



System Administration

SL1 version 10.1.4

Table of Contents

Introduction	8
Who Should Read This Manual?	8
What's In This Manual?	8
Requirements	9
Global Settings	10
Global Settings for API	11
Global Settings for Appliances	12
The Web Configuration Utility	14
Global Settings for Asset Automation	15
Global Settings for User Logins, Discovery, Data Collection, UI Features, and Expiration Warnings	16
Global Settings for Data Retention	29
Normalization and Roll-Up of Performance Data	32
Collection of Raw Data	32
Data Normalization and Rollup	34
Example	36
Storage of Raw and Rolled Up Data	36
Global Settings for Inbound Email and Outbound Email	36
Global Settings for Login Alert Messages	40
Global Settings for Password Reset Emails	42
Defining the Email Message for "I forgot my password"	42
Global Settings for System Thresholds	45
Global Settings for Interface Thresholds	49
Settings in Silo.Conf	57
Disabling the User Interface on a Database Server	67
Collector Groups	69
Installing, Configuring, and Licensing Data Collectors	70
Technical Information About Data Collectors	71
Duplicate IP Addresses	71
Open Ports	71
Viewing the List of Collector Groups	71
Creating a Collector Group	72
Editing a Collector Group	75
Collector Groups and Load Balancing	76
Tuning Collector Groups in the silo.conf File	77
Collector Affinity	79
Failover for Collector Groups for Component Devices	80
Collector Groups for Merged Devices	80
Creating a Collector Group for Data Storage Only	81
Deleting a Collector Group	82
Aligning the Collector Group for A Single Device	83
Aligning the Collector Group in a Device Template	84
Changing the Collector Group for One or More Devices	85
Managing the Host Files for a Collector Group	85
Processes for Collector Groups	86
Enabling and Disabling Concurrent PowerShell for Collector Groups	87
Enabling Concurrent PowerShell on All Collector Groups	87
Disabling Concurrent PowerShell on All Collector Groups	87
Enabling Concurrent PowerShell on a Specific Collector Group	88
Disabling Concurrent PowerShell on a Specific Collector Group	88
Daily Health Tasks	89

What is a Healthy SL1 System?	90
Monitoring System Events	93
Searching the System Logs	94
Deleting Entries from the System Logs	95
Monitoring System Processes	96
Viewing the List of System Processes	96
Searching and Filtering the List of System Processes	97
Monitoring the Status of Each Appliance	98
Monitoring User Actions and Events	99
Viewing the List of Audit Logs	100
Searching and Filtering the List of Audit Logs	101
Special Characters	102
Generating Reports on Audit Logs	105
Monitoring the Status of Data Collectors	106
Upgrading SL1	108
Reference: The System Updates Page	112
Reference: Viewing the List of Updates	112
Reference: Viewing the Log Files for Updates	114
Workflow	114
Planning	115
Scheduling Maintenance Windows	116
Updating SL1 Appliances to Oracle Linux	117
SL1 Release Prior to 8.1.1	117
SL1 Releases Prior to 8.10.0	117
Pre-Upgrade Best Practices	117
Backing Up Settings in the NextUI	118
Backing Up SSL Certificates	118
Setting the Timeout for PhoneHome Watchdog	118
Adjusting the Timeout for Staging and Deploying	119
SL1 8.14 and Later Releases	119
SL1 8.12 and Prior Releases	119
Running the System Status Script Before Upgrading	120
Running the System Status Script on SL1 8.14.0 and Later Releases	120
Running the System Status Script on SL1 8.12.x Releases	120
Running the System Status Script on SL1 8.10 and Prior Releases	121
Upgrading the SL1 Distributed Architecture on SL1 Versions 8.5.0 and Earlier	122
Disabling Automatic Staging	123
Downloading the Updates	124
Importing the Updates	124
Staging the Update	125
Running the Pre-Upgrade Check	127
All SL1 Appliances:	127
Active Database Server:	127
All Database Servers:	128
Administration Portal:	128
Data Collectors and Message Collectors:	128
Downloading and Running the Pre-Upgrade Check	128
Putting All SL1 Appliances in Maintenance Mode	129
Deploying the Update	130
Putting All SL1 Appliances Out of Maintenance Mode	131
Performing Deltaless Upgrades	132
Upgrading the SL1 Distributed Architecture on SL1 versions 8.6.0 and Later	132

Special Steps for SL1 8.12.0 and Earlier	133
Downloading the Update	134
Importing the Update	135
Staging the Update	136
Automatic Staging	136
Manually Staging an Update	137
Running the Pre-Upgrade Check for SL1 10.1 and Later	139
Running the Pre-Upgrade Check	139
CentOS 5 Failure	140
Collector Group Membership	140
Eligibility Failure	140
Enabled Failure	140
Free Disk-Space Failure	140
Host File Failure	141
Patch-Hook Ownership Failure	141
RPM Database Failure	141
RPM Package Failure	141
Running the Pre-Upgrade Check for SL1 8.14 and Earlier	142
All SL1 Appliances:	142
Active Database Server:	143
All Database Servers:	143
Administration Portal:	143
Data Collectors and Message Collectors:	144
Downloading and Running the Pre-Upgrade Check	144
Putting All SL1 Appliances in Maintenance Mode	145
Deploying the Update	145
Putting All SL1 Appliances Out of Maintenance Mode	147
Upgrading the Extended Architecture	148
Prerequisites	148
Resizing the Disks on the Compute Node	148
Upgrade Steps	149
Manual Steps for Updates to 8.4.x and Earlier Systems	151
Automatically Upgrading MariaDB with a Script	152
Manually Upgrading MariaDB	154
Download RPMs to SL1 Appliances	154
Manually Upgrade Two Database Servers Configured for High Availability or Disaster Recovery	155
Step 1: On the Secondary Database Server	155
Step 2: On the Primary Database Server	156
Step 3: On the Secondary Database Server	158
Manually Upgrade Three Database Servers Configured for High Availability and Disaster Recovery	159
Step 1: On the Secondary Database Server	159
Step 2: On the Primary Database Server	159
Step 3: On the Secondary Database Server	160
Step 4: On the Disaster Recovery Database Server	161
Manually Upgrading Standalone Database Servers, All-In-One Appliances, Data Collectors, and Message Collectors	162
Rebooting Appliances in the SL1 Distributed Stack	163
Rebooting the Administration Portal	163
Rebooting Multiple Administration Portals	163
Rebooting a Single Administration Portal	164
Rebooting Data Collectors and Message Collectors	164
Rebooting Data Collectors and Message Collectors from the Appliance Manager page	164

Rebooting Data Collectors and Message Collectors from the Command Line	165
Rebooting Standalone All-In-One Appliance and Standalone Database Server	165
Rebooting Two Database Servers Configured for Disaster Recovery	165
Rebooting Two Database Servers in a High Availability Cluster	166
Rebooting Three Database Servers Configured for High Availability and Disaster Recovery	167
Restoring Settings for NextUI	169
Restoring the SSL Certificate	169
Resetting the Timeout for PhoneHome Watchdog	169
Updating Default PowerPacks	170
Configuring Subscription Billing	172
Monitoring and Maintaining SL1	173
Monitoring and Managing User Access	174
Viewing Information about Each Access Session	174
Deleting a User's Session	175
Viewing Lockouts and Unlocking Lockouts	176
Global Settings for Lockouts	177
Audit Logs	177
Managing Scheduled Tasks	177
Viewing the List of Schedules	178
Enabling or Disabling One or More Schedules	179
Deleting One or More Schedules	180
Monitoring Overall System Usage and Statistics	180
Viewing an Overview of All Events	181
Viewing Events by Appliance and Event Source	183
Diagnostic Tools	187
ScienceLogic SL1 Self-Monitoring	188
Viewing Information About ScienceLogic Processes	188
Viewing the List of ScienceLogic Processes	189
Searching and Filtering the List of ScienceLogic Processes	191
Editing the Parameters of a ScienceLogic Process	192
Debugging a Process and Viewing Debug Logs	194
Viewing Information About Unhandled Exceptions	196
Viewing the List of Unhandled Exceptions	196
Searching and Filtering the list of Unhandled Exceptions	197
Saving the Unhandled Exception to the Local Computer	198
Viewing the Output of the System Status Script	198
Viewing the Database Tables on the Database Server	199
Accessing the Database Tool	199
Disabling Normalization, Re-Enabling Normalization, and Backfilling Raw Data	201
Enable Logging for Data Pull Storage Objects	204
Enable	204
Disable	204
Controlling Log Settings	205
Setting UI Developer Log Levels	205
Setting UI/REST MySQL Query Log Levels	205
Configuring Advanced Log Settings	206
Downloading Logs from the PHP Developer Logs page	206
Changing Passwords	207
Disabling phpMyAdmin	208
Changing the Password for the Default Account for the User Interface	209
Changing the Password for the Default Console User	210
Changing the Password for the Web Configuration Utility	210

Changing Database Passwords	211
Configuring a New Password on Collector Appliances	211
Editing Silo.Conf	213
Updating the master.system_settings_licenses Table	214
Changing the MySQL Root Password on Database Appliances	214
Recovering the Root MySQL Password	215
Recovering the MySQL SNMP User Account on Collector Appliance	216
Changing the IP Address of an SL1 Appliance	217
Changing the IP Address on an All-In-One Appliance	218
Step 1. Stop the EM7 Service	218
Step 2. Change the IP Address in the Configuration Files	219
Step 3. Change the IP Address in the /etc/hosts File	219
Step 4. Change the IP Address in the Network Interface Configuration File	219
Step 5. Update the IP Address in the MySQL Database	221
Step 6. Reboot the Appliance	221
Step 7. Change the Database Appliance IP Address in the Administration Portals	222
Changing the IP Address on a Database Server	224
Step 1. Stop the EM7 Service	224
Step 2. Change the IP Address in the Configuration Files	225
Step 3. Change the IP Address in the /etc/hosts File	225
Step 4. Change the IP Address in the Network Interface Configuration File	225
Step 5. Update the IP Address in the MySQL Database	227
Step 5a: For Database Servers Configured with PhoneHome	227
Step 5b For Clustered Database Appliances (using HA, DR, or HA+DR)	228
Step 6. Reboot the Appliance	230
Step 7. Change the Database Appliance IP Address in the Administration Portals, Data Collectors, and Message Collectors	230
Changing the IP Address on a Database Collector or Message Collector	233
Step 1. Stop the EM7 Service	233
Step 2. Change the IP Address in the Configuration Files	233
Step 3. Change the IP Address in the /etc/hosts File	233
Step 4. Change the IP Address in the Network Interface Configuration File	234
Step 5. Update the IP Address in the MySQL Database	235
Step 6. Reboot the Appliance	236
Step 7. Change the Database Appliance IP Address in the Administration Portals	236
Changing Name Servers on an SL1 Appliance	240
Changing Name Servers on an SL1 Appliance	240
Installing Additional RPMs on an SL1 Appliance	242
Backup Management	243
Overview	243
Creating a Backup Credential	243
Configuration Backups	244
Defining a Configuration Backup	245
Restoring a Configuration Backup	247
Full Backup	248
Defining a Full Backup	249
Restoring a Full Backup	251
Additional Configuration for Solaris NFS Mounts	253
Defining a DR Backup	253
Restoring a DR Backup	256
Subscription Licenses	258
Viewing the Subscription License Usage	259

Viewing Delivery Status	259
Manually Uploading License Usage to ScienceLogic	260
Downloading the Daily License Usage File	261
Manually Uploading the Daily License Usage File to ScienceLogic	262
Uploading the ScienceLogic Receipt	263
Data Retention Settings for Licensing	264
CAC Authentication	266
Prerequisites	267
Importing an SSL Certificate	268
Updating the ScienceLogic Configuration File	269
Defining the Client Certificate	269
Testing the Configuration	271
Installing an SSL Certificate	273
Certificates for ScienceLogic Servers	274
Requesting a Commercial SSL Certificate	274
Creating Your Own Certificate	275
Installing the Certificate on an SL1 Appliance	277
Authentication Profiles and Resources	279
Authentication Profiles	280
Viewing the List of Authentication Profiles	280
Filtering the List of Authentication Profiles	281
The "default" Authentication Profile	282
Creating an Authentication Profile	283
Editing an Authentication Profile	286
Deleting One or More Authentication Profiles	286
Authentication Resources	286
Viewing the List of Authentication Resources	287
Filtering the List of Authentication Resources	288
The "EM7 Internal" Resource	289
The Legacy Authentication Resources	289
Creating an LDAP/AD Authentication Resource	290
Creating an SSO Authentication Resource	296
Editing an Authentication Resource	301
Deleting an Authentication Resource	301
Managing Host Files	303
Viewing the List of Host Entries	304
Creating a New Host Entry	305
Editing a Host Entry	306
Using an Existing Host File Entry to Create a New Host File Entry (Save As)	308
Deleting One or More Host Entries	309

Chapter

1

Introduction

Overview

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

This chapter includes the following topics:

<i>Who Should Read This Manual?</i>	8
<i>What's In This Manual?</i>	8
<i>Requirements</i>	9

Who Should Read This Manual?

This manual is intended for System Administrators who must monitor and maintain the health of SL1.

This manual describes tasks in the **[System]** tab that are related to maintenance and monitoring of SL1. 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, health tasks, maintenance tasks, 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 SL1 appliances.



Chapter

2

Global Settings

Overview

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

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

This chapter includes the following topics:

<i>Global Settings for API</i>	11
<i>Global Settings for Appliances</i>	12
<i>Global Settings for Asset Automation</i>	15
<i>Global Settings for User Logins, Discovery, Data Collection, UI Features, and Expiration Warnings</i> ..	16
<i>Global Settings for Data Retention</i>	29
<i>Global Settings for Inbound Email and Outbound Email</i>	36
<i>Global Settings for Login Alert Messages</i>	40
<i>Global Settings for Password Reset Emails</i>	42
<i>Global Settings for System Thresholds</i>	45
<i>Global Settings for Interface Thresholds</i>	49
<i>Settings in Silo.Conf</i>	57
<i>Disabling the User Interface on a Database Server</i>	67

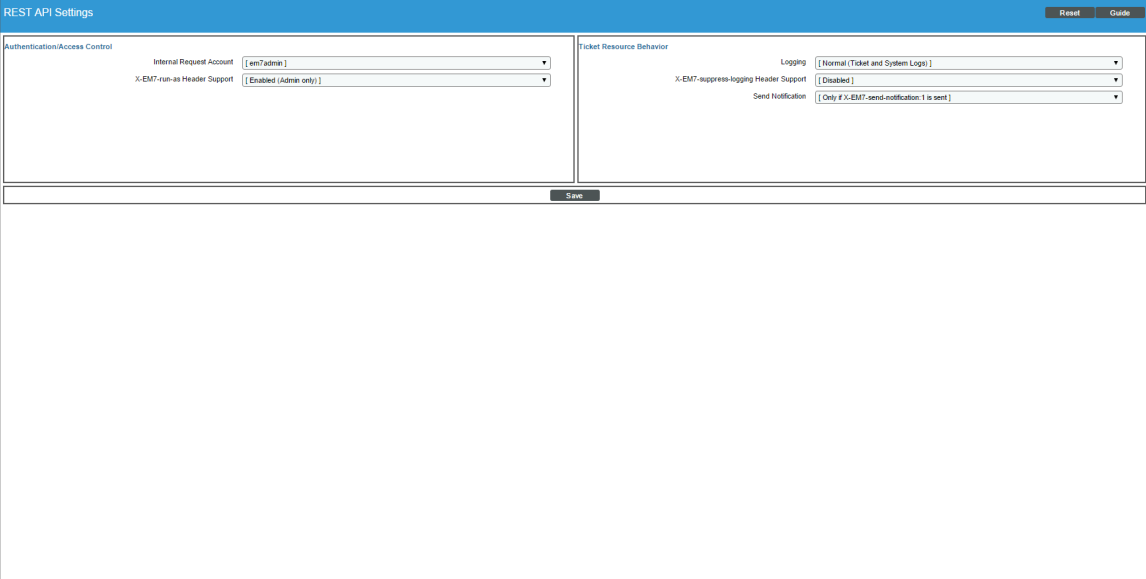
Global Settings for API

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).



The screenshot displays the 'REST API Settings' interface. It is divided into two main sections: 'Authentication/Access Control' and 'Ticket Resource Behavior'. In the 'Authentication/Access Control' section, there are two dropdown menus: 'Internal Request Account' set to '[em7Admin]' and 'X-EM7-run-as Header Support' set to '[Enabled (Admin only)]'. In the 'Ticket Resource Behavior' section, there are three dropdown menus: 'Logging' set to '[Normal (Ticket and System Logs)]', 'X-EM7-suppress-logging Header Support' set to '[Disabled]', and 'Send Notification' set to '[Only if X-EM7-send-notification: 1 is sent]'. A 'Save' button is located at the bottom center of the form area.



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 SL1 to make API requests without a password. For details on building such an API request, see the **ScienceLogic API** manual.
 - **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. For details on using this header, see the **ScienceLogic API** manual.
 - **Logging.** Specifies which logs SL1 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, SL1 will write the standard entry to the audit log that indicates a user performed a write-operation using the API. However, SL1 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, SL1 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, SL1 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 SL1 from writing to the ticket log for the ticket that was created or updated.
- **Send Notification**. When a ticket is created or updated, SL1 can automatically send notification emails to the ticket assignee and ticket watchers. This option specifies the conditions under which SL1 will send notification emails when tickets are created or updated using the API. Choices are:
 - *Only if X-EM7-send-notification:1 is sent*. SL1 will send notification emails for a ticket only when the X-EM7-send-notification header is set to 1. For details on using this header, see the manual **Using the ScienceLogic API**.
 - *Sent after every write operation*. SL1 will send notification emails for every API request that creates or updates a ticket.

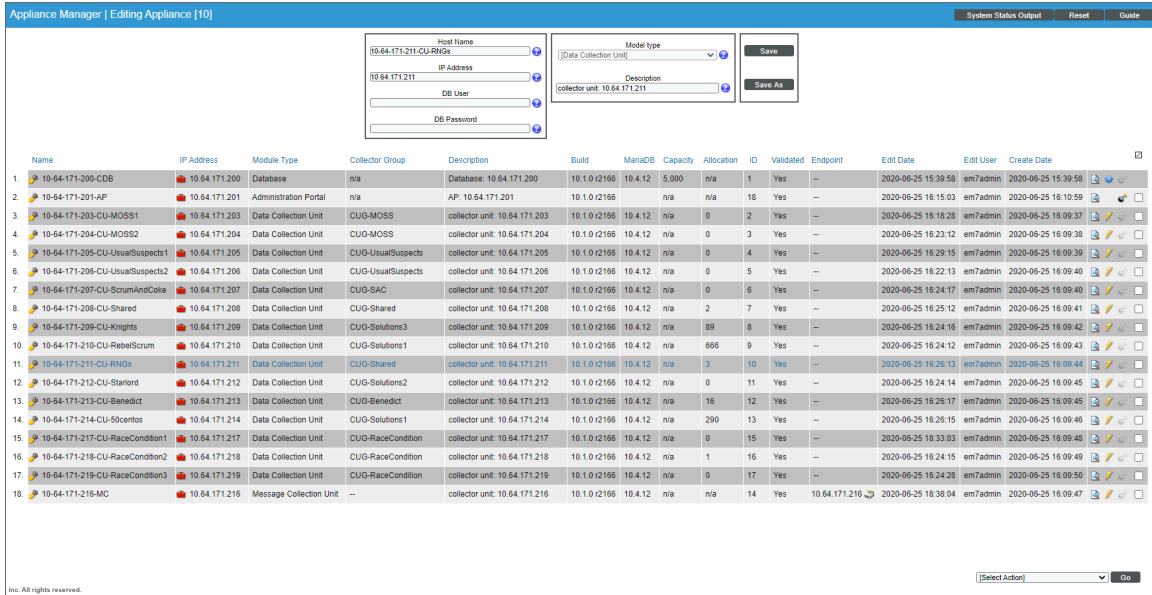
3. Click the **[Save]** button to save changes in this page.

Global Settings for Appliances

The **Appliance Manager** page (System > Settings > Appliances) allows you to view information, including license status, about each ScienceLogic appliance in your system.


From the **Appliance Manager** page, you can also access the Web Configuration Utility for each ScienceLogic appliance by clicking the toolbox icon () , or you can access the database administration tool for each Database Server or All-In-One Appliance by clicking the gear icon () .

For each Database Server, Data Collector, and Message Collector, you can view the click on the magnifying-glass icon () to view the output of the system status script for that appliance.



The screenshot shows the 'Appliance Manager | Editing Appliance [10]' interface. At the top, there are configuration fields for Host Name (10.64.171.211), IP Address (10.64.171.211), DB User, and DB Password. A 'Model type' dropdown is set to '[Data Collection Unit]'. Below these fields are 'Save' and 'Save As' buttons. The main area contains a table with columns: Name, IP Address, Module Type, Collector Group, Description, Build, MariaDB, Capacity, Allocation, ID, Validated, Endpoint, Edit Date, Edit User, and Create Date. The table lists 18 appliances, including Database, Administration Portal, and various Data Collection Units (CUG-MOSS, CUG-UsualSuspects, CUG-SAC, CUG-Shared, CUG-Solutions, CUG-RNGs, CUG-Stanford, CUG-Benedict, CUG-Solutions1, CUG-RaceCondition, CUG-RaceCondition2, CUG-RaceCondition3) and a Message Collection Unit (10-64-171-216-MC).

To edit information about a ScienceLogic appliance:

1. Go to the **Appliance Manager** page (System > Settings > Appliances).
2. Locate the ScienceLogic appliance you want to edit. Click its wrench icon () .
3. The fields in the top pane are populated with values from the selected ScienceLogic appliance.
4. You can edit one or more of the following fields:
 - **Host Name**. Name of the ScienceLogic appliance.
 - **IP Address**. Primary IP address for the ScienceLogic appliance

NOTE: For Data Collection Units that are part of a Phone Home configuration, ensure that the Primary IP address for the Data Collection Unit is its loopback IP.

- **Module Type**. This field is read-only. Possible values are:
 - All In One
 - 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*.

- **Description.** Description of the ScienceLogic appliance.

5. You can edit two optional fields for Data Collector or Message Collector.

- **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 SL1 System, you must define the **DB User** and **DB Password** for each Data Collector or Message Collector.

6. 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 SL1 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.

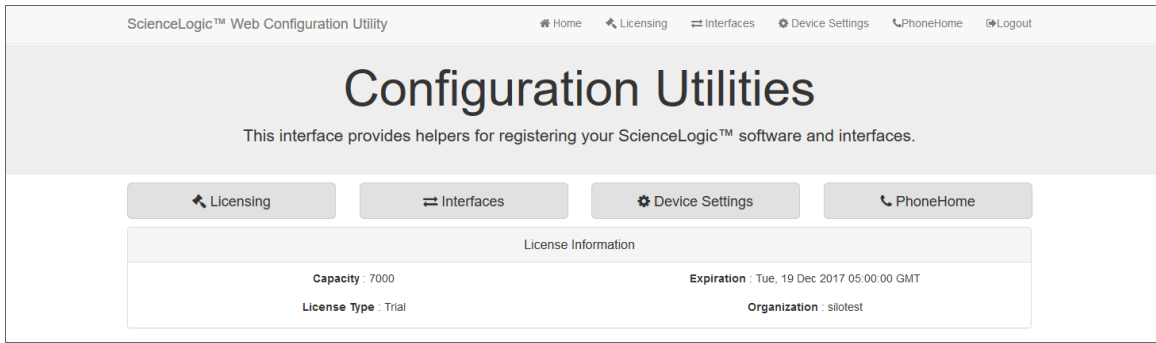
Perform the following steps to log in to the Web Configuration Utility:

1. You can log in to the Web Configuration Utility using any web browser supported by SL1. 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 of the appliance.

3. You will be prompted to type your username and password. Log in as **em7admin** with the appropriate password. The default password is **em7admin**. After logging in, the main **Configuration Utility** page appears:



4. In the **Configuration Utility**, you can license a SL1 appliance, configure interfaces, and edit settings for the SL1 appliance and the Database Server if applicable.
 - For details on using the **Configuration Utility** to license a SL1 appliance, see the manual *Installation and Initial Configuration*.
 - For details on using the 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.

Field Name	Data Source	Alert on Change	Auto Update	Edit User	Edit Date
1. Make	Dynamic App then Internal	No	Yes		2007-12-02 14:49:40
2. Model	Dynamic App then Internal	No	Yes		2007-12-02 14:49:40
3. Serial Number	Dynamic App	No	Yes		2007-12-02 14:49:40
4. Operating System	Dynamic App	No	Yes		2007-12-02 14:49:40
5. Host ID / SID	Dynamic App	No	Yes		2007-12-02 14:49:40
6. OS System Name	Dynamic App then Internal	No	Yes		2007-12-02 14:49:40
7. CPU Count	Dynamic App then Internal	No	Yes		2007-12-02 14:49:40
8. CPU Type/Make	Dynamic App then Internal	No	Yes		2010-01-12 14:49:40
9. CPU Speed	Dynamic App then Internal	No	Yes		2007-12-02 14:49:40
10. Firmware / BIOS Name	Dynamic App	No	Yes	em7admin	2011-01-12 17:17:12
11. Installed Memory	Dynamic App then Internal	No	Yes		2011-01-12 14:49:29
12. Hard Disk/Drive	N/A	No	---		2008-01-12 14:26:25
13. Asset Notes	N/A	No	---		2007-12-02 14:49:40
14. Asset Tag	Dynamic App	No	Yes		2007-12-02 14:49:40
15. Disk Array Size	N/A	No	---		2007-12-02 14:49:40
16. Disk Count	N/A	No	---		2007-12-02 14:49:40
17. Disk Size	N/A	No	---		2007-12-02 14:49:40
18. Method	N/A	No	---		2007-12-02 14:49:40
19. Annual Rate	N/A	No	---		2007-12-02 14:49:40
20. RFID Number	N/A	No	---		2007-12-02 14:49:40
21. Technician	N/A	No	---		2007-12-02 14:49:40
22. DNS Name	N/A	No	---		2007-12-02 14:49:40
23. DNS Domain Name	N/A	No	---		2007-12-02 14:49:40
24. Profile Memo	N/A	No	---		2007-12-02 14:49:40
25. Facility	N/A	No	---		2007-12-02 14:49:40
26. Floor	N/A	No	---		2007-12-02 14:49:40
27. Room	N/A	No	---		2007-12-02 14:49:40
28. Pile	N/A	No	---		2007-12-02 14:49:40
29. Panel	N/A	No	---		2007-12-02 14:49:40
30. Punch	N/A	No	---		2007-12-02 14:49:40

For each standard asset field, you can specify:

- Whether the field can be automatically populated by SL1 .
- Whether the field's value should be automatically updated by SL1 .
- Whether or not SL1 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 SL1 .

