, if applicable.

To change the IP address of a Database Server:

Step 1. Stop the EM7 Service

Before changing the IP address, you must stop the EM7 service. To stop the EM7 service:

1. Either go to the console of the Database Server 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 systemctl stop em7
```

Step 2. Change the IP Address in the silo.conf File

You must change the `ipaddress` value in the `silo.conf` configuration file. To do so:

1. Either go to the console of the Skylar One appliance or use SSH to access the server.
2. Enter the following at the command line:

```
sudo visilo
```

3. Change the following line in the `[LOCAL]` section of the file to specify the new IP address:

```
ipaddress = new_IP_address
```

4. Save and quit the file (`:wq`).

Step 3. Change the IP Address in the /etc/hosts File

If the `/etc/hosts` file contains an entry for the appliance, update the entry with the new IP address.

1. Either go to the console of the Skylar One appliance or use SSH to access the server.
2. Enter the following at the command line:

```
sudo vi /etc/hosts
```

3. If you see an IP address for the Database Server, change the IP address to the new IP address.

Step 4. Change the IP Address in the Network Interface Configuration File

NOTE: Be sure to set the IPADDR, PREFIX, GATEWAY and DNS# variables to the appropriate values for the new network. The PREFIX is the subnet mask in CIDR notation.

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 Database Server or use SSH to access the server.
2. Login as user **em7admin** with the appropriate password.
3. Enter the following at the command line:

```
sudo ifconfig
```

4. Your output will look like this:

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

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

5. Examine the output, find the interface that uses the old IP address, and note its name.

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

```
sudo vi /etc/sysconfig/network-scripts/ifcfg-interface name you noted  
in step #5
```

For example, from our output, we could enter:

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

7. The ifcfg file will look like this:

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

8. You can edit one or more of the following settings:

- **DNS1**=IP address of the DNS server that will be used by the Database Server.
- **IPADDR**=New IP address of the Database Server.
- **PREFIX**=netmask for the Database Server.
- **GATEWAY**=IP address of the network gateway that will be used by the Database Server.

9. Save your changes and exit the file (:wq)

10. At the command line, enter the following:

```
sudo service network restart
```

Step 5. Update the IP Address in the MySQL Database

In this step, you must set the new IP address in the `master.system_settings_licenses` table so that when the Database Server is restarted, Skylar One recognizes the new IP address as licensed.

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:

```
silo_mysql
```

4. At the mysql prompt, enter the following query:

```
UPDATE master.system_settings_licenses SET ip="[new IP address]"  
WHERE ip="[old IP address] LIMIT 1"
```

For example:

```
[em7admin@hostname ~]$ silo_mysql  
  
MariaDB [(none)]> UPDATE master.system_settings_licenses SET  
ip="192.168.10.22" WHERE ip="10.1.1.240";  
Query OK, 1 row affected (0.01 sec)  
Rows matched: 1 Changed: 1 Warnings: 0  
  
MariaDB [(none)]>
```

5. Enter "exit" to exit the MySQL session.

Step 5a: For Database Servers Configured with PhoneHome

If your Database Server is configured with PhoneHome, perform the following additional steps to change the IP address of the Database Server.

NOTE: If you add additional Database Servers with new IP addressees, wait for the Database Servers to migrate, and then the PhoneHome collectors will attempt to connect using the new Database Server IP addresses. After the databases reconnect, you can remove database entries with old IP addresses from the PhoneHome configuration.

1. Either go to the console of the Database Server or use SSH to access the server.
2. Enter the following at the command line:

```
phonehome status
```

NOTE: If you are planning to change the IP address of multiple Database Servers, you will want to update all of the relevant IP addresses in PhoneHome in this step.

3. The output will look like the following:

```
Phone Home Client configuration:  
Config Revisions: Device: 2 Destinations: 7 Global: 9  
  
Device Id Type      State  Status  Forwards  Name  
-----  
 15     collector  Enabled forwarded  Phone Home collector 15  
  
Device Id Type      State  Host/Ip    Port    Name  
-----  
 11     database  Enabled  192.168.2.2  7705  Phone Home database 11  
 12     database  Enabled  192.168.2.4  7705  Phone Home database 12  
 13     database  Enabled  192.168.2.6  7705  Phone Home database 13
```

4. Note the Device ID for the Database Server..
5. Run the `phonehome set` command to change the IP address for the device ID that corresponds to the Database Server. To do this, enter the following:

```
phonehome set <device_id> ip="<new_ip_address>"
```

where:

- `<device_id>` is the device ID you noted in step #4.
- `<new_ip_address>` is the new IP address.

For example

```
[root@<database_hostname username>]# phonehome set 11 ip="<new_ip_address>"
```

```
Reloading sshd configurations
```

6. For each Database Server that you want to change the IP address, perform step #5.

Step 5b For Clustered Database Appliances (using HA, DR, or HA+DR)

If your Database Servers are clustered for High Availability (HA), Disaster Recovery (DR), or both (HA+DR), to change the IP address of the Database Servers, you must also modify the clustering software configuration files, as described in this step.

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

WARNING: Changes to the running Cluster Resource Manager (CRM) configuration take effect immediately. ScienceLogic recommends that you wait to wait to change the virtual IP address until all of the Database Servers have been moved to its new location, if applicable.

2. You must edit the settings for the virtual IP for the cluster. At the command line, enter the following:

```
crm resource stop virtual_ip
crm resource param virtual_ip set ip <new_IP_address>
crm resource param virtual_ip set cidr_netmask <new_subnet_mask_in_
CIDR_notation>
crm resource start virtual_ip
```

3. In a High Availability configuration, the two Database Servers use two rings:

- ring0 defines the private interfaces that are connected directly to one another via crossover cable.
- ring1 defines the public interfaces that host the virtual IP and conduct the Skylar One related tasks.

To update these values in a High Availability configuration, you must edit the file /etc/corosync/corosync.conf..

4. Use a file editor like vi to edit /etc/corosync/corosync.conf.

5. You will see something like this:

```
nodelist {
  node {
    ring0_addr: 192.168.25.200
    ring1_addr: 10.1.20.25
    name: hardb1
    nodeid: 1
  }
  node {
    ring0_addr: 192.168.25.201
    ring1_addr: 10.1.20.26
    name: hardb2
    nodeid:2
  }
}
```

6. For the Database Server with the new IP address, edit the value for ring0_addr to match the new IP address. Save the file.

DRBD is the service that synchronizes the Database Servers in a High Availability or Disaster Recovery configuration. The DRBD file /etc/drbd.d/r0.res defines how data (on resource 0) is synchronized.

- In Skylar One configured for High Availability , DRBD uses the private IPs to synchronize high-availability data.
- In Skylar One configured for High Availability plus Disaster Recovery, DRBD uses the private interface to synchronize high-availability data and the virtual IP and the public IP addresses to synchronize data for disaster recovery.
- In Skylar One configured for Disaster Recovery, DRBD uses the virtual IP and the public IP addresses to synchronize data.=

7. Use a file editor like vi to edit /etc/drbd.d/r0.res. It will look something like this:

```
resource r0 {
    protocol A;
    device /dev/drdb1;
    stacked-on-top-of r0-L {
        address 127.0.0.1:7789;
        proxy on hadrdb1 hadrdb2 {
            inside 127.0.0.1:7790;
            outside 192.168.25.200:7788;
        }
    }
    on hadrdb2
        disk /dev/mapper/em7vg-db;
        address 127.0.0.1:7789;
        meta-disk internal;
        proxy on hadrdb3 {
            inside 127.0.0.1:7790;
            outside 192.168.25.201:7788;
        }
    }
}
```

8. Replace instances of the old IP address with the new IP address and save the file.
9. Shut down the Database Server .
10. Upon reboot of the Database Server, run a discovery session to rediscover the Database Server with its new IP address.

Step 6. Reboot the Appliance

Reboot the Database Server to apply all of the changes you made.

CAUTION: If you are migrating a High Availability cluster, shut down the secondary first, then the primary, so that Skylar One does not perform a failover. Restart up the primary first, and after it is up and running, restart on the secondary.

The system will boot up and will start the interface with the new IP address. Skylar One will start up and will learn from the database that the new IP address matches its configuration file and the value in the database table. Therefore, Skylar One will keep the current license for the appliance.

Step 7. Change the Database Appliance IP Address in the Administration Portals, Data Collectors, and Message Collectors

You must edit the configuration for each Skylar One node that communicates with the Database Server. To do so, perform the following steps on each Administration Portal, Data Collector, and Message Collector in your Skylar One system.

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.
4. Click the **[Device Settings]** button in the upper-right of the page. The **Settings** page appears.
5. In the **Settings** page, enter the following:
 - **Database IP Address**. The new IP address of the Database Server.
 - **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.**
6. Click the **[Save]** button. You may now log out of the Web Configuration Utility.

Step 8. Confirm the Change in Skylar One

After changing the IP address for your Skylar One appliance, go to the **Appliance Manager** page (System > Settings > Appliances) and confirm that the correct IP address for that appliance appears in the **IP Address** column. If it does not, you must *update it*.

Changing the IP Address on a Data Collector or Message Collector

You can change the IP address of a Data Collector or Message Collector in two ways:

- In the Web Configuration Utility
- From the command line of the Data Collector or Message Collector

Regardless of the method you choose, Skylar One will automatically update the IP address of the Data Collector or Message Collector in:

- /etc.silo.conf
- /etc/sysconfig/network-scripts/ifcfg-ens160
- /etc/hosts
- the master.system_settings_licenses table, in the ip column

The following sections explain the necessary steps.

Using the Web Configuration Utility to Change the IP Address of a Data Collector or Message Collector

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.
4. Select the **Interfaces** tab.
5. In the **Interfaces** page, select the interface that you want to edit.
6. Enter values in the following field:
 - **Interface IP Address.** Required. Enter the IP address for the bonded interface in standard IPv4, dotted-octet format.
7. The IP address for the Data Collector or Message Collector will be automatically updated in all the necessary configuration file and database tables.

Using the Command Line to Change the IP Address of a Data Collector or Message Collector

1. Either go to the console of the Data Collector or Message Collector or use SSH to access the server.
2. Enter the following at the command line:

```
sudo visilo
```

3. At the shell prompt, enter:

```
update_ip <interface_ID> <new_IP>
```

where:

- *interface_ID* is the name of the interface, usually ensn32 or ens160.
- *new_IP* is the new IP address to assign to the interface.

4. The IP address for the Data Collector or Message Collector will be automatically updated in all the necessary configuration file and database tables.

Confirming the IP Address Change on the Appliance Manager Page

After changing the IP address for your Skylar One appliance, go to the **Appliance Manager** page (System > Settings > Appliances) and confirm that the correct IP address appears for that appliance in the **IP Address** column.

If the **Appliance Manager** page does not display the new IP address, you can manually update the address. To do so:

1. Go to the **Appliance Manager** page (System > Settings > Appliances).
2. Click the wrench icon (🔧) of the appliance you want to update. Its values appear in the fields at the top of the page.
3. In the **IP Address** field, type the new IP address.
4. Click **[Save]**.

Chapter

11

Changing Domain Name Servers (DNS) and Host Names on a Skylar One Appliance

Overview

This chapter describes how to change domain name servers (DNS) and hostnames on a Skylar One (formerly SL1) appliance.

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>Changing Name Servers on a Skylar One Appliance</i>	280
<i>Changing Hostnames on a Skylar One Appliance</i>	281

Changing Name Servers on a Skylar One Appliance

Skylar One appliances include a local Domain Name Server (DNS) cache to improve performance. You can use the `/etc/dnsmasq-resolv.conf` file to change the DNS server information for the following Skylar One appliance types:

- Database Servers
- Data Collectors
- Message Collectors
- All-In-One Appliances

NOTE: You do not need to edit the `/etc/resolv.conf` file, which should have a single entry of
`nameserver 127.0.0.1`.

IMPORTANT: The following change procedure does not apply if you are running Skylar One on a cloud-based service like AWS or Azure. In that scenario, the DNS settings are provided automatically by DHCP, and configuration changes must be made in the DHCP provider's settings.

Additionally, the following procedure does not apply to [Administration Portals](#).

To update the name servers on most Skylar One appliance types:

1. Edit the `/etc/dnsmasq-resolv.conf` file by entering the following command:

```
sudo vi /etc/dnsmasq-resolv.conf
```

2. Change the `nameserver` entry in the file to the IP address of the new DNS or add new DNS entries to the file.
3. Save and quit to commit the changes. This change immediately causes the OS to use the new DNS, with no reboot or service restarts required. If you have multiple nameservers listed in the file, the system will try each entry in the list until it gets a response or runs out of nameservers.

Changing Name Servers on Administration Portals

For Skylar One Administration Portals, the Domain Name Server (DNS) server settings are configured at installation. You cannot adjust the DNS settings later through the Web Configurator. Instead, you must use the command line interface (CLI) to change the DNS server information. This action requires no downtime.

IMPORTANT: If you are running Skylar One on a cloud-based service like AWS or Azure, the following change procedures do not apply. The DNS settings are provided automatically by DHCP, and configuration changes must be made in the DHCP provider's settings.

To change the DNS settings for Skylar One Administration Portals:

1. Edit the `/etc/resolv.conf` file by entering the following command:

```
sudo vi /etc/resolv.conf
```

2. Change the `nameserver` entry to the IP address of the new DNS or add new DNS entries to the file.
3. Save and quit to commit the changes. This change immediately causes the OS to use the new DNS, with no reboot or service restarts required. If you have multiple nameservers listed in the file, the system will try each entry in the list until it gets a response or runs out of nameservers.

Next, add the DNS to the interface configuration file so that the change will persist if the network service is restarted or the Administration Portal is rebooted.

To add one or more domain name servers to the interface configuration file:

1. Either go to the console of the Administration Portal or use SSH to access it.
2. Log in as user `em7admin` with the appropriate password.
3. Determine the name of your primary interface (not the "lo" interface) by running the following command:

```
ip addr
```

4. Edit the corresponding interface configuration file in the `/etc/sysconfig/network-scripts` directory:

```
sudo vi /etc/sysconfig/network-scripts/ifcfg-{interface name}
```

5. Find the "DNS1" entry and change the IP address to the IP address of the new DNS.

NOTE: You can enter additional DNS servers and define them as DNS2, DNS3, and so on.

6. Save and quit to commit the changes.

Changing Hostnames on a Skylar One Appliance

To change hostnames on a Skylar One appliance:

1. Either go to the console of the Skylar One appliance or use SSH to access the server.
2. Set the hostname by running the following command:

```
sudo hostnamectl set-hostname <host.example.com>
```

3. Verify that the hostname was set:

```
hostnamectl

hostnamectl
  Static hostname: host.example.com
  Icon name: computer-vm
  Chassis: vm
  Machine ID: 59013338cba643ec8ed2ec9883dfddf4
  Boot ID: dde124fbc75c484ea7d71443277af659
  Virtualization: vmware
  Operating System: Oracle Linux Server 7.9
  CPE OS Name: cpe:/o:oracle:linux:7:9:server
  Kernel: Linux 3.10.0-1160.36.2.el7.x86_64
  Architecture: x86-64
```

The **em7** service will notice this change and update the **Appliances** page (System > Settings > Appliances) in a few minutes.

4. Edit the **/etc/hosts** file to replace the old hostname with the new one:

```
sudo vi /etc/hosts
```

5. Restart the **mysql** service to update the database **@@hostname**, which is how the command-line Message of the Day (MOTB) pulls the active Database Server:

```
sudo systemctl restart mysql
```

6. You can also restart the task manager service to get the changes picked up more quickly:

```
sudo systemctl restart em7.service
```

Chapter

12

Backup Management

Overview

This chapter describes how to prepare to back up your Skylar One (formerly SL1) system, to define and run your backups, and to restore from different backup types.

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>Types of Skylar One Backups</i>	285
<i>The Workflow for Backing Up and Restoring Skylar One</i>	288
<i>Planning for Skylar One Backups</i>	289
<i>Creating Backup Credentials</i>	289
<i>Configuring Backups</i>	291
<i>Enabling tmux</i>	298
<i>Mounting Backup Files</i>	299
<i>Restoring Backups</i>	302
<i>Unmounting Backup Files</i>	326
<i>Retaining Backups</i>	327
<i>Additional Configuration for Solaris NFS Remote Shares</i>	328

Types of Skylar One Backups

Skylar One allows you to define three types of backups for your system:

- **Configuration Backups**, which include the core database tables and files required to restore a Skylar One system. A configuration backup includes scope and policy information, but *not performance data, data collected using configuration Dynamic Applications, events, or logs*.
- **Full Backups**, which include all Skylar One databases and tables. A full backup does *not* back up the configuration files in `/etc/backup.conf`. As a result, PhoneHome configuration files are not backed up in a full backup; you can back up PhoneHome files and other configuration files in a configuration backup.
- **Disaster Recovery Backups**, which include a full backup of the disaster recovery database for Skylar One systems configured for disaster recovery (DR). The DR backup is similar to the full backup.

NOTE: If you have a large Skylar One system and large backup files, you have the option to [perform standard configuration or full backups on the disaster recovery Database Server](#). This is different than performing a disaster recovery backup.

The **Backup Management** page (System > Settings > Backup) lets you define your backups and run them on demand.

CAUTION: ScienceLogic does not support vmotion or VMware Snapshots for backups of data. For backup purposes, ScienceLogic supports only ScienceLogic backups to remote storage. Be advised that vmotion and VMware Snapshots can cause Skylar One outages.

Configuration Backups

A **configuration backup** stores a copy of the core database tables that are required to restore a Skylar One system. Configuration backups use the "MySQL dump" tool to create backups.

Note the following information about configuration backups:

- Skylar One automatically retains your last seven configuration backups, but will retain only one backup per day. If more than one configuration backup is created on the same day, Skylar One retains the most recent one.
- Skylar One can launch configuration backups automatically at the interval you specify.
- During configuration backup, the ScienceLogic database remains online.
- When the configuration backup starts, Skylar One creates a temporary mount point to your remote share.

- If you have a large system and very large backup files, you can use an alternative method to perform backups that reduces performance issues during backup. For more information, see the section [Performing Config Backups and Full Backups on a Disaster Recovery Database Server](#).

What Does a Configuration Backup Include?

A configuration backup includes:

- Scope and policy information, but *not performance data, data collected using configuration Dynamic Applications, or events*.
- By default, the following files are backed up during a configuration backup:
 - /etc/backup.conf
 - /etc/corosync/corosync.conf
 - /etc/drbd.d/r0.res
 - /etc/drbd-proxy.license
 - /etc/hosts
 - /etc/my.cnf.d/silo_mysql.cnf
 - /etc/nginx/*
 - /etc/phonehome/*
 - /etc/php-fpm.d/*.conf
 - /etc/postfix/main.cf
 - /etc/silo.conf
 - /etc/siteconfig/*
 - /etc/sl_vault/encryption_key
 - /etc/sl_vault/vault_conf.yml
 - /etc/ssh/*.key
 - /etc/ssh/*.pub
 - /etc/sysconfig/network-scripts/ifcfg-*
 - /etc/sysctl.d/*
 - /etc/systemd/system/mariadb.service.d/*.conf
 - /opt/em7/nextui/nextui.conf
 - /usr/libexec/postfix/main.cf
- All files and folders specified in /etc/backup.conf. If you have additional files that you want to include in configuration backups, you can include them in the file /etc/backup.conf. For more information, see the section on [Adding Files to Include in Configuration Backups](#).

- The following databases:
 - ***insight_agent***. Includes configuration information for the Skylar One agent.
 - ***master***. Includes system-level settings for Skylar One, Dynamic Application definitions and alignments, run book automation and action policies, monitoring policy definitions, and credentials.
 - ***master_access***. Includes user account information, access keys, and access hooks.
 - ***master_ap2***. Includes files from the new UI, including files from Business Services.
 - ***master_biz***. Includes asset information, dashboards, distribution lists, document templates, IT Service policy information, organization information, product SKU information, RSS feeds, ticketing information, and user preferences. By default, configuration backups do not include the *ticket_external_requests table* from the *master_biz* database.
 - ***master_custom***. Includes GUI customizations, dashboard widget definitions, and PowerPack files.
 - ***master_dev***. Includes information associated with device records, excluding performance data, data collected using configuration Dynamic Applications, events, or logs.
 - ***master_dns***. Includes DNS information.
 - ***master_events***. The configuration backup includes only the *event_suppressions* table from this database. This table stores event suppression settings.
 - ***master_filestore***. Includes information about files, PowerPacks, and notes. By default, configuration backups do not include the tables *metadata_system_package*, *metadata_system_patch*, *storage_system_package*, and *storage_system_patch*.
 - ***master_platform***. Includes information about ScienceLogic appliances.
 - ***master_reports***. Includes custom report definitions.
 - ***mysql***. Contains the configuration settings for the MariaDB database.
 - ***scheduler***. Includes all instances of scheduled items: reports, discovery sessions, etc.
 - ***sysinfo***. Contains the configuration settings for High Availability, Disaster Recovery, and PhoneHome Collectors.

Full Backups

A **full backup** creates a complete backup of the ScienceLogic database. Full backups use a built-in tool call MariaBackup.

Note the following information about full backups:

- Skylar One can launch full backups automatically at the interval you specify.
- During a full backup, the ScienceLogic database remains online.
- If you have a large system and very large backup files, you can use an alternative method to perform backups that reduces performance issues during backup. For more information, see the section [Performing Config Backups and Full Backups on a Disaster Recovery Database Server](#).

What Does a Full Backup Include?

A full backup includes all databases and tables.

A full backup does *not* back up the configuration files in `/etc/backup.conf`. As a result, PhoneHome configuration files are not backed up in a full backup; you can back up PhoneHome files and other configuration files in a configuration backup.

NOTE: For very large Skylar One systems, ScienceLogic recommends you use a SAN with snapshot technology to backup and restore data.

NOTE: If your Skylar One System uses AWS RDS (remote database), the **Full Backup** option is disabled.

Disaster Recovery Backups

Disaster recovery backups are similar to full backups, but are available only for Skylar One systems that are configured for disaster recovery (DR).

DR backups temporarily stop replication, mount the database, and run a full backup of the DR database. The process then re-enables replication and performs a partial resynchronization from the primary.

Note the following information about DR backups:

- DR backups use 'tar' to create a copy and compress the `/data.local/db` directory.
- During DR backup, the primary ScienceLogic database remains online.
- DR backups are not available on two-node high availability (HA) clusters.

What Does a Disaster Recovery Backup Include?

Disaster recovery backups include all configuration data, performance data, and log data.

NOTE: If your Skylar One System uses AWS RDS (remote database), the **DR Backup** option is disabled.

The Workflow for Backing Up and Restoring Skylar One

The workflow for backing up Skylar One is:

1. [Plan the backup](#).
2. [Create a backup credential](#).
3. [Configure the backup](#).

After the backup file has been created, should you ever need to restore that backup, the workflow is:

1. [Enable tmux](#), if necessary
2. [Mount the backup](#) (NFS and SMB shares only).
3. [Restore the backup](#).
4. [Unmount the backup](#) (NFS and SMB shares only).
5. [Retain the backup](#) (full and DR backups only).

Planning for Skylar One Backups

Before creating a backup of your Skylar One system, you must determine the following:

- The external system or service to which the backup will be stored. Your options are:
 - NFS file share
 - SMB file share
 - S3 storage service

NOTE: It is the responsibility of the system administrator for the external system to create any necessary entries in /etc/fstab required to allow the Skylar One system to access the share as root.

- The hostname or IP address of the system or service to which the backup will be stored.
- The directory to which you will write the backup.

Creating Backup Credentials

To configure a backup, you must include a credential that allows Skylar One to write to the external systems where you will store the backups. There are two types of credentials that you can create for this task:

- An [S3 Backup credential](#), if you are backing up your system to an S3 storage service
- A [Basic/Snippet credential](#), for all other backup scenarios

The sections below describe how to create both credential types.

Creating an S3 Backup Credential

You can use an S3 storage service to store configuration backups for Skylar One. To do so, you will need to create a credential that enables Skylar One to connect to the S3 service. Skylar One includes an **S3 Backup** credential type, which uses field names and terminology specific to S3 services, that you can use to connect with your S3 service.

NOTE: Skylar One supports the use of Amazon Web Services (AWS) or MinIO for S3 backup storage.

NOTE: Before creating an S3 backup credential, you must have the following information:

- The Access Key ID and Secret Access Key for the S3 account on which you want to store the backup
- The endpoint URL
- The encryption password and salt for the backup file

To define an S3 backup credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create S3 Backup Credential*. The **Create Credential** modal page appears.
3. Supply values in the following fields:
 - **Name.** Type a unique name for the credential. Can be any combination of alphanumeric characters, up to 64 characters.
 - **All Organizations.** Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the **What organization manages this service?** drop-down field to align the credential with those specific organizations. For more information about credentials and organizations, see the section on "Aligning Organizations with a Credential" in the **Discovery & Credentials** manual.
 - **Timeout (ms).** Type the time, in milliseconds, after which Skylar One will stop trying to communicate with the S3 storage service.
 - **Provider.** Select the S3 storage provider you want to use to store the backup. Choices are *Amazon Web Services (AWS) S3* and *Minio Object Storage*.
 - **Access Key ID.** Type the Access Key ID for the S3 account on which you want to store the backup.
 - **Secret Access Key.** Type the Secret Access Key for the S3 account on which you want to store the backup.
 - **Endpoint.** Type the URL of the S3 endpoint. The endpoint URL should not include the bucket name.
 - **Region.** Select the region of the S3 endpoint.

- **Bucket.** Type the name of the S3 bucket on which you want to store the backup.
- **Encryption Password.** Type the encryption password for the backup file.
- **Encryption Salt.** Type the encryption salt used to safeguard the backup file encryption password.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the section on "Using the Credential Tester Panel" in the *Discovery & Credentials* manual.

Creating a Basic/Snippet Backup Credential

If you are backing up your Skylar One system to a platform other than an Amazon Web Services S3 bucket, you must create a **Basic/Snippet** backup credential that allows Skylar One to write to the external backup systems.

To create a **Basic/Snippet** backup credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click **[Create New]** and select *Create Basic/Snippet Credential*. The **Create Credential** modal appears. If you are using the classic Skylar One user interface, go to the **Credential Management** page (System > Manage > Credentials), click the **[Actions]** button, and select *Create Basic/Snippet Credential* to use the **Credential Editor** modal.
3. Define values in the following fields:
 - **Name.** Name of the credential. Can be any combination of alphanumeric characters.
 - **Organization.** Select **All Organizations** or select an organization from the drop-down.
 - **Hostname/IP.** The hostname or IP address.
 - **Port.** This field is deprecated. Backups will not use this field.
 - **Timeout (ms).** This field is deprecated. Backups will not use this field.
 - **Username.** Username to use when connecting to the external system. If you are backing up to NFS-remote, this field is not required.
 - **Password.** Password to use when connecting to the external system. If you are backing up to NFS-remote, this field is not required.
4. Click **[Save & Close]**.

Configuring Backups

This section describes how to configure Skylar One backups.

Depending on the backup type, this process might include multiple steps:

1. *Include additional files or directories* (configuration backups only).
2. *View the included databases and tables* (configuration backups only).
3. Define the *configuration*, *full*, or *disaster recovery* backup options.

When you define the backup options, you must determine the external storage protocol and directory on which the backup will be stored. You have the following storage options:

- NFS mount
- SMB mount
- S3 storage service

NOTE: You cannot store backups on the Skylar One Database Server or All-In-One Appliance.

Adding Files to Include in Configuration Backups

All files and directories that are specified in **/etc/backup.conf** are included in configuration backups. If you have additional files or directories that you want to include in configuration backups, you can edit the **/etc/backup.conf** file to include them.

To add files or directories that you want to include in configuration backups:

1. Either go to the console of the Skylar One appliance or use SSH to access the server.
2. Enter the following at the command line:

```
sudo vi /etc/backup.conf
```

3. Move the cursor below the line that says, "Custom files can be added below this line".
4. Make the following updates based on what you want to include:
 - To add a file, enter the full directory path and filename.
 - To add a directory or folder, enter an asterisk (*) instead of a filename after the full directory path.
5. Save your changes and quit the file (`:wq`).

Defining a Configuration Backup

To define and schedule a configuration backup:

1. Go to the **Backup Management** page (System > Settings > Backup).
2. In the **Configuration Backup** pane, provide values in the following fields:
 - **Every**. Together with the *Interval* field, specifies how frequently Skylar One should automatically execute a full backup. Enter a number.

- **Interval.** Together with the **Every** field, you must specify how frequently Skylar One should automatically execute a configuration backup. Your choices are:
 - *Disabled.* Configuration backups are disabled. If you want to run on-demand configuration backups only, you do not need to change this setting.
 - *Day.* Skylar One will execute configuration backups daily as specified (for example, every 2 days).
 - *Week.* Skylar One will execute configuration backups weekly as specified (for example, every 2 weeks).
 - *Month.* Skylar One will execute configuration backups monthly as specified (for example, every 1 month).
- **Start Time / Date.** If you enabled configuration backups, you must specify the start time. You must specify the date on which you want configuration backups to begin, as well as the time of day you want the backups to run. For example, you might want to run configuration backups during a maintenance window late at night. Click the field to open a calendar and time selector.
- **Timezone.** Optional. Specify the timezone to use when running a backup. The default is UTC. Use the drop-down list to select the timezone.
- **Configuration Credentials.** Select the credential you created for Skylar One to use to connect to your external system (NFS, SMB, or S3). For more information, see [Creating a Backup Credential](#).
- **Configuration Protocol.** Specify the type of external system where the configuration backup will be stored. Choices are:
 - *NFS-Remote.* When you select this option, Skylar One stores the configuration backup on an NFS share. If you select the *NFS-remote* option, and your NFS share is hosted on a Solaris system, you must perform the steps listed in the [Additional Configuration for Solaris NFS Mounts](#) section.
 - *SMB-Remote.* When you select this option, Skylar One stores the configuration backup on an SMB share.
 - *S3.* When you select this option, Skylar One stores the configuration backup in an S3 bucket.
- **Configuration Subdirectory.** Specify a directory on the NFS or SMB share or S3 service in which you would like to store the configuration backup. When entering the subdirectory path, include the leading slash (""). ScienceLogic recommends creating and configuring a unique subdirectory for each platform that you are backing up to the share; do not back up multiple platforms to the exact same subfolder in a share. You will need this value as well as your [backup credential information](#) and any mount options that you specify in step 3 below if you [mount the backup](#) to an NFS or SMB share.
- **Configuration Backup on DR.** Select this checkbox if you want to perform configuration backups from the Disaster Recovery Database Server. This alternative method to backing up Skylar One can reduce performance issues during backup for users with large systems and very large backup files. For more information, see the section on [Performing Configuration and Full Backups on the DR Database Server](#).

NOTE: Skylar One will retain your two most recent configuration backups.

3. To access additional configuration options for this backup, click **[Advanced Settings]**. These advanced fields refer to the column names in the Skylar One Database:
 - **Configuration SMB Mount Options.** Corresponds to the backup_smb_mount_options database column.
 - **Configuration NFS Mount Options.** Corresponds to the backups_nfs_mount_options database column.
 - **Configuration CMD Options.** Corresponds to the backup_cmd_options database column.
 - **Configuration Comp CMD Options.** Corresponds to the comp_cmd_options database column.
 - **Configuration DB List (JSON).** Corresponds to the backup_db_list database column.
 - **Configuration CB Table Exclude (JSON).** Corresponds to the backup_cb_table_exclude database column.

NOTE: For more information about these advanced settings, contact ScienceLogic Support.

NOTE: You will need the values from the mount option fields as well as your *backup credential information* and the value you entered in the **Configuration Subdirectory** field in step 2 above when mounting an NFS or SMB share.

4. Click **[Save]**. Skylar One will execute the configuration backup at the specified interval, starting on the date you specified in the **Start Time / Date** field and using the time you specified in the **Start Time / Date** field.
5. To run the backup immediately, click the **[Backup Now]** button under **Configuration Backup** in the **Immediate Backup** section of the **Backup Management** page. Skylar One will run the backup immediately, as well as running the scheduled backup you configured in this procedure.

Defining a Full Backup

Full backups include all databases and tables. Skylar One automatically launches this backup at the frequency and time you specify.

To define and schedule a full backup:

1. Go to the **Backup Management** page (System > Settings > Backup).
2. In the **Full Backup** pane, provide values in the following fields:

- **Every.** Together with the *Interval* field, specifies how frequently Skylar One should automatically execute a full backup. Enter a number.
- **Interval.** Together with the *Every* field, specifies how frequently Skylar One should automatically execute a full backup. Your choices are:
 - *Disabled.* Full backups are disabled.
 - *Day.* Skylar One will execute full backups daily as specified (for example, every 2 days).
 - *Week.* Skylar One will execute full backups weekly as specified (for example, every 2 weeks).
 - *Month.* Skylar One will execute full backups monthly as specified (for example, every 1 month).
- **Start Time / Date.** If you enabled full backups, you must specify the start time. You must specify the date on which you want configuration backups to begin, as well as the time of day you want the backups to run. For example, you might want to run backups during a maintenance window late at night. Click the field to open a calendar and time selector.
- **Timezone.** Optional. Specify the timezone to use when running a backup. The default is UTC. Use the drop-down list to select the timezone.
- **Full Credentials.** Select the credential you created for Skylar One to use to connect to your external system (NFS, SMB, or S3). For more information, see [Creating a Backup Credential](#).
- **Full Protocol.** Specify the type of external system where the full backup will be stored. Choices are:
 - *NFS-Remote.* When you select this option, Skylar One stores the backup on an NFS share. If you select the *NFS-remote* option, and your NFS share is hosted on a Solaris system, you must perform the steps listed in the [Additional Configuration for Solaris NFS Mounts](#) section.
 - *SMB-Remote.* When you select this option, Skylar One stores the backup on an SMB share.
 - *S3.* When you select this option, Skylar One stores the backup in an S3 bucket.
- **Full Subdirectory.** Specify a directory on the NFS or SMB share or S3 service in which you would like to store the full backup. When entering the subdirectory path, include the leading slash ("/"). ScienceLogic recommends creating and configuring a unique subdirectory for each platform that you are backing up to the share; do not back up multiple platforms to the exact same subfolder in a share. You will need this value as well as your [backup credential information](#) and any mount options that you specify in step 3 below if you [mount the backup](#) to an NFS or SMB share.
- **Full Backup on DR.** Select this checkbox if you want to perform full backups from the Disaster Recovery Database Server. This alternative method to backing up Skylar One can reduce performance issues during backup for users with large systems and very large backup files. For more information, see the section on [Performing Configuration and Full Backups on the DR Database Server](#).

- **Full Retention Period.** Specify the number of full backups you want Skylar One to keep before deleting them. Skylar One will retain the number of backups you specify, plus one additional backup. For example, if you specify "2", Skylar One will retain your most recent backup plus the two backups that immediately precede it. If you keep the default value of "0", then Skylar One retains all of your backups.

NOTE: Retaining all or a large number of your backups might cause storage constraints in your backup share.

3. To access additional configuration options for this backup, click **[Advanced Settings]**. These advanced fields refer to the column names in the Skylar One Database:
 - **Full SMB Mount Options.** Corresponds to the backup_smb_mount_options database column.
 - **Full NFS Mount Options.** Corresponds to the backups_nfs_mount_options database column.
 - **Full CMD Options.** Corresponds to the backup_cmd_options database column.
 - **Full Comp CMD Options.** Corresponds to the comp_cmd_options database column.
 - **Custom mariabackup Options.** Lets you specify one or more custom MariaDB backup options. For details on these options, see <https://mariadb.com/kb/en/mariabackup-options/>.

NOTE: For more information about these advanced settings, contact ScienceLogic Support.

4. Click **[Save]**. Skylar One will execute the full backup at the frequency and time you specified in the **Every**, **Interval**, and **Start Time/Date** fields.
5. To run the backup immediately, click the **[Backup Now]** button under **Full Backup** in the **Immediate Backup** section of the **Backup Management** page. Skylar One will immediately run the backup and will still run the backup at the frequency and time you specified in the **Every**, **Interval**, and **Start Time/Date** fields.

Defining a Disaster Recover Backup

To define and schedule a DR backup:

1. Go to the **Backup Management** page (System > Settings > Backup).
2. In the **DR Backup** pane, provide values in the following fields:
 - **Every.** Together with the **Interval** field, specifies how frequently Skylar One should automatically execute a full backup. Enter a number.

- **Interval.** Together with the *Every* field, you must specify how frequently Skylar One should automatically execute a DR backup. Your choices are:
 - *Disabled.* DR backups are disabled. If you want to run on-demand DR backups only, you do not need to change this setting.
 - *Day.* Skylar One will execute DR backups daily as specified (for example, every 2 days).
 - *Week.* Skylar One will execute DR backups weekly as specified (for example, every 2 weeks).
 - *Month.* Skylar One will execute DR backups monthly as specified (for example, every 1 month).
- **Start Time/Date.** If you enabled DR backups, you must specify the start time. You must specify the date on which you want DR backups to begin, as well as the time of day you want the backups to run. For example, you might want to run DR backups during a maintenance window late at night. Click the field to open a calendar and time selector.
- **Timezone.** Optional. Specify the timezone to use when running a backup. The default is UTC. Use the drop-down list to select the timezone.
- **DR Credentials.** Select the credential you created for Skylar One to use to connect to your external system (NFS, SMB, or S3). For more information, see [Creating a Backup Credential](#).
- **DR Protocol.** Specifies where Skylar One should store the DR backups. Choices are:
 - *NFS-Remote.* When you select this option, Skylar One stores the DR backup on an NFS share. If you select the *NFS-remote* option, and your NFS mount is hosted on a Solaris system, you must perform the steps listed in the [Additional Configuration for Solaris NFS Mounts](#) section of this chapter.
 - *SMB-Remote.* When you select this option, Skylar One stores the DR backup on an SMB share.
 - *S3.* When you select this option, Skylar One stores the DR backup in an S3 bucket.
- **DR Subdirectory.** Specify a directory on the NFS or SMB share or S3 service in which you would like to store the DR backup. When entering the subdirectory path, include the leading slash ("/"). ScienceLogic recommends creating and configuring a unique subdirectory for each platform that you are backing up to the share; do not back up multiple platforms to the exact same subfolder in a share. You will need this value as well as your [backup credential information](#) and any mount options that you specify in step 3 below if you [mount the backup](#) to an NFS or SMB share.
- **DR Retention Period.** Specify the number of DR backups you want Skylar One to keep before deleting them. Skylar One will retain the number of backups you specify, plus one additional backup. For example, if you specify "2", Skylar One will retain your most recent backup plus the two backups that immediately precede it. If you keep the default value of "0", then Skylar One retains all of your backups.

NOTE: Retaining all or a large number of your backups might cause storage constraints in your backup share.

3. To access additional configuration options for this backup, click **[Advanced Settings]**. These advanced fields refer to the column names in the Skylar One Database:

- `backup_smb_mount_options`
- `backups_nfs_mount_options`
- `backup_cmd_options`
- `comp_cmd_options`

NOTE: For more information about these Advanced Settings, contact ScienceLogic Support.

4. Click **[Save]**. Skylar One will execute the DR backup at the frequency and time you specified in the **Backup Frequency** and **Start Time** fields.
5. To run the backup immediately, click the **[Backup Now]** button under **DR Backup**. Skylar One will immediately run the backup and will still run the backup at the frequency and time you specified in the **Backup Frequency** and **Start Time** fields.

Enabling tmux

Some of the subsequent processes required for mounting and restoring from backup files can take a long time to complete. For this reason, ScienceLogic typically recommends that you open a `tmux` session for these tasks. Before you can do so, however, the `tmux` utility might need to be enabled on your system, depending on your Skylar One version and deployment type.

The **`tmux` utility** is a terminal multiplexer that enables a number of terminals to be created, accessed, and controlled from a single screen, strengthens session-control mechanisms and aligns with industry-wide security practices.

Starting with Skylar One version 12.3.4, the `tmux` utility is disabled by default if you are on a non-STIG Skylar One deployment and access a Skylar One system using SSH.

NOTE: This is a change in behavior from Skylar One versions 12.2.1.1 through 12.3.3, where the `tmux` utility was enabled by default. For more information, see the [12.3.3 release notes](#).

If you are on a STIG-compliant Skylar One deployment, the `tmux` utility is enabled by default. ScienceLogic encourages non-STIG users enable the `tmux` utility as well.

If `tmux` is enabled, sessions are automatically locked after 15 minutes of idleness or if an unclean SSH disconnect or dropped SSH connection occurs. Upon login, Skylar One checks for and attaches any detached `tmux` session if it finds them; otherwise, it starts a new session.

The utility also facilitates advanced features like scroll-back buffering with search, built-in clipboarding, multiple sessions and panes, detaching or attaching sessions, and session supervision or sharing.

To enable the `tmux` utility in non-STIG deployments:

1. Either go to the console of the Skylar One appliance or use SSH to access the Skylar One appliance.
2. Open a shell session on the server.
3. Type the following at the command line to edit the **silo.conf** file:

```
sudo visilo
```

4. Change the following line in the **[OS_HARDENING]** section of the file to enable tmux:

```
TMUX = true
```

NOTE: If the **[OS_HARDENING]** heading does not already exist in the **silo.conf** file, you must add that immediately above the **TMUX = true** setting.

5. Save and quit the file (`:wq`).
6. Log out of Skylar One and then log back in. The tmux utility is now enabled.

For more information about tmux shortcuts and usage, see <https://tmuxcheatsheet.com/>.

NOTE: When using the command-line interface on a Skylar One appliance, the interface is meant for limited administrative purposes only.

TIP: When completing a task in a `tmux` session that will take longer than 15 minutes to complete, such as restoring from backup files, you should open a named `tmux` session to avoid the session timing out:

```
tmux new -t <session name>
```

Mounting Backup Files

If your backup is stored in an NFS or SMB share, then you must mount the backup before you can use it to restore your Skylar One system. When you are finished restoring from the backup, you must *unmount the backup files*. This section describes how to mount these two share types.

NOTE: If you store your backup files in an S3 bucket, you can skip this section.

Mounting NFS Shares

Before beginning this process, you should have the following information:

- The IP address that is defined in the *backup credential*
- Based on which type of backup files you are mounting, the mount directory that is defined in the **Configuration Subdirectory**, **Full Subdirectory**, or **DR Subdirectory** field on the **Backup Management** page (System > Settings > Backup)
- Based on which type of backup files you are mounting, the value from the **Configuration NFS Mount Options**, **Full NFS Mount Options**, or **DR NFS Mount Options** field, which you can access by clicking **[Advanced Settings]** on the **Backup Management** page (System > Settings > Backup)

NOTE: If you are running a version of Skylar One that does not enable the tmux utility by default, you must first [enable it](#).

To mount a remote NFS share:

1. Either go to the console of the Database Server where you want to restore the backup, or use SSH to log in to the Database Server, and log in as user **em7admin**.
2. Start a named `tmux` session, if necessary:

```
tmux new -t <session name>
```

where you replace `<session name>` with a name for the `tmux` session.

3. Execute the following command to mount the remote NFS share:

```
sudo /bin/mount -t nfs <IP from backup credential>:<subdirectory>
/mnt -o <NFS mount options>
```

where:

- `<IP from backup credential>` is the IP address you defined in your *backup credential*.
- `<subdirectory>` is the mount directory that is defined in the **Configuration Subdirectory**, **Full Subdirectory**, or **DR Subdirectory** field on the **Backup Management** page (System > Settings > Backup) to store the backup.
- `<NFS mount options>` is the value from the **Configuration NFS Mount Options**, **Full NFS Mount Options**, or **DR NFS Mount Options** field, which you can access by clicking **[Advanced Settings]** on the **Backup Management** page (System > Settings > Backup).

For example:

```
sudo /bin/mount -t nfs 10.64.70.63:/backups /mnt -o lookupcache=none
```

4. If you are prompted to provide a password, enter the password that you defined in your *backup credential*. If the command completes without an error, the NFS share has been mounted and will remain mounted until it is [manually unmounted](#).

Mounting SMB Shares

Before beginning this process, you should have the following information:

- The IP address and username that is defined in the *backup credential*
- Based on which type of backup files you are mounting, the mount directory that is defined in the **Configuration Subdirectory**, **Full Subdirectory**, or **DR Subdirectory** field on the **Backup Management** page (System > Settings > Backup)
- Based on which type of backup files you are mounting, the value from the **Configuration SMB Mount Options**, **Full SMB Mount Options**, or **DR SMB Mount Options** field, which you can access by clicking **[Advanced Settings]** on the **Backup Management** page (System > Settings > Backup)

NOTE: If you are running a version of Skylar One that does not enable the tmux utility by default, you must first [enable it](#).

To mount a remote SMB share:

1. Either go to the console of the Database Server where you want to restore the backup, or use SSH to log in to the Database Server, and log in as user **em7admin**.
2. Start a named `tmux` session, if necessary:

```
tmux new -t <session name>
```

where you replace `<session name>` with a name for the `tmux` session.

3. Execute the following command to mount the remote SMB share:

```
sudo /bin/mount -t cifs //<IP from backup credential>/<full
subdirectory> /mnt -o user=<username from backup credential>,<SMB
mount options>
```

where:

- `<IP from backup credential>` is the IP address you defined in your *backup credential*.
- `<username from backup credential>` is the username you defined in your *backup credential*.
- `<full subdirectory>` is the mount directory that is defined in the **Configuration Subdirectory**, **Full Subdirectory**, or **DR Subdirectory** field on the **Backup Management** page (System > Settings > Backup) to store the backup.
- `<SMB mount options>` is the value from the **Configuration SMB Mount Options**, **Full SMB Mount Options**, or **DR SMB Mount Options** field, which you can access by clicking **[Advanced Settings]** on the **Backup Management** page (System > Settings > Backup).

For example:

```
sudo /bin/mount -t cifs //10.64.70.48/Lab_Backups /mnt -o
user=administrator,iocharset=utf8,file_mode=0777,dir_
mode=0777,uid=600,gid=607,setuids,noperm,sec=ntlm,vers=1.0
```

4. If you are prompted to provide a password, enter the password that you defined in your *backup credential*. If the command completes without an error, the NFS share has been mounted and will remain mounted until it is *manually unmounted*.

Restoring Backups

This section describes how to restore the following backup types:

- *Configuration backup from an S3 bucket*
- *Configuration backup from a remote NFS or SMB share*
- *Full backup from an S3 bucket*
- *Full backup from a remote NFS or SMB share*
- *DR backup from an S3 bucket*
- *DR backup from a remote NFS or SMB share*

NOTE: To complete the restore steps, you must be familiar with how to edit a file using the vi text editor. If you need assistance with these steps, please contact ScienceLogic Support.

Restoring a Configuration Backup from an S3 Bucket

If you have performed configuration backups, you can restore your system from a configuration backup in the event of data corruption or other failure. The configuration backup file contains one SQL (.sql) file for each database that was included in the configuration backup. Configuration backups *do not* contain historical data from devices.

IMPORTANT: The following steps assume that the Database Server to which you are restoring the backup has not been previously configured and is on the same revision number as the Database Server that was used to create the backup file.

To restore a database using configuration backup files that were stored on an S3 storage service:

1. Either go to the console of the Database Server where you want to restore the backup, or use SSH to log in to the Database Server, then log in as user **em7admin**.

2. Ensure that the firstboot process has completed successfully by entering the following command:

```
tail /var/log/em7/em7_firstboot.log
```

The last few lines of the output should look similar to the following:

```
2025-02-05 21:42:00 All firstboot scripts completed successfully in  
772 seconds.
```

```
DISABLING FIRSTBOOT SCRIPT
```

If the output does not indicate that the firstboot process was completed successfully and the firstboot script is disabled, resolve those issues before proceeding.

3. Determine if the temporary extraction directory exists on the system already by running the following command:

```
ls /data.local/db/.hidden_tmp
```

If you receive an error such as "No such file or directory" or "Permission denied", continue to the next step. Otherwise, if the directory already exists, proceed to step 8.

4. To create a temporary extraction directory that can be used for extracting the backup files:

```
sudo mkdir /data.local/db/.hidden_tmp
```

NOTE: The directory might already exist. If you receive an error indicating that the directory already exists, continue with the following steps to ensure that it is configured properly.

5. Grant the **em7admin** user permissions for this directory:

```
sudo setfacl -R -m u:em7admin:rwx /data.local/db/.hidden_tmp
```

6. Set the default permissions on this directory:

```
sudo setfacl -d -R -m u:em7admin:rwx /data.local/db/.hidden_tmp
```

7. Exclude the directory from being scanned by MariaDB:

```
sudo sl1-config -q -y my mysqld ignore_db_dirs .hidden_tmp
```

NOTE: This command will overwrite any prior values. It is not expected that this option will be set to any other value.

8. After you have created the temporary extraction directory or confirmed that it already exists, ensure that it is empty by running the following command:

```
ls /data.local/db/.hidden_tmp/
```

If you see files in this directory, confirm that no one else is in the process of performing a restoration and then remove those files with the following command:

```
rm -r /data.local/db/.hidden_tmp/*
```

9. Run the following commands to set up the environment for restoration:

```
export RCLONE_CONFIG_BACKUP_BUCKET_ACL="private"
export RCLONE_CONFIG_BACKUP_CHUNK_SIZE="128Mi"
export RCLONE_CONFIG_BACKUP_TYPE="s3"
export RCLONE_CONFIG_BACKUP_UPLOAD_CONCURRENCY="4"
export RCLONE_CONFIG_BACKUP_PROVIDER="<--- Provider --->"
export RCLONE_CONFIG_BACKUP_ACCESS_KEY_ID="<--- Access Key ID --->"
export RCLONE_CONFIG_BACKUP_SECRET_ACCESS_KEY="<--- Secret Access Key
--->"
export RCLONE_CONFIG_BACKUP_ENDPOINT="<--- Endpoint URL --->"
export RCLONE_CONFIG_BACKUP_REGION="<--- Region --->"
export RCLONE_CONFIG_CRYPT_TYPE="crypt"
export RCLONE_CONFIG_CRYPT_DIRECTORY_NAME_ENCRYPTION="false"
export RCLONE_CONFIG_CRYPT_FILENAME_ENCRYPTION="off"
export RCLONE_CONFIG_CRYPT_REMOTE="backup:/<--- Bucket --->/<--- Folder --->"
export RCLONE_CONFIG_CRYPT_PASSWORD="'rclone obscure '<--- Encryption
Password --->'"
export RCLONE_CONFIG_CRYPT_PASSWORD2="'rclone obscure '<--- Encryption
Salt --->'"
```

10. Extract the archive from S3:

```
rclone cat crypt:<--- File Name without the .bin ---> | tar -xI pigz
-C /data.local/db/.hidden_tmp/
```

11. Your target directory (/data.local/db/.hidden_tmp) now includes a number of subdirectories. Run the following command to list the directories:

```
ls -l /data.local/db/.hidden_tmp/
```

You should see output similar to the following:

```
drwxr-x---+ 3 em7admin em7admin 20 Jul 31 22:57 data
drwxr-x---+ 11 em7admin em7admin 4096 Jul 31 22:57 etc
```

```
drwxr-x---+ 3 em7admin em7admin 17 Jul 31 22:57 opt  
drwxr-x---+ 3 em7admin em7admin 21 Jul 31 22:57 usr
```

The data subdirectory contains the SQL files needed to restore the database. The other subdirectories contain operating system files that are included in the backup.

12. Prior to starting the restoration process, stop all Skylar One services:

```
sudo siloctl stop
```

13. To restore operating system files, you must copy them back to their original locations.

CAUTION: ScienceLogic advises you not to bulk copy the operating system files back to their original locations. If you have reinstalled the system or are restoring to a different system, bulk copying the operating system files can result in the system becoming unusable.

14. A freshly installed Skylar One system will have most configuration files set to the recommended default values. If you are restoring on a freshly installed system, you need to restore the `silo.conf` file; otherwise, you can skip this step. To restore the `silo.conf` file:

```
sudo cp /data.local/db/.hidden_tmp/etc/silo.conf /etc/silo.conf
```

NOTE: If the IP address has changed, then after you restore the `silo.conf` file, you should use the `visilo` command to update the configuration to reflect the new IP address.

15. Review the other configuration files, comparing them to the files that you want to replace, and merge any changes in if needed.

16. To restore the database files, you must load the required SQL files. Run the following command to list all of the SQL files, replacing <numeric timestamp> with the numeric timestamp of when the backup was taken:

```
ls -l /data.local/db/.hidden_tmp/data/backup/remote<numeric timestamp>/staging/
```

NOTE: You can use tab completion by tying the path to the remote storage and then pressing the **[Tab]** key to complete the directory name.

Your output should be similar to:

```
-rw-r-----+ 1 em7admin em7admin 2262 Jul 30 21:17 insight_agent.sql
-rw-r-----+ 1 em7admin em7admin 138346 Jul 30 21:17 master_access.sql
-rw-r-----+ 1 em7admin em7admin 3036487 Jul 30 21:17 master_ap2.sql
-rw-r-----+ 1 em7admin em7admin 1839176 Jul 30 21:17 master_biz.sql
-rw-r-----+ 1 em7admin em7admin 46676500 Jul 30 21:17 master_custom.sql
-rw-r-----+ 1 em7admin em7admin 176017 Jul 30 21:17 master_dev.sql
-rw-r-----+ 1 em7admin em7admin 10849 Jul 30 21:17 master_dns.sql
-rw-r-----+ 1 em7admin em7admin 2142 Jul 30 21:17 master_events.event_suppressions.sql
-rw-r-----+ 1 em7admin em7admin 161861676 Jul 30 21:17 master_filestore.sql
-rw-r-----+ 1 em7admin em7admin 3569078 Jul 30 21:17 master_platform.sql
-rw-r-----+ 1 em7admin em7admin 5996388 Jul 30 21:17 master_reports.sql
-rw-r-----+ 1 em7admin em7admin 393919805 Jul 30 21:17 master.sql
-rw-r-----+ 1 em7admin em7admin 4630405 Jul 30 21:17 mysql.sql
-rw-r-----+ 1 em7admin em7admin 7475 Jul 30 21:17 phonehome.sql
-rw-r-----+ 1 em7admin em7admin 7257 Jul 30 21:17 scheduler.sql
-rw-r-----+ 1 em7admin em7admin 61479 Jul 30 21:17 sysinfo.sql
-rw-r-----+ 1 em7admin em7admin 2683 Jul 30 21:17 vault.sql
```

17. When restoring a newly installed system, you must restore the mysql.sql file first, as this will rewrite the privilege tables with the information needed to correctly restore the other files. To restore the mysql privilege tables, run the following commands, replacing `<numeric timestamp>` with the numeric timestamp of when the backup was taken:

```
sudo /bin/mysql mysql < /data.local/db/.hidden_
tmp/data/backup/remote<numeric timestamp>/staging/mysql.sql;
sudo /bin/mysql -e "FLUSH PRIVILEGES"
```

18. Confirm that the privilege tables were restored correctly:

```
silo_mysql -e "SELECT VERSION()"
```

The current version of MariaDB should appear in the output. If you receive an error message, this means the privilege tables were not restored or the `silo.conf` file for this backup was not restored. Resolve these issues before proceeding.

19. For all SQL databases **except** the MySQL database that was already restored, run the following command:

```
silo_mysql <name_of_database> < data.local/db/.hidden_
tmp/data/backup/remote<numeric timestamp>/staging/<name_of_
database>.sql
```

where:

- `<name_of_database>` is the name of the SQL database you are restoring.
- `<numeric timestamp>` is the numeric timestamp of when the backup was taken.

20. To restore all the databases that are included in the backup file, repeat the previous step for each `.sql` file except `mysql.sql`, which was already restored.
21. After you have restored all of the databases, re-license the Database Server using the standard licensing procedure. For details, see the section on *Licensing and Configuring an Appliance* in the *Installation and Initial Configuration* manual.
22. Start all Skylar One services:

```
sudo siloctl start
```

23. Clean up temporary files:

```
rm -rf /data.local/db/.hidden_tmp/*
```

Restoring a Configuration Backup from a Remote NFS or SMB Share

If you have performed configuration backups, you can restore your system from a configuration backup in the event of data corruption or other failure. The configuration backup file contains one SQL (`.sql`) file for each database that was included in the configuration backup. Configuration backups **do not** contain historical data from devices.

IMPORTANT: The following steps assume that the Database Server to which you are restoring the backup has not been previously configured and is on the same revision number as the Database Server that was used to create the backup file.

NOTE: Before beginning this procedure, ensure that your configuration backup files are already *mounted on the remote share*.

To restore a database using configuration backup files that were stored on a remote NFS or SMB share:

1. Either go to the console of the Database Server where you want to restore the backup, or use SSH to log in to the Database Server, then log in as user **em7admin**.
2. Ensure that the firstboot process has completed successfully by entering the following command:

```
tail /var/log/em7/em7_firstboot.log
```

The last few lines of the output should look similar to the following:

```
2025-02-05 21:42:00 All firstboot scripts completed successfully in  
772 seconds.
```

```
DISABLING FIRSTBOOT SCRIPT
```

If the output does not indicate that the firstboot process was completed successfully and the firstboot script is disabled, resolve those issues before proceeding.

3. Determine if the temporary extraction directory exists on the system already by running the following command:

```
ls /data.local/db/.hidden_tmp
```

If you receive an error such as "No such file or directory" or "Permission denied", continue to the next step. Otherwise, if the directory already exists, proceed to step 8.

4. To create a temporary extraction directory that can be used for extracting the backup files:

```
sudo mkdir /data.local/db/.hidden_tmp
```

NOTE: The directory might already exist. If you receive an error indicating that the directory already exists, continue with the following steps to ensure that it is configured properly.

5. Grant the **em7admin** user permissions for this directory:

```
sudo setfacl -R -m u:em7admin:rwx /data.local/db/.hidden_tmp
```

6. Set the default permissions on this directory:

```
sudo setfacl -d -R -m u:em7admin:rwx /data.local/db/.hidden_tmp
```

7. Exclude the directory from being scanned by MariaDB:

```
sudo sl1-config -q -y my mysqld ignore_db_dirs .hidden_tmp
```

NOTE: This command will overwrite any prior values. It is not expected that this option will be set to any other value.

8. After you have created the temporary extraction directory or confirmed that it already exists, ensure that it is empty by running the following command:

```
ls /data.local/db/.hidden_tmp/
```

If you see files in this directory, confirm that no one else is in the process of performing a restoration and then remove those files with the following command:

```
rm -r /data.local/db/.hidden_tmp/*
```

9. Execute the following command to extract the backup into your directory, where *<full path and filename for backup.tgz>* is the location and name of your backup file:

```
sudo tar -xI pigz -C /data.local/db/.hidden_tmp/ -f <full path and file name for backup.tgz>
```

For example:

```
tar -xI pigz -C /data.local/db/.hidden_tmp/ -f /mnt/db1_config_2021-02-01_21-00-00.tgz
```

NOTE: Depending on the permissions of your NFS or SMB share, the **em7admin** user might not have permissions to read this file. ScienceLogic recommends granting read access to the file using the following command:

```
chmod o+r <full path and file name for backup.tgz>
```

10. Your target directory (`/data.local/db/.hidden_tmp`) now includes a number of subdirectories. Run the following command to list the directories:

```
ls -l /data.local/db/.hidden_tmp/
```

You should see output similar to:

```
drwxr-x---+ 3 em7admin em7admin 20 Jul 31 22:57 data
drwxr-x---+ 11 em7admin em7admin 4096 Jul 31 22:57 etc
```

```
drwxr-x---+ 3 em7admin em7admin 17 Jul 31 22:57 opt
drwxr-x---+ 3 em7admin em7admin 21 Jul 31 22:57 usr
```

The data subdirectory contains the SQL files needed to restore the database. The other subdirectories contain operating system files that are included in the backup.

11. Prior to starting the restoration process, stop all Skylar One services:

```
sudo siloctl stop
```

12. To restore operating system files, you must copy them back to their original locations.

CAUTION: ScienceLogic advises you not to bulk copy the operating system files back to their original locations. If you have reinstalled the system or are restoring to a different system, bulk copying the operating system files can result in the system becoming unusable.

13. A freshly installed Skylar One system will have most configuration files set to the recommended default values. If you are restoring on a freshly installed system, you need to restore the `silo.conf` file; otherwise, you can skip this step. To restore the `silo.conf` file:

```
sudo cp /data.local/db/.hidden_tmp/etc/silo.conf /etc/silo.conf
```

NOTE: If the IP address has changed, then after you restore the `silo.conf` file, you should use the command `sudo visilo --no-validation` to update the configuration to reflect the new IP address.

14. Review the other configuration files, comparing them to the files that you want to replace, and merge any changes in if needed.

15. To restore the database files, you must load the required SQL files. Run the following command to list all of the SQL files, replacing <numeric timestamp> with the numeric timestamp of when the backup was taken:

```
ls -l /data.local/db/.hidden_tmp/data/backup/remote<numeric timestamp>/staging/
```

NOTE: You can use tab completion by tying the path to the remote storage and then pressing the **[Tab]** key to complete the directory name.

Your output should be similar to:

```
-rw-r-----+ 1 em7admin em7admin 2262 Jul 30 21:17 insight_agent.sql
-rw-r-----+ 1 em7admin em7admin 138346 Jul 30 21:17 master_access.sql
-rw-r-----+ 1 em7admin em7admin 3036487 Jul 30 21:17 master_ap2.sql
-rw-r-----+ 1 em7admin em7admin 1839176 Jul 30 21:17 master_biz.sql
-rw-r-----+ 1 em7admin em7admin 46676500 Jul 30 21:17 master_custom.sql
-rw-r-----+ 1 em7admin em7admin 176017 Jul 30 21:17 master_dev.sql
-rw-r-----+ 1 em7admin em7admin 10849 Jul 30 21:17 master_dns.sql
-rw-r-----+ 1 em7admin em7admin 2142 Jul 30 21:17 master_events.event_suppressions.sql
-rw-r-----+ 1 em7admin em7admin 161861676 Jul 30 21:17 master_filestore.sql
-rw-r-----+ 1 em7admin em7admin 3569078 Jul 30 21:17 master_platform.sql
-rw-r-----+ 1 em7admin em7admin 5996388 Jul 30 21:17 master_reports.sql
-rw-r-----+ 1 em7admin em7admin 393919805 Jul 30 21:17 master.sql
-rw-r-----+ 1 em7admin em7admin 4630405 Jul 30 21:17 mysql.sql
-rw-r-----+ 1 em7admin em7admin 7475 Jul 30 21:17 phonehome.sql
-rw-r-----+ 1 em7admin em7admin 7257 Jul 30 21:17 scheduler.sql
-rw-r-----+ 1 em7admin em7admin 61479 Jul 30 21:17 sysinfo.sql
-rw-r-----+ 1 em7admin em7admin 2683 Jul 30 21:17 vault.sql
```

16. When restoring a newly installed system, you must restore the mysql.sql file first, as this will rewrite the privilege tables with the information needed to correctly restore the other files. To restore the mysql privilege tables, run the following commands, replacing `<numeric timestamp>` with the numeric timestamp of when the backup was taken:

```
sudo -s /bin/mysql mysql < /data.local/db/.hidden_
tmp/data/backup/remote<numeric timestamp>/staging/mysql.sql;
sudo -s /bin/mysql -e "FLUSH PRIVILEGES"
```

17. Confirm that the privilege tables were restored correctly:

```
silo_mysql -e "SELECT VERSION()"
```

The current version of MariaDB should appear in the output. If you receive an error message, this means the privilege tables were not restored or the `silo.conf` file for this backup was not restored. Resolve these issues before proceeding.

18. For all SQL databases **except** the MySQL database that was already restored, run the following command:

```
silo_mysql <name_of_database> < /data.local/db/.hidden_
tmp/data/backup/remote<numeric timestamp>/staging/<name_of_
database>.sql
```

where:

- `<name_of_database>` is the name of the SQL database you are restoring.
- `<numeric timestamp>` is the numeric timestamp of when the backup was taken.

19. To restore all the databases that are included in the backup file, repeat the previous step for each `.sql` file except `mysql.sql`, which was already restored.
20. After you have restored all of the databases, re-license the Database Server using the standard licensing procedure. For details, see the section on *Licensing and Configuring an Appliance* in the *Installation and Initial Configuration* manual.
21. Start all Skylar One services:

```
sudo siloctl start
```

22. Unmount the backup directory from the share:

```
sudo umount /mnt
```

If the command completes without an error, you have successfully unmounted the share.

NOTE: If you receive an error, ensure that nothing is accessing the share and that your current path is not inside the share, then try again. To successfully unmount the share, nothing can be using it.

23. Clean up temporary files:

```
rm -rf /data.local/db/.hidden_tmp/*
```

Restoring a Full Backup from an S3 Bucket

NOTE: These steps assume that the Database Server to which you are restoring the backup has not been previously configured and is on the same platform revision number as the Database Server used to create the backup file.

To restore a Skylar One system using a full backup file from an S3 bucket:

1. Either go to the console of the Database Server where you want to restore the backup, or use SSH to log in to the Database Server.
2. Log in as user ***em7admin*** and then assume root privileges:

```
sudo -s
```

3. Execute the following commands:

WARNING: Executing this command will stop the database. Skylar One will not be operational until you complete the restore procedure.