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

Global Settings for User Logins, Discovery, Data Collection, UI Features, and Expiration Warnings

The **Behavior Settings** page (System > Settings > Behavior) allows you to define global parameters that affect:

- User Logins
- Discovery
- Data collection
- Settings that affect the display and behavior of the user interface
- Expiration warnings for asset warranties and SSL certificates

The parameters in the **Behavior Settings** page affect all pages, devices, and discovery functionality in SL1 . For most settings, you can define a one-time, manual override in the affected page. You can also override many of these settings per device. For example, you can define global parameters for nightly discovery in this page, but you can override these settings for a specific device on the **Settings** tab of the **Device Investigator** page or the **Device Properties** page (Devices > Device Manager > wrench icon) in the classic user interface.

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

1. Go to the **Behavior Settings** page (System > Settings > Behavior).

- In 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.

NOTE: Do not include a trailing forward slash ("/") at the end of the Interface URL. When SL1 generates URLs for tickets or events (for example, in email messages), the trailing forward slash will be automatically included.

- Force Secure HTTPS.** If enabled, forces users to use HTTPS (secure HTTP) instead of HTTP when users connect to the user interface.
- 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.
 - 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.

- **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:
 - *MD5 (Legacy)*
 - *SHA-512 (FIPS 140-2 Compliant)*
 - *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.
- **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 login attempts from the IP address will be denied.
 - *Lockout by Username and IP Address.* All login attempts by the username from the IP address will be denied.
 - *Lockout by Username (default).* All login attempts by the username will be denied.
 - *Disabled.* Lockouts are disabled.
- **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, SL1 will not allow another attempt for the number of seconds specified in this field. Choices are:
 - *Disabled.* SL1 does not enforce a delay between failed logins.
 - *1 Second.* After a failed login, SL1 will not allow another attempt for one second.
 - *2 seconds.* After a failed login, SL1 will not allow another attempt for two seconds.
 - *4 seconds.* After a failed login, SL1 will not allow another attempt for four seconds.
 - *8 seconds.* After a failed login, SL1 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 SL1, 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 see the 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

- **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 SL1, 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 see the 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.
- **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.
- **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 SL1 system. Can be up to 128 characters in length. This field is useful for companies or organizations with multiple SL1 systems. If a value is provided in this field, SL1 will include a "system identifier" value in each event generated by the current SL1 system. This allows users to easily determine the source SL1 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 SL1 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. Choices in this field are:
 - *1 min*. SL1 will poll the newly discovered interfaces every minute.
 - *5 mins*. SL1 will poll the newly discovered interfaces every five minutes. This is the default value for this field.
 - *10 mins*. SL1 will poll the newly discovered interfaces every 10 minutes.
 - *15 mins*. SL1 will poll the newly discovered interfaces every 15 minutes.
 - *30 mins*. SL1 will poll the newly discovered interfaces every 30 minutes.

- *60 mins.* SL1 will poll the newly discovered interfaces every 60 minutes.
- *120 mins.* SL1 will poll the newly discovered interfaces every 120 minutes.
- **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 > Discovery) does not work for DHCP devices, SL1 will automatically use the community string specified in this field.
- **Strip FQDN From Inbound Email Device Name.** In Events from Email policies, specifies how SL1 will match the regular expression for device name. Choices are:
 - *Enabled.* SL1 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), SL1 will align the event with the matching device with the highest Device ID.
 - *Disabled.* SL1 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 SL1 does not find an exact match in the incoming email, SL1 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**, SL1 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**), SL1 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, SL1 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, SL1 automatically sends email notifications to all watchers when a ticket is created, assigned, or updated.
 - *Disabled*. When you select this option, SL1 does not automatically send email notifications to all watchers when a ticket is created, assigned, or updated.
- **Prevent Browser Saved Credentials**. This checkbox specifies whether or not the user interface 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 SL1. This is a security measure, to prevent clickjacking 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.
- **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.

- **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). The user interface will display the following in the lower right of the page:

Previous Login: *yyyy-mm-dd hh-mm-ss from user's IP address.*

Failed Login: *yyyy-mm-dd hh-mm-ss from user's IP address.*

- **Ignore trap agent-addr varbind.** If you select this checkbox, SL1 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 > 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) will not display the *Create a Ticket* option. This field is unselected by default.
- **Enable CDP Topology**. If selected, SL1 will use Cisco Discovery Protocol (CDP) for each device that supports CDP. SL1 will 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 SL1, 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**. If selected, SL1 will use Link Layer Discovery Protocol (LLDP) for each device that supports LLDP. SL1 will then generate topology maps from the discovered LLDP relationships.
- **Enable Community String Indexing (VLAN Topology)**. If selected, SL1 will perform discovery of VLANs during topology collection. By default, this option is not selected because the SNMP requests used to discover VLANs might cause some types of hardware to erroneously reboot.
- **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 SL1. 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. SL1 also uses this default value to perform timezone conversions when no user timezone setting is available. For example, if SL1 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 SL1 from monitoring and reporting on NFS "shared" file systems. SL1 will monitor and report only on local file systems.
- **Port Polling Type**. Specifies how SL1 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. The options are:
 - *0. Model Device Only*. Discover if device is up and running and if so, discover device's make and model.
 - *1. Discover Dynamic Apps*. Discovery tool will search for Dynamic Applications associated with the device. Discovery will also perform *0. Model Device Only* discovery.

- 2. *Initial Population of Apps*. Discovery tool will retrieve subset of data from Dynamic Applications, to save time. Discovery tool will later retrieve full sets of data from each Dynamic Application. Discovery tool will also perform 1. *Discover Dynamic Apps* and 0. *Model Device Only*.
- 3. *Discover SSL Certificates*. Discovery tool will search for SSL certificates and retrieve SSL data. Discovery tool will also perform 2. *Initial Population of Apps*, 1. *Discover Dynamic Apps*, and 0. *Model Device Only*.
- 4. *Discover Open Ports*. Discovery tool will search for open ports. Discovery tool will also perform 3. *Discover SSL Certificates*, 2. *Initial Population of Apps*, 1. *Discover Dynamic Apps*, and 0. *Model Device Only*.

NOTE: If your system includes a firewall and you select option 4: *Discover Open Ports*, discovery may be blocked and/or may be taxing to your network.

- 5. *Advanced Port Discovery*. Discovery tool will search for open ports, using a faster TCP/IP connection method. Discovery tool will also perform 3. *Discover SSL Certificates*, 2. *Initial Population of Apps*, 1. *Discover Dynamic Apps*, and 0. *Model Device Only*.

NOTE: If your system includes a firewall and you select option 5: *Advanced Port Discovery*, some auto-discovered devices may remain in a pending state for some time after discovery. These devices will achieve a healthy status, but this might take several hours.

- **Rediscovery Scan Level (Nightly)**. Specifies the data to be gathered/updated each day during the nightly discovery process. The nightly discovery process will find any changes to previously discovered devices. The **Rediscovery Scan Level (Nightly)** field contains the same options as the **Initial Discovery Scan Level** field.
- **Discovery Scan Throttle**. Specifies the amount of time a discovery process should pause between each IP address in a discovery session. Pausing discovery processes between IP addresses 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.

NOTE: The **Discovery Scan Throttle** setting does not affect nightly auto discovery.

- **Port Scan All IPs**. Specifies whether SL1 should scan all IP addresses on a device for open ports. The choices are:
 - *0. Disabled*. SL1 will scan only the Admin Primary IP address (the IP address SL1 uses to communicate with the device) for open ports.
 - *1. Enabled*. SL1 will scan all discovered IP addresses for open ports.

NOTE: The *Port Scan All IPs* setting affects initial discovery, nightly auto discovery, and re-discovery.

- **Port Scan Timeout.** Length of time, in milliseconds, after which SL1 should stop trying to scan an IP address for open ports and begin scanning the next IP address (if applicable). Choices are between 60,000 and 1,800,000 milliseconds.

NOTE: The *Port Scan Timeout* setting affects initial discovery, nightly auto discovery, and re-discovery.

- **Restart Windows Services (Agent required).** Specifies whether SL1 should automatically restart failed Windows services that have been defined on the device with a startup type of "automatic". The choices are:
 - *0. Disabled.* SL1 will not automatically restart failed Windows services that have been defined on the device with a startup type of "automatic".
 - *1. Enabled.* SL1 will automatically restart failed Windows services that have been defined on the device with a startup type of "automatic".

NOTE: 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.

- **Hostname Precedence.** Specifies which name SL1 will use for each discovered device. Choices are:
 - *SNMP System Name.* Use the device name specified in the device's SNMP System MIB.
 - *DNS Reverse Lookup.* Use the device name specified in the device's reverse-lookup record.

NOTE: If *SNMP System Name* is selected, and SL1 cannot find an SNMP name for the device, SL1 will assign the name returned by the DNS Reverse Lookup. If SL1 cannot find a DNS Reverse Lookup name for the device, SL1 will use the device's Admin Primary IP address as the device name in SL1.

- **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 SL1, 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 SL1, 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}
 - {ifdescr}
 - {if_id}
 - {did}
 - {ifindex}
 - {ifphysaddress}
 - {iftype}
 - {ifspeed}
 - {ifhighspeed}
 - {ifoperstatus}
 - {ifadminstatus}
- **DNS Hostnames.** If SL1 will use the DNS Reverse Lookup name as the device name (see the description of the field **Hostname Precedence**), this field specifies whether SL1 will use the fully-qualified domain name or only the hostname for each discovered device. Choices are:
 - *Strip Device Name (Hostname).* SL1 will use only the device name as the DNS hostname for each device.
 - *Use Full Domain Name (FQDN).* SL1 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 SL1 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 SL1 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 SL1 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 SL1 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.
- **Topology Map Rendering.** This field lets you choose the rendering engine for the maps in SL1. You can choose from *HTML5*, which is used with the maps on the **Maps** page, or *Flash*, which was used with Classic Maps (Views). The default is *Flash*.
- **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 **Dashboards** manual.
- **Enable CBQoS Collection.** If selected, SL1 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 SL1 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, SL1 applies an interpolation function that spaces the data at regular time intervals. For example, suppose you have specified that SL1 should collect interface data every five minutes. However, due to network traffic across the Data Collectors, SL1 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, SL1 first applies an interpolation to the data to prepare the data for normalization.
- **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.

NOTE: Enabling the **Enhanced OID Translation** option might affect performance on large environments with a large number of traps.

- **Enable Concurrent SNMP Collection BETA.** 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.
- **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 SL1 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.

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

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).

The screenshot displays the 'Data Retention Settings' page. It is organized into three main sections: System Data Retention, Collection Data Retention, and Subscription Data Retention. Each section contains a list of data types with corresponding retention settings. The settings include a numerical value, a unit (days, months, or years), and a tooltip indicating the equivalent number of days. For example, '3 months*' is equivalent to 90 days. A 'Save' button is located at the bottom of each section. At the bottom of the page, a note states '* 1 month = 30 days'.

Section	Data Type	Retention	Equivalent
System Data Retention	Audit Logs	3 months*	90 days
	Event Logs	3 months*	90 days
	Access Logs	12 months*	360 days
	System Logs	31 days	31 days
	Collector Unit Data Buffer	2 days	2 days
Collection Data Retention	Raw Performance Data	7 days	7 days
	Hourly Rollup Performance Data	120 days	120 days
	Daily Rollup Performance Data	24 months*	720 days
	Configuration Data	7 days	7 days
	Journal Data	90 days	90 days
	Bandwidth Data	31 days	31 days
	Hourly Rollup Bandwidth Data	120 days	120 days
	Daily Rollup Bandwidth Data	24 months*	720 days
	Bandwidth Billing Data	24 months*	720 days
	Device Logs Age	90 days	90 days
	Device Logs Max	10000 records	10000 records
	Raw ITSM Data	31 days	31 days
	ITSM Service Metrics Data	30 days	30 days
	Hourly Rollup ITSM Service Metrics Data	120 days	120 days
	Daily Rollup ITSM Service Metrics Data	12 months*	360 days
ITSM Key Metrics Data	120 days	120 days	
Hourly Rollup ITSM Key Metrics Data	120 days	120 days	
Daily Rollup ITSM Key Metrics Data	24 months*	720 days	
Subscription Data Retention	Subscriber Device Configuration Data	0 months*	0 days
	Subscriber Device Usage Data	0 months*	0 days
	Subscriber System Configuration Data	3 months*	90 days
	Subscriber System Usage Data	3 months*	90 days
	Subscriber Device Type Data	3 months*	90 days
	Subscriber Daily Delivery Data	3 months*	90 days

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 SL1 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, SL1 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 120 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 120 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.
- **Hourly Rollup ITSM Service Metrics Data**. Number of days to retain hourly normalized values for metrics in IT Service policies. The default value is 120 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 120 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 365 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.
- **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 6 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 6 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 SL1 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 SL1 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. SL1 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 SL1, 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 SL1 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 SL1 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	Collects raw performance data from a device at the following intervals: <ul style="list-style-type: none"> • 1 minute • 2 minutes • 3 minutes

Collector	Collected Data and Intervals
	<ul style="list-style-type: none"> • 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 Poll Frequency field, in the Properties Editor page (System > Manage > Applications > Create or use the ).</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 Aggregation Frequency field, in the IT Service Editor page (Registry > IT Services > IT Service Manager > Create or use the ).</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

Collector	Collected Data and Intervals
	<ul style="list-style-type: none"> • 15 minutes • 30 minutes • 60 minutes • 120 minutes <p>You can specify the frequency at which SL1 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 select the  for the given interface).</p>
Additional Performance Data	SL1 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, SL1 collects this data every 5 minutes.

Data Normalization and Rollup

SL1 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 SL1 rolls up data, SL1 groups data into larger sets and calculates the average value for the larger set.

SL1 supports two types of rollup:

- **Hourly.** Groups and averages data that is collected at intervals of 60 minutes or less. SL1 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. SL1 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.

SL1 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

Frequency of Raw Collection	Rollup
Every 60 minutes	60 minutes, 24 hours
Every 120 or longer	24 hours

Before SL1 normalizes data, SL1 transforms the data. To transform the data, SL1 does the following:

- For bandwidth data and data from Dynamic Applications of type "Performance", SL1 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", SL1 evaluates presentation formulas. Counter metrics are first transformed into rates before evaluation.

NOTE: During the data transform steps, SL1 does not directly rollup the raw data in the database tables.

When SL1 rolls up data, SL1 must normalize that data, as follows:

NOTE: As a new piece of data is collected by SL1, the hourly normalization and daily normalization is calculated. SL1 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 SL1, 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, SL1 collects data about file system usage on the device named *my_device*. As each raw data point is collected, SL1 normalizes and rolls up the collected data for file system usage for *my_device*. SL1 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 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 SL1 should store raw data and rolled up and normalized data:

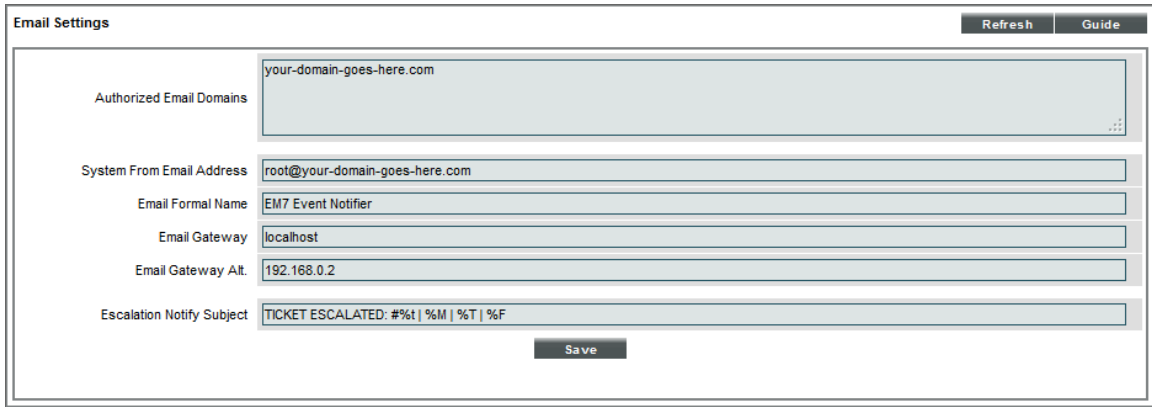
- You can define system-wide, default settings in the **Data Retention Settings** page (System > Settings > Data Retention). These settings apply to all collected data. However, you can override these system settings in the **Device Thresholds** page (Registry > Devices > Device Manager > wrench icon > Thresholds).
- 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 ). 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 SL1 will send and receive email. SL1 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).



The screenshot shows the 'Email Settings' page with the following fields and values:

Field	Value
Authorized Email Domains	your-domain-goes-here.com
System From Email Address	root@your-domain-goes-here.com
Email Formal Name	EM7 Event Notifier
Email Gateway	localhost
Email Gateway Alt.	192.168.0.2
Escalation Notify Subject	TICKET ESCALATED: #%t %M %T %F

Buttons: Refresh, Guide, Save

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 SL1. SL1 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 SL1 will send all outbound email.
- **Email Formal Name.** Name that will appear in the "from" field in email messages sent from SL1. This value can be any alphanumeric value, up to 64 characters in length.
- **Email Gateway.** IP address or fully-qualified name of SL1's SMTP Relay server. If SL1 is to send outgoing messages, this field must be defined. Examples of when SL1 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 SL1 SMTP Relay server is the same as the fully-qualified name of the Database Server or All-In-One Appliance.

If SL1 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), SL1 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.

- **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, SL1 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.
%B	Organization	Organization billing ID.
%b	Organization	Impacted organization.
%C	Organization	Organization CRM ID.

Variable	Source	Description
%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.
%7	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.

Variable	Source	Description
%q	Asset	Asset zone.
%R	Event Policy	Event policy cause/action text.
%r	System	Unique ID / name for the current SL1 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.
%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 SL1, 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 click the **[OK]** button to agree to the terms and conditions of use for that system.


To add a custom login alert message to SL1:

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:



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, SL1 will check the account information for that user. If the user's account information includes an email address, SL1 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. SL1 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, SL1 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 SL1 will not resend an email.

If the user's account information does not include an email address, SL1 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 SL1 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 SL1 when an end user selects the *I forgot my password* option from the **Login** page.

To define the email message sent by SL1 :

1. Go to the **Password Reset Email Editor** page (System > Settings > Password Reset Email).

The screenshot shows the 'Password Reset Email Editor' interface. It features a blue header with the title and a 'Guide' button. The main area contains three fields: 'Priority' (a dropdown menu set to '[High]'), 'Subject' (a text input field containing 'EM7 | %O [automated message]'), and 'Message' (a large text area containing the email body text: 'Hello %fn %n, Your password for account %A has been reset. Use the following link to log-in and choose a new password: %L'). A 'Save' button is located at the bottom center of the form.

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, SL1 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 SL1 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 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, and number of component devices.



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 **Device Thresholds** page (Devices > Device Manager > wrench icon > Thresholds). 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:

- **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, SL1 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), SL1 uses the value in this field if there is no differing value specified in the **Discovery Session Editor** page.

- During *re-discovery* (clicking the binocular icon () in the Device Properties page), SL1 will use the value in this field if there no value is specified in the **Device Thresholds** page (Devices > Device Manager > wrench icon > Thresholds) for the device.
 - During *nightly auto-discovery* (run automatically by SL1 every night, to update device information), SL1 uses the value in this field if no differing value is specified in the **Device Thresholds** page (Devices > Device Manager > wrench icon > Thresholds) for a device.
- **Maximum Allowed Interfaces.** Specifies the maximum number of interfaces per device. If a device exceeds this number of interfaces, SL1 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), SL1 uses the value in this field if there is no differing value specified in the **Discovery Session Editor** page.
 - During *re-discovery* (clicking the binocular icon () in the Device Properties page), SL1 will use the value in this field if there is no differing value is specified in the **Device Thresholds** page (Devices > Device Manager > wrench icon > Thresholds) for the device.
 - During *nightly auto-discovery* (run automatically by SL1 every night, to update device information), SL1 uses the value in this field if no differing value is specified in the **Device Thresholds** page (Devices > Device Manager > wrench icon > Thresholds) for a device.
- **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 SL1's ping (round-trip time divided by 2). The default value is 100 ms. When the latency threshold is exceeded, SL1 generates an event for that device.
- **System Availability.** During polling, SL1 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, SL1 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. SL1 polls component devices for availability at the frequency defined in the Dynamic ApplicationFor 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 SL1 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, SL1 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, SL1 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), SL1 does not generate events for that file system.

- **Rollover Percent.** For any collected data that uses a 32-bit counter, you can specify how SL1 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, SL1 assumes that the counter has been manually reset or that the two values were collected out of order. SL1 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 SL1 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}), SL1 concludes that the 32-bit counter rolled over.
- For example, if you specified "25" in this field, SL1 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, SL1 concludes that the 32-bit counter rolled over.
- When SL1 determines a counter has rolled over, SL1 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 SL1 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 SL1 records a counter value that is lower than the previously collected value and the platform has determined that the value is not a rollover, SL1:

- 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, SL1 concludes that the data was collected out of order.
- When SL1 determines a data point has been collected out of order, SL1 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, SL1 concludes that the 32-bit counter was manually reset to zero (0). SL1 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, SL1 concludes that the 64-bit counter was manually reset to zero (0). SL1 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 > Device Manager > 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 > Device Manager > wrench icon) for a device, this field specifies the percentage of packets that must be returned during an availability check for SL1 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**, SL1 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**, SL1 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 SL1 purges the component device. When a device is purged, SL1 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 **Device Thresholds** page 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 SL1 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", SL1 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 **Device Thresholds** page for the device.

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

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 SL1 should monitor an interface, SL1 will collect data about the interface and also monitor performance thresholds for the interface. SL1 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, SL1 will generate an event.

To define global thresholds for interfaces:

1. Go to **Interface Thresholds Defaults** page (System > Settings > Thresholds > Interface).

Metric	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	0.000	Disabled
Discards % In		

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**, **kpbs**, **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

Threshold	Default Value	Default Status
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
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**, **kpbs**, **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

Threshold	Default Value	Default Status
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
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

Threshold	Default Value	Default Status
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 > InboundErrors	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

Threshold	Default Value	Default Status
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 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

Threshold	Default Value	Default Status
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
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

Threshold	Default Value	Default Status
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
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, SL1 triggers an event.

- For thresholds that include the word *Falling*, when a value falls below the specified value, SL1 triggers an event.
- For thresholds that do not include the word *Rising* or *Falling*, when a value exceeds the specified value, SL1 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 SL1. 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 SL1. 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 SL1. 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 SL1 appliance has a configuration file called **silو.conf**, which contains configuration information about the appliance itself, such as the IP address, licensing information, and directory locations. The default settings in silو.conf are configured automatically when the appliance is installed. The following section describes how you can add additional, non-default settings to silو.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.

From the **Device Settings** page of the Web Configuration Utility, you can edit the **silو.conf** file and the following files:

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

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

To edit the silo.conf file:

1. [Log in to the Web Configuration Utility](#). The **Configuration Utilities** page appears.
2. Click the **[Device Settings]** button. The **Settings** page appears.

ScienceLogic™ Web Configuration Utility

Home Licensing Interfaces Device Settings PhoneHome Logout

Settings

Configure your appliance.

Web Configuration Username: em7admin

Web Config Password (change only):

Confirm Web Config Password:

Appliance Type: All in One

Save

Edit Files

chrony.conf silo.conf chrony.d/servers.conf

3. In the Edit Files section, click **silo.conf**. The Silo.conf Editor modal page appears:

Silo.conf Editor

```
[LOCAL]
rootdir = /opt/em7
vardir = /var/lib/em7
logdir = /var/log/em7
rundir = /run/em7
ipaddress = 10.2.14.27
dbdir = /data/db
dbserver = 127.0.0.1
dbport = 7706
dbuser = root
dbpasswd = em7admin
portcheck = /usr/bin/nmap
model_type = 1
eventmanager = internal,api,email,syslog,trap,dynamic
disable_itiil_compliance = 1

[SLADMIN]
sladmin_user = em7admin
sladmin_cred = Jz3j_-ghaKrTn0cFmzQRSSPA_IIIKECDAl8VU6Njvi828zrGgGec99R5yV1eAtn_

[CENTRAL]
dbserver = 127.0.0.2
dbport = 7706
dbuser = root
dbpasswd = em7admin
```

Save

4. For ISO installs of 10.1.0, the value for `dbpassword` is encrypted in the `silو.conf` file.

- To manually unencrypt this password:
 - Log in to the console of the Database Server or SSH to the Database Server.
 - Enter the following at the shell prompt:

```
/opt/em7/share/scripts/encrypt_decrypt_password.py --action decrypt
```

- To manually encrypt the password:
 - Log in to the console of the Database Server or SSH to the Database Server.
 - Enter the following at the shell prompt:

```
/opt/em7/share/scripts/encrypt_decrypt_password.py --action encrypt
```

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

- **store_timeout**. You can edit this setting in the `silو.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 `silو.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.

- **eventmanager**. You can edit this setting in the `silو.conf` file on each SL1 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 SL1 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 SL1 appliance the 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 SL1, 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, SL1 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_v1_trap_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 SL1 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 SL1 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 SL1, 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. SL1 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.Conf 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.Conf 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.Conf 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 wait timeout, in seconds.
- **write_timeout**. You can edit this setting in the `silو.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.

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

```
write_timeout=X
```

where:

- *X* is the write timeout, in seconds.
- **memory_limit**. You can edit this setting in the `silو.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 `silو.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).

- **message_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.

```
message_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 SL1 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 SL1 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]**. Adhoc 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 adhoc 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 adhoc report settings. Its default value is "second".
- *ui_php_timeout*. This setting controls the amount of time elapsed after which an inactive SL1 reports session will time out. Its default value is 1800.

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

NOTE: All changes to the silo.conf file are logged in the SL1 Database Server.

Disabling the User Interface on a Database Server

Database Servers are automatically configured to provide the user interface. If your SL1 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 vi /etc/siteconfig/firewalld-rich-rules.siteconfig
```

3. Add 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

A Collector Group is a group of SL1 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.

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.

This chapter will describe how to create and manage Collector Groups.

NOTE: If you are using a SL1 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.

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

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

This chapter includes the following topics:

<i>Installing, Configuring, and Licensing Data Collectors</i>	70
<i>Technical Information About Data Collectors</i>	71
<i>Duplicate IP Addresses</i>	71
<i>Open Ports</i>	71
<i>Viewing the List of Collector Groups</i>	71
<i>Creating a Collector Group</i>	72
<i>Editing a Collector Group</i>	75
<i>Collector Groups and Load Balancing</i>	76
<i>Tuning Collector Groups in the silo.conf File</i>	77
<i>Collector Affinity</i>	79
<i>Failover for Collector Groups for Component Devices</i>	80
<i>Collector Groups for Merged Devices</i>	80
<i>Creating a Collector Group for Data Storage Only</i>	81
<i>Deleting a Collector Group</i>	82
<i>Aligning the Collector Group for A Single Device</i>	83
<i>Aligning the Collector Group in a Device Template</i>	84
<i>Changing the Collector Group for One or More Devices</i>	85
<i>Managing the Host Files for a Collector Group</i>	85
<i>Processes for Collector Groups</i>	86
<i>Enabling and Disabling Concurrent PowerShell for Collector Groups</i>	87
<i>Enabling Concurrent PowerShell on All Collector Groups</i>	87
<i>Disabling Concurrent PowerShell on All Collector Groups</i>	87
<i>Enabling Concurrent PowerShell on a Specific Collector Group</i>	88
<i>Disabling Concurrent PowerShell on a Specific Collector Group</i>	88

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.

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

To view the list of already-defined Collector Groups in your SL1 system:

1. Go to the **Collector Group Management** page (System > Settings > Collector Groups).

Collector Group Management | Creating New Group

Collector Group Name:

Collector Follower:

Collector Selection:

Message Collector: 10-64-171-216-MC

Collectors Available For Follower:

Fallback Mode:

Follower Delay (minutes):

Fallback Delay (minutes):

Save

Collector Group Registry | Found 9 collector groups

Name	ID	# Collectors	Msg Collectors	# Devices	Edit User	Edit Date
CUG-Benedict	7	1		16	em7admin	2020-06-25 16:09:51
CUG-MOBS	1	2	10-64-171-216-MC	0	em7admin	2020-06-25 16:09:50
CUG-RaceCondition	8	3		1	em7admin	2020-06-25 16:50:19
CUG-SAC	3	1		0	em7admin	2020-06-25 16:09:50
CUG-Shared	6	2		125	em7admin	2020-06-25 16:09:51
CUG-Solutions1	5	2		956	em7admin	2020-06-25 16:39:04
CUG-Solutions2	4	1		0	em7admin	2020-06-26 15:09:49
CUG-Solutions3	9	1		89	em7admin	2020-06-26 15:10:15
CUG-UsualSuspects	2	2		0	em7admin	2020-06-25 16:09:50

Inc. All rights reserved.

2. The **Collector Group Registry** pane displays a list of all Collector Groups in your SL1 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 SL1 to each Collector Group.
- **# 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 SL1 system, you can define one or more Collector Groups. Each Collector Group must include at least one Data Collector.

To define a new Collector Group:

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.

The screenshot shows the 'Collector Group Management | Creating New Group' interface. The top pane contains a form with the following fields:

- Collector Group Name:** A text input field.
- Collector Follower:** A dropdown menu with options: 'Off (Maximize Manageable Devices)', 'Generate Alert On Collector Outage' (set to 'Yes'), and 'Enable Concurrent SNMP Collection BETA' (set to 'Use systemwide default').
- Collector Selection:** A list box containing '10-64-171-216-MC'.
- Collectors Available For Failover:** A dropdown menu set to 'N/A'.
- Failback Mode:** A dropdown menu set to 'Automatic'.
- Failover Delay (minutes):** A text input field set to '5'.
- Failback Delay (minutes):** A text input field set to '5'.
- Save:** A button.

The bottom pane shows the 'Collector Group Registry | Found 9 collector groups' table:

ID	Name	# Collectors	Mtu Collectors	# Devices	Exit User	Exit Date
1	CUG-Benedict	7	1	16	em7admin	2020-06-25 16:09:51
2	CUG-MOBS	1	2	0	em7admin	2020-06-25 16:09:50
3	CUG-RaceCondition	8	3	1	em7admin	2020-06-25 16:50:19
4	CUG-SAG	3	1	0	em7admin	2020-06-25 16:09:50
5	CUG-Shared	6	2	125	em7admin	2020-06-25 16:09:51
6	CUG-Solutions1	5	2	956	em7admin	2020-06-25 18:39:04
7	CUG-Solutions2	4	1	0	em7admin	2020-06-26 15:09:49
8	CUG-Solutions3	9	1	89	em7admin	2020-06-26 15:10:15
9	CUG-UsualSuspects	2	2	0	em7admin	2020-06-25 16:09:50

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.
- **Generate Alert on Collector Outage.** Specifies whether or not the platform should generate an event if a Data Collector has an outage.
- **Enable Concurrent SNMP Collection BETA.** 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).
- **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 that a single Message Collector can be used by multiple Collector Groups.

NOTE: 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, SL1 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.
 - SL1 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, SL1 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.

CAUTION: If the number of available Data Collectors is less than the value in the **Collectors Available For Failover** field, SL1 will not failover within the Collector Group. **SL1 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.** EM7 will generate a critical event.

- **Failback 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, SL1 will automatically redistribute data-collection tasks among the Collector Group, including the previously failed Data Collector.
 - *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 SL1 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.
 - **Failback Delay (minutes).** Applies only if you selected *On (Maximize Reliability)* in the **Collector Failover** field and *Automatic* in the **Failback Mode** field. Specifies the number of minutes SL1 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.
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](#) and the section on [aligning a device group with a Collector Group](#).

Editing a Collector Group

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:

1. Go to the **Collector Group Management** page (System > Settings > Collector Groups).

- In the **Collector Group Management** page, go to the **Collector Group Registry** pane at the bottom of the page.

The screenshot shows the 'Collector Group Management | Creating New Group' interface. The top pane contains a form with the following fields:

- Collector Group Name (text input)
- Collector Follower (dropdown menu)
- Generate Alert On Collector Outage (dropdown menu)
- Enable Concurrent SHAP Collection (BETA) (checkbox)
- Collector Selection (dropdown menu)
- Message Collector (dropdown menu)
- Collectors Available For Failover (dropdown menu)
- Failback Mode (dropdown menu)
- Follower Delay (minutes) (text input)
- Failback Delay (minutes) (text input)
- Save button

The bottom pane shows the 'Collector Group Registry | Found 9 collector groups' table:

Name	ID	# Collectors	My Collectors	# Devices	Edit User	Edit Date
CUG-Benedict	7	1		16	em7admin	2020-06-25 16:09:51
CUG-MOSS	1	2	10-64-171-216-MC	0	em7admin	2020-06-25 16:09:50
CUG-RaceCondition	8	3		1	em7admin	2020-06-25 16:50:19
CUG-SAC	3	1		0	em7admin	2020-06-25 16:09:50
CUG-Shared	5	2		125	em7admin	2020-06-25 16:09:51
CUG-Solutions1	5	2		956	em7admin	2020-06-25 16:39:04
CUG-Solutions2	4	1		0	em7admin	2020-06-26 15:09:49
CUG-Solutions3	9	1		89	em7admin	2020-06-26 15:10:15
CUG-UsualSuspects	2	2		0	em7admin	2020-06-25 16:09:50

- Find the Collector Group you want to edit. Click its wrench icon ().
- 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 [Creating a Collector Group](#).
- Click the **[Save]** button to save any changes to the Collector Group.

Collector Groups and Load Balancing


To perform initial discovery, SL1 uses a single, selected Data Collector from the Collector Group. This allows you to easily troubleshoot discovery if there are any problems.

After each discovered device is modeled (that is, after SL1 assigns a device ID and creates the device in the database), SL1 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.

SL1 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

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, SL1 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.

SL1 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 `silو.conf` File

With the addition of execution environments to SL1, SL1 sorts data collections in to a two-process-pool model.

SL1 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. SL1 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.Conf 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 PowerPacks that can be processed in parallel.	2
dynamic_collect_num_request_workers	The maximum number of request workers in each worker pool. In general, this value controls the number of collections within a PowerPack 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 requests workers.	200

NOTE: The database values for these parameters are "Null" by default, which specifies that SL1 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", SL1 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.

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.
- **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.
 - If the Dynamic Application will consume cache produced by a Dynamic Application aligned to a non-root device (for instance, a cluster device).
- **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. 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 > Device Manager > wrench icon) for either the physical device or the component device.
- From the **Actions** menu in the **Device Manager** page (Devices > Device Manager), 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 > Device Manager > wrench icon) for either the physical device or the component device, , select *Unmerge Devices* to unmerge devices.

- From the **Actions** menu in the the **Device Manager** page (Devices > Device Manager), 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, SL1 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

From the **Collector Group Management** page, you can create a **Virtual Collector Group** 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. **SL1 will stop collecting data from devices aligned with a Virtual Collector Group.**

To define a Virtual Collector Group:

1. Go to the **Collector Group Management** page (System > Settings > Collector Groups).

The screenshot shows the 'Collector Group Management | Creating New Group' interface. The top pane contains a form with the following fields:

- Collector Group Name:** A text input field.
- Collector Follower:** A dropdown menu with the option '[Off (Maximum Manageable Devices)]'.
- Generate Alert On Collector Outage:** A dropdown menu with the option '[Yes]'.
- Enable Concurrent SNMP Collection [BETA]:** A dropdown menu with the option '[Use systemwide default]'.
- Collector Selection:** A dropdown menu.
- Message Collector:** A dropdown menu with the value '10-64-171-216-MC'.
- Collectors Available For Failover:** A dropdown menu with the value 'N/A'.
- Failback Mode:** A dropdown menu with the option '[Automatic]'.
- Failover Delay (minutes):** A text input field with the value '5'.
- Failback Delay (minutes):** A text input field with the value '5'.
- Save:** A button.

The bottom pane shows the 'Collector Group Registry | Found 9 collector groups' table:

	Name	ID	# Collectors	Mtu Collectors	# Devices	Exit User	Exit Date	
1.	CUG-Benedict	7	1		16	em7admin	2020-06-25 16:09:51	
2.	CUG-MOBS	1	2	10-64-171-216-MC	0	em7admin	2020-06-25 16:09:50	
3.	CUG-RaceCondition	8	3		1	em7admin	2020-06-25 16:50:19	
4.	CUG-SAC	3	1		0	em7admin	2020-06-25 16:09:50	
5.	CUG-Shared	6	2		125	em7admin	2020-06-25 16:09:51	
6.	CUG-Solutions1	5	2		956	em7admin	2020-06-25 16:39:04	
7.	CUG-Solutions2	4	1		0	em7admin	2020-06-26 15:08:49	
8.	CUG-Solutions3	9	1		89	em7admin	2020-06-26 15:10:15	
9.	CUG-UsualSuspects	2	2		0	em7admin	2020-06-25 16:09:50	

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

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:

1. Go to the **Collector Group Management** page (System > Settings > Collector Groups).

The screenshot shows the 'Collector Group Management | Creating New Group' interface. It features a form with several sections: 'Collector Group Name' with a text input and a dropdown; 'Collector Fallback' with a dropdown and a checkbox for 'Off (Maximum Manageable Devices)'; 'Generate Alert On Collector Outage' with a dropdown set to 'Yes'; 'Enable Concurrent SNMP Collection (BETA)' with a dropdown set to 'Use systemwide default'; 'Collector Selection' with a dropdown menu; 'Message Collector' with a dropdown menu showing '10.64.171-216-MC'; 'Collectors Available For Fallback' with a dropdown; 'Fallback Mode' with a dropdown set to 'Automatic'; 'Fallback Delay (minutes)' with two input fields, both set to '5'; and a 'Save' button.

Below the form is the 'Collector Group Registry | Found 9 collector groups' table:


ID	Name	# Collectors	Msg Collectors	# Devices	Edit User	Edit Date
1	CUG-Benedict	7	1	16	em7admin	2020-06-25 16:09:51
2	CUG-MOSS	1	2	0	em7admin	2020-06-25 16:09:50
3	CUG-RaceCondition	8	3	1	em7admin	2020-06-25 16:50:19
4	CUG-SAC	3	1	0	em7admin	2020-06-25 16:09:50
5	CUG-Shared	6	2	125	em7admin	2020-06-25 16:09:51
6	CUG-Solutions1	5	2	956	em7admin	2020-06-25 16:39:04
7	CUG-Solutions2	4	1	0	em7admin	2020-06-25 15:08:49
8	CUG-Solutions3	9	1	89	em7admin	2020-06-26 15:10:15
9	CUG-UsualSuspects	2	2	0	em7admin	2020-06-25 16:09:50

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 bomb icon ().

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

- In the **Device Manager** page, find the device you want to edit. Click its wrench icon (). The **Device Properties** page appears:

Close	Properties	Thresholds	Collections	Monitors				
Schedule	Logs	Toolbox	Interfaces	Relationships	Tickets	Redirects	Notes	
Device Name	MAILSRV	Managed Type	Physical Device					
IP Address / ID	10.20.0.185 21	Category	Servers					
Class	Microsoft	Sub-Class	Windows 2000 Server					
Organization	System	Uptime	0 days, 00:00:00					
Collection Mode	Active	Collection Time	2014-06-18 14:15:00					
Description	Hardware: x86 Family 6 Model 11 Stepping 1 AT/AT COMPATIBLE - So	Group / Collector	CUG em7_ao_205					
Device Hostname								

Device Properties		Organization	Asset
		Actions	Reset
		Reset	Guide
Identification			
Device Name	MAILSRV	IP Address	[10.20.0.185 - verified]
Organization	[System]		
Monitoring & Management		Preferences	
Device Class	Microsoft Windows 2000 Server	Auto-Clear Events	<input checked="" type="checkbox"/>
SNMP Read/Write	[Cisco SNMPv2 - Example] [None]	Accept All Logs	<input checked="" type="checkbox"/>
Availability Port	[UDP] [161 - SNMP]	Daily Port Scans	<input checked="" type="checkbox"/>
Latency Port	[ICMP] [ICMP]	Auto-Update	<input checked="" type="checkbox"/>
Avail+Latency Alert	[Disable]	Scan All IPs	<input type="checkbox"/>
User Maintenance	[Disabled] [Maintenance Collection Enabled]	Dynamic Discovery	<input checked="" type="checkbox"/>
Collection	[Enabled] [CUG]	Preserve Hostname	<input checked="" type="checkbox"/>
Coll. Type	[Standard]	Disable Asset Update	<input type="checkbox"/>
Critical Ping	[Disabled]		
Dashboard	[None]		
Event Mask	[Group in blocks every 10 minutes]		
Save			

- In the **Device Properties** page, you can select a Collector Group from the **Collection** fields.
- 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 > Device Manager).

Device Name	IP Address	Device Gateway	Device Class Sub-class	DID	Organisation	Current State	Collection Group	Collection Status	SNMP Capable	SNMP Enable
10 Forward	10.20.0.195	Servers	NET-SNMP FreeBSD	54	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.108	10.20.0.108	Network Router	Cisco Systems 2501	72	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.123	10.20.0.123	Network Router	Cisco Systems 7206VXR	112	System	Minor	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.13	10.20.0.13	Unknown	Generic SNMP	107	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.135	10.20.0.135	Network Switches	Cisco Systems Catalyst 3509G-XL	131	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.141	10.20.0.141	Network Switches	Cisco Systems Catalyst WS-C6009-CatOS	118	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.146	10.20.0.146	Network Broadbar	Netopia Netopia 3346 v8 2r1	2	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.147	10.20.0.147	Network Broadbar	Netopia Netopia 3381 v8 0.10	175	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.148	10.20.0.148	Network Broadbar	Netopia Netopia (R3100, R4500, R7x00, R9	165	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.149	10.20.0.149	Network Broadbar	Netopia R7200-T	162	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.151	10.20.0.151	Unknown	Generic SNMP	141	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.160	10.20.0.160	Unknown	Generic SNMP	165	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.163	10.20.0.163	Unknown	Generic SNMP	164	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.175	10.20.0.175	Unknown	Generic SNMP	44	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.176	10.20.0.176	Unknown	Konica Corporation OEM	41	System	Minor	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.190	10.20.0.190	Unknown	Generic SNMP	56	System	Minor	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.191	10.20.0.191	Office Printers	Konica Minolta Fiery X3e Z2C-KM	57	System	Minor	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.201	10.20.0.201	Unknown	Generic SNMP	48	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.203	10.20.0.203	Unknown	Generic SNMP	52	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.209	10.20.0.209	Telephony	Quintum Tenor	53	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.222	10.20.0.222	Unknown	Generic SNMP	138	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.26	10.20.0.26	Unknown	Generic SNMP	171	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.52	10.20.0.52	Unknown	ASKEY Computer Corp. OEM	5	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.59	10.20.0.59	Unknown	Generic SNMP	3	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	
10.20.0.61	10.20.0.61	Unknown	Generic SNMP	84	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2	

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 SL1 system. When you create or edit an entry in the **Host File Entry Manager** page, SL1 automatically sends an update to every Data Collector in the specified Collector Group.

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


- The SL1 system does not reside in the end-customer's domain
- The SL1 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 SL1 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, SL1 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.

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.

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

```
INSERT INTO master.system_custom_config (`field`, `field_value`) VALUES ('enable_powershell_service', '0');
```

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).
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).
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', '0', '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).

Chapter



4

Daily Health Tasks

Overview

The tasks in this chapter help you monitor the health of your SL1 system. You can perform these tasks daily (or more frequently, if you require) to gather information about the overall status of your SL1 system.

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

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

This chapter includes the following topics:

What is a Healthy SL1 System?	90
Monitoring System Events	93
<i>Searching the System Logs</i>	94
<i>Deleting Entries from the System Logs</i>	95
Monitoring System Processes	96
<i>Viewing the List of System Processes</i>	96
<i>Searching and Filtering the List of System Processes</i>	97
Monitoring the Status of Each Appliance	98
Monitoring User Actions and Events	99
<i>Viewing the List of Audit Logs</i>	100
<i>Searching and Filtering the List of Audit Logs</i>	101
<i>Special Characters</i>	102
<i>Generating Reports on Audit Logs</i>	105

What is a Healthy SL1 System?

The following table presents a broad list of focus areas for SL1 health that are important to track for all SL1 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 SL1 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 SL1 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	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
		characteristics.		
Memory Consumption	Avoid swap usage	The MariaDB database will make use of available memory for caching purposes over time, but no SL1 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	Built-in HF rows-behind compared with HF object processing rate	Backlog time < 1 processing cycle

Focus Area	Check	Background	Specific Operation	Result
		should be completely processed 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.	SL1 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 SL1 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	Regular check of the system log	Daily maintenance tasks not being terminated due to an incomplete status.

Focus Area	Check	Background	Specific Operation	Result
		(scheduled nightly outside of core business hours) is to prune old data from the SL1 database, which is an essential activity for long term health.		
System Backup	Backups completing per schedule.	SL1 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.

Monitoring System Events

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

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

System Logs | Messages Found [1,317,344]

Purge Reset Guide

Search Message Search

Date	Module	Severity	Message
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-19: Non-digits in oid with object id: 8303, did:1336 in dynamic app:1148 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-19: Non-digits in oid with object id: 8305, did:1336 in dynamic app:1148 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-19: Non-digits in oid with object id: 8307, did:1336 in dynamic app:1148 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-19: Non-digits in oid with object id: 8308, did:1336 in dynamic app:1148 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-29: Non-digits in oid with object id: 8303, did:1331 in dynamic app:1148 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-29: Non-digits in oid with object id: 8307, did:1331 in dynamic app:1148 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-22: Non-digits in oid with object id: 8303, did:1330 in dynamic app:1148 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-20: Non-digits in oid with object id: 8331, did:1338 in dynamic app:1158 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-22: Non-digits in oid with object id: 8303, did:1330 in dynamic app:1148 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-20: Non-digits in oid with object id: 8332, did:1338 in dynamic app:1158 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-3: Non-digits in oid with object id: 8303, did:1325 in dynamic app:1148 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-22: Non-digits in oid with object id: 8305, did:1330 in dynamic app:1148 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-22: Non-digits in oid with object id: 8307, did:1330 in dynamic app:1148 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-22: Non-digits in oid with object id: 8307, did:1330 in dynamic app:1148 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-22: Non-digits in oid with object id: 8308, did:1330 in dynamic app:1148 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-22: Non-digits in oid with object id: 8310, did:1330 in dynamic app:1148 when updating performance value (val:2.666666666666667E-5)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-1: Non-digits in oid with object id: 8331, did:1328 in dynamic app:1158 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-3: Non-digits in oid with object id: 8307, did:1325 in dynamic app:1148 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-3: Non-digits in oid with object id: 8310, did:1325 in dynamic app:1148 when updating performance value (val:2.666666666666667E-5)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-1: Non-digits in oid with object id: 8332, did:1328 in dynamic app:1158 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-17: Non-digits in oid with object id: 8303, did:1337 in dynamic app:1148 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-17: Non-digits in oid with object id: 8307, did:1337 in dynamic app:1148 when updating performance value (val:None)
2012-03-19 11:12:43	em7_ao	Minor	11.PoolWorker-17: Non-digits in oid with object id: 8310, did:1337 in dynamic app:1148 when updating performance value (val:2.666666666666667E-5)
2012-03-19 11:11:58	em7_ao	Minor	11.PoolWorker-12: Non-digits in oid with object id: 8348, did:1312 in dynamic app:1163 when updating performance value (val:None)
2012-03-19 11:11:58	em7_ao	Minor	11.PoolWorker-12: Non-digits in oid with object id: 8349, did:1312 in dynamic app:1163 when updating performance value (val:None)

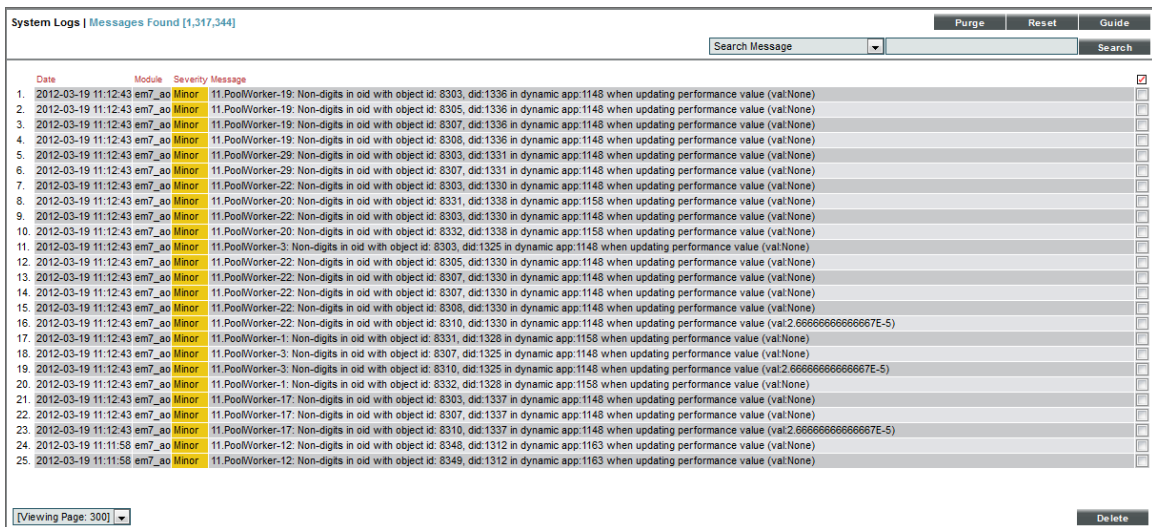
[Viewing Page: 300] Delete

2. In 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. To search the **System Logs**:

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



2. The search fields at the top of the **System Logs** page allows 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.

3. 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) allows you to view read-only information about the execution of SL1's system processes. System Processes gather, manipulate, and publish the data used in SL1. These system processes can be configured and debugged in the **Process Manager** page (System > Settings > Admin Processes).

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 SL1.
 - **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 SL1 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 SL1 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.

Searching and Filtering the List of System Processes

The **System Processes** page includes ten filters. You can filter the list of processes by one or multiple of the following parameters: appliance, process name, start time, end time, duration, frequency, percent, max instances, processed, and errors. Only processes that meet all the filter criteria will be displayed in the **System Processes** page.

You can filter by one or more of the following parameters. The list of system processes is dynamically updated as you select each filter.

- For eight of the filters, you must enter text to match against. The user interface will search for processes that match the text, including partial matches. Text matches are not case sensitive. You can use the following special characters in each filter:
 - `,` Specifies an "or" operation. For example:
 - "dell, micro" would match all values that contain the string "dell" OR the string "micro".
 - `!` Specifies a "not" operation. For example:
 - "!dell" would match all values that do not contain the string "dell".
- **Appliance**. You can enter text to match, including special characters, and the **System Processes** page will display only processes that have a matching appliance name.
- **Process**. You can enter text to match, including special characters, and the **System Processes** page will display only processes that have a matching process name.
- **ID**. You can enter text to match, and the **System Processes** page will display only processes that have a matching ID.
- **Start Time**. Only those processes that match all the previously selected fields and have the specified start date and time will be displayed. The choices are:
 - *All*. Display processes with all start dates and times.
 - *Last Minute*. Display only processes that started within the last minute.
 - *Last Hour*. Display only processes that started within the last hour.
 - *Last Day*. Display only processes that started within the last day.
 - *Last Week*. Display only processes that started within the last week.
 - *Last Month*. Display only processes that started within the last month.
 - *Last Year*. Display only processes that started within the last year.
- **End Time**. Only those processes that match all the previously selected fields and have the specified end date and time will be displayed. The choices are:
 - *All*. Display processes with all end dates and times.
 - *Last Minute*. Display only processes that ended within the last minute.
 - *Last Hour*. Display only processes that ended within the last hour.
 - *Last Day*. Display only processes that ended within the last day.

- *Last Week*. Display only processes that ended within the last week.
 - *Last Month*. Display only processes that ended within the last month.
 - *Last Year*. Display only processes that ended within the last year.
- **Duration**. You can enter text to match, including special characters, and the **System Processes** page will display only processes that have a matching duration.
 - **Frequency**. You can enter text to match, including special characters, and the **System Processes** page will display only processes that have a matching frequency.
 - **Percent**. You can enter text to match, including special characters, and the **System Processes** page will display only processes that have a matching percent.
 - **Instances**. This field is not currently in use. It is not recommended to filter the System Processes by this field.
 - **Max Instances**. You can enter text to match, including special characters, and the **System Processes** page will display only processes that have a matching number of Max Instances.
 - **Processed**. You can enter text to match, including special characters, and the **System Processes** page will display only processes that have a matching number of records processed.
 - **Errors**. You can enter text to match, including special characters, and the **System Processes** page will display only processes that have a matching number of errors.

Monitoring the Status of Each Appliance

The **Appliance Manager** page allows you to view information, including license status, about each appliance in your system.