```
siloctl stop --full
systemctl stop mariadb
rm -rf /data.local/db/*
mkdir /data.local/db/.tmp
cd /data.local/db/.tmp
export RCLONE_CONFIG_BACKUP_BUCKET_ACL="private"
export RCLONE_CONFIG_BACKUP_CHUNK_SIZE="128Mi"
export RCLONE_CONFIG_BACKUP_TYPE="s3"
export RCLONE_CONFIG_BACKUP_UPLOAD_CONCURRENCY="4"
export RCLONE_CONFIG_BACKUP_PROVIDER="<--- Provider --->"
export RCLONE_CONFIG_BACKUP_ACCESS_KEY_ID="<--- Access Key ID --->"
export RCLONE_CONFIG_BACKUP_SECRET_ACCESS_KEY="<--- Secret Access Key
--->"
export RCLONE_CONFIG_BACKUP_ENDPOINT="<--- Endpoint URL --->"
export RCLONE_CONFIG_BACKUP_REGION="<--- Region --->"
export RCLONE_CONFIG_CRYPT_TYPE="crypt"
export RCLONE_CONFIG_CRYPT_DIRECTORY_NAME_ENCRYPTION="false"
export RCLONE_CONFIG_CRYPT_FILENAME_ENCRYPTION="off"
export RCLONE_CONFIG_CRYPT_REMOTE="backup:/<--- Bucket --->/<---
Folder --->"
export RCLONE_CONFIG_CRYPT_PASSWORD="`rclone obscure '<--- Encryption
Password --->`"
export RCLONE_CONFIG_CRYPT_PASSWORD2="`rclone obscure '<---
Encryption Salt --->`"
time rclone cat crypt:<--- Full Backup File Name without the .bin ---
>|pigz -d|mbstream -vvv -x -C .

```

4. Execute the following command, where `<directory for data extraction>` is the directory you created in the previous step:

```
more /<directory for data extraction>/backup-my.cnf
```

5. Locate the line that looks like the following. Copy or write down the exact text that appears, such as:

```
innodb_data_file_path=ibdata1:354M;ibdata2:500M:autoextend:max:8925M
```

6. Execute the following command to edit the `/etc/my.cnf.d/silo_mysql.cnf` file:

```
vimmysql
```

7. Add the line you copied in Step 5 to the **mysql.siteconfig** file, such as:

```
innodb_data_file_path=ibdata1:354M;ibdata2:500M:autoextend:max:8925M
```

8. Save the file and exit the editor:

```
:wq
```

9. Execute the following command to build the updated configuration file:

```
vimmysql -f
```

10. Execute the following commands:

NOTE: The following process can take a long time to complete. To ensure the command is not interrupted, you should run the command in a **tmux** session. If you are running a version of Skylar One that does not enable the **tmux** utility by default, you must first [enable it](#).

To do so:

- Start a named **tmux** session:

```
tmux new -t <session name>
```

where you replace **<session name>** with a name for the **tmux** session.

- Execute the following commands:

```
mariabackup --prepare --use-memory=<80% of available memory> --  
target-dir=/data.local.db
```

```
mariabackup --move-back --force-non-empty-directories --target-  
dir
```

```
cd /data.local/db
```

```
rm -rf .tmp
```

```
chown -R mysql.mysql *
```

```
systemctl start mariadb
```

NOTE: These commands assume you have changed directories to the directory that contains the extracted backup files. If you want to run these commands from a location other than the location of the extracted backup files, then you must enter a value after "target-dir=" to point to the directory where the extracted files are located. To learn more about MariaBackup, see <https://mariadb.com/kb/en/full-backup-and-restore-with-mariabackup/>.

NOTE: Depending on the size of the backup, the mariabackup command might take a long time to complete.

11. Because a full database restoration requires the database username and password, you must check the MariaDB username and password that are stored in the **silo.conf** file.

NOTE: If the backup was taken from a Skylar One system installed from or upgraded to version 11.3.0 or later, the username will be **clientdbuser**.

If the backup was taken from a Skylar One system that was installed from or last upgraded to a version prior to 11.3.0, the username will be **root**.

- To check the username and password in the **silo.conf** file, enter the following command:

```
visilo
```

- Check (and if necessary, update) the values in the **dbuser**, **dbpasswd**, and **ap_pass** fields in both the **[LOCAL]** and **[CENTRAL]** sections.
- Save and close the file:

```
:wq
```

NOTE: Upon saving, visilo will validate that the MariaDB passwords work. If the passwords fail, ensure that you are entering the correct ones. If the passwords are unknown, you must perform the password recovery and reset procedure described in the section *Changing the MariaDB Password on Skylar One Appliances*.

12. Restart Skylar One:

```
silootl start --full
```

13. Re-license the Database Server using the standard licensing procedure. For details, see the section on "Licensing and Configuring an Appliance" in the *Installation and Initial Configuration* manual.

NOTE: The `start` process from step 12 will fail and restart until your Skylar One system is licensed. This is expected behavior.

Restoring a Full Backup from a Remote NFS or SMB Share

IMPORTANT: The following steps assume that the Database Server to which you are restoring the backup has not been previously configured and is on the same revision number as the Database Server that was used to create the backup file.

To restore a Skylar One system using a full backup file from an a remote NFS or SMB share:

1. Either go to the console of the Database Server where you want to restore the backup, or use SSH to log in to the Database Server, and log in as user `em7admin`.
2. Stop MariaDB on the Database Server to which you are restoring the backup and remove the existing files in the `/data.local/db` directory that is used by MariaDB:
 - Execute the following command:

```
sudo siloctl stop --full
```

WARNING: Executing this command will stop the database. Skylar One will not be operational until you complete the restore procedure.

- Wait for all services to stop, and then run:

```
sudo systemctl stop mariadb
```
- After MariaDB stops, run:

```
sudo sh -c 'rm -rf /data.local/db/*'
```

3. On the Database Server to which you are restoring the backup, extract the database files from the backup files into the `/data.local/db` directory.

NOTE: The following process can take a long time to complete. To ensure the command is not interrupted, you should run the command in a `tmux` session. If you are running a version of Skylar One that does not enable the `tmux` utility by default, you must first [enable it](#).

- Start a named tmux session:

```
tmux new -t <session name>
```

where you replace *<session name>* with a name for the tmux session.

- Mount the remote NFS or SMB share.
 - If you are mounting a remote NFS share, run the following command:

```
sudo /bin/mount -t nfs <IP from backup credential>:<full
subdirectory> /mnt -o <full NFS mount options>
```

where:

- <IP from backup credential> is the IP address you defined in your backup credential.
- <full subdirectory> is the directory on a remote NFS, SMB, or S3 mount that you specified in the Full Subdirectory field on the Backup Management page (System > Settings > Backup) to store the backup.
- <full NFS mount options> is the value from the Full NFS Mount Options field, which you can access by clicking Advanced Settings on the Backup Management page (System > Settings > Backup).

For example:

```
sudo /bin/mount -t nfs 10.64.70.63:/backups /mnt -o
lookupcache=none
```

- If you are mounting a remote SMB share, run the following command:

```
sudo /bin/mount -t cifs //<IP from backup credential>/<full
subdirectory> /mnt -o user=<username from backup
credential>,<full SMB mount options>
```

where:

- <IP from backup credential> is the IP address you defined in your backup credential.
- <full subdirectory> is the directory on a remote NFS, SMB, or S3 mount that you specified in the Full Subdirectory field on the Backup Management page (System > Settings > Backup) to store the backup.
- <username from backup credential> is the username you defined in your backup credential.
- <full SMB mount options> is the value from the **Full SMB Mount Options** field, which you can access by clicking **Advanced Settings** on the **Backup Management** page (System > Settings > Backup).

For example:

```
sudo /bin/mount -t cifs //10.64.70.48/Lab_Backups /mnt -o
user=administrator,iocharset=utf8,file_mode=0777,dir_
mode=0777,uid=600,gid=607,setuids,noperm,sec=ntlm,vers=1.0
```

NOTE: After entering one of the above commands, if you are prompted to provide a password, enter the password that you defined in your *backup credential*.

- To extract the backup file into your directory, run the following command:

```
sudo sh -c 'pigz -dc <full path and file name for backup.gz> |  
mbstream -x -C /data.local/db/'
```

where *<full path and filename for backup.gz>* is the location and name of your backup file.

For example:

```
sudo sh -c 'pigz -dc /mnt/db1_full_2021-02-01_21-00-00.gz |  
mbstream -x -C /data.local/db/'
```

TIP: Optionally, you can create a new `tmux` session to continue with other tasks while the backup file unpacks in the background:

- Press "Ctrl + b", type ":new", and then press "Enter". This opens a new `tmux` session.
- You can continue to monitor the progress in the original session. After the `mbstream` command is complete, run the following command:

```
watch -n1 'ps -ef |grep "[m]bstream";'
```

- You can use "Ctrl +B" and then type "s" to show all `tmux` sessions and switch between them.

4. Update the MariaDB configuration:

- Look up the maximum size limit for the tablespace file from the last time the backup was taken:

```
sudo grep "autoextend:max" /data.local/db/backup-my.cnf
```

- Copy or write down the output exactly as it appears. You will need this output in a later step. For example:

```
innodb_data_file_  
path=ibdata1:354M;ibdata2:500M:autoextend:max:8925M
```

- Execute the following command to open the MariaDB configuration:

```
sudo vimysql
```

- In the editor, find a line that is similar to the "innodb_data_file_path" output from the example above and replace it with the line you copied or wrote down from that step.

- Save the file and exit the editor:

```
:wq
```

5. Ensure that the database files are point-in-time consistent:

```
sudo mariabackup --prepare --use-memory=<80% of available memory> --
target-dir=/data.local/db
```

Depending on the size of the backup, the `mariabackup` command might take a long time to complete.

6. Ensure you have proper ownership on the files:

```
sudo chown -R mysql:mysql /data.local/db/*
```

7. Start MariaDB on the target Database Server:

```
sudo systemctl start mariadb
```

8. Because a full database restoration requires the database username and password, you must check the MariaDB username and password that are stored in the **silo.conf** file.

If the backup was taken from a Skylar One system installed from or upgraded to version 11.3.0 or later, the username will be **clientdbuser**.

If the backup was taken from a Skylar One system that was installed from or last upgraded to a version prior to 11.3.0, the username will be **root**.

- To check the username and password in the **silo.conf** file, enter the following command:

```
sudo visilo
```

- If needed, update the values in the `dbuser`, `dbpasswd`, and `ap_pass` fields in both the `[LOCAL]` and `[CENTRAL]` sections.
- Save and close the file:

```
:wq
```

Upon saving, `visilo` will validate that the MariaDB passwords work. If the passwords fail, ensure that you are entering the correct ones. If the passwords are unknown, you must perform the password recovery and reset procedure described in the section [Changing the MariaDB Password on Skylar One Appliances](#).

9. Restart Skylar One:

```
sudo siloctl start --full
```

10. Re-license the Database Server using the standard licensing procedure. For details, see the section on **Licensing and Configuring an Appliance** in the **Installation and Initial Configuration** manual. The process from step 9 will fail and restart until your Skylar One system is licensed. This is expected behavior.

Restoring a DR Backup from an S3 Bucket

NOTE: These steps assume that the Database Server to which you are restoring the backup has not been previously configured and is on the same platform revision number as the Database Server used to create the backup file.

NOTE: To complete these steps, you must be familiar with how to edit a file using the vi text editor. If you need assistance with these steps, please contact ScienceLogic Support.

To restore a Database Server using a DR backup file from an S3 bucket, perform the following steps:

1. Either go to the console of the Database Server where you want to restore the backup or use SSH to access the Database Server.
2. Log in as user **em7admin** and sudo to the root account:

```
sudo -s
```

NOTE: The following process can take a long time to complete. To ensure the command is not interrupted, you should run the command in a `tmux` session. If you are running a version of Skylar One that does not enable the `tmux` utility by default, you must first [enable it](#).

3. Start a named `tmux` session:

```
tmux new -t <session name>
```

where you replace `<session name>` with a name for the `tmux` session.

4. Execute the following commands:

WARNING: Executing this command will stop the database. Skylar One will not be operational until you complete the restore procedure.

```
sudoctl stop --full
systemctl stop mariadb
rm -rf /data/db/*
cd /data/db
export RCLONE_CONFIG_BACKUP_BUCKET_ACL="private"
export RCLONE_CONFIG_BACKUP_CHUNK_SIZE="128Mi"
export RCLONE_CONFIG_BACKUP_TYPE="s3"
```

```

export RCLONE_CONFIG_BACKUP_UPLOAD_CONCURRENCY="4"
export RCLONE_CONFIG_BACKUP_PROVIDER="--- Provider --->"
export RCLONE_CONFIG_BACKUP_ACCESS_KEY_ID="--- Access Key ID --->"
export RCLONE_CONFIG_BACKUP_SECRET_ACCESS_KEY="--- Secret Access Key
--->"
export RCLONE_CONFIG_BACKUP_ENDPOINT="--- Endpoint URL --->"
export RCLONE_CONFIG_BACKUP_REGION="--- Region --->"
export RCLONE_CONFIG_CRYPT_TYPE="crypt"
export RCLONE_CONFIG_CRYPT_DIRECTORY_NAME_ENCRYPTION="false"
export RCLONE_CONFIG_CRYPT_FILENAME_ENCRYPTION="off"
export RCLONE_CONFIG_CRYPT_REMOTE="backup:/--- Bucket --->/---
```

Folder --->"

```

export RCLONE_CONFIG_CRYPT_PASSWORD="`rclone obscure '--- Encryption
Password --->'`"
export RCLONE_CONFIG_CRYPT_PASSWORD2="`rclone obscure '--- Encryption Salt --->'`"
rclone cat crypt:<--- Full Backup File Name without the .bin ---
>|pigz -d|mbstream -vvv -x -C .
```

5. Execute the following command, where `<directory for data extraction>` is the directory you created in the previous step:

```
more /<directory for data extraction>/backup-my.cnf
```

6. Locate the line that looks like the following. Copy or write down the exact text that appears, such as:

```
innodb_data_file_path=ibdata1:354M;ibdata2:500M:autoextend:max:8925M
```

7. Execute the following command to edit the `/etc/my.cnf.d/silo_mysql.cnf` file:

```
vimmysql
```

8. Add the line you copied in Step 5 to the `mysql.siteconfig` file, such as:

```
innodb_data_file_path=ibdata1:354M;ibdata2:500M:autoextend:max:8925M
```

9. Save the file and exit the editor:

```
:wq
```

10. Execute the following command to build the updated configuration file:

```
vimmysql -f
```

11. Execute the following commands:

```
mariabackup --prepare --use-memory=<80% of available memory> --target-dir=/data.local.db

mariabackup --move-back --force-non-empty-directories --target-dir

cd /data.local/db

rm -rf .tmp

chown -R mysql.mysql *

systemctl start mariadb
```

NOTE: These commands assume you have changed directories to the directory that contains the extracted backup files. If you want to run these commands from a location other than the location of the extracted backup files, then you must enter a value after "target-dir=" to point to the directory where the extracted files are located. To learn more about MariaBackup, see <https://mariadb.com/kb/en/full-backup-and-restore-with-mariabackup/>.

NOTE: Depending on the size of the backup, the mariabackup command might take a long time to complete.

12. Because a full database restoration requires the database username and password, you must check the MariaDB username and password that are stored in the **silo.conf** file.

NOTE: If the backup was taken from a Skylar One system installed from or upgraded to version 11.3.0 or later, the username will be **clientdbuser**.

If the backup was taken from a Skylar One system that was installed from or last upgraded to a version prior to 11.3.0, the username will be **root**.

- To check the username and password in the **silo.conf** file, enter the following command:

```
visilo
```

- Check (and if necessary, update) the values in the `dbuser`, `dbpasswd`, and `ap_pass` fields in both the [LOCAL] and [CENTRAL] sections.
- Save and close the file:

```
:wq
```

NOTE: Upon saving, visilo will validate that the MariaDB passwords work. If the passwords fail, ensure that you are entering the correct ones. If the passwords are unknown, you must perform the password recovery and reset procedure described in the section [Changing the MariaDB Password on Skylar One Appliances](#).

13. Execute the following command to restart Skylar One and the database:

```
siloctl start --full
```

14. Re-license the Database Server using the standard licensing procedure. For details, see the section on "Licensing and Configuring an Appliance" in the **Installation and Initial Configuration** manual.

NOTE: The process from step 12 will fail and restart until your Skylar One system is licensed. This is expected behavior.

Restoring a DR Backup from a Remote NFS or SMB Share

To restore a Database Server using a DR backup file, perform the following steps:

NOTE: These steps assume that the Database Server has not been previously configured.

1. The Database Server to which you are restoring the backup must be at the same revision number as the Database Server that created the backup file.
2. Either go to the console of the Database Server where you want to restore the backup or use SSH to access the Database Server.
3. Log in as user **em7admin** and sudo to the root account:

```
sudo -s
```

4. Execute the following commands:

WARNING: Executing this command will stop the database. Skylar One will not be operational until you complete the restore procedure.

```
siloctl stop --full
```

```
systemctl stop mariadb
```

```
rm -rf /data/db/*
```

5. Execute the following commands, substituting the full pathname of your backup file:

```
cd /data/db
```

```
pigz -dc <full path and name to backup file.tgz> | tar xvf -
```

```
mv /data/db/data/db/* .
```

```
rm -rf /data/db/data
```

```
cp /data/db/etc/my.cnf.d/silo_mysql.cnf /root/silo_mysql.bak
```

```
rm -rf /data/db/etc
```

```
chown -R mysql:mysql /data/db/*
```

6. Execute the following commands to restart Skylar One and the database:

```
siloctl start --full
```

```
systemctl start mariadb
```

Unmounting Backup Files

If your backup was stored in an NFS or SMB share, then you must unmount the backup when you are finished accessing it. This section describes how to unmount these two share types.

NOTE: If you store your backup files in an S3 bucket, you can skip this section.

To unmount the NFS or SMB share:

1. Either go to the console of the Database Server where you restored the backup, or use SSH to log in to the Database Server, and log in as user **em7admin**.
2. Verify that the share is mounted by running the following command:

```
mountpoint /mnt
```

If the share is mounted, the command will return `"/mnt is a mountpoint"`.

3. To unmount the share, run the following command:

```
sudo umount /mnt
```

If the command completes without an error, you have successfully unmounted the share.

NOTE: If you receive an error, ensure that nothing is accessing the share and that your current path is not inside the share, then try again. To successfully unmount the share, nothing can be using it.

Retaining Backups

This section describes how to retain full and DR backups.

Retaining Full Backups

NOTE: This section applies only to users who are running a version of Skylar One prior to 11.2.0. If you are running Skylar One 11.2.0 or later, you can set the retention value in the **Full Retention Period** field on the **Backup Management** page (System > Settings > Backup). For more information, see the section on [Defining a Full Backup](#).

NOTE: The **Database Tool** page is available only in versions of Skylar One prior to 12.2.1 and displays only for users that have sufficient permissions to access the page.

To specify the number of full backups to retain in Skylar One versions prior to 11.2.0:

1. Go to the **Database Tool** page (System > Tools > DB Tool).
2. In the SQL Query field, enter the following:

```
UPDATE master.system_settings_backup SET backup_retention = <number  
of backups to retain> WHERE id = 2
```

3. For example, the following command would retain five full backups, the last four plus the current backup:

```
UPDATE master.system_settings_backup SET backup_retention = 4 WHERE  
id = 2
```

4. Click **[Go]**. Skylar One will create an entry in `/var/log/em7/silo.log` when a backup is deleted.

Retaining DR Backups

NOTE: This section applies only to users who are running a version of Skylar One prior to 11.2.0. If you are running Skylar One 11.2.0 or later, you can set the retention value in the **DR Retention Period** field on the **Backup Management** page (System > Settings > Backup). For more information, see the section on [Defining a Disaster Recovery Backup](#).

NOTE: The **Database Tool** page is available only in versions of Skylar One prior to 12.2.1 and displays only for users that have sufficient permissions to access the page.

To specify the number of DR backups to retain in Skylar One versions prior to 11.2.0:

1. Go to the **Database Tool** page (System > Tools > DB Tool).
2. In the SQL Query field, enter the following:

```
UPDATE master.system_settings_backup SET backup_retention = <number  
of backups to retain> WHERE id = 3
```

3. For example, the following command would retain five DR backups, the last four and the current backup:

```
UPDATE master.system_settings_backup SET backup_retention = 4 WHERE  
id = 3
```

4. Click **[Go]**. Skylar One will create an entry in `/var/log/em7/silo.log` when a backup is deleted.

Additional Configuration for Solaris NFS Remote Shares

To use the **NFS-remote** backup protocol with an NFS share hosted on a Solaris system, you must configure the Solaris system to allow the backup process to change file ownership permissions. To do this:

- In **/etc/dfs/dfstab** on the Solaris system, you must specify that the fully-qualified domain name of the Database Server or All-In-One Appliance can access the NFS file system as root. For example:

```
share -F nfs -o sec=sys,root=database.scienceologic.local -d
"ScienceLogic Backup Share" /export/home/backup
```

- In **/etc/default/nfs** on the Solaris system, include the line `NFSMAPID_DOMAIN=<domain of Database Server or All-In-One Appliance>`. For example:

```
NFSMAPID_DOMAIN=ScienceLogic.local
```

You can test this configuration by mounting the NFS file system from the console of your Skylar One appliance, creating a new file on the file system using the "touch" command, and then executing the command "ls -la". If the Solaris system is configured correctly, the output of the ls command will indicate that the new file was created and is owned by the "root" user.

Performing Configuration and Full Backups on the DR Database Server

Users with large systems and very large backup files can use an alternative method to back up Skylar One: Performing backups from the disaster recovery Database Server. The benefit to this alternative method is that it reduces performance issues during the backup procedure.

It is important to note the differences between a disaster recovery backup and a configuration or full backup on the disaster recovery Database Server:

- A *disaster recovery backup* is a method for **backing up the Disaster Recovery system**.
- A configuration or full backup performed on the disaster recovery Database Server is an alternative method to **backing up Skylar One**.

Unlike the disaster recovery backup, which uses tar to make a compressed copy of the /data/db directory, performing a configuration backup from a disaster recovery Database Server uses the MySQL Dump tool and backs up the same data as described in the section on [Configuration Backups](#).

Performing a full backup from a disaster recovery Database Server uses the MariaBackup tool and backs up the same data as described in the section on [Full Backups](#).

NOTE: ScienceLogic recommends that you use the **Full Backup On DR** option instead of the **Full Backup > [Backup Now]** button in the **Immediate Backup** section. The **Full Backup On DR** option includes additional checks by the MariaBackup tool to detect database consistency instead of just copying the content of **/data.local/db** without any additional checks.

During a configuration backup or full backup from a disaster recovery Database Server:

- The disaster recovery Database Server does not appear active to any other Skylar One appliances.
- No applications or services can connect to the disaster recovery Database Server.

To perform a full backup on the disaster recovery Database Server:

1. Follow the instructions in the section on [Creating a Credential](#).
2. Follow the instructions in the section on [Configuring Backups](#).
3. Depending on the type of backup you are defining, select the **Configuration Backup on DR** or **Full Backup on DR** checkbox.
4. Complete the remaining fields on that page as needed, and then click **[Save]**.
5. If necessary, follow the instructions in the section on [Mounting Backup Files](#).
6. To restore the backup, follow the instructions in the section on [Restoring Backups](#).

NOTE: After you define configuration or full backups on the Disaster Recovery Database Server, you should disable any other existing backups. However, you can continue to use standard DR Backup to back up your Disaster Recovery system.

Configuration Backup on a Disaster Recovery Database Server

To perform a configuration backup on the Disaster Recovery Database Server:

1. Go to the **Backup Management** page (System > Settings > Backup).
2. In the **Configuration Backup** pane, select the **Configuration Backup on DR** checkbox.
3. Complete the remaining fields on that page as needed, and then click **[Save]**.
4. Follow the instructions in the section on [Creating a Credential](#).
5. Following the instructions in the section on [Defining a Configuration Backup](#).
6. To restore the backup, follow the steps in the section [Restoring a Configuration Backup](#).

NOTE: If you enabled configuration backups on the Disaster Recovery Database Server, you should disable standard configuration backups.

Full Backup on a Disaster Recovery Database Server

To perform a full backup on the Disaster Recovery Database Server:

1. Go to the **Backup Management** page (System > Settings > Backup).
2. In the **Full Backup** pane, select the **Full Backup on DR** checkbox.
3. Complete the remaining fields on that page as needed, and then click **[Save]**.
4. Follow the instructions in the section on [Creating a Credential](#).
5. Following the instructions in the section on [Defining a Full Backup](#).
6. To restore the backup, follow the steps in the section [Restoring a Full Backup](#).

NOTE: If you enabled full backups on the Disaster Recovery Database Server, you should disable standard full backups.

Chapter

13

Viewing License Data

Overview

This chapter describes license data for 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 (...).

Viewing License Information

The **License Information** modal enables you to:

- View a list of all third-party licenses that are aligned with Skylar One
- Search for specific licenses
- View the full text of each license

The **Skylar One License Info** page lists all of the licenses that are aligned with your Skylar One system.

To view license information:

1. Go to the **Skylar One License Info** page (Misc > Skylar One License Info).

NOTE: To view this page in the classic Skylar One user interface, click the Toolbox button in the upper-right of the Skylar One browser session and then select *License Information*.

2. On the **Skylar One License Information** page, you can do the following:

- To view any license in its entirety, click its right-arrow icon. When you do, the icon becomes a down-arrow, and the full license information appears.
- To view all of the licenses in their entirety, click the **Expand All** link.
- To view only the condensed information for each license, click the **Collapse All** link.
- To search for a specific license, type part or all of its name in the search box in the upper-right of the page and then press the **Enter** key.

Chapter

14

Subscription Data

Overview

If you have a subscription license, you can use the **Subscription Usage** page (Manage > Subscription Usage) to:

- View current or historical subscription license usage graphs for Skylar One (formerly SL1)
- Download subscription license usage data for manual upload to the ScienceLogic billing server
- Upload a receipt from the ScienceLogic billing server

NOTE: This page is not available in the classic user interface.

If your Skylar One system is configured to communicate with the ScienceLogic billing server, usage data will be sent automatically from your Skylar One system to the ScienceLogic billing server once a day. After the ScienceLogic billing server receives the usage data, Skylar One will automatically mark the license usage file as delivered.

If your Skylar One system is not configured to communicate with the ScienceLogic billing server or if the connection to the ScienceLogic billing server fails, you can manually upload usage data to the ScienceLogic billing server.

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>Ensuring Accurate Data</i>	335
<i>Viewing Subscription Usage</i>	335
<i>Viewing Delivery Status</i>	339
<i>Manually Uploading License Usage to ScienceLogic</i>	340
<i>Data Retention Settings for Licensing</i>	341

Ensuring Accurate Data

If the reports described in this chapter do not appear to be correct, validate your configuration by opening the **Subscription Usage** page (Manage > Subscription Usage) and clicking the **[License Data Delivery Status]** button to view the **License Data** modal. For secure or air-gapped systems, the latest "Summary Date" value will be within the last 48 hours.

Viewing Subscription Usage

IMPORTANT: ScienceLogic uses the data collected for each device to assign a license to the device, and you are billed based on the number of device licenses you consume. The rules that govern how licenses are assigned are based on your subscription with ScienceLogic. Certain types of devices do not count against your device license usage in Skylar One. For more information, see the [Non-billable Devices](#) section.

The **Subscription Usage** page (Manage > Subscription Usage) provides a visual overview of your device license usage, including historical and current information. The information available in each panel of the page is described in this section and is accessible from the drop down menu in the top right of the page.


Device Name	Organization	Class
localhost	System	ScienceLogic, Inc. EM7 All-In-One
test_by_device_application_group	System	NET-SNMP Linux
test_by_device_class	System	OpenStack Cloud Physical Service
test_by_device_class	System	OpenStack Cloud Physical Service
test_by_inactive_device	System	Virtual Device Content Verification

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.


Current License Usage Totals

The bar chart in the upper-left corner of the **Subscription Usage** page shows the current total consumption of device licenses by license type. Each bar represents a license type. You can hover over a bar to see the number of licenses currently being consumed for that type. In the example below, we have four license types. We see that we are currently using 12 Ratio licenses.



Historical License Usage Totals

The line graph in the upper-right corner of the **Subscription Usage** page shows subscription license usage over time by license type. You can hover over a data point to see how many licenses of a given type were in use at that point in time.



A Date Range selector lets you narrow your view by using one of the following:

- All
- Current Quarter
- Last Quarter
- Current & Last Quarter
- Last Year

Whereas the "Current License Usages Totals" bar chart showed us that the most used license type currently is the Ratio license, the historical graph shows us that the Non-Billable license type is the most used over the given time range.

Current License Totals

The Current License usage table in the lower half of the **Subscription Usage** page lets you view information about your current subscription license usage in one of two ways:

- **Current License Usage by Device.** Choose "Current License Usage" from the drop-down list. Provides a list of all devices and the licenses they are using. You can search the table or filter on Device Name, Class, Organization, Category, or License Type.

Current License Usage by Device				
Current License Usage	Type to search device licenses			
	ORGANIZATION	CLASS	CATEGORY	LICENSE TYPE
Current License Usage	System	NET-SNMP Linux	Servers	Test Inactive
Current License Counts	System	Virtual Device Content Verification	Virtual	Test Inactive
Current License Category	System	Virtual Device Content Verification	Virtual	Test Inactive
test_by_inactive_device	System	Virtual Device Content Verification	Virtual	Test Inactive

NOTE: You can click on a device to open the **Device Investigator** page for the device. For more information, see *Using the Device Investigator* in the online help.

TIP: If you are looking for a specific device license or licenses, click the gear icon (⚙) to the right of the **Search** field and select *Advanced*. For more information, see the "Performing an Advanced Search" topic in the *Introduction to Skylar One* manual.

- **Current License Counts by Device Class.** Choose "Current License Counts" from the drop-down list. Provides a list of all device classes that are associated with devices in the system, the device category, the total licenses currently consumed by devices of that device class, and a breakdown count for each license type for the device class. You can filter this table by device Class or Category.

Current License Counts by Device Class						
Current License Counts	Type to search device class licenses					
CLASS	CATEGORY	TOTAL	INFRASTRUCTURE	NON-BILLABLE	RATIO	STORAGE
AWS Account	Cloud.Account	1	1	0	0	0
AWS Service	Cloud.Service	1	0	1	0	0
Cisco Systems CRS-1 1...	Network.Router	1	0	0	1	0
Citrix NetScaler	Network.Services	1	0	0	1	0
Dell EMC Unisphere fo...	Storage.Management	1	0	0	0	1
Dell EMC Unity LUN	Storage.LUN	2	0	0	0	2
Dell EMC Unity Storag...	Storage.Pool	1	0	0	0	1

- **Current License Counts by Device Category.** Choose "Current License Category" from the drop-down list. Provides the Total number of licensed categories, their Test status, and number of Tests by Category.

Current License Counts by Device Category			
Current License Category	Type to search device category licenses		
Category	Total	Test Inactive	Test by Category
Servers	1	1	0
Virtual	9	7	2

TIP: The "Current License Category" category information also auto-populates and can be displayed in the "Current License Totals" bar chart above.

Viewing Delivery Status

The **License Data Delivery Status** modal displays the status of one or more daily license usage files. To view the **License Data Delivery Status** page:

NOTE: This feature is only available for secure users with systems that do not send data to the billing server daily. It is not populated for standard online users.

1. Go to the **Subscription Usage** page (Manage > Subscription Usage).
2. In the **Current License Usage by Device** section, click the **[License Data Delivery Status]** button.
3. The **License Data Delivery Status** modal appears and displays a list of daily license usage files. For each daily license usage file, the **License Data Delivery Status** page displays the following:
 - **Summary Date.** Date associated with the daily license-usage file.
 - **Delivery Status.** Possible values are:
 - "0" (zero). File has not been uploaded to the ScienceLogic billing server.
 - "1" (one). File has been uploaded to the ScienceLogic billing server and may be deleted from the Skylar One system by the automated maintenance process.
 - **Summary Size.** Size of the daily license usage file.

Manually Uploading License Usage to ScienceLogic

If your Skylar One system is configured to communicate with ScienceLogic, usage data will automatically be sent to the ScienceLogic billing server once a day. After the ScienceLogic billing server receives the usage data, Skylar One will automatically mark the license usage file as delivered.

If your Skylar One system is not configured to communicate with ScienceLogic or if the connection to the ScienceLogic billing server fails:

- You can use the **License Data Delivery Status** page to manually download the daily license-usage file.
- You can then log in to the ScienceLogic billing server and manually upload the daily license-usage file.
- You can then use the **License Data Delivery Status** page to upload the ScienceLogic "receipt" to your Skylar One system, allowing Skylar One to mark the license usage file as delivered.
- License usage files will not be deleted from your system until they are delivered.

Downloading the Daily License Usage File

If your Skylar One system is not configured to communicate with ScienceLogic or if the connection to the ScienceLogic billing server fails, you can use the **License Data Delivery Status** page to manually download the daily license usage file. You can then log in to the ScienceLogic Licensing and Billing server and manually upload the daily license usage file.

To download the daily license-usage file using the **License Data Delivery Status** page:

1. Go to the **Subscription Usage** page (Manage > Subscription Usage).
2. Click the **[License Data Delivery Status]** button.
3. Select one or more daily license usage files to download to your local computer, and then click the **[Download]** button.

NOTE: If the download size exceeds 50 MB, the **[Download]** button is disabled.

4. The daily license usage file is saved to your local computer. The downloaded file is usually named "license_data.json.gz".

Manually Uploading the Daily License Usage File to ScienceLogic

After downloading the daily license usage file to your local computer, you can manually upload the file to the ScienceLogic billing server:

1. Log in to the ScienceLogic billing system.
2. Go to the **Subscription Data** page (Preferences > Account > Subscription Billing).
3. In the **Subscription Data** page, go to the **Subscription Data Update** pane. Use the **[Browse]** button to find the daily license-usage file that you downloaded to your local computer.
4. Click the **[Get Update]** button to upload the daily license-usage file to the ScienceLogic server.
5. The ScienceLogic server will provide a "receipt" file for you to download. This file is usually called "status_updated.json.gz". You must upload this receipt to your Skylar One system.

Uploading the ScienceLogic Receipt

After uploading the daily license usage file to the ScienceLogic Support Center, your customer success manager will reply with a "receipt" file for you to download. This file is usually called "status_updated.json.gz".

You must upload this "receipt" file to your Skylar One system to inform your Skylar One system that the upload was successful and that the Skylar One system may delete the daily license usage file.

To upload the "receipt" file:

1. Go to the **Subscription Usage** page (Manage > Subscription Usage).
2. Click the **[License Data Delivery Status]** button.
3. In the **Status Update File** field, click **[Choose File]** and browse to locate the "receipt" file.
4. Click the **[Upload]** button to upload the "receipt" file to your Skylar One system.

Data Retention Settings for Licensing

The **Data Retention Settings** page contains settings for subscribers.

To adjust these settings:

1. Go to the **Data Retention Settings** page (System > Settings > Data Retention).
2. The following sliders appear under the **Subscription Data Retention** heading:
 - **Subscriber Device Configuration Data**. View Only. This feature is no longer populated but historical data will remain. For users with a subscriber license. Number of months to retain the files and database tables that contain configuration information for a device. Default value is two months.
 - **Subscriber Device Usage Data**. View Only. This feature is no longer populated but historical data will remain. For users with a subscriber license. Number of months to retain information on total number of events and total number of tickets. Default value is two months.
 - **Subscriber System Configuration Data**. View Only. This feature is no longer populated but historical data will remain. For users with a subscriber license. Number of months to retain the files and database tables that contain configuration information for the Skylar One system. Default value is three months.

- **Subscriber System Usage Data.** View Only. This feature is no longer populated but historical data will remain. For users with a subscriber license. Number of months to retain information on total number of events and total number of tickets. Default value is three months.
- **Subscriber Device Type Data.** For users with a subscriber license. Number of months to retain the files and database tables that map each device to a device category, as per your subscriber license. Default value is three months. This data populates the license usage report and the license usage user interface reference in **Current License Totals** in the [Viewing Subscription Usage](#) section.
- **Subscriber Daily Delivery Data.** For users with a subscriber license. Number of months to retain the "crunched" license usage data that is calculated each day using the Subscriber Device Configuration Data, Subscriber System Configuration Data, Subscriber System Usage Data, and Subscriber Device Type Data. Skylar One will not prune data that has not yet been delivered to the ScienceLogic Licensing and Billing server. Default value is three months. This feature is available only for secure/air-gapped systems.

The **Data Retention Settings** page contains settings for feature data retention as well. This data is applicable to all systems.

To adjust these settings:

1. Go to the **Data Retention Settings** page (System > Settings > Data Retention).
2. The following sliders appear under the **Feature Data Retention** heading:
 - **Feature Configuration Data.** Feature data is information about how Skylar One is being used, including configuration data and performance metrics. Number of months to retain raw feature configuration values. Default value is six months.
 - **Feature Performance Data.** Feature data is information about how Skylar One is being used, including configuration data and performance metrics. Number of months to retain raw performance metric values. Default value is six months.
 - **Feature Performance Data Aggregation Daily.** For users with a subscriber license. Number of months to retain the files and database tables that contain aggregated daily-performance values. Default value is six months.
 - **Feature Performance Data Aggregation Hourly.** For users with a subscriber license. Number of months to retain the files and database tables that contain hourly-performance values. Default value is six months.

Feature Data Retention

Feature Configuration Data	 6 months* <small>[Current: 6 months*]</small>
Feature Performance Data	 6 months* <small>[Current: 6 months*]</small>
Feature Performance Data Aggregation Daily	 6 months* <small>[Current: 6 months*]</small>
Feature Performance Data Aggregation Hourly	 6 months* <small>[Current: 6 months*]</small>

Save

* 1 month = 30 days

Chapter

15

CAC Authentication

Overview

This chapter describes how Skylar One (formerly SL1) supports Common Access Card (CAC) authentication. The **Client Certificate & CAC Authentication** page (System > Settings > Authentication > Admin SSL Certificates, or System > Settings > Authentication > SSL Certificates in the classic Skylar One user interface) allows you to define a check for SSL certificate that controls whether the login page is displayed to the end user.

This feature is primarily used to authenticate Common Access Card (CAC) users against a Department of Defense (DoD) issued server-side certificate; however, based on your business needs, this feature can also be used with your own client/server certificates.

NOTE: You can use CAC authentication to log in to either the default Skylar One user interface ("AP2") or the classic Skylar One user interface. Follow the steps described in this chapter to configure your CAC authentication, regardless of which user interface you use.

NOTE: Currently, Skylar One does not support client-side certificate authentication for login to the console, either through SSH or through a keyboard connected to the appliance.

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>Using CAC Authentication</i>	345
<i>Prerequisites</i>	346
<i>Importing SSL Certificates</i>	347
<i>Extracting the Common Name from a Certificate for Authentication</i>	348
<i>Defining the Client Certificate Chain</i>	349
<i>Verifying SSL Certificate File Import and Resolving Issues</i>	351
<i>Clearing the Skylar One Cache and Restarting NGINX</i>	352
<i>Testing the Configuration</i>	352
<i>Troubleshooting CAC Authentication</i>	353
<i>Accessing the Appliance without CAC Authentication</i>	355
<i>Special Circumstance: Multiple Levels of Intermediate Certificates</i>	355

Using CAC Authentication

Skylar One supports CAC authentication. The [Client Certificate & CAC Authentication](#) page allows you to define a check for SSL certificate that controls whether the login page is displayed to the end user. This feature is primarily used to authenticate Common Access Card (CAC) users against a Department of Defense (DoD) issued server-side certificate; however, based on your business needs, this feature can also be used with your own client/server certificates.

The CAC is a United States DoD smartcard issued as standard identification for Active Duty Military personnel, reserve personnel, civilian employees, and eligible contractor personnel. A User Principal Name (UPN) is recommended, and in some instances required, when using CAC.

CAC provides applications with a more secure way to authenticate the identity of a user, application, or device. However, even if a user authenticates with a certificate, it does not mean that they user is authorized to access the requested data. For more information on authentication and authorization, see the DoD documentation on authentication and authorization for DoD web servers.

DoD has implemented an external interoperability strategy for secure information sharing with external partners. Some DoD industry partners have implemented corporate PKIs, and others have obtained certificates from approved commercial PKIs. Some DoD international allied and coalition partners also have established PKIs to issue certificates to their personnel. Systems and applications with user populations that hold approved external credentials should be configured to accept those credentials rather than requiring the users to obtain Common Access Cards (CACs) or External Certification Authority (ECA) certificates. For the complete list of DoD approved external PKIs and interoperability tools, see the DoD documentation on interoperability.

DoD policy requires that external credentials have an assurance level of medium hardware or higher, so systems accepting external credentials must have an assurance level enforcement capability. Depending on technology, this can be accomplished through use of the Interoperability Root CAs (IRCAs) or implementation of a local certificate policy object identifier (OID) filtering solution such as the DoD PKE Trust Anchor Constraints Tools (TACT). For a complete list of approved partner OIDs, see the DoD documentation on the approved assurance levels from external partner PKIs.

Systems and applications typically have configuration properties that control security settings related to PKI functionality. Security settings should be configured to support all desired PKI functions and comply with DoD authentication policy.

Skylar One allows you to configure appliances that provide the user interface (Administration Portal, All-In-One Appliance, or the Database Server) for use with DoD certificates or your own certificates.

The CAC is used as the user's authentication to Skylar One. If the Authentication Profile (System > Settings > Authentication > Profiles) contains both the "CAC/Client Cert" and "EM7 Login Page" credential sources *and* a CAC is not presented or is invalid, then the ScienceLogic login page is presented to the end user.

NOTE: You can use CAC authentication to log in to either the default Skylar One user interface ("AP2") or the classic Skylar One user interface. Follow the steps described in this chapter to configure your CAC authentication, regardless of which user interface you use.

NOTE: Currently, Skylar One does not support client-side certificate authentication for login to the console, either through SSH or through a keyboard connected to the appliance.

Prerequisites

To use client certificate authentication with Skylar One, you must first meet the following requirements:

1. Organizations must be created and configured. For more information, see "Creating and Editing Organizations" in the *Organizations and Users* manual.
2. An LDAP or AD Credential must be configured with a Service Account that has the appropriate permissions to query AD, typically read access.
3. Create one or more User Policies if you will use Skylar One authentication. You do not need to configure user policies if you are using Active Directory (AD) or LDAP. For more information, see "Creating a User Policy" in the *Organizations and Users* manual.
4. If you are using LDAP or Active Directory as your user store, you must configure this as your Authentication Resource before setting up your Authentication Profile. For more information, see *Authentication Resources*.
5. Configure an Authentication Profile for CAC authentication. When setting up your Authentication Profile for CAC, align the "CAC/Client Cert" credential with the profile as the first credential source. You can align the EM7 Login Page as a secondary credential source for administrator access, but this is not required. For more information, see *Creating an Authentication Profile*.
6. Configure an emergency account ("break glass" account) for the Database Server. Because CAC will work only with the Database Server's DNS name, an emergency account ensures that the em7admin account is used only as a last resort.
7. Go to the **Authentication Profiles** page (System > Settings > Authentication > Profiles). Select the wrench icon (🔧) for the default profile. In the **Authentication Profile Editor** page, in the **Aligned Credential Sources** field, delete any existing CAC/Client Cert credentials.
8. Your users must have either:
 - Valid CACs with valid client-side certificates already loaded onto the cards, or
 - Valid client-side certificates installed in their web browser.
9. If CACs are used, the browser through which the user logs on to the user interface must be able to read security certificates from the cards.
10. The Administration Portal, All-In-One Appliance, or the Database Server will request a certificate from the CAC or client web server only when the appliance uses HTTPS. In Skylar One 12.2.0 and above, the use of HTTPS is enforced by default. In Skylar One versions prior to 12.2.0, you must go to the **Behavior Settings** page (System > Settings > Behavior), and select the **Force Secure HTTPS** setting checkbox.

NOTE: In Skylar One 12.2.0 and above, the **Force Secure HTTPS** setting does not appear as an option on the **Behavior Settings** page.

11. On the **SSL Certificates** page (System > Settings > Authentication > Admin SSL Certificates, or System > Settings > Authentication > SSL Certificates in the classic Skylar One user interface), you must install the certificate chain in PEM format on the Administration Portal, All-In-One Appliance, or the Database Server. A certificate chain usually includes a root CA certificate and an intermediate certificate. Your organization might require multiple intermediate certificates to provide access to all users. To learn more about importing a certificate, see the section [Importing an SSL Certificate](#).

NOTE: If you want to extract part of the Common Name to customize the username that is displayed in Skylar One after CAC authentication, you can edit the ScienceLogic configuration file to customize the displayed username. You do not need to do this if you are using the msUPN. For more information, see [Extracting the Common Name from a Certificate for Authentication](#).

12. In the **Client Certificate & CAC Authentication** page (System > Settings > Authentication > CAC Client Cert Auth, or System > Settings > Authentication > CAC/ClientCert Auth in the classic Skylar One user interface), you must configure the server-side certificate and test it against your client-side certificate. For more information, see [Defining the Client Certificate Chain](#).

Importing SSL Certificates

Secure Sockets Layer, or SSL, is a protocol for securely transmitting data via the Internet. SSL uses a private key to encrypt data to be transferred over an Internet connection. In Skylar One, you can import server-side SSL certificate files, including DoD certificate files used in CAC authentication, to the Administration Portal, All-In-One Appliance, or the Database Server.

Note the following:

- You must have one root certificate and one certificate for each intermediate authority in the client certificate chain. If you have users with CACs issued by different intermediate authorities, you must import SSL certificates for all possible client authentication chains into Skylar One.
- All SSL certificates must be in PEM format.
- You can test your SSL certificate files by using the following command, where <certificate_file_name> is the full name of the certificate file:

```
openssl x509 -text -noout -in <certificate_file_name>
```

TIP: It is a best practice to check each certificate file before attempting to import the file. If you encounter an error, resolve that error before you continue.

To import an SSL certificate for CAC authentication:

1. Go to the **SSL Certificates** page (System > Settings > Authentication > Admin SSL Certificates, or System > Settings > Authentication > SSL Certificates in the classic Skylar One user interface).
2. In the **SSL Certificates** page, click the **[Actions]** menu. Select **Import PEM Certificate File**. The **Import Certificate File (PEM format)** modal appears.
3. In the **Import Certificate File (PEM format)** modal, enter the following:
 - **Description**. Description of the certificate.
 - **CA File**. Browse for the server-side certificate file on your local computer.
4. Click the **[Save]** button to load the certificate to the Administration Portal, All-In-One Appliance, or the Database Server.
5. Repeat these steps for each certificate file you want to import. When finished, verify that all of your certificates appear in the listing shown on the SSL Certificates page (System > Settings > Authentication > Admin SSL Certificates, or System > Settings > Authentication > SSL Certificates in the classic Skylar One user interface).

TIP: A best practice is to make note of the value in the Hash field shown for each certificate and verify that the hash values match the symlink files in the /var/lib/em7/certs directory on the appliance after completing the configuration of your client certificate chain. You will use these hash values in [Verifying SSL Certificate File Import and Resolving Issues](#).

Extracting the Common Name from a Certificate for Authentication

By default, the certificate configuration file (em7_certificate.conf) is configured to display the full common name (CN) of the CAC user as the username in Skylar One after authentication. If this meets your requirements, then you do not need to update the configuration file and can skip this section.

NOTE: If you are using the Microsoft User Principal Name (MS UPN) in your certificates, you do not need to make any edits in the configuration file.

However, if you require that Skylar One use only a portion of the CN, then you can edit the certificate configuration file to parse out a username from the CN in the certificate.

For example, in some instances you might want to use an employee's ID number as the username. To do that, you must edit the Nginx configuration file.

To do so:

1. Log in to the console of the Skylar One appliance as the root user.
2. Navigate to the directory `/etc/nginx/conf.d/`:

```
cd /etc/nginx/conf.d/
```

3. Open the file em7_certificate.conf with a text editor like vi:

```
vi em7_certificate.conf
```

4. Modify the file to extract the CN from the full Distinguished Name (DN) found in the certificate based on how you want to map the username to an LDAP system or how you want the usernames to look if you are using Skylar One internal as the backend of your authentication configuration.

This is the default configuration of the file:

```
# Create the Username for EM7 to use from the Certificate
# Default: Pull the Common Name from the DN.
map $ssl_client_s_dn $ssl_client_username {
  ~/?CN=(?<CN>[^/,]+) $CN;
}
```

Modify the string to extract the name. The following is a regular expression that extracts the CN from the full DN found in the certificate:

```
map $ssl_client_s_dn $ssl_client_username { ~/CN=[A-Z\.]+(?<num>[0-9]+)
  $num; }
```

5. Save and quit (:wq) the file.

Defining the Client Certificate Chain

After importing your SSL Certificates, you must consolidate the SSL PEM certificates into a combined file (emt_combined.crt). On the CAC/ClientCert Auth menu, select all of the desired SSL PEM certificates. After saving, Skylar One will update the em7_combined.crt file with all of the selected SSL PEM certificates. Skylar One will then use only the selected PEM certificates for validating and authenticating users.

You can also define some custom settings for client-side certificate authentication. You can define error messages that are displayed to the end user if authentication fails. Optionally, you can also define IP addresses in this modal for which the user interface will not perform certificate authentication, if you have not already created an Authentication Profile for this purpose. See [Accessing the Appliance without CAC Authentication](#) for more information.

When authentication is successful, the user interface displays the **ScienceLogic Login** page to the user.

To define the authentication settings:

1. Access the user interface with your CAC or a browser with your client-side certificate installed.
2. Go to the **Client Certificate & CAC Authentication** page (System > Settings > Authentication > CAC Client Cert Auth, or System > Settings > Authentication > CAC/ClientCert Auth in the classic Skylar One user interface).
3. Supply a value in each of the following fields:

- **Root CA Certificates.** Select *all* root and intermediate certificates that make up the chain from a list of certificates installed on the **SSL Certificates** page (System > Settings > Authentication > Admin SSL Certificates, or System > Settings > Authentication > SSL Certificates in the classic Skylar One user interface). Your client-side certificate will be authenticated against the selected server-side root and intermediate certificates.
- **Auth Failure Message.** Enter text for the error message that appears to users if authentication fails.

CAUTION: You cannot save your authentication settings until you enter text in the "Auth Failure Message" field.

- **Ignore Networks.** In this field, you can enter a list of networks and hosts from which certificate authentication *is not required*. During each login, the platform will compare the client's IP address to the list entered in this field. If the client's IP address is included in this field, Skylar One will not require certificate authentication from that client.

NOTE: If you are using Authentication Profiles to configure access from specific resources from which certificate authentication is not required, you do not need to use the *Ignore Networks* field. For more information, see [Accessing the Appliance without CAC Authentication](#).

- In the *Ignore Networks* field, you can enter one or more IP addresses, each separated by a new-line character (press the [**<Enter>**] key).
- In the list of IPs to ignore, you can enter only the first octet, only the first and second octet, only the first, second, and third octet, or all four octets. Skylar One will interpret the entry as if the rightmost octet is followed by * (asterisk).

For example:

- 192.168.10.142 will allow a single host to log in to the user interface without certificate authentication
- 192 behaves the same as entering 192*. This will allow all hosts included in 192.0.0.1 through 192.254.254.254 to log in to the user interface without certificate authentication
- 192.168.10.24 behaves the same as entering 192.168.10.24*. This will allow all hosts 192.168.10.24, 192.168.10.240, 192.168.10.241, 192.168.10.242, 192.168.10.243, 192.168.10.244, 192.168.10.245, 192.168.10.246, 192.168.10.247, 192.168.10.248, and 192.168.10.249

4. Click the **[Save]** button to save your settings. The user interface displays the message:

Settings Saved Successfully. Configuration must be tested in order to take effect.

CAUTION: Do not click the Test link at this time.

Verifying SSL Certificate File Import and Resolving Issues

After you have imported your SSL certificates and configured your client certificate chain, it is important to verify the your certificate files were imported correctly and are valid in Skylar One.

To verify that your SSL certificate files were imported correctly:

1. Either go to the console of the Skylar One appliance where you imported the SSL certificates, or use SSH to log in.
2. Navigate to the /var/lib/em7/certs directory. At the shell prompt, enter:

```
ls -l
```

3. Review the list of hash symlink files in the directory and compare them to the list of certificates on the SSL Certificates page. Ensure that the hash values shown in Skylar One match the hash symlink files. Note that the hash symlink in the /var/lib/em7/certs directory (in blue text) for a certificate file is appended with ".0", as shown in the image below.

```
[root@aio-169-161 em7admin]# ll /var/lib/em7/certs/
total 140
lrwxrwxrwx. 1 s-em7-http s-em7-core 21 Jun 25 03:41 4f5db21f.0 -> DoD_.pem
lrwxrwxrwx. 1 s-em7-http s-em7-core 21 Jun 25 03:41 60085f15.0 -> DoD_.pem
lrwxrwxrwx. 1 s-em7-http s-em7-core 21 Jun 25 03:41 9cf5f371.0 -> DoD_.pem
lrwxrwxrwx. 1 s-em7-http s-em7-core 21 Jun 25 03:41 d7fc5635.0 -> DoD_.pem
-rw-r--r--. 1 s-em7-http s-em7-core 1716 Jun 25 03:35 [REDACTED].pem
-rw-r--r--. 1 s-em7-http s-em7-core 1716 Jun 25 03:36 [REDACTED].pem
-rw-r--r--. 1 s-em7-http s-em7-core 1753 Jun 25 03:36 [REDACTED].pem
-rw-r--r--. 1 s-em7-http s-em7-core 1753 Jun 25 03:37 [REDACTED].pem
-rw-r--r--. 1 s-em7-http s-em7-core 1753 Jun 25 03:37 [REDACTED].pem
-rw-r--r--. 1 s-em7-http s-em7-core 1269 Jun 25 03:35 [REDACTED].pem
lrwxrwxrwx. 1 s-em7-http s-em7-core 22 Jun 25 03:41 ebe73690.0 -> DoD_.3.pem
lrwxrwxrwx. 1 s-em7-http s-em7-core 21 Jun 25 03:41 ec465775.0 -> DoD_.pem
-rw-rw-r--. 1 s-em7-core s-em7-core 9960 Jun 25 03:41 em7_combined.crt
-rw-rw-r--. 1 s-em7-core s-em7-core 1489 Jun 2 00:35 em7_default.crt
-rw-rw-r--. 1 s-em7-core s-em7-core 11972 Jun 2 00:35 em7_import_dodeca2.cac
-rw-rw-r--. 1 s-em7-core s-em7-core 10485 Jun 2 00:35 em7_import_dodeca.cac
-rw-rw-r--. 1 s-em7-core s-em7-core 5302 Jun 2 00:35 em7_import_rel3_dodroot_1024.cac
-rw-rw-r--. 1 s-em7-core s-em7-core 66452 Jun 2 00:35 em7_import_rel3_dodroot_2048.cac
-rw-r--r--. 1 s-em7-http s-em7-core 0 Jun 25 03:41 f9b9dee864d0b27dda9dde4bbdfb9cf7.sync
[root@aio-169-161 em7admin]#
```

All of the following must be true. If any of these are not true, then the certificate file was not imported and saved correctly in Skylar One:

- One hash symlink file should exist in the directory for each of the imported certificate files.
- The file size of the "em7_combined.crt" file is equal to the combined file sizes of all of the certificate (.pem) files. (The file "em7_combined.crt" is not equal to the "em7_default.crt" file.)

- When you view the contents of the "em7_combined.crt" file using `cat` or similar command, the file is the concatenation of all of the certificate (.pem) files. NGINX references the "em7_combined.crt" file as the file containing the client certificate chain.

4. If any of the conditions listed above are not true, then the certificate file was not imported and saved correctly in Skylar One. To resolve the problem:
 - a. Log in to the Skylar One appliance user interface.
 - b. Go to the **Client Certificate & CAC Authentication** page (System > Settings > Authentication > CAC Client Cert Auth, or System > Settings > Authentication > CAC/ClientCert Auth in the classic Skylar One user interface).
 - c. Edit the *Auth Failure Message* field. Make a change to the section.
 - d. Click **[Save]**.
 - e. Repeat steps 1-3 to verify your certificate files.

Clearing the Skylar One Cache and Restarting NGINX

Before you proceed to testing the configuration, you must clear the Skylar One cache, restart nginx, and close any browsers you have open. This will ensure the best outcome when testing.

1. Log in to the user interface on the Skylar One appliance.
2. Click on the **[Toolbox]** button (≡) and choose Misc > *Clear Skylar One System Cache*.
3. Log out of the Skylar One appliance.
4. Either go to the console of the Skylar One appliance where you imported the SSL certificates, or use SSH to log in to the appliance.
5. Run the following command to restart nginx:

```
sudo systemctl restart nginx
```

6. Close any open browsers that have been used to access the appliance.

Testing the Configuration

After you define the certificate authentication settings, you must test your client-side certificate against the server-side certificate you selected in the **Root CA Certificates** field. Testing your configuration is required to prevent an incorrect configuration from preventing administrator access to the user interface. If the test is successful, the certificate authentication settings will be applied. If the test is unsuccessful, the certificate authentication settings will not be applied.

To test certificate authentication settings:

1. With your CAC inserted in the reader, access the user interface of your Skylar One appliance using the IP address or domain name defined in the **AP Hostname Pattern** field of the CAC Authentication Profile (System > Settings > Authentication > Profiles). Log in with an administrator account.

2. Go to the **Client Certificate & CAC Authentication** page (System > Settings > Authentication > CAC Client Cert Auth, or System > Settings > Authentication > CAC/ClientCert Auth in the classic Skylar One user interface).
3. After defining the certificate, you will see the following message at the top of the pane:

Configuration must be tested in order to take effect: TEST.

4. Click the **TEST** link. Skylar One will attempt to authenticate your client-side certificate against the selected server-side certificate.
5. If the test authentication is successful, Skylar One will display the following message at the top of the pane and end users with the appropriate client certificate or CAC can now access the user interface using client certificate authentication:

Configuration verified and enabled.

6. A new field, **Client Cert / CAC Auth**, appears with a default value of *Allowed*. Do not edit this field.
7. Set the *Certificate User Field* to "Common Name" (default) or "MS UPN".

NOTE: If you are using LDAP or Active Directory (AD) for user authentication, set this field to "MS UPN".

8. Select the **[Save]** button to save the setting in the **Client Cert / CAC Auth field**.
9. If the test authentication is unsuccessful, the user interface will display the following message at the top of the pane. The settings will not be applied, and client certificate authentication will not be used until the problem is corrected:

ERROR: configuration was not successfully tested with CAC or Client Certificate.

NOTE: If you experience the error above, double-check the following: verify there are not any simple mistakes by reviewing any information you manually entered; check to see if there is a mismatch between the certificate chain installed in nginx versus what the browser uses; make sure the cert file names do not contain spaces or blanks in the file name.

Troubleshooting CAC Authentication

There are a few common issues you might experience while testing CAC authentication. If your test is unsuccessful, review the following troubleshooting steps.

Failed to Identify Personal Identity Verification (PIV) Card

If you receive the "Failed to identify PIV card" message, verify the following:

- All root and intermediate certificates have been uploaded in a PEM format.
- All root and intermediate certificates have not expired and are configured properly.
- The client certificate has not expired.
- Customer username information is in the Microsoft User Principal Name (MS UPN and not the Common Name (CN)).

Failed CAC Authentication After Disaster Recovery (DR) Failover

If your CAC authentication testing fails after DR failover, verify the following:

- The DR node domain has been added to the Auth Profile for pattern matching.
- The DR has all of the required certificates. (You might be required to manually upload and save the certificates again.)

NOTE: For more information, see the chapter on "Disaster Recovery with Two Appliances" in the *High Availability and Disaster Recovery Configuration* manual.

Failed CAC Authentication After Setting Up High Availability (HA)

Assuming you have deployed a Skylar One distributed system with one Database Server and two or more Administration Portals and your CAC authentication testing fails after setting up HA, you must ensure the following:

- The first Administration Portal where you configured CAC is working and you authenticated with CAC before verifying that the second or third Administration Portal actually works. You should not have the CAC/Client Cert in the aligned credential source on the default profile but in a new profile created for CAC only.

NOTE: ScienceLogic suggests having at least two profiles (a default profile and a CAC profile). You should enter an AP Hostname Pattern on the CAC profile but keep the AP Hostname Pattern blank on the default profile.

- Upon successful CAC login on the first Administration Portal, you will notice that any login attempts to the second or third Administration Portal will fail. For this, you need to first verify that the content of the `/var/lib/em7/certs` contains the PEM files that are identical to the first Administration Portal. You must also ensure that the hash files representing your PEM files are identical to the first Administration Portal and that the combined file is identical to the first Administration Portal.
- Once verified, restart nginx on the second and third Administration Portals and ensure that nginx is running correctly.

- Verify that you can log in with CAC from the second or third Administration Portals as you have done with the first Administration Portal.

NOTE: For more information, see the chapter on "High Availability with Two Appliances" in the *High Availability and Disaster Recovery Configuration* manual.

Accessing the Appliance without CAC Authentication

In certain circumstances, you might need to access your Skylar One appliance without using CAC authentication. For example, the following are some reasons you might want to use another authentication type:

- For use during initial setup
- For appliance access when a certificate has expired
- For maintenance or administrator accounts
- For certain internal networks that will not require certificate authentication

You can configure the appliance to accept a login in these cases in two ways:

- By configuring an Authentication Profile to use an alternative authentication resource (for example, EM7 Internal) for certain networks or hosts. For more information, see the chapter on "Authentication Profiles" in the *System Administration* manual.
- By using the Ignore Networks field on the **Client Certificate & CAC Authentication** page (System > Settings > Authentication > CAC Client Cert Auth, or System > Settings > Authentication > CAC/ClientCert Auth in the classic Skylar One user interface). For more information, see [Defining the Client Certificate Chain](#).

Special Circumstance: Multiple Levels of Intermediate Certificates

By default, Skylar One is configured to handle the typical certificate hierarchy, which comprises three levels: root, intermediate, and client certificates. This represents a depth of 2 from the root to the client certificate. If your organization will use CAC authentication in which you have multiple levels of intermediate certificates in the hierarchy, you will need to change this setting (`ssl_verify_depth`) as described in the procedure below.

To update the value of `ssl_verify_depth`:

1. Log in to the console of the ScienceLogic appliance as the root user.
2. Navigate to the directory `/etc/nginx/conf.d/`:

```
cd /etc/nginx/conf.d/
```

3. Open the file em7ngx_web_ui.conf with a text editor like vi:

```
vi em7ngx_web_ui.conf
```

4. Edit the ssl_verify_depth value to be the depth from client certificate to the root certificate (for example, 3):

```
ssl_verify_depth 3;
```

5. Save and quit (:wq) the file.

Chapter

16

Authentication Profiles and Resources

Overview

This chapter describes the following topics:

- **Authentication Profiles**. Policies that align user accounts with one or more types of authentication.
- **Authentication Resources**. Configuration policies that describe how Skylar One (formerly SL1) should communicate with a user store.

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:

Authentication Profiles	358
Authentication Resources	362

Authentication Profiles

Authentication profiles are policies that align user accounts with one or more types of authentication:

- **Alignment by pattern matching.** Skylar One examines the URL or IP address that a user enters in a browser to connect to an Administration Portal, Database Server, or All-In-One Appliance. If the URL or IP address matches the criteria specified in an authentication profile, Skylar One will automatically use the matching profile to perform user authentication.
- **Credential Source.** Specifies from where Skylar One should extract the username and password or certificate to be authenticated. These credentials are passed to Skylar One through HTTP. Skylar One then passes the credentials to each authentication resource specified in the authentication profile (for example, CAC/Client Cert). The authentication resources communicate with user stores that can authenticate the credentials entered by a user.
- **Authentication Resource.** Specifies the connector to use to communicate with the user store, the credential to use to connect to the user store (if applicable), such as your Active Directory server, and the URLs to examine during authentication. Authentication Resource also maps attributes from the user's account in the user store to fields in the ScienceLogic user account. For details on creating an authentication resource, see the section on [Authentication Resources](#).

NOTE: If you will be using Single Sign-On (SSO) as your method of authentication, your SSO resource must be placed in its own Authentication Profile, since it will take priority over any other authentication method defined. If you have multiple SSO resources, each *must* be in its own profile.

Viewing the List of Authentication Profiles

To view a list of all authentication profiles in Skylar One:

1. Go to the **Authentication Profiles** page (System > Settings > Authentication > Profiles).
2. The following information is displayed about each authentication profile:
 - **Profile Name.** Name of the authentication profile.
 - **Access.** Indicates whether the authentication profile is shared with all organizations or is private.

NOTE: The **Access** column displays only for Administrator user accounts and user accounts assigned to the System organization.

- **ID.** Unique numeric ID, automatically assigned by Skylar One to each authentication profile.
- **Hostname Pattern.** This field is used to match the URL or IP address that a user enters in a browser to connect to an Administration Portal, Database Server, or All-In-One Appliance. If the URL or IP address matches the value in this field, Skylar One applies the authentication profile to the user for the current session.
- **Priority Order.** If your Skylar One System includes multiple authentication profiles, Skylar One evaluates the authentication profiles in priority order, ascending. This column displays the priority order value for the authentication profiles, where 0 (zero) is the highest priority.
- **Edited By.** The user who created or last edited the authentication profile.
- **Last Edited.** Date and time the authentication profile was created or last edited.

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

The "default" Authentication Profile

Skylar One includes a *default* authentication profile, for which the following rules apply:

- You cannot delete the *default* profile.
- If an **AP Hostname Pattern** fails to match all the other authentication profiles, Skylar One applies the *default* authentication profile.
- In ISO systems, initially the *default* profile is pre-configured to allow ScienceLogic administrators to log in via *CAC/Client Certificate*, *HTTP Auth*, or the Skylar One *Login Page* and the authentication resource *Internal*. This allows administrators to log in and perform initial configuration on the Skylar One system. After initial configuration, administrators can edit the *default* profile as best fits their organization.

- In patched systems, the *default* profile is included in the patch and is pre-configured to allow ScienceLogic administrators to log in via the ScienceLogic login page and the authentication resource *Internal*. It allows credentials via *CAC/Client Certificate*, *HTTP Auth*, or the Skylar One *Login Page*. It can also use the legacy authentication resources *SSO (legacy)*, *LDAP/AD (legacy)*, and *Internal*.

Creating an Authentication Profile

To create a new authentication profile:

1. Go to the **Authentication Profiles** page (System > Settings > Authentication > Profiles).
2. Click the **[Create]** button. The **Authentication Profile Editor** modal appears.
3. Enter values in the following fields:
 - **Name**. Name of the authentication profile.
 - **Sharing Permission**. Indicates if the authentication profile is shared or private. Choices are:
 - *Shared with all organizations*. The authentication profile is shared with users across all organizations.
 - *Private (visible to System organization only)*. The authentication profile is private to only user accounts assigned to the System organization.

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

- **Priority Order**. If your Skylar One System includes multiple authentication profiles, Skylar One evaluates the authentication profiles in ascending priority order. Skylar One will apply the authentication profile that matches the hostname or IP in the current URL AND has the lowest value in the **Priority Order** field.
- **Pattern Type**. Specifies how Skylar One will evaluate the value in the **AP Hostname Pattern** field. Choices are:
 - *Wildcard*. Skylar One will perform a text match, with wildcard characters (asterisks).
 - *Regex*. Skylar One will use regular expressions to compare the **AP Hostname Pattern** to the current session information.
- **AP Hostname Pattern**. This field is used to match the URL or IP address that a user enters in a browser to connect to an Administration Portal, Database Server, or All-In-One Appliance. If the URL or IP address matches the value in this field, Skylar One applies the authentication profile to the user for the current session.
 - For example, if you specify "*" (asterisk), any IP address or URL will match. Skylar One will then apply this authentication profile to every session on an Administration Portal, Database Server, or All-In-One Appliance.

- If you enter "192.168.38.235", Skylar One will apply the authentication profile to each session on an Administration Portal, Database Server, or All-In-One Appliance where the user enters "192.168.38.235" into the browser.
- If you enter "*.*sciencelogic.local*", Skylar One will apply the authentication profile to each session on an Administration Portal, Database Server, or All-In-One Appliance where the user enters a URL ending with ".*sciencelogic.local*" into the browser.

NOTE: Do not include underscores (_) in the **AP Hostname Pattern** field. URLs with underscores are not considered valid in Skylar One authentication profiles.

- **Available Credential Sources.** This field tells Skylar One how to retrieve the user's credentials from the HTTP request to Skylar One. To align a credential source with the authentication profile, highlight the credential source and click the right-arrow button. You can select zero, one, or multiple credential sources for the authentication profile. Initially, this pane displays a list of all the credential sources:

NOTE: If you will be using CAC authentication, align the CAC/Client Cert credential source. If this is your primary method of logging in to Skylar One, align CAC/Client Cert as the number one credential source. ScienceLogic recommends having EM7 Login Page aligned, as well, for administrator or maintenance access.

- *CAC/Client Cert.* Skylar One will retrieve a certificate from the HTTP request.
- *HTTP Auth.* Skylar One will retrieve a user name and password from the HTTP request.
- *Login Page.* Skylar One will retrieve a user name and password from the ScienceLogic login page fields.

NOTE: If you are using Single Sign-On (SSO) authentication, the **Available Credential Sources** field is ignored. You do not have to align a credential source because credentials are submitted directly to an Identity Provider (IdP) instead of Skylar One.

- **Aligned Credentials Sources.** This field displays the list of credential sources that have been aligned with the authentication profile. The authentication profile will examine each credential source in the order in which it appears in this list. When the authentication profile finds the user's credential, the authentication profile stops examining any remaining credential sources in the list.

- **Available Authentication Resources.** This field tells Skylar One which authentication resources to use to authenticate the retrieved credentials. To align an authentication resource with the authentication profile, highlight the authentication resource and click the right-arrow button. You must select at least one authentication resource (but can select more than one). For details on creating an authentication resource, see the section on [Authentication Resources](#).
- **Aligned Authentication Resources.** This field displays the list of authentication resources that have been aligned with the authentication profile. The authentication profile will examine each authentication resource in the order in which it appears in this list. When an authentication resource successfully authenticates the user, the authentication profile stops executing any remaining authentication resources in the list.

4. Click the **[Save]** button to save your changes to the new authentication profile.

Editing an Authentication Profile

The **Authentication Profiles** page allows you to edit an existing authentication profile. To do so:

1. Go to the **Authentication Profiles** page (System > Settings > Authentication > Profiles).
2. Find the authentication profile that you want to edit. Click its wrench icon (🔧).
3. The **Authentication Profile Editor** modal page appears. In this page, you can edit the value of one or more fields.
4. Click the **[Save]** button to save your changes to the authentication profile.

Deleting One or More Authentication Profiles

The **Authentication Profiles** page allows you to delete one or more authentication profiles from Skylar One. To do so:

1. Go to the **Authentication Profiles** page (System > Settings > Authentication > Profiles).
2. Select the checkbox of each authentication profile that you want to delete.
3. Click the **Select Actions** menu (in the lower right), select *DELETE Authentication Profile*, and then click the **[Go]** button. The selected authentication profiles will be deleted.

NOTE: You cannot delete the *default* authentication profile.

Authentication Resources

An authentication resource is a configuration policy that describes how Skylar One should communicate with a user store. An authentication resource specifies the connector to use to communicate with the user store, the credential to use to connect to the user store (if applicable), and the URLs to examine during authentication. An authentication resource also maps attributes from the user's account in the user store to fields in the ScienceLogic user account.

Viewing the List of Authentication Resources

The **Authentication Resource Manager** page displays a list of all authentication resources in the Skylar One System.

To view the list of authentication resources :

1. Go to the **Authentication Resource Manager** page (System > Settings > Authentication > Resources).
2. The following information is displayed about each authentication resource:

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

- **Resource Name.** Name of the authentication resource.
- **Access.** Indicates whether the authentication resource is shared with all organizations or is private.

NOTE: The **Access** column displays only for Administrator user accounts and user accounts assigned to the System organization.

- **ID.** Unique numeric ID, automatically assigned by Skylar One to each authentication resource.
- **Type.** Specifies the user store that is associated with the resource. Possible types are:
 - *Internal.* The authentication resource communicates and passes information to and from the ScienceLogic Database.
 - *LDAP/AD.* The authentication resource communicates and passes information to and from an LDAP server or Active Directory server.
 - *SSO.* The authentication resource communicates and passes information to and from a SAML Identity Provider (IdP) or Service Provider (SP).
- **Connector.** The software that allows communication between the authentication resource and the user store. Possible connectors are:
 - *Internal.* Software that communicates with the ScienceLogic Database.
 - *LDAP/AD.* Software that communicates with an LDAP server or Active Directory server.

- *LDAP/AD - Legacy*. Software that communicates with an LDAP server or Active Directory server for ScienceLogic servers that were configured prior to version 7.8 of Skylar One. Skylar One systems that were upgraded to version 7.8 using patches can continue to use the same authentication methods without making changes to user accounts or the LDAP server or Active Directory server.
- *OneLogin*. Software that communicates with a SAML Identity Provider (IdP).
- *SimpleSAML - Legacy*. Software that communicates with a SAML Identity Provider (IdP) and Service Provider (SP) for ScienceLogic servers that were configured prior to version 7.8 of Skylar One. Skylar One systems that were upgraded to version 7.8 using patches can continue to use the same authentication methods without making changes to user accounts, the SAML configuration, or the SSO provider.
- **Edited By**. The user who created or last edited the authentication resource.
- **Last Edited**. Date the time the authentication resource was created or last edited.

The "Internal" Resource

The *Internal* resource allows you to access the user store in the ScienceLogic database.

- By default, each Skylar One System includes the *Internal* authentication resource.
- You cannot create an *Internal* authentication resource.
- You cannot edit or delete the *Internal* authentication resource included with your Skylar One System.
- Each Skylar One System can include only one the *Internal* authentication resource.

Creating an LDAP/AD Authentication Resource

The **LDAP/AD Auth Resource Editor** page allows you to define an authentication resource for use with an LDAP/AD user store. An LDAP/AD authentication resource specifies the connector (communication software) to use to communicate with the LDAP/AD user store and the credential to use to connect to the user store. An LDAP/AD authentication resource can also map attributes from the user's LDAP/AD account to fields in the ScienceLogic user account.

ScienceLogic administrators can use LDAP or Active Directory to authenticate ScienceLogic users. There are two ways to use LDAP or Active Directory authentication with Skylar One:

- You can configure Skylar One to automatically create user accounts for existing LDAP or Active Directory users and then always use LDAP or Active Directory to authenticate those users when they log in to Skylar One.
- You can use LDAP or Active Directory to authenticate one or more ScienceLogic users when they log in to Skylar One.

To create an LDAP/AD authentication resource:

1. Go to the **Authentication Resource Manager** page (System > Settings > Authentication > Resources).
2. Click the **[Actions]** menu and then select *Create LDAP/AD Resource*. The **LDAP/AD Auth Resource Editor** modal page appears.
3. Enter values in the following fields:

Basic Settings

- **Name.** Name of the LDAP/AD authentication resource.
- **Sharing Permission.** Indicates if the authentication resource is shared or private. Choices are:
 - *Shared with all organizations.* The authentication resource is shared with users across all organizations.
 - *Private (visible to System organization only).* The authentication resource is private to only user accounts assigned to the System organization.

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

- **Read Credential.** Credential that allows Skylar One to read data from an LDAP or Active Directory server. Select from a list of all LDAP and Active Directory credentials to which you have access. If this field has been set to a credential to which you do not have access, this field will display the value *Restricted Credential*. If you set this field to a different credential, the entry for *Restricted Credential* will be removed from the field; you will not be able to re-align the field with the *Restricted Credential*.
- **Write Credential.** Credential that allows Skylar One to write data to an LDAP or Active Directory server. Select from a list of all LDAP and Active Directory credentials to which you have access. If this field has been set to a credential to which you do not have access, this field will display the value *Restricted Credential*. If you set this field to a different credential, the entry for *Restricted Credential* will be removed from the field; you will not be able to re-align the field with the *Restricted Credential*.
- **User Name Suffix.** Optional field. Because Skylar One can authenticate against multiple LDAP or Active Directory servers, there is a risk of collision among user names. In this field, you can enter a string to append to the user name to minimize the risk of collision. For example:
 - Suppose we entered *@ad.local* in this field.
 - Suppose the next LDAP/AD user logs in to Skylar One with the user name *bishopbrennan*.
 - Skylar One will log that user in as *bishopbrennan@ad.local*.

NOTE: A best practice to avoid collisions is to use email addresses as user names.

- **User Display Name.** Select what name to display from the following options:
 - *disable.* Uses the current default behavior, which displays the user's username in the Skylar One user interface and logs.
 - *email address.* Displays the user's email address in the Skylar One user interface and logs.

- *user principal name*. Displays the value from the UPN field on this page in the Skylar One user interface and logs.
- **UPN**. "User principal name." If you select *user principal name* in the **User Display Name** field, then the value from this field displays in the Skylar One user interface and audit logs. Enter one of the following:
 - *email address*. Displays the user's email address in the Skylar One user interface and audit logs.
 - *user principal name*. Displays the value from the UPN field on this page in the Skylar One user interface and audit logs.

NOTE: Your organization membership(s) might affect the list of credentials you can see in the **Read Credential** field and the **Write Credential** field. For details, see the **Discovery & Credentials** manual.

- **Search Filter**. Specifies where to find the user's account information in LDAP or Active Directory. You must tell Skylar One where to find the LDAP or AD attribute that maps to the user's account name in Skylar One.

For example, an LDAP user might use his/her uid value to log in to Skylar One. In the ScienceLogic account, that uid value will then become the user's **Account Login Name**.

You can use the following variables in the search filter:

- [%u]. ScienceLogic login name.
- %e. Email address.
- An example search filter for LDAP might be:

```
(& (objectClass=person) (uid=%u) )
```

This says to search in the object class called "person" for the uid that matches the ScienceLogic login name (entered when the user logs in to Skylar One and then stored in the variable %u).

- An example search filter for Active Directory might be:

```
(sAMAccountName=%u)
```

This says to search for the samaccountname attribute that matches the ScienceLogic login name (entered when the user logs in to Skylar One and then stored in the variable %u).

- For more information on the syntax of LDAP and AD search filters, see [RFC 4515](#).

- **Sync directory values on login.** If an LDAP or AD administrator makes changes to an LDAP or AD account, Skylar One will automatically retrieve those updates and apply them to the user's account in Skylar One (in the **Account Properties** page) the next time the user logs in to Skylar One. (For more information about user account properties, see the *Organizations & Users* manual.)
- **Sync account values to directory on save.** If a ScienceLogic administrator made changes to the ScienceLogic account, Skylar One will automatically write those changes to the user's account in LDAP or Active Directory.

NOTE: The *Sync account values to directory on save* option requires a write credential.

Attribute Mapping

If you have configured Skylar One to automatically create ScienceLogic accounts for LDAP or AD users, these fields specify the LDAP or AD attribute value that will be automatically inserted into each field in each user's **Account Properties** page. (For more information about user account properties, see the *Organizations & Users* manual.)

Skylar One automatically populates as many of these fields as possible. You can edit or delete the default values provided by Skylar One. For example, Skylar One automatically inserts the value of the LDAP/AD attribute "sn" (surname) into the **Last Name** field in each user's **Account Properties** page.

NOTE: Skylar One requires that the LDAP or AD attribute name that you specify in each field uses **all lower-case characters**.

- **First Name.** Specifies the LDAP or AD attribute value that will be automatically inserted into the **First Name** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the LDAP/AD attribute "givenname" into this field.
- **Last Name.** Specifies the LDAP or AD attribute value that will be automatically inserted into the **Last Name** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the LDAP/AD attribute "sn" into this field.
- **Title.** Specifies the LDAP or AD attribute value that will be automatically inserted into the **Title** field in each user's **Account Properties** page.
- **Department.** Specifies the LDAP or AD attribute value that will be automatically inserted into the **Department** field in each user's **Account Properties** page.
- **Phone.** Specifies the LDAP or AD attribute value that will be automatically inserted into the **Phone** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the LDAP/AD attribute "telephonenumber" into this field.
- **Fax.** Specifies the LDAP or AD attribute value that will be automatically inserted into the **Fax** field in each user's **Account Properties** page.

- **Mobile**. Specifies the LDAP or AD attribute value that will be automatically inserted into the **Mobile** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the LDAP/AD attribute "mobile" into this field.
- **Pager**. Specifies the LDAP or AD attribute value that will be automatically inserted into the **Pager** field in each user's **Account Properties** page.
- **Primary Email**. Specifies the LDAP or AD attribute value that will be automatically inserted into the **Primary Email** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the LDAP/AD attribute "mail" into this field.
- **Secondary Email**. Specifies the LDAP or AD attribute value that will be automatically inserted into the **Secondary Email** field in each user's **Account Properties** page.
- **Street Address**. Specifies the LDAP or AD attribute value that will be automatically inserted into the **Street Address** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the LDAP/AD attribute "streetaddress" into this field.
- **Suite/Building**. Specifies the LDAP or AD attribute value that will be automatically inserted into the **Suite/Building** field in each user's **Account Properties** page.
- **City**. Specifies the LDAP or AD attribute value that will be automatically inserted into the **City** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the LDAP/AD attribute "l" into this field.
- **State**. Specifies the LDAP or AD attribute value that will be automatically inserted into the **State** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the LDAP/AD attribute "st" into this field.
- **Postal Code**. Specifies the LDAP or AD attribute value that will be automatically inserted into the **Postal Code** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the LDAP/AD attribute "postalcode" into this field.
- **Country**. Specifies the LDAP or AD attribute value that will be automatically inserted into the **Country** field in each user's **Account Properties** page.
- **Organization**. Specifies the LDAP or AD attribute value that will be used to automatically define the **Primary Organization** field in each user's **Account Permissions** page. You must also specify one of the following:
 - *directory attribute specifies organization ID*. If selected, the attribute in the **Organization** field specifies an organization ID.
 - *directory attribute specifies organization name*. If selected, the attribute in the **Organization** field specifies an organization name.
 - *directory attribute specifies organization CRM ID*. If selected, the attribute in the **Organization** field specifies the CRM ID of an organization.

NOTE: To use Attribute Mapping for **Organization**, your LDAP/AD schema must include an attribute that maps to ScienceLogic Organization names, Organization IDs, or Organization CRM IDs.

NOTE: When you create a new LDAP/AD user, you must align a user policy with that user. If the aligned user policy specifies an organization for the user, the value from the user policy will overwrite the value from Attribute Mapping.

User Policy Alignment

- **Type.** Specifies whether Skylar One should automatically create ScienceLogic accounts for each LDAP or Active Directory user in the search base (which is specified in the credential), whether Skylar One should simply use LDAP or Active Directory to authenticate one or more users, or whether Skylar One will refuse to authenticate specific users. Choices are:
 - *Do not authenticate new users from directory.* Only those users who have an account already created in Skylar One can log in to Skylar One. However, if one or more users' **Account Permissions** page specifies *LDAP/Active Directory* in the **Authentication Method** field, Skylar One will authenticate those users with either LDAP or Active Directory, using the settings and credentials specified in this page.
 - *Static policy alignment.* If an LDAP or AD user logs in to Skylar One using the LDAP or AD attribute specified in the **Search Filter** field, Skylar One will automatically create an account for that user. Skylar One will use **one user policy** (specified in the **Policy** field) to create all imported LDAP or AD user accounts. Skylar One will also use the settings and credentials specified in this page when creating the account.
 - *Dynamic policy alignment.* If an LDAP or AD user logs in to Skylar One using the LDAP or AD attribute specified in the **Search Filter** field, Skylar One will automatically create an account for that user. Skylar One will **choose from among multiple user policies** to create imported LDAP or AD user accounts. For example, some imported user accounts might use "user policy A"; other imported user accounts might use "user policy B". Skylar One will also use the settings and credentials specified in this page when creating the account.

NOTE: If you selected *Static policy alignment* in the **Type** field, you must supply a value in the **Policy** field:

- **Policy.** Specifies the user policy to use to automatically create a ScienceLogic account for each LDAP or AD user. Select from a list of all user policies that specify *LDAP/Active Directory* in the **Authentication Method** field.

NOTE: If you selected *Dynamic policy alignment* in the **Type** field, you must supply values in the **Attribute**, **Value**, and **Policy** fields.

- **Attribute.** Specifies the LDAP or AD attribute you want to use to differentiate imported user accounts. For example, you could select the attribute "department" and then assign different user policies to import user accounts from different departments. You can also use this field to exclude LDAP or AD accounts for which you do not want to create a ScienceLogic account.
- **Value.** Specifies the LDAP or AD attribute value. That is, you specify one of the possible values for the attribute (specified in the **Attribute** field). Skylar One will compare the value in this field to the retrieved value for the **Attribute**.
- **Policy.** Choose one of the following:
 - *Do Not Authenticate.* If selected, if the retrieved value of the specified **Attribute** matches the value in the **Value** field, the user is not authenticated. This setting applies to new users for whom LDAP or Active Directory would have to create a new account in Skylar One and for users who already have an account in Skylar One.
 - *the policy you want to associate with that value.* Select from a list of all user policies that specify *LDAP /Active Directory* in the **Authentication Method** field.
 - For example, suppose you specified "department" in the **Attribute** field. Suppose that the "department" attribute could have two possible values: "Sales" or "NOC".
 - Suppose you created two user policies. One user policy, called "Sales User Policy", includes the appropriate ticket queues and access keys for Sales personnel. Another user policy, called "NOC User Policy", include the appropriate ticket queues and access keys for NOC personnel.
 - In one of the **Value** fields, you could specify "Sales". In the corresponding **Policy** field, you could then specify "Sales User Policy".
 - In the next **Value** field, you could specify "NOC". In the corresponding **Policy** field, you could specify "NOC User Policy".
 - After defining these two **Value** fields and corresponding **Policy** fields, user accounts from the Sales department would be imported into Skylar One using the Sales User Policy. User accounts from the NOC department would be imported into Skylar One using the NOC User Policy.
- To define additional **Value** and **Policy** fields, click on the plus-sign icon (+).

4. Click the **[Save]** button to save your changes to the new authentication resource.

Creating an SSO Authentication Resource

The **SSO Auth Resource Editor** page allows you to define an authentication resource for use with a SAML IdP. An SSO authentication resource specifies the connector (communication software) to use to communicate with the SAML IdP and the URLs to use to send and retrieve information from the SAML IdP. An SSO authentication resource can also maps attributes from the user's SSO account to fields in the ScienceLogic user account.

ScienceLogic administrators can use SSO to authenticate ScienceLogic users. There are two ways to use SSO authentication with Skylar One:

- You can configure Skylar One to automatically create user accounts for existing SSO users and then always use SSO to authenticate those users when they log in to Skylar One.
- You can use SSO to authenticate one or more ScienceLogic users when they log in to Skylar One.

To create an SSO authentication resource:

1. Go to the **Authentication Resource Manager** page (System > Settings > Authentication > Resources).
2. Click the **[Actions]** menu and then select *Create SSO Resource*. The **SSO Auth Resource Editor** page appears.
3. Enter values in the following fields:

Basic Settings

- **Name.** Name of the SSO authentication resource.
- **IdP Entity ID.** Globally unique name used as a SAML identifier configured on the IdP, usually in the format of an absolute URL.
- **IdP Cert Fingerprint.** The SHA1 certificate fingerprint, provided by the identity provider or service provider. Note that this field is not the serial number of the certificate.

NOTE: If you supply the IdP certificate when you configure the SSO Authentication Resource, the IdP certificate fingerprint is not required and will not be used for IdP response validation. Instead, the full certificate that you provide in the **IdP Certificate** field will be used.

- **IdP Certificate.** To ensure that communication between the IdP and Skylar One is signed, type the full, PEM-encoded certificate from the IdP.
- **Sharing Permission.** Indicates if the authentication resource is shared or private. Choices are:
 - *Shared with all organizations.* The authentication resource is shared with users across all organizations.
 - *Private (visible to System organization only).* The authentication resource is private to only user accounts assigned to the System organization.

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

- **User Name Suffix.** Optional field. If you don't supply a value in this field, Skylar One retrieves the SAML **NameID** attribute and uses that value as the ScienceLogic username.
 - You can supply the variable **%u** in this field, and the Skylar One retrieves the SAML **NameID** attribute and uses that value as the ScienceLogic user name.
 - You can supply the value **%attribute_name%**, where attribute name is a SAML attribute other than **NameID**. Skylar One will use the value of the attribute as the ScienceLogic user name.

- Because a user can authenticate against multiple SSO servers, there is a risk of collision among user names. In this field, you can enter a string to append to the ScienceLogic user name to minimize risk of collision. For example:
 - You can enter a string, with no SAML attribute specified. When you don't specify a SAML attribute in this field, Skylar One will retrieve the SAML **NameID** attribute and append the string you specify in this field.

Suppose we entered `@scienelogic.local` in this field.

Suppose the next SSO user logs in to Skylar One with the SAML **NameID** of `bishopbrennan`.

Skylar One will log in that user as `bishopbrennan@scienelogic.local`.

- You can enter one or more SAML attribute names, surrounded by percent signs (%), with text preceding it and/or text appended. Skylar One will retrieve the value of the SAML attribute and use that value plus any preceding text or appended text as the the ScienceLogic user name.

Suppose we entered `%sn%-external` in this field.

Suppose the next SSO user logs in to Skylar One with their SAML **sn** (last name) attribute of `krilly`

Skylar One will log in that user as `krilly-external`.

NOTE: A best practice to avoid collisions is to use email addresses as user names.

- IdP SSO URL.** The URL to which Skylar One will send login requests to the IdP. This field must contain an absolute URL.
- IdP SLS URL.** Optional field. If you want each user to be automatically logged out of the IdP when that user logs out of Skylar One, enter the URL to which Skylar One will post the logout request to the IdP. If you leave this field blank, a user can log out of Skylar One without automatically logging out of the IdP.
- Sync directory values on login.** If an SSO administrator makes changes to an SSO account, Skylar One will automatically retrieve those updates and apply them to the user's account in the **Account Properties** page the next time the user logs in to Skylar One. (For more information about user account properties, see the *Organizations & Users* manual.)
- Signing Options.** Specifies whether digital signing is required for communication between the IdP and Skylar One. Choices are:
 - Disable.** No digital signature is required.
 - IdP Response.** Messages from the IDP to Skylar One must be signed. Skylar One will use the value in the **IdP Certificate** field to validate the signature.

- *SP Request and IdP Response.* Messages from the IdP to Skylar One must be signed. Skylar One will use the value in the **IdP Certificate** field to validate the signature. Messages from Skylar One to the IdP must also be signed.
- **Strict Mode.** If you selected *IdP Response* or *SP Request and IdP Response* in the Signing Options field, this field is automatically set to *enable*. This field enforces validation of the SAML response and its attributes. As a best practice, disable this field while initially configuring Skylar One and the IdP. As a best practice, enable this field for production use.
- **Integrated Windows Auth.** If you are using Active Directory Federation Services (ADFS) as your IdP, select *Enable* in this field.

Attribute Mapping

If you have configured Skylar One to automatically create ScienceLogic accounts for SSO users, these fields specify the SAML attribute value that will be automatically inserted into each field in each user's **Account Properties** page. (For more information about user account properties, see the *Organizations & Users* manual.)

Skylar One automatically populates as many of these fields as possible. You can edit or delete the default values provided by Skylar One. For example, Skylar One automatically inserts the value of the SAML attribute "sn" (surname) into the **Last Name** field in each user's **Account Properties** page.

NOTE: Skylar One requires that the SAML attribute name that you specify in each field uses all lowercase characters.

- **First Name.** Specifies the SAML attribute value that will be automatically inserted into the **First Name** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the SAML attribute "givenname" into this field.
- **Last Name.** Specifies the SAML attribute value that will be automatically inserted into the **Last Name** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the SAML attribute "sn" into this field.
- **Title.** Specifies the SAML attribute value that will be automatically inserted into the **Title** field in each user's **Account Properties** page.
- **Department.** Specifies the SAML attribute value that will be automatically inserted into the **Department** field in each user's **Account Properties** page.
- **Phone.** Specifies the SAML attribute value that will be automatically inserted into the **Phone** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the SAML attribute "telephonenumber" into this field.
- **Fax.** Specifies the SAML attribute value that will be automatically inserted into the **Fax** field in each user's **Account Properties** page.
- **Mobile.** Specifies the SAML attribute value that will be automatically inserted into the **Mobile** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the SAML attribute "mobile" into this field.
- **Pager.** Specifies the SAML attribute value that will be automatically inserted into the **Pager** field in each user's **Account Properties** page.

- **Primary Email.** Specifies the SAML attribute value that will be automatically inserted into the **Primary Email** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the SAML attribute "mail" into this field.
- **Secondary Email.** Specifies the SAML attribute value that will be automatically inserted into the **Secondary Email** field in each user's **Account Properties** page.
- **Street Address.** Specifies the SAML attribute value that will be automatically inserted into the **Street Address** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the SAML attribute "streetaddress" into this field.
- **Suite/Building.** Specifies the SAML attribute value that will be automatically inserted into the **Suite/Building** field in each user's **Account Properties** page.
- **City.** Specifies the SAML attribute value that will be automatically inserted into the **City** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the SAML attribute "l" into this field.
- **State.** Specifies the SAML attribute value that will be automatically inserted into the **State** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the SAML attribute "st" into this field.
- **Postal Code.** Specifies the SAML attribute value that will be automatically inserted into the **Postal Code** field in each user's **Account Properties** page. By default, Skylar One inserts the value of the SAML attribute "postalcode" into this field.
- **Country.** Specifies the SAML attribute value that will be automatically inserted into the **Country** field in each user's **Account Properties** page.
- **Organization.** Specifies the SAML attribute value that will be used to automatically define the **Primary Organization** field in each user's **Account Permissions** page. You must also specify one of the following:
 - *directory attribute specifies organization ID.* The attribute in the **Organization** field specifies an organization ID.
 - *directory attribute specifies organization name.* The attribute in the **Organization** field specifies an organization name.
 - *directory attribute specifies organization CRM ID.* The attribute in the **Organization** field specifies the CRM ID of an organization.

NOTE: To use Attribute Mapping for **Organization**, your SAML schema must include an attribute that maps to All-In-One Appliance Organization names, Organization IDs, or Organization CRM IDs.

NOTE: When you create a new SSO user, you must align a user policy with that user. If the aligned user policy specifies an organization for the user, the value from the user policy will overwrite the value from Attribute Mapping.

User Policy Alignment

- **Type**. Specifies whether Skylar One should automatically create ScienceLogic accounts for each SSO user, whether Skylar One should simply use SSO to authenticate one or more users, or whether Skylar One will refuse to authenticate specific users. Choices are:
 - *Do not authenticate new users*. Only those users who have an account already created in Skylar One can log in to Skylar One, which will authenticate those users with SSO using the settings specified in this page.
 - *Static policy alignment*. If an SSO user tries to access Skylar One, Skylar One will automatically create an account for that user. Skylar One will use **one user policy** (specified in the **Policy** field) to create the imported SSO user accounts for this authentication resource. Skylar One will also use the settings specified in this page when creating the account.
 - *Dynamic policy alignment*. If an SSO user tries to access Skylar One, Skylar One will automatically create an account for that user. Skylar One will choose from among **multiple user policies** to create imported SSO user accounts for this authentication resource. For example, some imported user accounts might use "user policy A"; other imported user accounts might use "user policy B". Skylar One will also use the settings specified in this page when creating the account.

NOTE: If you selected *Static policy alignment* in the **Type** field, you must supply a value in the **Policy** field.

- **Policy**. Specifies the user policy to use to automatically create a ScienceLogic account for each SSO user. Select from a list of all user policies.

NOTE: If you selected *Dynamic policy alignment* in the **Type** field, you must supply values in the **Attribute**, **Value**, and **Policy** fields.

- **Attribute**. Specifies the SAML attribute you want to use to differentiate imported user accounts. For example, you could select the attribute *department* and then assign different user policies to import user accounts from different departments. You can also use this field to exclude SSO accounts for which you **do not want to allow authentication**.
- **Value**. Specifies the SAML attribute value. That is, you specify one of the possible values for the attribute (specified in the **Attribute** field). Skylar One will compare the value in this field to the retrieved value for the **Attribute**.
- **Policy**. Choose one of the following:
 - *Do Not Authenticate*. If the retrieved value of the specified **Attribute** matches the value in the **Value** field, the user is not authenticated. This setting applies to new users for whom SSO would have to create a new account in Skylar One and for users who already have an account in Skylar One.

- the policy you want to associate with that value. Select from a list of all user policies that specify SSO in the **Authentication Method** field.
 - For example, suppose you specified **department** in the **Attribute** field. Suppose that the **department** attribute could have two possible values: *Sales* or *NOC*.
 - Suppose you created two user policies. One user policy, called *Sales User Policy*, includes the appropriate ticket queues and access keys for Sales personnel. Another user policy, called *NOC User Policy*, includes the appropriate ticket queues and access keys for NOC personnel.
 - In one of the **Value** fields, you could specify *Sales*. In the corresponding **Policy** field, you could then specify *Sales User Policy*.
 - You could then click on the plus-sign icon (+) and add another **Value** field and another **Policy** field.
 - In the next **Value** field, you could specify *NOC*. In the corresponding **Policy** field, you could specify *NOC User Policy*.
 - After defining these two **Value** fields and the corresponding **Policy** fields, user accounts from the *Sales* department would be imported into Skylar One using the *Sales User Policy*.
 - User accounts from the *NOC* department would be imported into Skylar One using the *NOC User Policy*.
 - To define additional **Value** and **Policy** fields, click on the plus-sign icon (+).

4. Click the **[Save]** button to save your changes to the new authentication resource.

Editing an Authentication Resource

The **Authentication Resource Manager** page allows you to edit an existing authentication resource. To do so:

1. Go to the **Authentication Resource Manager** page (System > Settings > Authentication > Resources).
2. Find the authentication resource that you want to edit. Click its wrench icon (🔧).
 - For LDAP/AD Resources, the **LDAP/AD Auth Resource Editor** page appears. In this page, you can edit the values for one or more fields. For more information, see the [Creating an LDAP/AD Authentication Resource](#) section.
 - For SSO Resources, SSO Auth Resource Editor page appears. In this page, you can edit the values for one or more fields. For more information, see the [Creating an SSO Authentication Resource](#) section.
3. Click the **[Save]** button to save your changes to the authentication resource.

Deleting an Authentication Resource

The **Authentication Resource Manager** page allows you to delete one or more authentication resources from Skylar One. To do so:

1. Go to the **Authentication Resource Manager** page (System > Settings > Authentication > Resources).
2. Select the checkbox of each authentication resource that you want to delete.
3. Click the **Select Actions** menu (in the lower right), select *DELETE Authentication Resource*, and then click the **[Go]** button. The selected authentication resources will be deleted.

NOTE: You cannot delete the *Internal* authentication resource.

Chapter

17

Installing a TLS (SSL) Certificate

Overview

This chapter describes how to acquire and install a TLS (SSL) certificate for your Skylar One (formerly SL1) system.

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>Using TLS Certificates</i>	379
<i>Certificates for ScienceLogic Servers</i>	379
<i>Requesting a Commercial TLS Certificate</i>	379
<i>Creating Your Own Certificate</i>	381
<i>Installing the Certificate on a Skylar One Appliance</i>	383

Using TLS Certificates

TLS is an acronym for Transport Layer Security, a cryptographic protocol that provides security for network communications to ensure privacy, integrity, and authenticity. It is a more secure and modern successor to the now-deprecated Secure Sockets Layer (**SSL**) protocol.

Both protocols encrypt data and authenticate users to facilitate the secure exchange of data. Usually, URLs that include "HTTPS" are using TLS for security.

To implement TLS, a certificate resides on the web server and is used to encrypt the data and to identify the website. The certificate contains information about the certificate holder, the domain for which the certificate was issued, the name of the Certificate Authority who issued the certificate, and the root and the country in which the certificate was issued.

There are two ways to acquire a TLS certificate:

- You can acquire a certificate from a vendor (called a "certificate authority"), such as VeriSign or GeoTrust.
- You can "self-sign" your own certificate. Using available tools (both open source and proprietary), you can create and sign your own SSL certificate instead of purchasing one from a certificate authority.

Skylar One includes a self-signed certificate from ScienceLogic. Self-signed certificates can trigger a warning message in some browsers. For this reason, some customers might prefer to purchase a TLS certificate from a certificate authority and install the certificate on one or more servers.

Certificates for ScienceLogic Servers

Each Skylar One appliance includes a self-signed certificate from ScienceLogic.

Each Skylar One appliance uses the Nginx web server and OpenSSL.

If you want to use commercial TLS certificates with Skylar One, you must acquire certificates for the following Skylar One appliances:

- For each Administration Portal, Database Server, or All-In-One Appliance you must acquire at least one certificate for the standard user interface and the Configuration Utility.
- For each Data Collector, you must acquire one certificate, for use with the Configuration Utility.
- For each Message Collector, you must acquire one certificate, for use with the Configuration Utility.

Requesting a Commercial TLS Certificate

To acquire a commercial TLS certificate, you must first create a private key and then use the private key to create a Certificate Signing Request (CSR). You must then send the CSR to a Certificate Authority (CA). Some well-known CAs are VeriSign, GeoTrust, Thawte, GoDaddy, and Comodo. The CA will charge you a fee and send you a certificate for use with your private key.

To create a CSR, perform the following on each Skylar One appliance.

1. Either go to the console of the Skylar One appliance or use SSH to access the server. Open a shell session on the Skylar One appliance. Log in as "em7admin". Then, do one of the following:
 - If you are running a Security Technical Implementation Guide (STIG) deployment of Skylar One, complete steps 2 and 3.
 - If you are running a non-STIG deployment of Skylar One, complete steps 4 and 5.
2. To generate a private key for the server in your STIG deployment of Skylar One, enter the following at the shell prompt:

```
sudo openssl genrsa -out <keyname>_unencrypted.key 4096
```

where:

- <keyname> is a name for the private key. For example, you might want to name the private key for an administration portal *adminport.key*.

NOTE: Make sure the file is **not** named *silossi.key*, which is the name of the pre-existing ScienceLogic, self-signed certificate file.

3. Encrypt this generated key by entering the following:

```
sudo openssl pkcs8 -topk8 -v2 aes256 -in <keyname>_unencrypted.key -  
out <keyname>.key
```

After encrypting the key, proceed to step 6.

4. To generate a private key for the server in your non-STIG deployment of Skylar One, enter the following at the shell prompt:

```
sudo openssl genrsa -aes256 -out <keyname>.key 4096
```

where:

- <keyname> is a name for the private key. For example, you might want to name the private key for an administration portal *adminport.key*.

NOTE: Make sure the file is **not** named *silossi.key*, which is the name of the pre-existing ScienceLogic, self-signed certificate file.

5. Enter a passphrase for the key when prompted.

TIP: A best practice is to make a backup copy of the key file and the passphrase and store both in a secure location.

6. Remove the passphrase from the key before generating a Certificate Signing Request (CSR). To do this, enter the following command at the shell prompt, inserting the keyname you used where indicated:

```
sudo openssl rsa -in <keyname>.key -out <keyname>.key.insecure
```

7. Create a Certificate Signing Request (CSR) for the private key you created in the previous steps. To do this, enter the following command at the shell prompt:

```
sudo openssl req -new -key <keyname>.key.insecure -out <keyname>.csr
```

where:

- `<keyname>` is a name for the CSR for the specific server. For example, you might want to name the private key for an administration portal `adminimport.key` and name the CSR for that key `adminimport.csr`.

NOTE: Make sure the keyname is **not** `silossi.key`. This is the name of the pre-existing ScienceLogic, self-signed certificate file.

8. Enter the demographic information for your key.
 - Enter a two-letter Country Name (for example, US).
 - Enter your State or Province full name (for example, Virginia).
 - Enter your Locality Name or city (for example, Reston).
 - Enter your Organization Name or company (for example, ScienceLogic).
 - Enter the Common Name, that is, your server's hostname (for example, `myhost.scienceologic.com`).
 - Enter your Email Address. This is where you want communication from the Certificate Authority to be sent.
9. Send the .csr file you generated to a Certificate Authority. The Certificate Authority will provide details on how to send the .csr file. The Certificate Authority will then send you a .crt file. The .crt file is the public key that matches your private key for the Skylar One appliance. Some Certificate Authorities, e.g. GoDaddy, might use an intermediate certificate to sign the provided certificate. If an intermediate certificate is used, the Certificate Authority will provide a bundle of chained certificates in a second .crt file.

Creating Your Own Certificate

There are two reasons you might create your own TLS certificate:

- If your organization is a root Certificate Authority (for example, some departments of the United States government), you can create your own private key and public key for each ScienceLogic server.
- If your security requirements permit a self-signed certificate, you can create your own private key and public key for each Skylar One appliance.

NOTE: Remember to create key pairs for all for each Skylar One appliance in your Skylar One system, and also remember to create two key pairs for each Administration Portal in your Skylar One system.

NOTE: If your organization is a Certificate Authority, see your organization's internal documentation on creating a certificate for NGINX.

To create a self-signed certificate:

1. Either go to the console of the Skylar One appliance or use SSH to open a shell session on the Skylar One appliance.
2. Log in as an administrator (such as em7admin).
3. Generate a private key for the server. To do this, enter the following at the shell prompt:

```
sudo openssl genrsa -aes256 -out <keyname>.key 4096
```

where `<keyname>` is a name for the private key. For example, you might want to name the private key for an Administration Portal `adminport.key`.

CAUTION: Make sure the file is **not** named `silossl.key`. This is the name of the pre-existing ScienceLogic self-signed certificate file.

4. Enter a passphrase for the key when prompted.

TIP: A best practice is to make a backup copy of the key file and the passphrase and store both in a secure location.

5. Remove the passphrase from the key before you continue. To do this, enter the following command at the shell prompt, inserting the keyname you used where indicated:

```
sudo openssl rsa -in <keyname>.key -out <keyname>.key.insecure
```

6. Create a self-signed certificate based on the private key you generated in the previous steps. To do this, enter the following at the shell prompt:

```
sudo openssl req -new -x509 -nodes -sha1 -days 365 -key <keyname>.key  
-out <keyname>.crt
```

where:

- `<keyname>.key` is the private key for the Skylar One appliance .
- `<keyname>.crt` is the public key (certificate) for the Skylar One appliance.

For example, you might want to name the private key for an Administration Portal `adminport.key`, and name the certificate file for that key `adminport.crt`. The resulting `.crt` file is the public key that matches your private key for the Skylar One appliance.

CAUTION: Make sure the files are **not** named `siloss1.crt` and `siloss1.key`. These are the names of the pre-existing ScienceLogic self-signed certificate files.

7. Copy your private key and certificate files to `/etc/nginx`.
8. **On Collectors.** Add the private key and certificate file to each Collector for the Configuration Utility. To do this, add the names of the new `.key` and `.crt` files to the following files:

```
/etc/nginx/conf.d/em7ngx_web_ui.conf
```

```
/etc/nginx/conf.d/em7ngx_em7proxy_web_ui.conf
```

9. **On the Administration Portal, Database Server, or All-in-One Appliance.** Add the private key and certificate file for the user interface. To do this, add the names of the new `.key` and `.crt` files to the following files:

```
/etc/nginx/conf.d/em7ngx_web_ui.conf
```

```
/etc/nginx/conf.d/em7ngx_em7proxy_web_ui.conf
```

10. Restart the Web Configuration Utility and web server by entering the following command:

```
sudo systemctl restart nginx
```

Installing the Certificate on a Skylar One Appliance

ScienceLogic does not provide support for third party certificates. Be advised that installing a new TLS certificate can affect the operation of TLS services.

Most certificate authorities provide support and resources on installing and enabling their certificates in Nginx web servers. If you have questions, please refer to your Certificate Authority.

WARNING: The following steps will stop and restart the Skylar One appliance and temporarily make the Administration Portal site unavailable. Confirm with your System Administrator that you are permitted to restart the ScienceLogic Web Service.

NOTE: These instructions assume that you are familiar with the Linux shell and the "vi" editor.

To install a commercial TLS certificate on a Skylar One appliance, perform the following:

1. Acquire a certificate from a certificate authority.
2. Copy the certificate files (*.key and all *.crt files) to a server that can access the Skylar One appliance via SFTP.

NOTE: Make sure the files are **not** named *silossl.crt* and *silossl.key*. These are the names of the pre-existing ScienceLogic self-signed certificate files.

3. Use SFTP or SCP to copy the .crt file(s) and the .key file to the Skylar One appliance in the `/etc/nginx` directory.
4. Either go to the console of the Skylar One appliance or use SSH to access the server. Open a shell session on the Skylar One appliance. Log in as "em7admin".
5. If an intermediate certificate has been used to sign the certificate file, execute the following commands to combine the server certificate and the bundle of chained certificates provided by the Certificate Authority, entering the server certificate name, bundle name, and combined certificate name where indicated:

```
cd /etc/nginx
```

```
cat <server certificate name>.crt <bundle name>.crt > <combined  
certificate name>.crt
```

Use the combined .crt file name when updating the nginx configuration.

6. For each appliance, edit the following files to configure the certificate for the Configuration Utility:
 - `/etc/nginx/conf.d/em7webconfig.conf`
 - `/etc/nginx/conf.d/em7_sladmin.conf`
 - Edit the following lines, removing references to *silossl.crt* and *silossl.key* and replacing them with the names of the new .key and .crt files:

```
ssl_certificate /etc/nginx/<name of .crt file>;
```

```
ssl_certificate_key /etc/nginx/<name of .key file>;
```

7. In addition, for each Administration Portal, Database Server, and All-In-One Appliance, you must also edit the following files to configure the certificate for the user interface:

- `/etc/nginx/conf.d/em7nginx_web_ui.conf`
- `/etc/nginx/conf.d/em7nginx_em7proxy_web_ui.conf`
- Edit the following lines, removing references to *silossl.pem* and *silossl.key* and replacing them with the names of the new key files:

```
ssl_certificate /etc/nginx/<name of .crt file>;
```

```
ssl_certificate_key /etc/nginx/<name of .key file>;
```

8. Next, you will need to restart the webconfig and webserver. To do this, execute the following command:

- For all appliances, enter:

```
sudo systemctl restart nginx
```

9. To test the certificate, open a browser session and connect to the Administration Portal, Database Server, or All-In-One Appliance using HTTPS.

- From the Administration Portal, go to the **Appliance Manager** page (System > Settings > Appliances).
- Select the toolbox icon (🔧) for each server. Notice that the URL for the Configuration Utility includes https.

Chapter

18

Managing Host Files

Overview

The **Host File Entry Manager** page allows you to edit and manage host files for all of the Data Collectors from a single page in the Skylar One (formerly SL1) system. When you create or edit an entry in the **Host File Entry Manager** page, Skylar One automatically sends an update to every Data Collector in the specified Collector Group.

The **Host File Entry Manager** page is helpful when:

- The Skylar One system does not reside in the end-customer's domain
- The Skylar One system does not have line-of-sight to an end-customer's DNS service
- A customer's DNS service cannot resolve a host name for a device that the Skylar One system monitors

You can create host file entries for each device managed by the Skylar One system. You can create duplicate host file entries, one for each Collection Group, to ensure that all Collection Groups can resolve all host names for monitored devices.

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:

Viewing the List of Host Entries	387
Creating a New Host Entry	387
Editing a Host Entry	388

Using an Existing Host File Entry to Create a New Host File Entry (Save As)	388
Deleting One or More Host Entries	389

Viewing the List of Host Entries

To view the list of host entries, perform the following steps:

1. Go to the **Host File Entry Manager** page (System > Customize > Host Files).
2. The **Host File Entry Manager** page displays the following about each host entry:
 - **IP Address**. The IP address to resolve with the host name.
 - **Hostnames and Aliases**. The host name to align with the specified IP address. You can also include a space-delimited list of aliases for the host name.
 - **Description**. Description of the host entry.
 - **Organization**. Organization associated with the host.
 - **CUG**. The Collector Group to which Skylar One will send the host entry. The host entry will be added to the host file on each Data Collection Server in the Collector Group.
 - **Edited By**. User who created or last edited the host entry.
 - **Last Edit**. Date the host entry was created or last edited.

Creating a New Host Entry

To create a host file entry:

1. Go to the **Host File Entry Manager** page (System > Customize > Host Files).
2. Click the **[Action]** menu and choose **Create New Entry**. The **Create New Host File Entry** modal page appears.
3. In the **Create New Host File Entry** modal page, supply values in the following fields:
 - **IP Address**. The IP address to resolve with the hostname.

NOTE: Server hostnames should be aligned to external IP addresses when supporting Network Address Translation (NAT) environments.

- **Hostnames and Aliases**. The hostname to align with the specified IP address. You can also include a space-delimited list of aliases for the host name.
- **Description**. Description of the host entry. This field is not written to the host file. This field is for administrators to use when managing host file entries.

- **Organization.** Organization associated with the host. You can select from a list of all existing organizations. This field is not written to the host file. This field is for administrators to use when managing host file entries. For example, a service provider could assign each customer its own organization and then use this field to manage host file entries for each customer.

4. Click the **[Save]** button to save the new host entry.

Editing a Host Entry

To edit a host entry, perform the following steps:

1. Go to the **Host File Entry Manager** page (System > Customize > Host Files).
2. Click the wrench icon (🔧) for the host file entry you want to edit. The **Editing Host File Entry** modal page appears, populated with values from the selected host file entry.
3. In the **Editing Host File Entry** modal page, you can edit one or more of the following fields:
 - **IP Address.** The IP address to resolve with the host name.
 - **Hostnames and Aliases.** The host name to align with the specified IP address. You can also include a space-delimited list of aliases for the host name.
 - **Description.** Description of the host entry. This field is not written to the host file. This field is for administrators to use when managing host file entries.
 - **Organization.** Organization associated with the host. You can select from a list of all existing organizations. This field is not written to the host file. This field is for administrators to use when managing host file entries. For example, a service provider could assign each customer its own organization and then use this field to manage host file entries for each customer.
4. Click the **[Save]** button to save your changes.

Using an Existing Host File Entry to Create a New Host File Entry (Save As)

To create a new host entry, using an existing host entry as the template:

1. Go to the **Host File Entry Manager** page (System > Customize > Host Files).
2. Click the wrench icon (🔧) for the host file entry you want to edit. The **Editing Host File Entry** modal page appears, populated with values from the selected host file entry.
3. In the **Editing Host File Entry** modal page, you can edit one or more of the following fields:
 - **IP Address.** The IP address to resolve with the host name.
 - **Hostnames and Aliases.** The host name to align with the specified IP address. You can also include a space-delimited list of aliases for the host name.
 - **Description.** Description of the host entry. This field is not written to the host file. This field is for administrators to use when managing host file entries.

- **Organization.** Organization associated with the host. You can select from a list of all existing organizations. This field is not written to the host file. This field is for administrators to use when managing host file entries. For example, a service provider could assign each customer its own organization and then use this field to manage host file entries for each customer.

4. Click the **[Save As]** button to save your changes as a new host file entry. A pop-up message appears, asking if you want to save your edits as a new entry. Click the **[OK]** button.

Deleting One or More Host Entries

To delete one or more host entries, perform the following steps:

1. Go to the **Host File Entry Manager** page (System > Customize > Host Files).
2. Select the checkbox for each host file entry you want to delete.
3. Click the **Select Action** field in the lower right, then select *DELETE Host File Entry*. Click the **[Go]** button.

© 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