From the **Appliance Manager** page, you can also access the Web Configuration Utility for each appliance and the database administration tool for the Database Server.



To view the **Appliance Manager** page:

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

The screenshot shows the 'Appliance Manager' interface. At the top right are 'Reset' and 'Guide' buttons. The main area contains a form with fields for 'Host Name', 'IP Address', 'Model type' (a dropdown menu currently showing 'Database [2]'), and 'Description'. There are 'Save' and 'Save As' buttons to the right of the form. Below the form is a table listing existing appliances.

	Name	IP Address	Module Type	Collector Group	Description	Build	Capacity	Allocation	ID	Validated	Edit Date	Edit User	Create Date	
1.	em7_db	10.0.9.90	Database	n/a	EM7 Database	18872	500	n/a	1	Yes	2012-03-20 19:08:08	em7admin	2012-03-20 18:55:16	
2.	em7_ap	10.0.9.50	Administration Portal	n/a	Application Server: 10.0.9.50	18872	n/a	n/a	3	Yes	2012-03-20 19:07:14	em7admin	2012-03-20 19:07:14	
3.	em7_cu1	10.0.9.54	Data Collection Unit	CUG1		18872	n/a	352	6	Yes	2012-03-20 19:10:05	em7admin	2012-03-20 19:10:04	
4.	em7_cu2	10.0.9.56	Data Collection Unit	CUG2		18872	n/a	148	5	Yes	2012-03-20 19:10:05	em7admin	2012-03-20 19:10:02	
5.	em7_mc	10.0.9.53	Message Collection Unit	--		18872	n/a	n/a	7	Yes	2012-03-20 19:10:20	em7admin	2012-03-20 19:10:05	
6.	em7_is	10.0.9.58	Integration Server	n/a	Integration Server: 10.0.9.58	18872	n/a	n/a	2	Yes	2012-03-20 19:06:18	em7admin	2012-03-20 19:06:18	

2. The **Appliance Manager** page displays the following about each appliance:

- **Name**. Name of the appliance.
 - **IP Address**. Primary IP address for the appliance.
 - **Module Type**. Type of appliance. Choices are:
 - All-In-One Appliance
 - Database Server
 - Administration Portal
 - Data Collector
 - Message Collector
 - Integration Server
 - **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.
 - **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.
 - **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.
3. For Database Servers, you can click the gear icon () to access the phpMyAdmin interface for the Database Server. In this interface, you can view all the database tables on the Database Server.
 4. 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.

Monitoring User Actions and Events

The **Audit Logs** page provides an audit trail for SL1. The **Audit Logs** page displays a record of actions in SL1 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 logs in to SL1
- Organization name changes
- The addition, editing, or deletion of elements in SL1

2. The **Audit Logs** page displays all actions that are performed by users or managed elements in SL1. For each action, the **Audit Logs** page displays:
 - **Date/Time**. Date and time the action occurred and the log entry was created.
 - **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."
 - **Organization**. Organization associated with the action.
 - **Message**. Text of the log entry.

Searching and Filtering the List of Audit Logs

The Filter-While-You-Type fields appear as a row of blank fields at the top of the list. These fields allow you to filter the items that appear in the list.

The list is dynamically updated as you select each filter. For each filter, you must make a selection from a drop-down menu or type text to match against. SL1 will search for entries that match the text, including partial matches. Text matches are not case-sensitive, and you can use special characters in each text field.

By default, the cursor is placed in the first Filter-While-You-Type field. You can use the <Tab> key or your mouse to move your cursor through the fields.

You can filter by one or more of the following parameters. Only items that meet all of the filter criteria are displayed on the page.

The following describes each filter on the **Audit Logs** page:

- **Date/Time**. Only those audit logs that have the specified creation date will be displayed. The choices are:
 - *All*. Display all audit logs that match the other filters.
 - *Last Minute*. Display only audit logs that have been created within the last minute.
 - *Last Hour*. Display only audit logs that have been created within the last hour.
 - *Last Day*. Display only audit logs that have been created within the last day.
 - *Last Week*. Display only audit logs that have been created within the last week.
 - *Last Month*. Display only audit logs that have been created within the last month.
 - *Last Year*. Display only audit logs that have been created within the last year.
- **Source**. You can enter text to match, including special characters, and the Audit Logs page will display only audit logs that have a matching source.
- **Organization**. You can enter text to match, including special characters, and the Audit Logs page will display only audit logs that have a matching organization.
- **Message**. You can enter text to match, including special characters, and the Audit Logs page will display only audit logs that have a matching message.

Special Characters

You can include the following special characters to filter by each column except those that display date and time:

NOTE: When searching for a string, SL1 will match substrings by default, even if you do not include any special characters. For example, searching for "hel" will match both "hello" and "helicopter". When searching for a numeric value, SL1 will not match a substring unless you use a special character.

String and Numeric

- , (comma). Specifies an "OR" operation. Works for string and numeric values. For example:
"dell, micro" matches all values that contain the string "dell" OR the string "micro".
- & (ampersand). Specifies an "AND" operation. Works for string and numeric values. For example:
"dell & micro" matches all values that contain both the string "dell" AND the string "micro", in any order.
- ! (exclamation point). Specifies a "not" operation. Works for string and numeric values. For example:
"!dell" matches all values that do not contain the string "dell".
"! ^ micro" would match all values that do not start with "micro".
"!fer\$" would match all values that do not end with "fer".
"! ^ \$" would match all values that are not null.
"! ^ " would match null values.
"!\$" would match null values.
"!*" would match null values.
"happy, !dell" would match values that contain "happy" OR values that do not contain "dell".

NOTE: You can also use the "!" character in combination with the arithmetic special characters (min-max, >, <, >=, <=, =) described below.

- * (asterisk). Specifies a "match zero or more" operation. Works for string and numeric values. For a string, matches any string that matches the text before and after the asterisk. For a number, matches any number that contains the text. For example:

"hel*er" would match "helpers" and "helicopter" but not "hello".

"325*" would match "325", "32561", and "325000".

"*000" would match "1000", "25000", and "10500000".

- ? (question mark). Specifies "match any one character". Works for string and numeric values. For example:

"l?ver" would match the strings "oliver", "levers", and "lover", but not "believer".

"135?" would match the numbers "1350", "1354", and "1359", but not "135" or "13502"

String

- ^ (caret). For strings only. Specifies "match the beginning". Matches any string that begins with the specified string. For example:

"^sci" would match "scientific" and "sciencelogic", but not "conscious".

"^happy\$" would match only the string "happy", with no characters before or after.

"!^micro" would match all values that do not start with "micro".

"!^\$" would match all values that are not null.

"!^" would match null values.

- \$ (dollar sign). For strings only. Specifies "match the ending". Matches any string that ends with the specified string. For example:

"ter\$" would match the string "renter" but not the string "terrific".

"^happy\$" would match only the string "happy", with no characters before or after.

"!fer\$" would match all values that do not end with "fer".

"!^\$" would match all values that are not null.

"!\$" would match null values.

NOTE: You can use both ^ and \$ if you want to match an entire string and only that string. For example, "^tern\$" would match the strings "tern" or "Tern" or "TERN"; it would not match the strings "terne" or "cistern".

Numeric

- min-max. Matches numeric values only. Specifies any value between the minimum value and the maximum value, including the minimum and the maximum. For example:

"1-5" would match 1, 2, 3, 4, and 5.

- - (dash). Matches numeric values only. A "half open" range. Specifies values including the minimum and greater or including the maximum and lesser. For example:

"1-" matches 1 and greater. So would match 1, 2, 6, 345, etc.

"-5" matches 5 and less. So would match 5, 3, 1, 0, etc.

- > (greater than). Matches numeric values only. Specifies any value "greater than". For example:

">7" would match all values greater than 7.

- < (less than). Matches numeric values only. Specifies any value "less than". For example:

"<12" would match all values less than 12.

- >= (greater than or equal to). Matches numeric values only. Specifies any value "greater than or equal to". For example:

">=7" would match all values 7 and greater.

- <= (less than or equal to). Matches numeric values only. Specifies any value "less than or equal to". For example:

"<=12" would match all values 12 and less.

- = (equal). Matches numeric values only. For numeric values, allows you to match a negative value. For example:

"=-5" would match "-5" instead of being evaluated as the "half open range" as described above.

Additional Examples

- "aio\$". Matches only text that ends with "aio".
- "^shu". Matches only text that begins with "shu".
- "^silo\$". Matches only the text "silo", with no characters before or after.
- "!silo". Matches only text that does not contain the characters "silo".
- "!^silo". Matches only text that does not start with "silo".
- "!O\$". Matches only text that does not end with "O".
- "!^silo\$". Matches only text that is not the exact text "silo", with no characters before or after.
- "!"^". Matches null values, typically represented as "--" in most pages.
- "!"\$". Matches null values, typically represented as "--" in most pages.

- "!"^\$". Matches all text that is not null.
- "silo, !aggr". Matches text that contains the characters "silo" and also text that does not contain "aggr".
- "silo, 02, !aggr". Matches text that contains "silo" and also text that contains "02" and also text that does not contain "aggr".
- "silo, 02, !aggr, !01". Matches text that contains "silo" and also text that contains "02" and also text that does not contain "aggr" and also text that does not contain "01".
- "^s*i!*o\$". Matches text that contains the letter "s", "i", "l", "o", in that order. Other letters might lie between these letters. For example "sXiXo" would match.
- "!^s*i!*o\$". Matches all text that does not that contains the letter "s", "i", "l", "o", in that order. Other letters might lie between these letters. For example "sXiXo" would not match.
- "!vol&!silo". Matches text that does not contain "vol" AND also does not contain "silo". For example, "volume" would match, because it contains "vol" but not "silo".
- "!vol&02". Matches text that does not contain "vol" AND also contains "02". For example, "happy02" would match, because it does not contain "vol" and it does contain "02".
- "aggr, !vol&02". Matches text that contains "aggr" OR text that does not contain "vol" AND also contains "02".
- "aggr, !vol&!infra". Matches text that contains "aggr" OR text that does not contain "vol" AND does not contain "infra".
- "*". Matches all text.
- "!*". Matches null values, typically represented as "--" in most pages.
- "silo". Matches text that contains "silo".
- "!silo". Matches text that does not contain "silo".
- "!^silo\$". Matches all text except the text "silo", with no characters before or after.
- "-3,7-8,11,24,50-". Matches numbers 1, 2, 3, 7, 8, 11, 24, 50, and all numbers greater than 50.
- "-3,7-8,11,24,50-,a". Matches numbers 1, 2, 3, 7, 8, 11, 24, 50, and all numbers greater than 50, and text that includes "a".
- "?n". Matches text that contains any single character and the character "n". For example, this string would match "an", "bn", "cn", "1n", and "2n".
- "n*SAN". Matches text the contains "n", zero or any number of any characters and then "SAN". For example, the string would match "nSAN", and "nhamburgerSAN".
- "^?n*SAN\$". Matches text that begins with any single character, is following by "n", and then zero or any number of any characters, and ends in "SAN".

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, SL1 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]**.

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).

Collector Status										Refresh	Guide
	Collector Name	Collector ID	Collector Address	Group ID	Group name	Last State Change	Collector State	Sync State			
1	MOSS_PATCH_MC	6	10.2.3.9	--	--	--	Available [0]	In Sync [0]			
2	MOSS_PATCH_CUG2	4	10.2.3.8	1	CUG2	2015-08-05 10:43:09	Available [0]	In Sync [0]			
3	MOSS_PATCH_CUG1	5	10.2.3.7	2	CUG1	2015-10-20 14:18:37	Available [0]	In Sync [0]			
4	MOSS_CUG3	10	10.2.3.12	5	CUG3	2015-10-20 13:56:28	Available [0]	In Sync [0]			

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 SL1.
- **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.
- **Sync State.** Specifies whether the Data Collector or Message Collector is in synch with the latest configuration data on the Database Server.

Chapter



5

Upgrading SL1

Overview

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

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

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

This chapter includes the following topics:

<i>Reference: The System Updates Page</i>	112
<i>Reference: Viewing the List of Updates</i>	112
<i>Reference: Viewing the Log Files for Updates</i>	114
<i>Workflow</i>	114
<i>Planning</i>	115
<i>Scheduling Maintenance Windows</i>	116
<i>Updating SL1 Appliances to Oracle Linux</i>	117
<i>Pre-Upgrade Best Practices</i>	117
<i>Backing Up Settings in the NextUI</i>	118
<i>Backing Up SSL Certificates</i>	118
<i>Setting the Timeout for PhoneHome Watchdog</i>	118
<i>Adjusting the Timeout for Staging and Deploying</i>	119

<i>Running the System Status Script Before Upgrading</i>	120
<i>Upgrading the SL1 Distributed Architecture on SL1 Versions 8.5.0 and Earlier</i>	122
<i>Upgrading the SL1 Distributed Architecture on SL1 versions 8.6.0 and Later</i>	132
<i>Upgrading the Extended Architecture</i>	148
<i>Manual Steps for Updates to 8.4.x and Earlier Systems</i>	151
<i>Automatically Upgrading MariaDB with a Script</i>	152
<i>Manually Upgrading MariaDB</i>	154
<i>Rebooting Appliances in the SL1 Distributed Stack</i>	163
<i>Restoring Settings for NextUI</i>	169
<i>Restoring the SSL Certificate</i>	169
<i>Resetting the Timeout for PhoneHome Watchdog</i>	169
<i>Updating Default PowerPacks</i>	170
<i>Configuring Subscription Billing</i>	172
<i>Reference: The System Updates Page</i>	112
<i>Reference: Viewing the List of Updates</i>	112
<i>Reference: Viewing the Log Files for Updates</i>	114
<i>Workflow</i>	114
<i>Planning</i>	115
<i>Scheduling Maintenance Windows</i>	116
<i>Updating SL1 Appliances to Oracle Linux</i>	117
<i>SL1 Release Prior to 8.1.1</i>	117
<i>SL1 Releases Prior to 8.10.0</i>	117
<i>Pre-Upgrade Best Practices</i>	117
<i>Backing Up Settings in the NextUI</i>	118
<i>Backing Up SSL Certificates</i>	118
<i>Setting the Timeout for PhoneHome Watchdog</i>	118
<i>Adjusting the Timeout for Staging and Deploying</i>	119
<i>SL1 8.14 and Later Releases</i>	119
<i>SL1 8.12 and Prior Releases</i>	119
<i>Running the System Status Script Before Upgrading</i>	120
<i>Running the System Status Script on SL1 8.14.0 and Later Releases</i>	120
<i>Running the System Status Script on SL1 8.12.x Releases</i>	120
<i>Running the System Status Script on SL1 8.10 and Prior Releases</i>	121
<i>Upgrading the SL1 Distributed Architecture on SL1 Versions 8.5.0 and Earlier</i>	122
<i>Disabling Automatic Staging</i>	123

<i>Downloading the Updates</i>	124
<i>Importing the Updates</i>	124
<i>Staging the Update</i>	125
<i>Running the Pre-Upgrade Check</i>	127
All SL1 Appliances:	127
Active Database Server:	127
All Database Servers:	128
Administration Portal:	128
Data Collectors and Message Collectors:	128
Downloading and Running the Pre-Upgrade Check	128
<i>Putting All SL1 Appliances in Maintenance Mode</i>	129
<i>Deploying the Update</i>	130
<i>Putting All SL1 Appliances Out of Maintenance Mode</i>	131
<i>Performing Deltaless Upgrades</i>	132
Upgrading the SL1 Distributed Architecture on SL1 versions 8.6.0 and Later	132
<i>Special Steps for SL1 8.12.0 and Earlier</i>	133
<i>Downloading the Update</i>	134
<i>Importing the Update</i>	135
<i>Staging the Update</i>	136
Automatic Staging	136
Manually Staging an Update	137
<i>Running the Pre-Upgrade Check for SL1 10.1 and Later</i>	139
Running the Pre-Upgrade Check	139
CentOS 5 Failure	140
<i>Running the Pre-Upgrade Check for SL1 8.14 and Earlier</i>	142
All SL1 Appliances:	142
Active Database Server:	143
All Database Servers:	143
Administration Portal:	143
Data Collectors and Message Collectors:	144
Downloading and Running the Pre-Upgrade Check	144
<i>Putting All SL1 Appliances in Maintenance Mode</i>	145
<i>Deploying the Update</i>	145
<i>Putting All SL1 Appliances Out of Maintenance Mode</i>	147
Upgrading the Extended Architecture	148
Prerequisites	148

Resizing the Disks on the Compute Node	148
Upgrade Steps	149
Manual Steps for Updates to 8.4.x and Earlier Systems	151
Automatically Upgrading MariaDB with a Script	152
Manually Upgrading MariaDB	154
Download RPMs to SL1 Appliances	154
Manually Upgrade Two Database Servers Configured for High Availability or Disaster Recovery	155
Step 1: On the Secondary Database Server	155
Step 2: On the Primary Database Server	156
Step 3: On the Secondary Database Server	158
Manually Upgrade Three Database Servers Configured for High Availability and Disaster Recovery	159
Step 1: On the Secondary Database Server	159
Step 2: On the Primary Database Server	159
Step 3: On the Secondary Database Server	160
Step 4: On the Disaster Recovery Database Server	161
Manually Upgrading Standalone Database Servers, All-In-One Appliances, Data Collectors, and Message Collectors	162
Rebooting Appliances in the SL1 Distributed Stack	163
Rebooting the Administration Portal	163
Rebooting Multiple Administration Portals	163
Rebooting a Single Administration Portal	164
Rebooting Data Collectors and Message Collectors	164
Rebooting Data Collectors and Message Collectors from the Appliance Manager page	164
Rebooting Data Collectors and Message Collectors from the Command Line	165
Rebooting Standalone All-In-One Appliance and Standalone Database Server	165
Rebooting Two Database Servers Configured for Disaster Recovery	165
Rebooting Two Database Servers in a High Availability Cluster	166
Rebooting Three Database Servers Configured for High Availability and Disaster Recovery	167
Restoring Settings for NextUI	169
Restoring the SSL Certificate	169
Resetting the Timeout for PhoneHome Watchdog	169
Updating Default PowerPacks	170
Configuring Subscription Billing	172

Reference: The System Updates Page

The **System Updates** page (System > Tools > Updates) allows you to update the software on your SL1 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 SL1 system, the SL1 system can automatically **stage** the software update. Staging is when the software is copied to each ScienceLogic appliance. Staging allows SL1 to simultaneously apply the software changes to each ScienceLogic appliance, regardless of the speed of the connection to each ScienceLogic appliance.

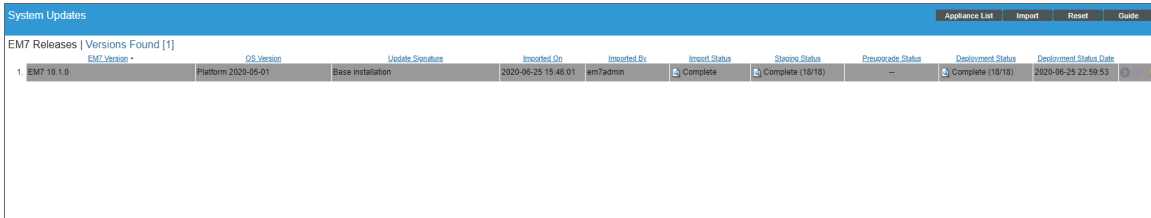
You can allow the SL1 system to automatically stage the software or you can manually stage the software.

After the software update is staged, you can 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 SL1 systems, users might have to create a Collector Group for a single Data Collector before applying updates.

Reference: Viewing the List of Updates

The **System Updates** page (System > Tools > Updates) displays the following about each update:



The screenshot shows the 'System Updates' page with a table of updates. The table has columns for EM7 Version, OS Version, Update Signature, Imported On, Imported By, Import Status, Staging Status, Prepopulate Status, Deployment Status, and Deployment Status Date. A single update is listed with the following details:


EM7 Version	OS Version	Update Signature	Imported On	Imported By	Import Status	Staging Status	Prepopulate Status	Deployment Status	Deployment Status Date
EM7 10.1.0	Platform 2020-09-01	Base Installation	2020-09-25 15:46:01	em7admin	Complete	Complete (18/18)	--	Complete (18/18)	2020-09-25 22:59:53


TIP: To sort the list of update files, 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 **Deployed Status Date** column sorts by descending order on the first click; to sort by ascending order, click the column heading again.

- **EM7 Version**. Name and version number for the software update.
- **OS Version**. Name and number of the platform OS update.

- **Update Signature.** Name of the entity that released the update and type of update. Usually "ScienceLogic Official Release".
- **Imported On.** Date and time the software update was loaded onto the SL1 system.
- **Imported By.** Name of the ScienceLogic user who loaded the software update onto the SL1 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 SL1 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 SL1 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.
- **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 SL1 appliances.
 - *Scheduled.* The SL1 system 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.
 - *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: If you did not select **Auto Stage** during import, the **Staging Status** column will include an asterisk (*) until you manually stage the update.

- **Preupgrade Status.** You can run the pre-upgrade check after importing and staging an update but before deploying the update. The pre-upgrade check will ensure that all criteria are met before deploying.
 - If you want to run the pre-upgrade check, select the purple checkmark for the selected row.
 - The possible values in this field are *In Progress* or *Complete*.
 - Clicking on the magnifying-glass icon () in this column displays the output of the pre-upgrade check.
 - If a pre-upgrade criterion fails, the **[Deploy]** button will be disabled for the selected row.

- If an appliance fails the pre-upgrade criteria, you can view the output from the system status script for each failed appliance. Go to the Appliance Manager page (System > Settings > Appliances), find the appliance that failed, and click on the magnifying-glass icon ().
- For details see [the section on the Pre-Upgrade Check](#) .
- **Deployment Status.** Specifies the current deployment state. Possible values are:
 - --. No deployment request is active, and software has not been deployed on any SL1 appliances.
 - *Scheduled.* The SL1 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 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.
- **Deployment Status Date.** Specifies the date and time the software update was last deployed.

Reference: Viewing the Log Files for Updates

From the **System Updates** page, you can view a log file that displays the history of the software update. To view this log file:

1. Go to the **System Updates** page (System > Tools > Updates).
2. In the **System Updates** page, find the software update for which you want to view the log files. Go to its **Import Status** column, **Staging Status** column, or **Deployment Status** column and click the log icon (.
3. The appropriate log page appears. In this modal page, each log entry displays:
 - Information about the status of the software update and its related actions.
 - For each action, the name and IP address of the appliance where the action occurred
 - The date and time each action occurred.

Workflow

The following sections describe the steps to plan and deploy an SL1 update.

If you would like assistance planning an upgrade path that minimizes downtime, please contact your Customer Success Manager.

The workflow for upgrading SL1 is:

1. [Planning](#)
2. [Scheduling maintenance windows](#)
3. (If necessary) [updating SL1 Appliances to a versions of SL1 that uses RedHat Linux \(SL1 8.1.1 and later\)](#).
4. [Pre-upgrade best practices for SL1](#)
5. [Backing Up Custom Settings in the NextUI](#)
6. [Backing Up SSL Certificates](#)
7. [Setting the Timeout for PhoneHome Watchdog](#)
8. [Adjusting Timeout for slow connections](#)
9. [Running the system status script](#) on the Database Server or All-In-One before upgrading.
10. (If necessary) [Upgrading the SL1 Distributed Architecture on SL1 Versions 8.5.0 and Earlier](#) using the System Update tool (System > Tools > Updates).
11. [Upgrading the SL1 Distributed Architecture on SL1 Versions 8.6.0 and Later](#) using the System Update tool (System > Tools > Updates).
12. [Removing SL1 Appliances from Maintenance Mode](#)
13. (Optional) [Upgrading the Extended Architecture](#)
14. (As needed) [Manual updates for systems upgrading from 8.4.x and earlier](#)
15. (As needed) [Upgrading MariaDB](#)

CAUTION: Refer to the release notes for your current release to determine if you must upgrade MariaDB.

16. (As needed) [Rebooting SL1 Appliances](#)

CAUTION: Refer to the release notes for your current release to determine if you must reboot all SL1 appliances after upgrading.

17. [Restoring Custom Settings in the NextUI](#)
18. [Restoring SSL Certificates](#)
19. [Resetting the Timeout for PhoneHome Watchdog](#)
20. [Updating PowerPacks](#)
21. (One-time) [Configure Subscription Billing](#). For details, see the **Subscription Billing** manual.

Planning

Before upgrading SL1 , 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?
2. Read the Known Issues for the release at <https://support.sciencelogic.com/s/topic/0TO0z000000E6w7GAC>
3. Identify all integrations and third-party applications that access the SL1 database or manipulate data on SL1. 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 custom work (PowerPacks, Run Book Automations, Event policies, Dashboard widgets) and ensure that it is backed up so you can restore it if necessary.

Scheduling Maintenance Windows

Upgrading SL1 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 SL1 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 SL1 appliances in maintenance mode. The SL1 system will not be available during this procedure. Update both the SL1 Distributed systems and SL1 Extended systems (if applicable).

CAUTION: If you are upgrading from a version of SL1 prior to 8.6.0, you will have to perform the first two bullets twice, once to upgrade to 8.6.0 and then again to upgrade to the latest update. You will need an additional maintenance window for the extra deploy step.

- **Update MariaDB (if required).** The SL1 system will not be available during this procedure.

CAUTION: Refer to the release notes for your current release to determine if you must upgrade MariaDB.

- **Reboot Appliances (if required).** Individual SL1 appliances will not be available during these procedures.

CAUTION: Refer to the release notes for your current release to determine if you must reboot all SL1 appliances after upgrading.

- Identify activities and customers that will be affected by maintenance windows and schedule and inform appropriately.

Updating SL1 Appliances to Oracle Linux

SL1 Release Prior to 8.1.1

SL1 8.1.1 included a complete update of the operating system for each SL1 appliance, from CentOS 5.11 to Oracle Linux. Major operating system components, including the database, web server, and High Availability/Disaster Recovery packages have been updated or replaced by new industry-standard packages.

When upgrading from a version prior to 8.1.1, each appliance must be migrated to 8.9.0 and the Oracle Linux 7.5 operating system.

ScienceLogic strongly suggests that you contact Customer Support or your Customer Success Manager to plan your migration from CentOS (versions of SL1 prior to 8.1.1) to the latest release.

SL1 Releases Prior to 8.10.0

Please note that SL1 8.10.0 and later versions do not support "mixed-mode", where the Database Servers and Administration Portals are running Oracle Linux (SL1 versions 8.1.1 and later) and the Data Collectors and Message Collectors are running CentOS (SL1 version prior to 8.1.1). If your system uses "mixed-mode", you must reinstall the Data Collectors and Message Collectors with an SL1 ISO that matches the SL1 version running on the Database Servers or Administration Portals before upgrading to the current versions of SL1.

ScienceLogic strongly suggests that you contact Customer Support or your Customer Success Manager to plan your migration from mixed mode to SL1 versions 8.10.0 and later

Pre-Upgrade Best Practices

Before you upgrade, check the following:

- Before you install an SL1 upgrade, ScienceLogic recommends reviewing 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.sciencelogic.com/s/system-requirements>.
- Before installing an SL1 upgrade, ScienceLogic recommends that you verify that recent backups are available for your system.
- Ensure that each SL1 appliance has a valid license.

- To apply updates to an existing Data Collector, that Data Collector must be a member of a Collector Group. In some SL1 systems, users might have to create a Collector Group for a single Data Collector.
- Ensure that each Data Collector is "available" to the Database Server. To check, see the **Collector Status** page (System > Monitor > Collector Status).

Backing Up Settings in the NextUI

To save any custom settings in the NextUI:

1. Login to the console of the Database Server or SSH to the Database Server.
2. Open a shell session.
3. Enter the following at the shell prompt:

```
cp /opt/em7/nextui/nextui.env /opt/em7/nextui/nextui.env.backup
```

Backing Up SSL Certificates

To backup your SSL Certificates:

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

```
cp /etc/nginx/siloss1.key /etc/nginx/siloss1.key.bak
cp /etc/nginx/siloss1.pem /etc/nginx/siloss1.pem.bak
```

4. Repeat these steps on each Database Server in your SL1 system.

Setting the Timeout for PhoneHome Watchdog

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 SL1 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.

Adjusting the Timeout for Staging and Deploying

If you have slow connections between SL1 appliances, you can adjust the timeout values for staging and deploying upgrades.

SL1 8.14 and Later Releases

To do this on an SL1 system running 8.14.0 or later:

1. Log in to the console of the Database Server or SSH to the Database Server.
2. Open a shell session.
3. Enter 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 SL1 appliance. The default value is 1800 seconds (30 minutes). You can increase this value for slow connections.

4. Enter 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 SL1 appliance. The default value is 3600 seconds (1 hour). You can increase this value for slow connections.

SL1 8.12 and Prior Releases

To do this on SL1 systems running versions prior to 8.14.0:

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

```
silosql -e "UPDATE master.system_settings_patcher SET value=<timeout_in_
seconds> WHERE param='staging_wait_time'"
```

- where `<timeout_in_seconds>` is the timeout value, in seconds, for staging for each SL1 appliance. The default value is 1800 seconds (30 minutes). You can increase this value for slow connections.

4. Enter the following at the shell prompt:

```
silo_mysql -e "UPDATE master.system_settings_patcher SET value=<timeout_in_seconds> WHERE param='deploy_wait_time'"
```

- `<timeout_in_seconds>` is the timeout value, in seconds, for deploying to each SL1 appliance. The default value is 3600 seconds (1 hour). You can increase this value for slow connections.


Running the System Status Script Before Upgrading

SL1 includes a script, `system_status.sh`, that provides diagnostic data for each appliance in your SL1 system.

Running the System Status Script on SL1 8.14.0 and Later Releases

If you are running SL1 version SL 8.14.0 or later, SL1 automatically runs the system status script every 15 minutes on each appliance in your SL1 system.

For SL1 systems running 8.140 or later, ScienceLogic recommends that you view the output from the system status script before upgrading:

1. In SL1, go to the **Appliance Manager** page (System > Settings > Appliances).
2. Locate the SL1 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 appliance in your SL1 system.
6. To get the very latest status before upgrading, [manually run the system status script](#) on each Database Server or All-In-One Appliance.

Running the System Status Script on SL1 8.12.x Releases

If you are running SL1 version 8.12.0 or later, SL1 includes the system status script already installed in the `/opt/em7/bin` directory. However, there is no user interface option in the **Appliance Manager** page (System > Settings > Appliances) until SL1 version 8.14.0.

Before upgrading from any 8.12.x or later version, ScienceLogic recommends that you manually run the system status script on each Database Server or All-In-One Appliance, to get the latest results.

To run the system status script on each SL1 appliance:

1. Either go to the console of the SL1 appliance or use SSH to access the SL1 appliance.
2. Open a shell session on the server.
3. Navigate to `/opt/em7/bin`:


```
cd /opt/em7/bin
```

4. At the shell prompt, enter the following:

```
sudo system_status.sh > /tmp/status
```

5. Enter the root password.
6. Navigate to /tmp/status to view the results.

```
cd /tmp  
cat status
```

7. If the output includes errors and you need help fixing them, contact ScienceLogic Customer Support to fix the errors before upgrading.

Running the System Status Script on SL1 8.10 and Prior Releases

If you are running an SL1 release prior to 8.12.0, you can download and run the system status script before upgrading. You should run this script on each SL1 appliance.

To download and run the system_status.sh script:

1. Go to https://docs.sciencelogic.com/system_status/system_status.sh to download the system status script. Save the file to a local computer.
2. Using WinSCP or another file-transfer utility, copy the file system_status.sh to a directory on the SL1 appliance.
3. Go to the console of the SL1 appliance or use SSH to access the SL1 appliance.
4. Navigate to the directory that includes the system_status.sh file.
5. Enter the following to make the script executable:

```
chmod +x system_status.sh.
```

6. To execute the script with the default options:

```
sudo ./system_status.sh
```

7. You can use the help option to view all the options:

```
./system_status.sh --help
```

You should see something like this:

```
-l [--logcollect] = gather logs into a single file  
-c [--crmreport] = create CRM report for HA/DR issues  
-f [--fix] = fix mode (fix anything marked [fixable])  
-m [--mysql-tuning] = a MySQL configuration tuning script (formerly known as  
"Gold Changes")  
-n [--no-upgrade] = Do not upgrade or prompt to upgrade, even if an upgrade is  
available  
-o [--outage] = outage investigator  
-p [--ptstalk] = run pt-stalk  
-h [--help] = help message (what you're reading now)  
-u [--autoupdate] = if an update is available, upgrade without prompting  
-v [--version] = show version number
```

Here is a sample output of the script run without options:

```
System Status v3.11
Tue Jul 17 15:39:43 UTC 2019
Using latest revision
Results marked with [KB ####] refer to a KB article number you can reference at
https://support.sciencelogic.com/s/article/####
Database Specs:
SL1 Architecture: AIO
SL1 Specs:
SL1 Release: EM7 8.10.0 [build 1023]
ISO Release: 8.1.1
Managed Devices: 2 (2 active)
Managed Components: 0 (0 active)
Managed Interfaces: 14 (3 active)
OS Specs:
OS Release: Oracle Linux Server release 7.5
Hostname: ha2
Active IPs:
- 192.168.##.##/22
Licensed IP: 192.168.##.##
Appliance Model: AIO
Appliance Type: Virtual - VMware Virtual Platform
MySQL Version: 10.1.36
Number of Cores: 2
Total RAM: 6GB
Disk Size: 60GB
- Disk size below minimum for 2 active devices, should be a minimum of 80GB
/data.local/db size: 27G
5 Minute Load Avg: 2.74
Memory Utilization: 65%
Other Information:
Disabled Processes [KB 1277]:
- Enterprise Database: Asset Record Maintenance
- Enterprise Database: Subscription Usage Crunch
Checking for errors...
- Devlogging ON [KB 1049]
- More than three kernels installed. [KB 1294]
- DRBD active, but DRBD SNMP extensions for monitoring not found [KB 1278]
[fixable]

All Done!
```

8. If the output includes errors and you need help fixing them, contact ScienceLogic Customer Support to fix the errors before upgrading.
9. Repeat these steps on each SL1 appliance in your SL1 system.

Upgrading the SL1 Distributed Architecture on SL1 Versions 8.5.0 and Earlier

To use the steps in this section, you must be running SL1 version 8.1.1 or later. If you are running a previous version, **ScienceLogic strongly suggests that you contact Customer Support or your Customer Success Manager to plan your migration from CentOS (versions of SL1 prior to 8.1.1) to the latest release.**

NOTE: The upgrade process might include importing multiple upgrade files. You must wait until an update file has imported successfully (i.e. the Import Status column displays Complete) before importing the next update file.

Any distributed system running 8.6.0 or later can be upgraded by importing, staging, and deploying a single update file. After you upgrade to 8.6.0, you can use the delta-less upgrade feature.

NOTE: The 8.4.2 release changed the firewall on all appliances from iptables to firewalld. If you have added a custom firewall rule, such as a non-standard port for Phone Home Collectors, these rules must be migrated before upgrading to an 8.4.2 or later release. Please contact ScienceLogic Support for more information.

Upgrading the Distributed Stack for SL1 includes the following steps:

- [Disabling automatic staging](#)
- [Downloading the Updates](#)
- [Importing the Updates](#)
- [Staging the Updates](#)
- [Running the Pre-Upgrade Check](#)
- [Putting All SL1 Appliances in Maintenance Mode](#)
- [Deploying the Update](#)
- [Putting All SL1 Appliances Out of Maintenance Mode](#)

CAUTION: SL1 versions 8.10.0 and later do not support Data Collectors and Message Collectors running the CentOS operating system. If your system includes Data Collectors and Message Collectors running the CentOS operating system, contact your Customer Success Manager for details on upgrading Data Collectors and Message Collectors to Oracle Linux before upgrading.

Disabling Automatic Staging

For systems running an SL1 version prior to 8.12.0, go to the **System Updates** page and disable automatic staging (System > Tools > Updates > Actions > Disable automatic staging).

If you have previously used manual staging, 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 Updates

NOTE: These steps do not affect the performance of SL1. ScienceLogic recommends that you perform these steps at least 3 days before upgrading.

To download updates for previous SL1 software versions that have reached their End of Life date and are no longer supported by ScienceLogic, contact ScienceLogic Support or a designated Customer Success Manager to get the update files.

Store the update files in a location that you can use to upload files to the SL1 system

Importing the Updates

NOTE: These steps do not affect the performance of SL1. ScienceLogic recommends that you perform these steps at least 3 days before upgrading.


If your system is not currently running a recent release, the upgrade process includes importing multiple update files. You must wait until an update file has imported successfully (i.e. the **Import Status** column displays *Complete* in both the **EM7 Releases** pane and the **ScienceLogic OS** pane) before importing the next update file.

Although you must import all the update files between your current release and SL1 8.6.0, you must stage only the final update.

To import an update:

1. Go to the **System Updates** page (System > Tools > Updates)

System Updates									
Installation mode: [Upgrade only]									
Actions: Import Reset Guide									
EM7 Releases Versions Found [2]									
Release Version	Update Signature	Imported On	Imported By	Import Status	Staging Status	Deployment	Deployment Status	Deployment Status Date	
EM7 8.2.0 jenkins_EM7_G3_8.2.0 r4269	ScienceLogic Internal QA Build	2016-12-04 16:39:35	em7admin	Missing Key	N/A	--	New/Unscheduled	--	
EM7 8.2.0	Base installation	2016-11-14 16:43:54	em7admin	Complete	N/A	Full (1/1)	New/Unscheduled	--	
ScienceLogic OS Updates Found [2]									
Release Version	Update Signature	Imported On	Imported By	Import Status	Staging Status	Deployment	Deployment Status	Deployment Status Date	
Platform 2016-12-01	ScienceLogic Internal QA Build	2016-12-04 16:39:35	em7admin	Missing Key	N/A	--	New/Unscheduled	--	
Platform 2016-11-11	Base installation	2016-11-14 16:43:54	em7admin	Complete	N/A	Outdated	New/Unscheduled	--	

2. Select the **[Import]** button.
3. In the **Import a new update** modal page, navigate to the software file for the update and select it. Select the **[Import]** button.
4. The **Import Status** column displays the status of the import. Possible values are:
 - *Processing*. Software update is currently being imported by the SL1 system.
 - *Complete*. Software update has been imported successfully.
 - *Failed*. Import has failed due to an unexpected condition. Contact ScienceLogic Support for assistance.
 - *Missing Base*. The SL1 system cannot import this software update until another software package has been imported. Check the log for a message stating which software package needs to be imported.
5. The update file will load an update in both the **EM7 Releases** pane and the **ScienceLogic OS** pane. You must wait until the update file has imported successfully (i.e. the **Import Status** column displays *Complete*) before importing the next update file.
6. For more detailed information about importing, in the **Import Status** column, select the log icon ().
7. Repeat these steps for each update file between your current version and 8.6.0

Staging the Update

NOTE: These steps do not affect the performance of SL1. ScienceLogic recommends that you perform these steps at least 3 days before upgrading.


When you manually stage a software update, SL1 checks the status of the software on each SL1 appliance. The platform then stages the software update only to those SL1 appliances that have not yet been staged for this software update.

You must stage only the latest update. For example, if you downloaded updates for 8.4.0, 8.5.0, and 8.6.0, you must stage only the update for 8.6.0.

To manually stage a software update:

1. Go to the **System Updates** page (System > Tools > Updates

System Updates									
Installation mode: [Upgrade only]									
Actions Import Reset Guide									
EM7 Releases Versions Found [2]									
Release Version	Update Signature	Imported On	Imported By	Import Status	Staging Status	Deployment	Deployment Status	Deployment Status Date	
1 EM7 8.2.0(jenkins_EM7_OS_8.2.0 r4269)	ScienceLogic Internal QA Build	2016-12-04 16:39:35	em7admin	Missing Key	N/A	--	New/Unscheduled	--	
2 EM7 8.2.0	Base installation	2016-11-14 16:43:54	em7admin	Complete	N/A	Full (1/1)	New/Unscheduled	--	
ScienceLogic OS Updates Found [2]									
Release Version	Update Signature	Imported On	Imported By	Import Status	Staging Status	Deployment	Deployment Status	Deployment Status Date	
1 Platform 2016-12-01	ScienceLogic Internal QA Build	2016-12-04 16:39:35	em7admin	Missing Key	N/A	--	New/Unscheduled	--	
2 Platform 2016-11-11	Base installation	2016-11-14 16:43:54	em7admin	Complete	N/A	Outdated	New/Unscheduled	--	

2. In the **EM7 Releases** pane, find the software update for EM7 8.6.0 . Select its checkbox.
3. In the **ScienceLogic OS** pane, find the latest OS update. Select its checkbox.
4. In the **Select Actions** field in the lower right, select *Stage Update (prepare for installation)*. Click the **[Go]** button.
5. The **Staging Status** column displays the status of the staging process. Possible values are:
 - -- Software is not currently being staged, but one or more SL1 appliances can install it.
 - N/A. This update cannot be installed in the current Installation mode.
 - *Obsolete*. The current update is not the latest or has already been installed. If the Installation Mode is set to Upgrade Only, you cannot install an obsolete update.
 - *Failed*. An unexpected error occurred in the staging process. Contact ScienceLogic Support.
 - *In progress*. Staging is in progress but has not completed.
 - *Complete*. Staging has completed and all appliances are ready to deploy the software.
6. For more detailed information about staging, in the **Staging Status** column, select the log icon ().
7. SL1 will copy the software update to each SL1 appliance that has not yet been staged.

Running the Pre-Upgrade Check

NOTE: These steps do not affect the performance of SL1. ScienceLogic recommends that you perform these steps at least 3 days before upgrading.

The pre-upgrade check examines the following:

All SL1 Appliances:

The pre-upgrade-checkv2.sh script checks the following on all SL1 appliances:

- Hostname and hosts file
- Uptime and average CPU load
- SL1 version
- OS versions
- MariaDB version
- Current disk utilization
- Is post update running?
- Checks yum config for configured proxy
- Check for packages installed outside of official SL1 update
- Gathers system status

Active Database Server:

The pre-upgrade-checkv2.sh script checks the following on all active Database Servers:

- Identifies Data Collectors that are not members of a Collector Group
- Lists all active SL1 appliances
- Lists all SL1 appliances that are disabled or ineligible for upgrade
- Checks for active staging schedules
- Check for active deployment schedules
- Are backups configured and successful?
- Checks pruner status
- Checks SIGTERMs for past 24 hours
- Compares number of tables on disk to the table definition cache
- Lists the largest 10 databases by size
- Searches for long running queries
- Shows status of used mysql connections and variables

- Lists mysql config files
- Gathers pt-diststats at 15 second interval three times
- Checks for system status wrapper bug
- Ensure that silo.log ownership is correct

All Database Servers:

The pre-upgrade-checkv2.sh script checks the following on all Database Servers:

- Checks HA/DR status
- Checks for crm bans
- Are custom SSL certs named silo default?
- Gathers output of top
- Is this SL1 Appliance memory swapping?

Administration Portal:

The pre-upgrade-checkv2.sh script checks the following on all Administration Portals:

- Checks web processes
- Are custom SSL certs named silo default?

Data Collectors and Message Collectors:

The pre-upgrade-checkv2.sh script checks the following on all active Data Collectors and Message Collectors

- Checks running python processes

Downloading and Running the Pre-Upgrade Check

To download and run the pre-upgrade check:

1. Contact ScienceLogic Customer Support to get the file. Save the file to a local computer.
2. Use WinSCP or another file-transfer utility to copy the file pre-upgrade-check-v2.sh to the directory /home/em7admin of each SL1 appliance.
3. Go to the console of each SL1 appliance or open a SSH session to each SL1 appliance.
4. Navigate to the directory /home/em7admin.
5. At the shell prompt, enter the following to make the script executable:

```
chmod +x pre-upgrade-check-v2.sh.
```

6. To execute the script with the default options:

```
sudo ./pre-upgrade-check-v2.sh -<appliance_type> > /tmp/preupgrade-check-v2-  
`hostname -s`. `date "+%Y%m%d"` 2>&1
```


where:

appliance_type is one of the following:

-d = Database Appliance

-o = All-In-One

-c = Data Collector

-m = Message Collector

-a = Admin Portal

7. To view the output from the script, view the file `/tmp/pre-upgrade-check-v2- <current hostname> . <today's date>`
8. If the output includes errors and you need help fixing them, contact ScienceLogic Customer Support to fix the errors before upgrading.
9. Perform these steps on each SL1 appliance.

Putting All SL1 Appliances in 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 SL1 appliances in maintenance mode. This will prevent spurious error messages and events during the deployment.

To enable user maintenance mode for all the SL1 appliances in your SL1 system:

1. Go to the Appliance Manager page (System > Settings > Appliances). Note the list of SL1 appliances in your system.
2. Go to the **Device Manager** page (Registry > Devices > Device Manager):

Device Manager Devices Found (176)										Actions	Report	Reset	Guide	
Device Name	IP Address	Device Category	Device Class / Subclass	OID	Organization	Current State	Collection State	Collection State	SNMP Credential	SNMP Version				
1. 10-Forward	10.20.0.195	Servers	NET-SNMP FreeBSD	54	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
2. 10.20.0.108	10.20.0.108	Network Router	Cisco Systems 2501	72	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
3. 10.20.0.123	10.20.0.123	Network Router	Cisco Systems 7206VXR	112	System	Minor	CUG1	Active	Cisco SNMPv2 - Exa V2					
4. 10.20.0.13	10.20.0.13	Unknown	Generic SNMP	107	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
5. 10.20.0.135	10.20.0.135	Network Switches Cisco Systems	Catalyst 3560G-XL	131	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
6. 10.20.0.141	10.20.0.141	Network Switches Cisco Systems	Catalyst WS-C6509-CatOS	110	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
7. 10.20.0.146	10.20.0.146	Network Broadbar Netopia	Netopia 3346 v8.2f1	2	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
8. 10.20.0.147	10.20.0.147	Network Broadbar Netopia	Netopia 3381 v8.0.10	175	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
9. 10.20.0.148	10.20.0.148	Network Broadbar Netopia	Netopia (R)3100, R450x, R700d, R9	163	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
10. 10.20.0.149	10.20.0.149	Network Broadbar Netopia	R7220CT	162	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
11. 10.20.0.151	10.20.0.151	Unknown	Generic SNMP	141	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
12. 10.20.0.160	10.20.0.160	Unknown	Generic SNMP	165	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
13. 10.20.0.163	10.20.0.163	Unknown	Generic SNMP	164	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
14. 10.20.0.175	10.20.0.175	Unknown	Generic SNMP	44	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
15. 10.20.0.176	10.20.0.176	Unknown	Konica Corporation OEM	41	System	Minor	CUG1	Active	Cisco SNMPv2 - Exa V2					
16. 10.20.0.190	10.20.0.190	Unknown	Generic SNMP	56	System	Minor	CUG1	Active	Cisco SNMPv2 - Exa V2					
17. 10.20.0.191	10.20.0.191	Office Printers	Konica Minolta Fiery X3e 22C-KM	57	System	Minor	CUG1	Active	Cisco SNMPv2 - Exa V2					
18. 10.20.0.201	10.20.0.201	Unknown	Generic SNMP	48	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
19. 10.20.0.208	10.20.0.208	Unknown	Generic SNMP	52	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
20. 10.20.0.209	10.20.0.209	Telephony	Quantum Tenor	53	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
21. 10.20.0.222	10.20.0.222	Unknown	Generic SNMP	138	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
22. 10.20.0.26	10.20.0.26	Unknown	Generic SNMP	171	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
23. 10.20.0.52	10.20.0.52	Unknown	ASKEY Computer Corp. OEM	6	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
24. 10.20.0.59	10.20.0.59	Unknown	Generic SNMP	3	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
25. 10.20.0.61	10.20.0.61	Unknown	Generic SNMP	84	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					

- In the **Device Manager** page, select the checkbox for each SL1 appliance in your SL1 system. This includes both primary and secondary Database Servers.
- In the **Select Action** drop-down list, select the following:
 - 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.
- Click the **[Go]** button.

Deploying the Update

NOTE: ScienceLogic recommends that you perform these steps during a maintenance window.

When you deploy a software update, SL1 checks the status of the software on each SL1 appliance. SL1 can deploy the software update to those SL1 appliances that have been successfully staged.

If SL1 is still staging the software update when you select the lightning bolt icon (⚡), SL1 will wait until staging has completed before deploying the updates to each SL1 appliance.

You must deploy only the latest update. For example, if you downloaded updates for 8.4.0, 8.5.0, and 8.6.0, you must deploy only the update for 8.6.0.

To deploy a software update:

1. Go to the **System Updates** page (System > Tools > Updates)

System Updates									
Installation mode: [Upgrade only]									
Actions Import Reset Guide									
EM7 Releases Versions Found [2]									
Release Version	Update Signature	Imported On	Imported By	Import Status	Staging Status	Deployment	Deployment Status	Deployment Status Date	
1. EM7 8.2.0 jenkins_EM7_G3_8.2.0 r4269	ScienceLogic Internal QA Build	2016-12-04 16:39:35	em7admin	Missing Key	N/A	--	New/Unscheduled	--	
2. EM7 8.2.0	Base installation	2016-11-14 16:43:54	em7admin	Complete	N/A	Full (1/1)	New/Unscheduled	--	
ScienceLogic OS Updates Found [2]									
Release Version	Update Signature	Imported On	Imported By	Import Status	Staging Status	Deployment	Deployment Status	Deployment Status Date	
1. Platform 2016-12-01	ScienceLogic Internal QA Build	2016-12-04 16:39:35	em7admin	Missing Key	N/A	--	New/Unscheduled	--	
2. Platform 2016-11-11	Base installation	2016-11-14 16:43:54	em7admin	Complete	N/A	Outdated	New/Unscheduled	--	

2. Ensure that the software update for EM7 8.6.0 has a **Staging Status** of *Complete*.
3. In the **EM7 Releases** pane, find the software update for EM7 8.6.0 . Select its lightning bolt icon ().
4. In the **ScienceLogic OS** pane, find the latest OS update. Select its lightning bolt icon ().
5. The software update will be deployed to those SL1 appliances in your system that have been successfully staged.
6. The **Deployment Status** column displays the status of the deployment. Possible values are:
 - *Full*. Software was fully deployed to all SL1 appliances.
 - *Partial*. Software was deployed to some but not all SL1 appliances.
 - *Outdated*. Software is not the latest version. A later version has been deployed.
 - *--*. Software has not been deployed on any SL1 appliances.
7. For more detailed information about deployment, in the **Deployment Status** column, select the log icon ().
8. SL1 will install the software update on each SL1 appliance .

Putting All SL1 Appliances Out of Maintenance Mode

To disable user maintenance mode for all the SL1 appliances in your SL1 system:

1. Go to the **Appliance Manager** page. Note the list of SL1 appliances in your system.
2. Go to the **Device Manager** page (Registry > Devices > Device Manager):

Device Manager Devices Found (176)										Actions	Report	Reset	Guide	
Device Name	IP Address	Device Category	Device Class / Subclass	OID	Organization	Current State	Collection Group	Collection State	SNMP Credential	SNMP Version				
1. 10-Forward	10.20.0.195	Servers	NET-SNMP FreeBSD	54	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
2. 10.20.0.108	10.20.0.108	Network.Router	Cisco Systems 2501	72	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
3. 10.20.0.123	10.20.0.123	Network.Router	Cisco Systems 7206VXR	112	System	Minor	CUG1	Active	Cisco SNMPv2 - Exa V2					
4. 10.20.0.13	10.20.0.13	Unknown	Generic SNMP	107	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
5. 10.20.0.135	10.20.0.135	Network.Switches Cisco Systems	Catalyst 3560G-XL	131	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
6. 10.20.0.141	10.20.0.141	Network.Switches Cisco Systems	Catalyst WS-C6509-CatOS	110	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
7. 10.20.0.146	10.20.0.146	Network Broadbat Netopia	Netopia 3346 v8.2f1	2	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
8. 10.20.0.147	10.20.0.147	Network Broadbat Netopia	Netopia 3381 v8.0.10	175	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
9. 10.20.0.148	10.20.0.148	Network Broadbat Netopia	Netopia (R)100, R450x, R700d, R9	163	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
10. 10.20.0.149	10.20.0.149	Network Broadbat Netopia	RT220CT	162	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
11. 10.20.0.151	10.20.0.151	Unknown	Generic SNMP	141	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
12. 10.20.0.160	10.20.0.160	Unknown	Generic SNMP	165	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
13. 10.20.0.163	10.20.0.163	Unknown	Generic SNMP	164	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
14. 10.20.0.175	10.20.0.175	Unknown	Generic SNMP	44	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
15. 10.20.0.176	10.20.0.176	Unknown	Konica Corporation OEM	41	System	Minor	CUG1	Active	Cisco SNMPv2 - Exa V2					
16. 10.20.0.190	10.20.0.190	Unknown	Generic SNMP	56	System	Minor	CUG1	Active	Cisco SNMPv2 - Exa V2					
17. 10.20.0.191	10.20.0.191	Office Printers	Konica Minolta Fiery X3e 22C-KM	57	System	Minor	CUG1	Active	Cisco SNMPv2 - Exa V2					
18. 10.20.0.201	10.20.0.201	Unknown	Generic SNMP	48	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
19. 10.20.0.208	10.20.0.208	Unknown	Generic SNMP	52	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
20. 10.20.0.209	10.20.0.209	Telephony	Quintum Tenor	53	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
21. 10.20.0.222	10.20.0.222	Unknown	Generic SNMP	138	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
22. 10.20.0.26	10.20.0.26	Unknown	Generic SNMP	171	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
23. 10.20.0.52	10.20.0.52	Unknown	ASKEY Computer Corp. OEM	6	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
24. 10.20.0.59	10.20.0.59	Unknown	Generic SNMP	3	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					
25. 10.20.0.61	10.20.0.61	Unknown	Generic SNMP	84	System	Healthy	CUG1	Active	Cisco SNMPv2 - Exa V2					

- In the **Device Manager** page, select the checkbox for each SL1 appliance in your SL1 system.
- In the **Select Action** drop-down list, select the following:
 - Change User Maintenance Mode: Disabled.** This option disables user maintenance mode for the selected devices.
- Click the **[Go]** button.

Performing Deltaless Upgrades

After performing the steps to upgrade to SL1 8.6.0, continue to the [Upgrading the SL1 Distributed Architecture on SL1 versions 8.6.0 and Later](#) to upgrade to the latest SL1 release.

Upgrading the SL1 Distributed Architecture on SL1 versions 8.6.0 and Later

As of 8.12.0, SL1 uses an improved System Updates tool (System > Tools > Updates).

As of 8.12.1.3, SL1 uses delta-less upgrades, meaning that you can import a single file and upgrade to the latest version.

NOTE: : You must be running SL1 8.6.0 or later to use delta-less updates.

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

Upgrading the Distributed Stack for SL1 includes the following steps:

- For systems prior to 8.12.0, [disabling automatic staging](#)
- [Downloading the Update](#)
- [Importing the Update](#)
- [Staging the Update](#)
- [Running the Pre-Upgrade Check](#)
- [Putting All SL1 Appliances in Maintenance Mode](#)
- [Deploying the Update](#)
- [Putting All SL1 Appliances Out of Maintenance Mode](#)

Special Steps for SL1 8.12.0 and Earlier

For systems running an SL1 version prior to 8.12.0, go to the **System Updates** page and disable automatic staging (System > Tools > Updates > Actions > Disable automatic staging).

If you have previously used manual staging, 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.

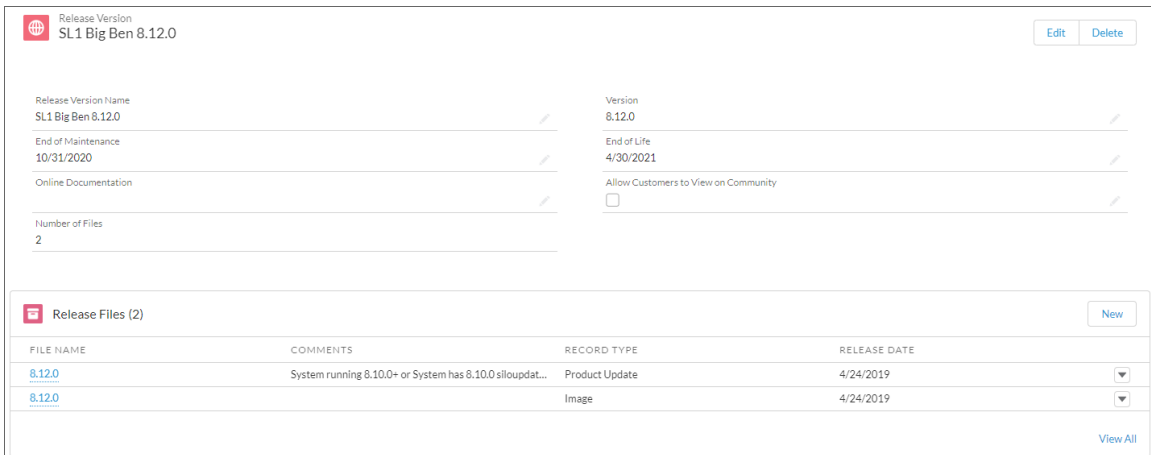
Downloading the Update

NOTE: To download updates for previous SL1 software versions that have reached their End of Life date and are no longer supported by ScienceLogic, contact ScienceLogic Support or a designated Customer Success Manager to get the update files.

Before you can load a patch or update onto your instance of the SL1 system, you must first download the patch or update to your local computer. To do this:

NOTE: These steps do not affect the performance of SL1. ScienceLogic recommends that you perform these steps at least 3 days before upgrading.

1. Log in to <https://support.sciencelogic.com>. Use your ScienceLogic customer account and password to access this site.
2. Select the [**Product Downloads**] button, select the **Product Downloads** menu, and choose *Platform*.
3. Find the release you are interested in and click its name.



Release Version
SL1 Big Ben 8.12.0

Edit Delete

Release Version Name SL1 Big Ben 8.12.0	Version 8.12.0
End of Maintenance 10/31/2020	End of Life 4/30/2021
Online Documentation	Allow Customers to View on Community <input type="checkbox"/>
Number of Files 2	

Release Files (2)

FILE NAME	COMMENTS	RECORD TYPE	RELEASE DATE
8.12.0	System running 8.10.0+ or System has 8.10.0 siloupdat...	Product Update	4/24/2019
8.12.0		Image	4/24/2019

View All


4. In the **Release Version** article, click on the link for the release image or release patch you want to download. Scroll to the bottom of the page.
5. Under **Files**, select the link for the file you want to download. The file is then downloaded to your local computer.

Importing the Update


To import a product update on to your SL1 system:

NOTE: These steps do not affect the performance of SL1. ScienceLogic recommends that you perform these steps at least 3 days before upgrading.

1. Make sure that you can navigate to the patch file.
2. In SL1, go to the **System Updates** page (System > Tools > Updates).
3. In the **System Updates** page, click the **[Import]** button.



4. 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 SL1 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.
5. Click the **[Import]** button.
6. 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 SL1 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 SL1 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.
7. The update file or patch file is imported to SL1 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 SL1 system, you must **stage** the software update. During staging, the SL1 system copies the software update to each ScienceLogic applianceGlobal Manager system. Staging allows SL1 to simultaneously apply the software changes to each ScienceLogic applianceGlobal Manager system, regardless of the speed of the connection to each ScienceLogic applianceGlobal Manager system. The SL1 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: These steps do not affect the performance of SL1. ScienceLogic recommends that you perform these steps at least 3 days before upgrading.

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


Automatic Staging

To enable automatic staging:

1. In SL1, 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 SL1 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 stage compared to the total number of ScienceLogic appliances.


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 ().

5. The *Staging Status* column can have one of the following statuses:


- --. No staging request is active and software has not been staged on any SL1 appliances.
- *Scheduled*. The SL1 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.


To disable automatic staging:

1. In SL1, 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 **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.

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 ().

Manually Staging an Update

You can manually stage a software update to one or more SL1 appliances.

For example, 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 SL1 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, SL1 checks the status of the software updated on each ScienceLogic appliance. SL1 then stages the software update **only to those SL1 appliances that have not yet been staged** for this software update.

To manually stage a software update:

1. In SL1, go to the **System Updates** page (System > Tools > Updates).

EM7 Version	OS Version	Update Signature	Imported On	Imported By	Import Status	Staging Status	Development Status	Development Status Date
EM7 8.12.0	Platform 2019-04-09	ScienceLogic Official Release	2019-04-19 05:13:55	em7admin	Complete	Complete (4/4)	Complete (4/4)	2019-04-19 09:44:10
EM7 8.10.0	Platform 2018-11-27	Base installation	2018-12-17 23:42:02	em7admin	Complete	Outdated	Outdated	2019-04-19 09:15:45

2. In the **System Updates** page, find the software update you want to stage. Select its staging icon ().
3. The software update will be copied to each ScienceLogic appliance that has not yet been staged.
4. The *Staging Status* column will display the number of ScienceLogic appliances that have been successfully stage compared to the total number of ScienceLogic appliances.

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 ().

5. The *Staging Status* column can have one of the following statuses:
 - *--*. No staging request is active and software has not been staged on any SL1 appliances.
 - *Scheduled*. The SL1 system 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.
- *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.

Running the Pre-Upgrade Check for SL1 10.1 and Later

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: These steps do not affect the performance of SL1. ScienceLogic recommends that you perform these steps at least 3 days before upgrading.

The pre-upgrade check examines the following:


- Is each SL1 Appliance eligible to be updates?
- Are updates enable on each SL1 Appliance?
- Are any of the SL1 Appliances running CentOS 5?
- Is this hostfile on each SL1 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?

Running the Pre-Upgrade Check

To run a pre-upgrade check:

1. Go to the **System Updates** page (System > Tools > Updates).

System Updates									
EM7 Releases Versions Found [1]									
EM7 Version	OS Version	Update Signature	Imported On	Imported By	Import Status	Staging Status	Preupgrade Status	Deployment Status	Deployment Status Date
1. EM7 10.1.0	Platform 2020-05-01	Base installation	2020-06-25 15:46:01	em7admin	Complete	Complete (10/10)	...	Complete (10/10)	2020-06-25 22:59:53

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.

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 SL1 appliance is not licensed or the license has expired
- The SL1 appliance cannot be reached over the network
- The Data Collector has failed over
- The SL1 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 SL1 appliances are enabled for patching.

However, if you have used a command-line tool to exclude an SL1 appliance from updates, the pre-upgrade check will fail. To fix this error, include the SL1 appliance for updates.

Free Disk-Space Failure

This test checks the root partition and requires 1 GB of free disk space. If the root partition does not have 1 GB of free disk space, the pre-upgrade check will fail.

If the root partition does not have 1 GB 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 for 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.

Follow these steps on this page 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.

```
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:

```
sudo package-cleanup --dupes
```

3. If there are duplicate packages:

```
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 SL1 ISO or patch bundle
3. If the package is not part of the SL1 ISO or patch bundle, uninstall the package.

Unmet dependencies

You'll have to reset the staging status of the appliance and stage it again. Contact ScienceLogic Customer Success for help in resetting the staging status.

Running the Pre-Upgrade Check for SL1 8.14 and Earlier

NOTE: These steps do not affect the performance of SL1. ScienceLogic recommends that you perform these steps at least 3 days before upgrading.

The pre-upgrade check examines the following:

All SL1 Appliances:

The pre-upgrade-checkv2.sh script checks the following on all SL1 appliances:

- Hostname and hosts file
- Uptime and average CPU load
- SL1 version
- OS versions
- MariaDB version
- Current disk utilization
- Is post update running?

- Checks yum config for configured proxy
- Check for packages installed outside of official SL1 update
- Gathers system status

Active Database Server:

The pre-upgrade-checkv2.sh script checks the following on all active Database Servers:

- Identifies Data Collectors that are not members of a Collector Group
- Lists all active SL1 appliances
- Lists all SL1 appliances that are disabled or ineligible for upgrade
- Checks for active staging schedules
- Check for active deployment schedules
- Are backups configured and successful?
- Checks pruner status
- Checks SIGTERMs for past 24 hours
- Compares number of tables on disk to the table definition cache
- Lists the largest 10 databases by size
- Searches for long running queries
- Shows status of used mysql connections and variables
- Lists mysql config files
- Gathers pt-diststats at 15 second interval three times
- Checks for system status wrapper bug
- Ensure that silo.log ownership is correct

All Database Servers:

The pre-upgrade-checkv2.sh script checks the following on all Database Servers:

- Checks HA/DR status
- Checks for crm bans
- Are custom SSL certs named silo default?
- Gathers output of top
- Is this SL1 Appliance memory swapping?

Administration Portal:

The pre-upgrade-checkv2.sh script checks the following on all Administration Portals:

- Checks web processes
- Are custom SSL certs named silo default?

Data Collectors and Message Collectors:

The pre-upgrade-checkv2.sh script checks the following on all active Data Collectors and Message Collectors

- Checks running python processes

Downloading and Running the Pre-Upgrade Check

To download and run the pre-upgrade check:

1. Contact ScienceLogic Customer Support to get the file. Save the file to a local computer.
2. Use WinSCP or another file-transfer utility to copy the file pre-upgrade-check-v2.sh to the directory /home/em7admin of each SL1 appliance.
3. Go to the console of each SL1 appliance or open a SSH session to each SL1 appliance.
4. Navigate to the directory /home/em7admin.
5. At the shell prompt, enter the following to make the script executable:

```
chmod +x pre-upgrade-check-v2.sh.
```

6. To execute the script with the default options:

```
sudo ./pre-upgrade-check-v2.sh -<appliance_type> > /tmp/preupgrade-check-v2-  
`hostname -s`. `date "+%Y%m%d" ` 2>&1
```

where:

appliance_type is one of the following:

-d = Database Appliance

-o = All-In-One

-c = Data Collector

-m = Message Collector

-a = Admin Portal

7. To view the output from the script, view the file /tmp/pre-upgrade-check-v2-<current hostname>.<today's date>
8. If the output includes errors and you need help fixing them, contact ScienceLogic Customer Support to fix the errors before upgrading.
9. Perform these steps on each SL1 appliance.

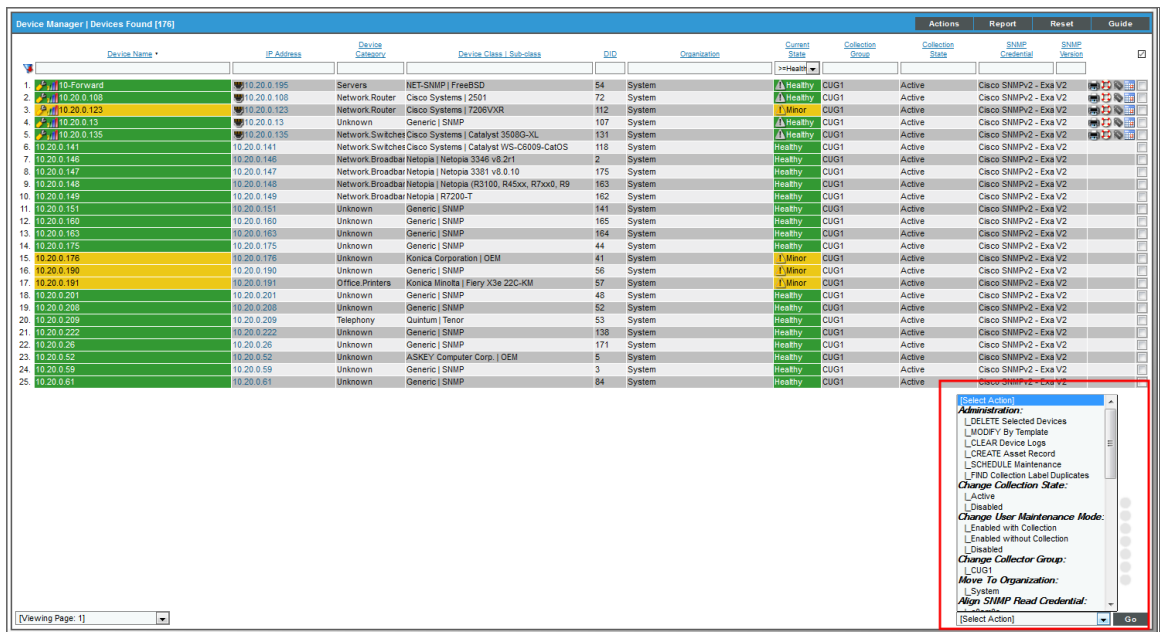
Putting All SL1 Appliances in Maintenance Mode

NOTE: ScienceLogic recommends that you perform these steps during a maintenance window.

Immediately before deploying a software update, ScienceLogic recommends that you put all SL1 appliances in maintenance mode. This will prevent spurious error messages and events during the deployment.

To enable user maintenance mode for all the SL1 appliances in your SL1 system:

1. Go to the Appliance Manager page (System > Settings > Appliances). Note the list of SL1 appliances in your system.
2. Go to the **Device Manager** page (Registry > Devices > Device Manager):



3. In the **Device Manager** page, select the checkbox for each SL1 appliance in your SL1 system. This includes both primary and secondary Database Servers.
4. In the **Select Action** drop-down list, select the following:
 - *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.
5. Click the **[Go]** button.

Deploying the Update

During deployment, avoid the following tasks:


- Running integrations and third-party applications that access the SL1 database or manipulate data on SL1
- Running discovery sessions
- Running nightly discovery
- Bringing HA/DR out of maintenance mode
- Adding new SL1 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


After you have imported and staged an update to SL1, run the pre-upgrade check, and fixed all errors found by the pre-upgrade check, you can either immediately deploy the update or deploy the update at a later time.

When you deploy an update, the update is installed on all appliances that have already been staged.

NOTE: When you deploy an update, SL1 checks to ensure that you have already deployed all required updates. If you have not, SL1 will generate an error message specifying the updates you must deploy before continuing with the current update.


To deploy a software update on your 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.

NOTE: If SL1 is still staging the patch when you click the lightning-bolt icon () , SL1 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 SL1 system that have already been staged. If one or more appliances in your SL1 system have been successfully staged, SL1 will deploy the update to those appliances.
5. During deployment, the Deployment Status column can have one of the following statuses:
 - --. No deployment request is active, and software has not been deployed on any SL1 appliances.
 - *Scheduled*. The SL1 system is aware of the deployment request and is preparing for deployment.

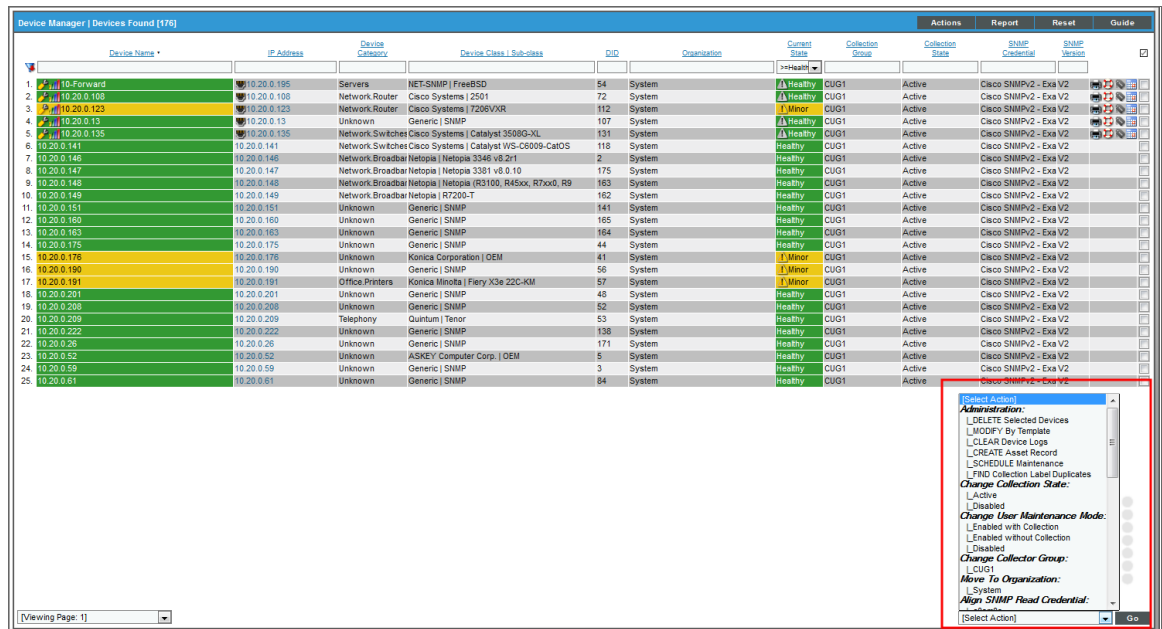
- *In Progress*. Deployment is in progress but has not completed.
- *Complete*. Deployment has completed, and all appliances are updated.
- *Incomplete*. Deployment has completed, and one or more, but not all, appliances are updated.
- *Canceled*. User manually canceled the deployment.
- *Outdated*. The current update is not the latest or has already been installed.
- *Failed*. An unexpected error occurred in the deployment process. Contact ScienceLogic Support.

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 ().

Putting All SL1 Appliances Out of Maintenance Mode

To disable user maintenance mode for all the SL1 appliances in your SL1 system:

1. Go to the **Appliance Manager** page. Note the list of SL1 appliances in your system.
2. Go to the **Device Manager** page (Registry > Devices > Device Manager):



3. In the **Device Manager** page, select the checkbox for each SL1 appliance in your SL1 system.
4. In the **Select Action** drop-down list, select the following:
 - *Change User Maintenance Mode: Disabled*. This option disables user maintenance mode for the selected devices.
5. Click the **[Go]** button.

NOTE: If your SL1 system includes the SL1 Extended Architecture, perform the steps in the section [Upgrading the Extended Architecture](#).

CAUTION: Refer to the release notes for your current release to determine if you must [upgrade MariaDB](#) after upgrading.

CAUTION: Refer to the release notes for your current release to determine if you must [reboot all SL1 appliances](#) after upgrading.

Upgrading the Extended Architecture

To upgrade the SL1 Extended Architecture, follow these procedures:

- [Resize the Disks on the Compute Node](#) (if the disks on each Compute Node are smaller than 350GB)
- [Upgrade Steps](#)

Prerequisites

- Unlike for 8.14 releases, for 10.1.0 and later releases, 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.
- Unlike for 8.14 releases, for 10.1.0 and later releases, 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 new nodes before upgrading your existing nodes to 10.1.0.
- Ensure that all nodes in the SL1 Extended Architecture can access the internet.

NOTE: To perform the upgrade, you must have a ScienceLogic customer account that allows you access to the artifact repository page on the Customer Portal. For details, contact your Customer Success Manager. To verify your access, go to <https://sciencelogic.jfrog.io/sciencelogic/webapp/#/profile>

Resizing the Disks on the Compute Node

The Kafka Messaging service requires additional disk space on each Compute Node. Before upgrading to 10.1.0, ensure that each disk on each existing Compute Node in the Compute Node cluster is at least 350GB.

If each disk on each existing Compute Node is not at least 350GB, perform the following steps on each Compute Node:

1. Resize the hard disk via your hypervisor to at least 350GB.
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 vgs | grep "VG Size"
```

Upgrade Steps

To upgrade from 8.14.x to 10.1.x:

1. Either go to the console of the Management Node or 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. Back up the following files:

- `/home/em7admin/sl1x-deploy/sl1x-inv.yml`
- `/home/em7admin/sl1x-deploy/output-files/cluster.yml`
- `/home/em7admin/sl1x-deploy/output-files/cluster.rkestate`
- `/home/em7admin/sl1x-deploy/output-files/kube_config_cluster.yml`

NOTE: ScienceLogic recommends that you back up these files at regular intervals.

4. Run the following command to enter the docker container:

```
docker-compose -f docker-compose.external.yml run --rm deploy shell
```

5. Run the following command and note the value of `HOSTS`:

```
kubectl get ing responder-ingress
```

6. Stop the ingress of SL1 Agent data and Collector Pipeline data:

```
helm delete --purge sl1-streamer
helm delete --purge sls-api-storeconfig
helm delete --purge avail-store
helm delete --purge da-postprocessing-service
```

7. Delete failed charts:

```
helm ls | awk '/FAILED/'
```

If the above command results in any output, run the following command:

```
helm delete --purge $(helm ls | awk '/FAILED/ { print $1 }')
```

8. Exit the ansible shell session:

```
exit
```

9. Monitor the queues until they are drained:

```
check https://<HOSTS>/api/queues/list/?api_key=asdfQ345sdf
```

where:

`HOSTS` is the value from step #5.

Refresh that page until all queues are at 0.

10. Navigate to the `s11x-deploy` directory and then download and extract the latest templates. To do this, enter the following at the shell prompt:

```
curl -u <username> -k https://sciencelogic.jfrog.io/sciencelogic/docker-compose-  
local/s11x-deploy.tar.gz | tar zxv - {docker-compose.external,s11x-inv-  
template}.yaml
```

When prompted, enter the API key.

11. Rename old inventory file:

```
mv s11x-inv.yaml s11x-inv.yaml.8.14
```

12. Copy the inventory template file to the name `s11x-inv.yaml`. To do this :

```
cp s11x-inv-template.yaml s11x-inv.yaml
```

13. Edit the file ***s11x-inv.yaml*** to match your SL1 Extended system. At the shell prompt ,enter:

```
vi s11x-inv.yaml
```

CAUTION: Do not remove colons when editing this file.

NOTE: For details on editing the `s11x-inv.yaml` file, see the ***Installation*** manual.

14. Save your changes and exit the file (`:wq`).

15. If you have added one or more new Compute Nodes or the Storage Manager node, enter the following at the shell prompt:

```
docker-compose -f docker-compose.external.yaml run --rm deploy ssh-keys --ask-pass
```

When prompted, enter the System Password that you entered in the ISO menu.

16. To update the SL1 Extended system, enter the following at the shell prompt:

```
docker-compose -f docker-compose.external.yaml run --rm deploy s11x
```

Manual Steps for Updates to 8.4.x and Earlier Systems

If you upgraded from an 8.4.x system or earlier, you must manually apply the following changes to every Message Collector and All-In-One Appliance in your SL1 system:

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

```
sudo vi /etc/siteconfig/siloconf.siteconfig
```

4. Locate the following line:

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

NOTE: On an All-In-One Appliance, this line will include additional entries in the comma-delimited list.

5. Add ",agent" to the end of the line. The line should now look like this:

```
eventmanager = syslog,trap,internal,agent
```

6. Save the file and exit vi (:wq).

7. At the command line, enter the following command to rebuild the configuration file:

```
sudo /opt/em7/share/scripts/generate-silo-conf.py > /etc/silo.conf
```

8.4.2 included a new feature for Ticketing. The Note Editor for tickets includes a drop-down menu where the user can specify that the note be saved as Plain Text or HTML. Ticket notes created in the API can also be saved as Plain Text or HTML. HTML is the default format for ticket notes in both the Note Editor and the API.

This new feature required a change to the database schema that will be performed immediately after a system is upgraded to 8.4.2 or later for the first time. If your system has not been upgraded to 8.4.2 or later, this schema change will be performed after upgrading to 8.6.0.

During the post-patch process, all existing ticket notes will be migrated to the new schema in batches. During the migration, all ticket notes will be unavailable. The user interface and API will not display ticket notes. Ticket notes cannot be viewed or updated during the post-patch process. On a system that included 2,000,000 ticket notes, this process took approximately 6 hours.

If you require access to all ticket notes immediately after installing an 8.4.2 or later release, contact ScienceLogic Customer Support for details on manually updating the database schema before you upgrade to 8.4.2 or later.

Automatically Upgrading MariaDB with a Script

NOTE: Refer to the release notes for your current release to determine if you must upgrade MariaDB. Not every SL1 update requires an upgrade of for MariaDB.

SL1 will automatically update MariaDB-client, MariaDB-common, and MariaDB-shared RPMs but will not update the MariaDB Server RPM. You must update the MariaDB Server RPM.

SL1 10.1.0 and later releases include the ***module_upgrade_mariadb*** script to automatically upgrade MariaDB server.

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`. SL1 will save these custom settings and apply them after you upgrade MariaDB.

The **`module_upgrade_mariadb`** script:

- Upgrades the following SL1 appliances:
 - All Database Servers
 - All-In-One Appliances
 - Data Collectors
 - Message Collectors
- Upgrade 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 servers 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`

To upgrade MariaDB, perform the following:

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:

```
sudo /opt/em7/bin/module_upgrade_mariadb -m all
```

3. To see all the options for the **`module_upgrade_mariadb`** script, enter the following at the shell prompt:

```
/opt/em7/bin/module_upgrade_mariadb -h
```

```
Usage: module_upgrade_mariadb -v <mariadb-server-version> -m <module_id> [-t|--test] [-y|--assumeeyes] [-h|--help]
```

4. The script includes these options:
 - `-m` parameter specifies the SL1 appliances that you want to upgrade. You can specify:
 - `-m <mid1, mid2...midN>` provides a comma-separated module IDs.
 - `-m all` : upgrade all appliances (Database Servers, All-In-One Appliances, Data Collectors, and Message Collectors).

- `-m all -db` : upgrade all Database Servers.
- `-m all-cu` : upgrade all Data Collectors and Message Collectors.
- `-t` parameter specifies not to upgrade but instead to run a test of the upgrade script.
- `-y` parameter specifies to automatically enter "yes" at all prompts.
- `-s` parameter specifies to ignore errors in the MySQL configuration files and proceed with the upgrade.
- `-p` parameter specifies the number of Data Collectors that you want to upgrade simultaneously. Database Servers will be upgraded one at a time. Possible values are 1 - 20. The default value is 1.
 - `-p number_of_modules` is the number of Data Collectors to upgrade simultaneously. Values are 1 - 20. The default value is 1.

5. To view the status of the automatic upgrade, enter the following:

```
monitor_upgrade_mariadb
```

Manually Upgrading MariaDB

NOTE: Refer to the release notes for your current release to determine if you must upgrade MariaDB. Not every SL1 update requires an upgrade of for MariaDB.

If you prefer to upgrade MariaDB manually, the following sections describe how to upgrade the MariaDB server for different SL1 appliance types and architectures.

When you update MariaDB, you must update the following SL1 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`. SL1 will save these custom settings and apply them after you upgrade MariaDB.

Download RPMs to SL1 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 SL1 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:

```
curl -Lk -u "<SL1 admin user name>:<SL1 admin password>"
http://localhost:10080/MariaDB-server-`rpm -q --queryformat '%{VERSION}-%
{RELEASE}.'%{ARCH}' MariaDB-client`.rpm > /tmp/MariaDB-server.rpm
curl -Lk -u "<SL1 admin user name>:<SL1 admin password>"
http://localhost:10080/`repoquery --repoid em7-os --queryformat '%{NAME}-%
{VERSION}-%{RELEASE}.'%{ARCH}'.rpm' galera-4` > /tmp/galera-4.rpm
```
3. To Navigate to /tmp directory, enter the following:

```
cd /tmp
```
4. Use SCP or another secure copy program to copy these files to the /tmp directory on each on each Database Server, All-In-One Appliance, Data Collector, and Message Collector:
 - MariaDDB-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 SL1 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, enter:

```
sudo -s
```
3. When prompted, enter the administrator password.
4. At the shell prompt, enter:

```
coro config
```
5. The following prompt appears:

- 1) Enable 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. To determine the current installed version of the RPMs:

```
sudo rpm -qa ^MariaDB-*
```

2. To stop SL1 and MariaDB, enter the following at the shell prompt:

```
sudo systemctl stop em7
sudo systemctl stop mariadb.service
```

3. To stop the MySQL resource, enter the following:

```
sudo crm resource stop mysql
```

4. To save the current enabled state for mariadb.service, enter the following:

```
export MSRVR=`sudo systemctl is-enabled mariadb.service`
```

5. To remove the older version of MariaDB-server, enter the following:

```
sudo rpm --nodeps -ev MariaDB-server
```

6. To replace the Galera package and install the new MariaDB-server package, enter the following:

```
sudo yum --disablerepo=* swap -- remove galera -- install /tmp/galera-4.rpm
sudo yum --disablerepo=* install /tmp/MariaDB-server.rpm
```

7. To remove incompatible backup packages, enter the following:

```
sudo yum remove percona-xtrabackup
```

NOTE: If the "yum remove" command fails, it means that the package does not exist on the SL1 appliance. You can ignore the error message.

8. To regenerate the configuration file for MariaDB, enter the following:

```
sudo /opt/em7/share/scripts/generate-my-conf.py -o /etc/my.cnf.d/silo_mysql.cnf
```

9. To re-start MariaDB, enter the following:

```
sudo systemctl daemon-reload
sudo systemctl start mariadb
```

10. To restart the MySQL resource:

```
sudo crm resource start mysql
```

11. To restore the mariadb.service enabled state, enter the following:

```
sudo systemctl ${MSRV::-1} mariadb.service
```

12. To upgrade the internal configuration for the database, enter the following:

```
sudo mysql_upgrade -u root -p
```

13. To restart the em7 service:

```
sudo systemctl start em7
```

```
sudo rpm -qa ^MariaDB-*
```

14. To stop SL1 and MariaDB, enter the following at the shell prompt:

```
sudo systemctl stop em7
sudo systemctl stop mariadb.service
```

15. To stop the MySQL resource, enter the following:

```
sudo crm resource stop mysql
```

16. To save the current enabled state for mariadb.service, enter the following:

```
export MSRV=`sudo systemctl is-enabled mariadb.service`
```

17. To remove the older version of MariaDB-server, enter the following:

```
sudo rpm --nodeps -ev MariaDB-server
```

18. To replace the Galera package and install the new MariaDB-server package, enter the following:

```
sudo yum --disablerepo=* swap -- remove galera -- install /tmp/galera-4.rpm
sudo yum --disablerepo=* install /tmp/MariaDB-server.rpm
```

19. To remove incompatible backup packages, enter the following:

```
sudo yum remove percona-xtrabackup
```

NOTE: If the "yum remove" command fails, it means that the package does not exist on the SL1 appliance. You can ignore the error message.

20. To regenerate the configuration file for MariaDB, enter the following:

```
sudo /opt/em7/share/scripts/generate-my-conf.py -o /etc/my.cnf.d/silo_mysql.cnf
```

21. To re-start MariaDB, enter the following:

```
sudo systemctl daemon-reload
sudo systemctl start mariadb
```

22. To restart the MySQL resource:

```
sudo crm resource start mysql
```

23. To restore the mariadb.service enabled state, enter the following:

```
sudo systemctl ${MSRV::-1} mariadb.service
```

24. To upgrade the internal configuration for the database, enter the following:

```
sudo mysql_upgrade -u root -p
```

25. To restart the em7 service:

```
sudo systemctl start em7
```

Step 3: On the Secondary Database Server

1. To determine the current installed version of the RPMs:

```
sudo rpm -qa ^MariaDB-*
```

2. To save the current enabled state for mariadb.service, enter the following:

```
export MSRVR=`sudo systemctl is-enabled mariadb.service`
```

3. To remove the older version of MariaDB-server, enter the following:

```
sudo rpm --nodeps -ev MariaDB-server
```

4. To replace the Galera package and install the new MariaDB-server package, enter the following:

```
sudo yum --disablerepo=* swap -- remove galera -- install /tmp/galera-4.rpm  
sudo yum --disablerepo=* install /tmp/MariaDB-server.rpm
```

5. To remove incompatible backup packages, enter the following:

```
sudo yum remove percona-xtrabackup
```

NOTE: If the "yum remove" command fails, it means that the package does not exist on the SL1 appliance. You can ignore the error message.

6. To regenerate the configuration file for MariaDB, enter the following:

```
sudo /opt/em7/share/scripts/generate-my-conf.py -o /etc/my.cnf.d/silo_mysql.cnf
```

7. To restore the mariadb.service enabled state, enter the following:

```
sudo systemctl ${MSRVR::-1} mariadb.service
```

8. To take the secondary Database Server out of maintenance mode, enter the following at the shell prompt:

```
sudo -s
```

9. When prompted, enter the administrator password.

10. At the shell prompt, enter:

```
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".

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, enter:

```
sudo -s
```
3. When prompted, enter the administrator password.
4. At the shell prompt, enter:

```
coro config
```
5. The following prompt appears:
 - 1) Enable 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. To determine the current installed version of the RPMs:

```
sudo rpm -qa ^MariaDB-*
```
2. To stop SL1 and MariaDB, enter the following at the shell prompt:

```
sudo systemctl stop em7  
sudo systemctl stop mariadb.service
```
3. To stop the MySQL resource, enter the following:

```
sudo crm resource stop mysql
```

4. To save the current enabled state for mariadb.service, enter the following:

```
export MSRV=`sudo systemctl is-enabled mariadb.service`
```
5. To remove the older version of MariaDB-server, enter the following:

```
sudo rpm --nodeps -ev MariaDB-server
```
6. To replace the Galera package and install the new MariaDB-server package, enter the following:

```
sudo yum --disablerepo=* swap -- remove galera -- install /tmp/galera-4.rpm  
sudo yum --disablerepo=* install /tmp/MariaDB-server.rpm
```
7. To remove incompatible backup packages, enter the following:

```
sudo yum remove percona-xtrabackup
```

NOTE: If the "yum remove" command fails, it means that the package does not exist on the SL1 appliance. You can ignore the error message.

8. To regenerate the configuration file for MariaDB, enter the following:

```
sudo /opt/em7/share/scripts/generate-my-conf.py -o /etc/my.cnf.d/silo_mysql.cnf
```
9. To re-start MariaDB, enter the following:

```
sudo systemctl daemon-reload  
sudo systemctl start mariadb
```
10. To restart the MySQL resource:

```
sudo crm resource start mysql
```
11. To restore the mariadb.service enabled state, enter the following:

```
sudo systemctl ${MSRV::-1} mariadb.service
```
12. To upgrade the internal configuration for the database, enter the following:

```
sudo mysql_upgrade -u root -p
```
13. To restart the em7 service:

```
sudo systemctl start em7
```

Step 3: On the Secondary Database Server

1. To save the current enabled state for mariadb.service, enter the following:

```
export MSRV=`sudo systemctl is-enabled mariadb.service`
```
2. To remove the older version of MariaDB-server, enter the following:

```
sudo rpm --nodeps -ev MariaDB-server
```
3. To replace the Galera package and install the new MariaDB-server package, enter the following:

```
sudo yum --disablerepo=* swap -- remove galera -- install /tmp/galera-4.rpm
```



```
sudo yum --disablerepo=* install /tmp/MariaDB-server.rpm
```

4. To remove incompatible backup packages, enter the following:

```
sudo yum remove percona-xtrabackup
```

NOTE: If the "yum remove" command fails, it means that the package does not exist on the SL1 appliance. You can ignore the error message.

5. To regenerate the configuration file for MariaDB, enter the following:

```
sudo /opt/em7/share/scripts/generate-my-conf.py -o /etc/my.cnf.d/silo_mysql.cnf
```

6. To restore the mariadb.service enabled state, enter the following:

```
sudo systemctl ${MSRV::-1} mariadb.service
```

7. At the shell prompt, enter:

```
sudo -s
```

8. When prompted, enter the administrator password.

9. At the shell prompt, enter:

```
coro config
```

10. The following prompt appears:

```
1) Disable Maintenance
2) Option Disabled
3) Promote DRBD
4) Stop Pacemaker
5) Resource Status
6) Quit
```

11. Enter "1".

Step 4: On the Disaster Recovery Database Server

1. Open an SSH session to the Disaster Recovery Database Server.

2. At the shell prompt, enter:

```
sudo -s
```

3. When prompted, enter the administrator password.

4. To save the current enabled state for mariadb.service, enter the following:

```
export MSRV=`sudo systemctl is-enabled mariadb.service`
```

5. To remove the older version of MariaDB-server, enter the following:

```
sudo rpm --nodeps -ev MariaDB-server
```

6. To replace the Galera package and install the new MariaDB-server package, enter the following:

```
sudo yum --disablerepo=* swap -- remove galera -- install /tmp/galera-4.rpm
sudo yum --disablerepo=* install /tmp/MariaDB-server.rpm
```

7. To remove incompatible backup packages, enter the following:

```
sudo yum remove percona-xtrabackup
```

NOTE: If the "yum remove" command fails, it means that the package does not exist on the SL1 appliance. You can ignore the error message.

8. To regenerate the configuration file for MariaDB, enter the following:

```
sudo /opt/em7/share/scripts/generate-my-conf.py -o /etc/my.cnf.d/silo_mysql.cnf
```

9. To restore the mariadb.service enabled state, enter the following:

```
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 SL1 appliance.
2. To stop SL1 and mariadb, enter the following at the shell prompt:

```
sudo systemctl stop em7
sudo systemctl stop mariadb.service
```

3. To save the current enabled state for the mariadb.service, enter the following:

```
export MSRV=`sudo systemctl is-enabled mariadb.service`
```

4. To remove the older version of MariaDB-server, enter the following:

```
sudo rpm --nodeps -ev MariaDB-server
```

5. To replace the Galera package and install the new MariaDB-server package, enter the following:

```
sudo yum --disablerepo=* swap -- remove galera -- install /tmp/galera-4.rpm
sudo yum --disablerepo=* install /tmp/MariaDB-server.rpm
```

6. To remove incompatible backup packages, enter the following:

```
sudo yum remove percona-xtrabackup
```

NOTE: If the "yum remove" command fails, it means that the package does not exist on the SL1 appliance. You can ignore the error message.

7. To regenerate the configuration file for MariaDB, enter the following:

```
sudo /opt/em7/share/scripts/generate-my-conf.py -o /etc/my.cnf.d/silo_mysql.cnf
```

8. To re-start MariaDB, enter the following:

```
sudo systemctl daemon-reload
sudo systemctl start mariadb
```

9. To restore the mariadb.service enabled state, enter the following:

```
sudo systemctl ${MSRV::-1} mariadb.service
```

10. To upgrade the internal configuration for the database, enter the following:

```
sudo mysql_upgrade -u root -p
```

11. To restart the em7 service:

```
sudo systemctl start em7
```

12. 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 SL1 Distributed Stack

NOTE: Refer to the release notes for your current release to determine if you must reboot all SL1 appliances. Not every SL1 update requires rebooting.

When an upgrade requires a reboot, use the steps listed in this section to reboot all SL1 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 SL1 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 SL1 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 SL1 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 SL1 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 SL1 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.

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.

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 SL1 appliance.
2. Log in as **em7admin** with the appropriate password.
3. Execute the following commands on the SL1 appliance to pause the system and shutdown MariaDB.

```
sudo touch /tmp/.proc_mgr_pause  
sudo systemctl stop mariadb
```

4. Execute the following command on the appliance to reboot the SL1 appliance:

```
sudo reboot
```

5. After the SL1 appliance has rebooted, either go to the console or use SSH to access the SL1 appliance.
6. Log in as **em7admin** with the appropriate password.
7. Execute the following command on the appliance to un-pause the system:

```
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. First, you should check the status of both Database Servers. To do this, enter the following at the shell prompt:

```
cat /proc/drbd
```

4. Your output will look like this:

```
1: cs:Connected ro:Primary/Secondary 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
```

5. Execute the following commands on the **primary** Database Server to pause the system and shutdown MariaDB. Enter the password for the em7admin user when prompted:

```
sudo touch /tmp/.proc_mgr_pause
sudo systemctl stop pacemaker
sudo systemctl stop mariadb
```

6. Execute the following command on the **primary** Database Server to reboot the appliance:

```
sudo reboot
```

7. After the primary appliance has rebooted, log in to the console of the **primary** Database Server again.

8. Execute the following commands on the **primary** Database Server:

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

9. Enter the password for the em7admin user and confirm the command when prompted.

10. Log in to the **secondary** Database Server as the em7admin user using the console or SSH.

11. Execute the following command on the **secondary** Database Server to reboot the appliance:

```
sudo reboot
```

12. 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. First, you should check the status of both Database Servers. To do this, enter the following at the shell prompt:

```
cat /proc/drbd
```

4. 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.

5. Execute the following command on the **secondary** Database Server to stop the cluster service:

```
sudo systemctl stop pacemaker
```

6. Enter the password for the em7admin user when prompted.

7. Either go to the console of the **primary** Database Server or use SSH to access the **primary** Database Server.
8. Log in as **em7admin** with the appropriate password.
9. Execute the following commands on the **primary** Database Server to pause the system and stop the cluster service. Enter the password for the em7admin user when prompted:

```
sudo touch /tmp/.proc_mgr_pause
sudo systemctl stop pacemaker
sudo systemctl stop mariadb
```

10. Execute the following command on the **primary** Database Server to reboot the appliance:

```
sudo reboot
```

11. 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.

12. Log in as **em7admin** with the appropriate password.

13. Execute the following command on the **primary** Database Server:

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

14. Enter the password for the em7admin user and confirm the command when prompted.

15. Either go to the console of the **secondary** Database Server or use SSH to access the **secondary** Database Server.

16. Log in as **em7admin** with the appropriate password.

17. Execute the following command on the **secondary** Database Server:

```
sudo reboot
```

18. 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. First, you should check the status of both Database Servers in the HA cluster. To do this, enter the following at the shell prompt:

```
cat /proc/drbd
```

4. 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.

5. To stop the cluster service, execute the following command on the **secondary** Database Server in the HA cluster:

```
sudo systemctl stop pacemaker
```
6. Enter the password for the em7admin user when prompted.
7. 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.
8. Log in as **em7admin** with the appropriate password.
9. To pause the system and stop the cluster service, execute the following commands on the **primary** Database Server in the HA cluster :

```
sudo touch /tmp/.proc_mgr_pause
sudo systemctl stop pacemaker
sudo systemctl stop mariadb
```
10. Enter the password for the em7admin user when prompted
11. To reboot the **primary** Database Server in the HA cluster, the following command on the **primary** Database Server in the HA cluster:

```
sudo reboot
```
12. 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.
13. Execute the following command on the **primary** Database Server in the HA cluster:

```
sudo rm /tmp/.proc_mgr_pause
```
14. Enter the password for the em7admin user and confirm the command when prompted.
15. 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.
16. Log in as **em7admin** with the appropriate password.
17. Execute the following command on the **secondary** Database Server in the HA cluster to reboot the appliance:

```
sudo reboot
```
18. Enter the password for the em7admin user when prompted.
19. Either go to the console of the Database Server for Disaster Recovery or use SSH to access the Database Server for Disaster Recovery.
20. Log in as **em7admin** with the appropriate password.
21. Execute the following command on the Database Server for Disaster Recovery to reboot the appliance:


```
sudo reboot
```

22. Enter the password for the em7admin user when prompted.

Restoring Settings for NextUI

To restore the backup of the custom settings in the NextUI:

1. Login to the console of the Database Server or SSH to the Database Server.
2. Open a shell session.
3. Enter the following at the shell prompt:

```
cp /opt/em7/nextui/nextui.env.backup /opt/em7/nextui/nextui.env
```

Restoring the SSL Certificate

To restore your SSL Certificates:

1. Login to the console of the Database Server or SSH to the Database Server.
2. 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 SL1 system.

Resetting the Timeout for PhoneHome Watchdog

You can manually re-set 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. At the command line, type the following:

```
phonehome watchdog view
```

3. You should see something like the following:

```
Current settings:
```

```
autosync: yes  
interval: 120  
state: enabled  
autoreconnect: yes  
timeoutcount: 1
```

```
check: default
```

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

5. To change the settings for SL1 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;
```

6. Repeat these steps on each Data Collector.


7. Repeat these steps on each , Message Collector.

8. Repeat these steps on each Database Server.

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 SL1. 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, SL1 will automatically install the PowerPack. If a PowerPack is included in an update and is currently installed on your system, SL1 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).

PowerPack™ Manager | PowerPacks™ Found [178] Actions Reset Guide

PowerPack™ Name	Version	Update	Publisher	Areas	Events	Classes	Reports	Widgets	Dashbds	Actions	Creds	Tckt Tol	DashSL1	Dev Tol	ID	Edited By	Last Edited	
1. a simple DCM app	3	No	ScienceLogic, Inc.	4	1										169	em7admin	2019-07-25 14:15:22	<input checked="" type="checkbox"/>
2. Microsoft PowerPack Test Cred:	1.7	No	ScienceLogic, Inc.								21				177	em7admin	2019-07-26 15:16:28	<input checked="" type="checkbox"/>
3. SNMP Dynamic Application Ref:	1.0	No	ScienceLogic, Inc.	5											179	em7admin	2019-07-26 15:56:25	<input type="checkbox"/>
4. 2-level DCM app test	1.1	No	ScienceLogic, Inc.	4	1	2									186	em7admin	2019-07-30 08:28:38	<input type="checkbox"/>
5. 3Com Base Pack	1	No	ScienceLogic, Inc.			298									30	em7admin	2019-07-25 11:17:02	<input type="checkbox"/>
6. Alcatel-Lucent Base Pack	1	No	ScienceLogic, Inc.			33									130	em7admin	2019-07-25 11:18:45	<input type="checkbox"/>
7. Atheon Base Pack	1.3	No	ScienceLogic, Inc.	3	5	6									114	em7admin	2019-07-25 11:18:27	<input type="checkbox"/>
8. Amazon Web Services	113	No	ScienceLogic, Inc.	211	222	149	13		13	5	3				41	em7admin	2019-07-25 11:17:23	<input type="checkbox"/>
9. APC Base Pack	7.3.6	No	ScienceLogic, Inc.	9	27	30									48	em7admin	2019-07-25 11:17:32	<input type="checkbox"/>
10. Application Management Pack	1	No	ScienceLogic, Inc.		24										5	em7admin	2019-07-25 11:16:39	<input type="checkbox"/>
11. Arista Networks: Base	1.0	No	ScienceLogic, Inc.			65									120	em7admin	2019-07-25 11:18:35	<input type="checkbox"/>
12. Aruba Base Pack	1.1	No	ScienceLogic, Inc.	3											146	em7admin	2019-07-25 11:18:58	<input type="checkbox"/>
13. AskEM7 Query Widgets	7.2.3	No	ScienceLogic, Inc.					3							34	em7admin	2019-07-25 11:17:07	<input type="checkbox"/>
14. Attachmate Base Pack	1.1	No	ScienceLogic, Inc.			13									56	em7admin	2019-07-25 11:17:35	<input type="checkbox"/>
15. Avaya Base Pack	7.3.5	No	ScienceLogic, Inc.			43									92	em7admin	2019-07-25 11:17:58	<input type="checkbox"/>
16. Avocent ACS Pack	7.3.6	No	ScienceLogic, Inc.	4											147	em7admin	2019-07-25 11:18:58	<input type="checkbox"/>
17. Avocent Base Pack	1.1	No	ScienceLogic, Inc.	3		6									54	em7admin	2019-07-25 11:17:35	<input type="checkbox"/>
18. Blue Coat Base Pack	1.2	No	ScienceLogic, Inc.	1	2	10									16	em7admin	2019-07-25 11:17:35	<input type="checkbox"/>
19. BlueCat Base Pack	1.1	No	ScienceLogic, Inc.	2											29	em7admin	2019-07-25 11:17:35	<input type="checkbox"/>

[Select Action]
Administration:
Update PowerPack(s)
Delete PowerPack(s)
[Select Action] Go

- In the **PowerPack Manager** page, select the checkbox for each PowerPack you want to install.
- In the **Select Action** drop-down field (in the lower right), choose *Update PowerPack(s)*.
- SL1 will display the following message before updating the PowerPack(s):

Update the selected PowerPacks?

NOTE: Any customizations to items contained in updated PowerPacks will be overwritten by the version contained within the more recently imported PowerPack file.

Click the **[OK]** button to continue the installation. Click the **[Cancel]** button to quit the update.

- 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 automatically installed in your SL1 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: The **Enable Selective PowerPack Field Protection** field on the **Behavior Settings** page (System > Settings > Behavior) affects how updates behave. If the **Enable Selective PowerPack Field Protection** checkbox is selected, certain fields in Event Policies, Dynamic Applications, and Device Classes will **not** be updated.

Configuring Subscription Billing

If your SL1 system is configured to communicate with the ScienceLogic billing server, usage data will be sent automatically from your SL1 system to the ScienceLogic billing server once a day. After the ScienceLogic billing server receives the usage data, SL1 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 SL1 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 SL1 systems, the value of **Summary Date** should be within the past 48 hours.
- For SL1 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

Monitoring and Maintaining SL1

Overview

This chapter describes how to manage user access, manage scheduled tasks, and more.

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

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

This chapter includes the following topics:

<i>Monitoring and Managing User Access</i>	174
<i>Managing Scheduled Tasks</i>	177
<i>Monitoring Overall System Usage and Statistics</i>	180
<i>Viewing an Overview of All Events</i>	181
<i>Viewing Events by Appliance and Event Source</i>	183

Monitoring and Managing User Access

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).

The screenshot shows the 'Access Sessions | Log Entries Found [125]' page. The table has the following columns: User Account, Last Address, State, Login Time, Last Hit Time, Logout Time, Session Duration, and Session ID. The table contains 26 rows of data, each representing a user session. The 'State' column indicates the session's status, such as 'Logged In', 'Expired', or 'Logged Out'. The 'Session Duration' column shows the length of the session in various units (seconds, minutes, hours). The 'Session ID' column provides a unique identifier for each session.

User Account	Last Address	State	Login Time	Last Hit Time	Logout Time	Session Duration	Session ID
em7Admin	10.128.38.60	Logged In	2017-05-02 16:27:15	2017-05-02 16:27:15	--	14 secs	840m9f8aaqabv20ajgkcor7
em7Admin	10.128.38.59	Logged In	2017-05-02 12:12:53	2017-05-02 12:12:53	--	4 hrs 14 mins 36 secs	2dd1euomog805oiz2t757dk35
em7Admin	10.128.38.67	Logged Out	2017-05-01 13:13:11	2017-05-01 16:50:55	2017-05-01 16:50:55	3 hrs 37 mins 44 secs	qq3rd490g3fnaamig7alf3d6
em7Admin	10.128.38.67	Logged Out	2017-05-01 10:30:09	2017-05-01 11:44:24	2017-05-01 11:44:24	1 hr 14 mins 15 secs	bfmp2aio6b6h63f0og18b6e
em7Admin	10.128.38.69	Expired	2017-04-28 15:17:59	2017-04-28 15:17:59	2017-04-28 15:17:59	0 secs	g0e41fmbaa2fuo3bh0cqm4
em7Admin	10.128.39.251	Expired	2017-04-26 15:47:57	2017-04-26 15:47:57	2017-04-26 15:47:57	0 secs	jny23accr95o8fman9h6f530g2
em7Admin	10.128.38.102	Expired	2017-04-25 09:03:21	2017-04-25 09:03:21	2017-04-25 09:03:21	0 secs	y8qas10j9o4h0d94pp0t1
em7Admin	10.128.39.224	Logged Out	2017-04-21 12:48:52	2017-04-21 15:10:40	2017-04-21 15:10:40	2 hrs 21 mins 48 secs	jybcu3ng3burda7jgnfadng0
em7Admin	10.128.39.216	Expired	2017-04-21 13:02:34	2017-04-21 13:02:34	2017-04-21 13:02:34	0 secs	5s0k0tecs7o75ab1glofh2o1
em7Admin	10.128.38.68	Logged Out	2017-04-20 11:46:49	2017-04-20 16:02:13	2017-04-20 16:02:13	4 hrs 15 mins 24 secs	980p0p0p0ncas88h0g7g4a0
em7Admin	10.128.39.182	Expired	2017-04-20 15:31:21	2017-04-20 15:31:21	2017-04-20 15:31:21	0 secs	0p0p0p0f61qg2sv4hnhg0f91
em7Admin	10.128.38.68	Expired	2017-04-20 11:46:48	2017-04-20 11:46:48	2017-04-20 11:46:48	0 secs	ja0n1a74d5ccsrah1a9muc0533
em7Admin	10.128.38.91	Expired	2017-04-19 12:56:38	2017-04-19 12:56:38	2017-04-19 12:56:38	0 secs	780ns7cmg0v238gbofnqno2
em7Admin	10.128.38.43	Logged Out	2017-04-18 09:59:26	2017-04-18 16:17:04	2017-04-18 16:17:04	6 hrs 17 mins 38 secs	amh0g7032ozm3mu50pigh6
em7Admin	10.128.38.53	Expired	2017-04-18 15:22:25	2017-04-18 15:22:25	2017-04-18 15:22:25	0 secs	ikx2k1kqp1q0j7u0cg5bp15
em7Admin	10.128.38.53	Expired	2017-04-18 13:45:34	2017-04-18 15:22:25	2017-04-18 15:22:25	1 hr 38 mins 51 secs	ikx2k1kqp1q0j7u0cg5bp15
em7Admin	10.128.38.53	Expired	2017-04-18 11:49:16	2017-04-18 13:45:34	2017-04-18 13:45:34	1 hr 56 mins 18 secs	ikx2k1kqp1q0j7u0cg5bp15
em7Admin	10.128.39.235	Expired	2017-04-17 12:57:57	2017-04-18 11:49:16	2017-04-18 11:49:16	22 hrs 51 mins 19 secs	ikx2k1kqp1q0j7u0cg5bp15
em7Admin	10.128.39.231	Logged Out	2017-04-17 16:28:25	2017-04-17 17:19:34	2017-04-17 17:19:34	51 mins 9 secs	08uaq2rs1ructq9h09v4305
em7Admin	10.128.39.219	Expired	2017-04-14 16:13:11	2017-04-14 16:13:11	2017-04-14 16:13:11	0 secs	gib0h904b05m0gbrko2u2gm3
em7Admin	10.128.39.228	Expired	2017-04-14 10:26:50	2017-04-14 11:30:54	2017-04-14 11:30:54	1 hr 4 mins 4 secs	gib0q5vfm58m0oub7363no261
em7Admin	10.128.39.228	Expired	2017-04-14 11:30:54	2017-04-14 11:30:54	2017-04-14 11:30:54	0 secs	gib0q5vfm58m0oub7363no261
em7Admin	10.128.39.219	Expired	2017-04-14 10:43:12	2017-04-14 10:43:12	2017-04-14 10:43:12	0 secs	u7hg4p0f0r1f5jg84k3aba0
em7Admin	10.128.38.43	Logged Out	2017-04-13 16:28:48	2017-04-13 16:41:18	2017-04-13 16:41:18	12 mins 30 secs	7nb23no6b0o3a3h0spk83o1
em7Admin	10.128.39.142	Expired	2017-04-13 12:22:58	2017-04-13 12:22:58	2017-04-13 12:22:58	0 secs	cg1p0p0p0f0r1f5jg84k3aba0
em7Admin	10.128.39.175	Logged Out	2017-04-12 10:29:16	2017-04-12 16:58:43	2017-04-12 16:58:43	6 hrs 29 mins 27 secs	oHfpmcaaqvqrmjs1c7jron4

2. For each session, the **Access Sessions** page displays:

- **User Account.** Username of person logging in to the user interface.
- **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.
- **Session ID**. Unique numeric ID assigned to each user session.

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).

User Account	Last Address	State	Login Time	Last-Hit Time	Logout Time	Session Duration	Session ID
em7admin	10.128.38.60	Logged In	2017-05-02 16:27:15	2017-05-02 16:27:15	--	14 secs	64d9m9f8aaqab20aigokiorc7
em7admin	10.128.38.59	Logged In	2017-05-02 12:12:53	2017-05-02 12:12:53	--	4 hrs 14 mins 36 secs	2d91kumog80bz2k2757dk35
em7admin	10.128.38.67	Logged Out	2017-05-01 13:13:11	2017-05-01 16:50:55	2017-05-01 16:50:55	3 hrs 37 mins 44 secs	qg3rd490g3fnaannmg7alf36
em7admin	10.128.38.67	Logged Out	2017-05-01 10:30:09	2017-05-01 11:44:24	2017-05-01 11:44:24	1 hr 14 mins 15 secs	bf8np2a10w8b6390ag168a4
em7admin	10.128.38.69	Expired	2017-04-28 15:17:59	2017-04-28 15:17:59	2017-04-28 15:17:59	0 secs	gg8ee4fmbaa26u30h3bqan4
em7admin	10.128.38.251	Expired	2017-04-26 15:47:57	2017-04-26 15:47:57	2017-04-26 15:47:57	0 secs	jm23accr5e6fman8h64538g2
em7admin	10.128.38.102	Expired	2017-04-25 09:03:21	2017-04-25 09:03:21	2017-04-25 09:03:21	0 secs	rj8qa10h6pkh9d94ppp1
em7admin	10.128.39.224	Logged Out	2017-04-21 12:48:52	2017-04-21 15:10:40	2017-04-21 15:10:40	2 hrs 21 mins 48 secs	yx0u3np3ibourd7jgnfadnbg0
em7admin	10.128.39.216	Expired	2017-04-21 13:02:34	2017-04-21 13:02:34	2017-04-21 13:02:34	0 secs	50k0t8ca7o75a01glof82b1
em7admin	10.128.38.68	Logged Out	2017-04-20 11:46:49	2017-04-20 16:02:13	2017-04-20 16:02:13	4 hrs 15 mins 24 secs	98l8p8uphunc888p87g4a0
em7admin	10.128.38.182	Expired	2017-04-20 15:31:21	2017-04-20 15:31:21	2017-04-20 15:31:21	0 secs	0p0p8h1q1q23vo4hnhg8f8f1
em7admin	10.128.38.68	Expired	2017-04-20 11:46:48	2017-04-20 11:46:48	2017-04-20 11:46:48	0 secs	ja0n1a74d5ccsrsh1aqmucc633
em7admin	10.128.38.91	Expired	2017-04-19 12:56:38	2017-04-19 12:56:38	2017-04-19 12:56:38	0 secs	780ns7cmg0a238g8ofnqno2
em7admin	10.128.38.43	Logged Out	2017-04-18 09:59:28	2017-04-18 16:17:04	2017-04-18 16:17:04	6 hrs 17 mins 38 secs	ah0p070320a3m50spjg86
em7admin	10.128.38.53	Expired	2017-04-18 15:22:25	2017-04-18 15:22:25	2017-04-18 15:22:25	0 secs	kk2k1kqp1p0q7u0cg5pp5
em7admin	10.128.38.53	Expired	2017-04-18 13:45:34	2017-04-18 15:22:25	2017-04-18 15:22:25	1 hr 36 mins 51 secs	kk2k1kqp1p0q7u0cg5pp5
em7admin	10.128.38.53	Expired	2017-04-18 11:48:16	2017-04-18 13:45:34	2017-04-18 13:45:34	1 hr 56 mins 18 secs	kk2k1kqp1p0q7u0cg5pp5
em7admin	10.128.39.235	Expired	2017-04-17 12:57:57	2017-04-18 11:49:16	2017-04-18 11:49:16	22 hrs 51 mins 19 secs	kk2k1kqp1p0q7u0cg5pp5
em7admin	10.128.39.231	Logged Out	2017-04-17 16:28:25	2017-04-17 17:19:34	2017-04-17 17:19:34	51 mins 9 secs	08uaq2rs1uctqhh9k4305
em7admin	10.128.39.219	Expired	2017-04-14 16:13:11	2017-04-14 16:13:11	2017-04-14 16:13:11	0 secs	g80h00405m0gk0a2u2gm3
em7admin	10.128.39.228	Expired	2017-04-14 10:30:50	2017-04-14 11:30:54	2017-04-14 11:30:54	1 hr 4 mins 4 secs	g80c0fmd58m0cubh7363nc261
em7admin	10.128.39.228	Expired	2017-04-14 11:30:54	2017-04-14 11:30:54	2017-04-14 11:30:54	0 secs	g80c0fmd58m0cubh7363nc261
em7admin	10.128.39.219	Expired	2017-04-14 10:43:12	2017-04-14 10:43:12	2017-04-14 10:43:12	0 secs	u7hg4p8f8f85g84r2abag0
em7admin	10.128.38.43	Logged Out	2017-04-13 16:28:48	2017-04-13 16:41:18	2017-04-13 16:41:18	12 mins 30 secs	7n523ncv0600h34hesp8d301
em7admin	10.128.39.142	Expired	2017-04-13 12:22:58	2017-04-13 12:22:58	2017-04-13 12:22:58	0 secs	cg8tpav0pk8d296mq04fus07
em7admin	10.128.39.175	Logged Out	2017-04-12 10:29:16	2017-04-12 16:58:43	2017-04-12 16:58:43	6 hrs 29 mins 27 secs	o8f8mcaajvgqjms1c7jron4

2. In the **Access Sessions** page, find the session you want to end. Click 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*.

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.



Attempt Account	From Address	Attempt Time	Trys
1. ttesterson	192.168.37.12	2012-03-20 18:59:00	3

Copyright © 2003 - 2012 ScienceLogic, Inc. All rights reserved.

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 bomb icon ().

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](#).

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*.)

Viewing the List of Schedules

The **Schedule Manager** page (Registry > Schedules > Schedule Manager) displays the following about each schedule:

Schedule Manager Schedules Found [18]														
Schedule Summary *	Schedule Description	Event ID	sch_id	Context	Timezone	Start Time	Duration	Recurrence Interval	End Date	Last Run	Owner	Organization	Visibility	Enabled
1. SAC Daily Discovery Maint	SAC Daily Discovery Maint	175	0	Discovery	America/Chicago	2017-10-01 08:00:00	30 minute	Every 2 Days	2018-10-01 08:00:00	--	AutoRegUser	System	World	Yes
2. SAC Daily Report Maint	SAC Daily Asset List Report	158	54	Reports	America/New_York	2017-10-01 08:00:00	30 minute	Every 1 Day	2018-10-01 08:00:00	--	AutoRegUser	System	World	Yes
3. SAC Hourly Ticket Maint	SAC Hourly Ticket Maintenance	155	40	Tickets	America/New_York	2017-10-01 08:00:00	30 minute	Every 24 Hours	2018-10-01 08:00:00	--	AutoRegUser	System	World	Yes
4. SAC One Time Dev Maint	SAC One Time Device Maint	165	66	Devices	America/New_York	2017-10-01 08:00:00	30 minute	--	--	--	AutoRegUser	System	World	Yes
5. SAC One Time Dev Maint	SAC One Time Device Maint	153	62	Devices	America/New_York	2017-10-01 08:00:00	30 minute	--	--	--	AutoRegUser	System	World	Yes
6. SAC Weekly IT Service Maint	SAC Weekly IT Service Sch	160	0	IT Services	America/New_York	2017-10-01 08:00:00	30 minute	Every 1 Week	2018-10-01 08:00:00	--	AutoRegUser	System	World	Yes
7. sch_1	sch_1_admin-system-world	166	0	Discovery	America/New_York	2017-10-01 08:00:00	30 minute	--	--	--	em7admin	System	World	Yes
8. sch_2	sch_2_admin-system-org	167	0	Discovery	America/New_York	2017-10-01 08:00:00	30 minute	--	--	--	em7admin	System	Organza	Yes
9. sch_4	sch_4_admin-system-world	169	0	Discovery	America/New_York	2017-10-01 08:00:00	30 minute	--	--	--	sac_user1	SAC_Scheduler_T	World	Yes
10. sch_5	sch_5_admin-system-org	170	0	Discovery	America/New_York	2017-10-01 08:00:00	30 minute	--	--	--	sac_user1	SAC_Scheduler_T	Organza	Yes
11. sch_7	sch_7_admin-system-world	172	0	Discovery	America/New_York	2017-10-01 08:00:00	30 minute	--	--	--	sac_user2	System	World	Yes
12. sch_8	sch_8_admin-system-org	173	0	Discovery	America/New_York	2017-10-01 08:00:00	30 minute	--	--	--	sac_user2	System	Organza	Yes
13. Schum Fu Pandas Discover	--	32	0	Discovery	America/New_York	2016-07-25 23:00:00	30 minute	Every 1 Day	--	2017-01-17 23:30:00	em7admin	System	World	Yes
14. Scrummy Bears Discovery	--	33	0	Discovery	America/New_York	2016-07-25 23:00:00	30 minute	Every 1 Day	--	2017-01-17 23:30:00	em7admin	System	World	Yes
15. System Patch Install - versio	--	34	862	Patches	UTC	2016-07-26 21:15:00	30 minute	Every 0 Minutes	--	2016-07-26 21:45:00	em7admin	System	World	Yes
16. System Patch Install - versio	--	50	1714	Patches	UTC	2016-11-11 18:33:00	30 minute	Every 0 Minutes	--	2016-10-11 19:03:00	em7admin	System	World	Yes
17. System Patch Install - versio	--	58	2169	Patches	UTC	2016-11-09 21:13:00	30 minute	Every 0 Minutes	--	2016-11-09 21:43:00	em7admin	System	World	Yes
18. System Patch Install - versio	--	125	2530	Patches	UTC	2016-12-11 19:28:00	30 minute	Every 0 Minutes	--	2016-12-11 19:58:00	em7admin	System	World	Yes

- **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. SL1 automatically creates this ID for each scheduled process.
- **sch id.** Displays a unique, numeric ID for the schedule. SL1 automatically creates this ID for each schedule.
- **Context.** Displays the area of SL1 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).

The screenshot shows the 'Schedule Manager | Schedules Found [18]' page. It features a table with columns: Schedule Summary, Schedule Description, Event ID, sch_id, Context, Timezone, Start Time, Duration, Recurrence Interval, End Date, Last Run, Owner, Organization, Visibility, and Enabled. The table lists 18 scheduled tasks, including 'SAC Daily Discovery Maint', 'SAC Daily Report Maint', 'SAC Hourly Ticket Maint', etc. A red box highlights a dropdown menu in the bottom right corner with the following options: Administration, DELETE Schedules, ENABLE Schedules, and DISABLE Schedules. The 'DISABLE Schedules' option is currently selected.

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).

The screenshot shows the 'Schedule Manager | Schedules Found [18]' page. It contains a table with columns: Schedule Summary, Schedule Description, Event ID, ssh_id, Context, Timezone, Start Time, Duration, Recurrence Interval, End Date, Last Run, Owner, Organization, Visibility, and Enabled. The table lists 18 entries, including 'SAC Daily Discovery Maint', 'SAC Daily Report Maint', 'SAC Hourly Ticket Maint', etc. A dropdown menu is open over the table, showing options: Administration, Enable Schedules, and Disable Schedules. The 'Go' button is also visible.

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.

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 SL1
- A pie graph showing the percent of the total data-collection load handled by each Data Collector or **Collector Group**
- 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.

NOTE: The pie graph does not appear for All-In-One Appliances.

To view the **System Usage** page:

1. Go to the **System Usage** page (System > Monitor > System Usage).
2. The **System Usage** page appears:

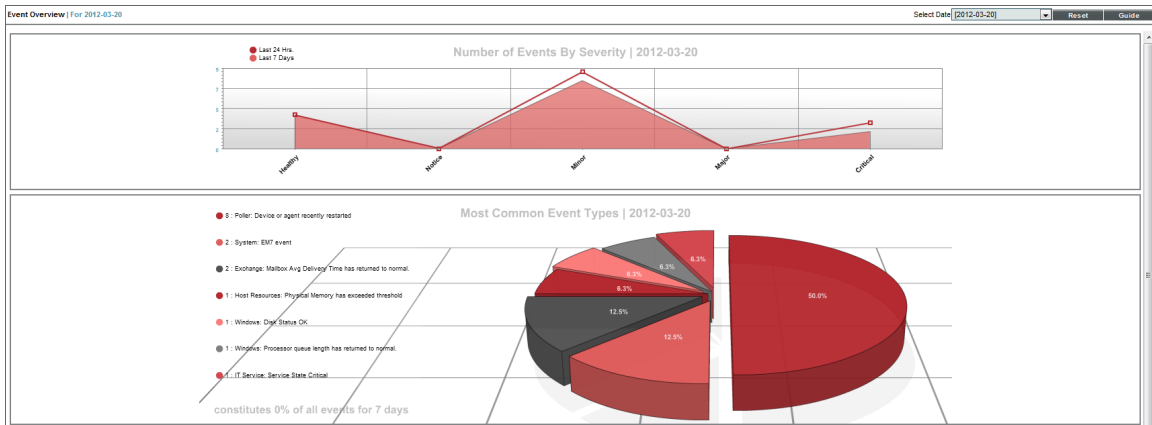


Viewing an Overview of All Events


The **Event Overview** page provides a graphical overview of all events in SL1.

To view the **Event Overview** page:


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




2. 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.



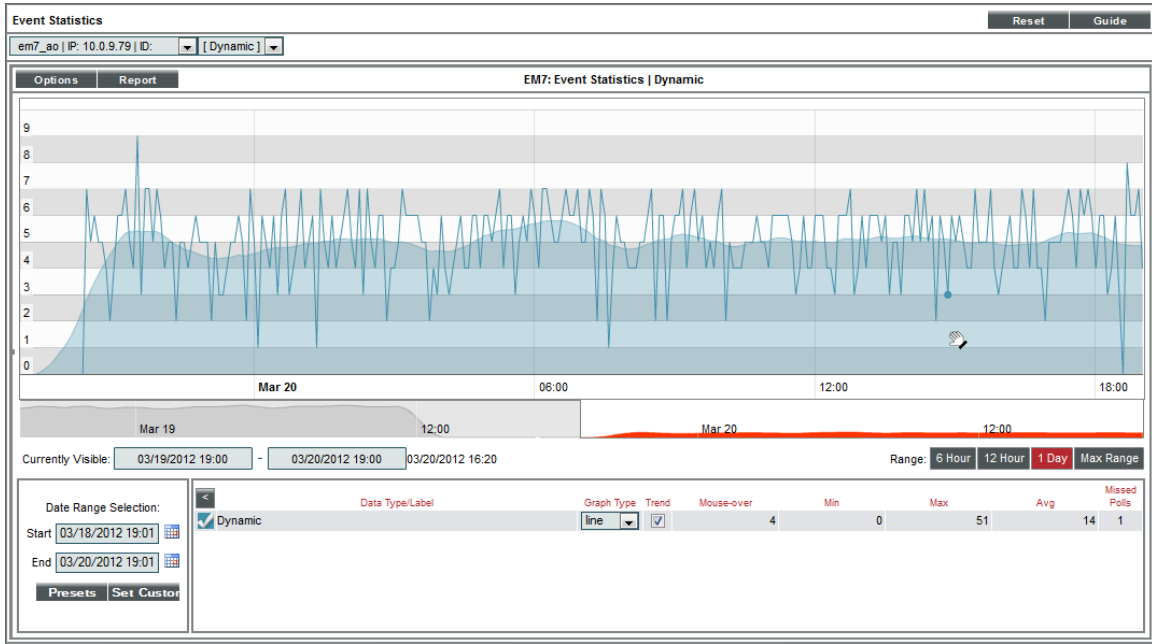
- **Mean Time-to-Resolution.** This bar graph displays the number of events generated in the last 24 hours, 7 days, 14 days, and 30 days, and their average resolution time.
 - The y-axis displays the number of events.
 - The x-axis displays the time span. There is a bar for 24 hours, 7 days, 14 days, 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.

To generate the report:

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



2. In the **Event Statistics** page, supply values in the following fields:

- **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 field in the upper right, select from the list of event types. The choices are:
 - *Syslog*. Event was generated from standard system log generated by device.
 - *Internal*. Event was generated by SL1.
 - *Trap*. Event was generated by an SNMP trap.
 - *Dynamic*. Event was generated by a monitoring application running on the device.
 - *API*. The event was generated by an external API.
 - *Email*. The event was generated by an incoming email.

3. 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

Diagnostic Tools

Overview

This chapter describes some diagnostic tools for troubleshooting and diagnosing problems in SL1.

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

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

This chapter includes the following topics:

ScienceLogic SL1 Self-Monitoring	188
Viewing Information About ScienceLogic Processes	188
<i>Viewing the List of ScienceLogic Processes</i>	189
<i>Searching and Filtering the List of ScienceLogic Processes</i>	191
<i>Editing the Parameters of a ScienceLogic Process</i>	192
Debugging a Process and Viewing Debug Logs	194
Viewing Information About Unhandled Exceptions	196
<i>Viewing the List of Unhandled Exceptions</i>	196
<i>Searching and Filtering the list of Unhandled Exceptions</i>	197
<i>Saving the Unhandled Exception to the Local Computer</i>	198
Viewing the Output of the System Status Script	198
Viewing the Database Tables on the Database Server	199
<i>Accessing the Database Tool</i>	199

<i>Disabling Normalization, Re-Enabling Normalization, and Backfilling Raw Data</i>	201
<i>Enable Logging for Data Pull Storage Objects</i>	204
Enable	204
Disable	204
<i>Controlling Log Settings</i>	205
<i>Setting UI Developer Log Levels</i>	205
<i>Setting UI/REST MySQL Query Log Levels</i>	205
<i>Configuring Advanced Log Settings</i>	206
<i>Downloading Logs from the PHP Developer Logs page</i>	206

ScienceLogic SL1 Self-Monitoring

SL1 includes the *ScienceLogic Support Pack PowerPack*, which monitors the SL1 platform and processes. This PowerPack includes the following Dynamic Applications:

- Support: Config Push
- Support: Configuration Parser (MY CNF)
- Support: Database Sizes
- Support: Datapull Configuration
- Support: Datapull Stats
- Support: DRDB Proxy State
- Support: File System
- Support: InnoDB Size
- Support: Maintenance Tracking
- Support: MySQL Performance
- Support: Platform Statistics
- Support: PT-DiskStats
- Support: Rows Behind
- Support: ScienceLogic Configuration

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 SL1.

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 SL1 is running.

SL1 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.

The **Process Manager** page allows you to view and edit the parameters of system processes.

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).

Process Name	Program File	Frequency	Runtime Offset	Async Throttle	Batch Factor	Time Factor	Run Length	State	Debug	ID	Edited By	Edit Date
Application & Report Server: Remote diagnostic	em7_httpd_admin	0	--	--	--	--	--	Enabled	Disabled	54	em7admin	2009-06-29 14:06:59
Application & Report Server: Scheduled Report Runner	scheduled_report_run.py	-1	--	25	--	15	15	Enabled	Disabled	53	em7admin	2009-07-14 12:20:00
Application & Report Server: Secure	em7_httpd	0	--	--	--	--	--	Enabled	Disabled	53	em7admin	2009-06-29 14:06:59
Application & Report Server: Standard	em7_httpd	0	--	--	--	--	--	Enabled	Disabled	52	em7admin	2009-06-29 14:06:59
Data Collection: Async Dynamic App Collection	async_dynamic_collect.py	-1	--	2	--	15	15	Enabled	Disabled	129	em7admin	2010-03-04 11:53:41
Data Collection: Availability	availability_collect.py	5	2	--	30	5	30	Enabled	Disabled	10	em7admin	2009-06-29 14:06:59
Data Collection: CDP Collection	cdp_collect.py	120	0	--	30	0	120	Enabled	Disabled	33	em7admin	2010-03-28 10:37:39
Data Collection: Critical Availability	em7_cavald	0	--	--	--	--	--	Enabled	Disabled	47	em7admin	2009-06-29 14:06:59
Data Collection: Critical Port	em7_polc	0	--	--	--	--	--	Enabled	Disabled	48	em7admin	2009-06-29 14:06:59
Data Collection: DNS Policy Monitoring	dns_collect.py	5	2	--	30	5	30	Enabled	Disabled	29	em7admin	2009-06-29 14:06:59
Data Collection: Dynamic App	dynamic_collect.py	1	0	--	20	15	16	Enabled	Disabled	11	em7admin	2009-06-29 14:06:59
Data Collection: Dynamic Refresh	dynamic_check.py	1440	200	--	30	0	1440	Enabled	Disabled	28	em7admin	2009-06-29 14:06:59
Data Collection: E-Mail Round-Trip	email_rt_collect.py	5	3	--	30	0	5	Enabled	Disabled	30	em7admin	2009-06-29 14:06:59
Data Collection: Filesystem statistics	filesystem_stats_collect.py	5	0	--	30	0	5	Enabled	Disabled	32	em7admin	2009-06-29 14:06:59
Data Collection: Host Filesystem Inventory	filesystem_inventory_collect.py	120	44	--	30	0	120	Enabled	Disabled	31	em7admin	2009-06-29 14:06:59
Data Collection: Interface Bandwidth	if_collect.py	1	0	--	20	10	11	Enabled	Disabled	12	em7admin	2009-06-29 14:06:59
Data Collection: L3 Topology Collection	l3topology_collect.py	120	90	--	30	1	240	Enabled	Disabled	34	em7admin	2010-03-28 10:37:39
Data Collection: OS Process	process_collect.py	120	0	--	20	0	120	Enabled	Disabled	14	em7admin	2009-06-29 14:06:59
Data Collection: OS Process Check	process_check.py	5	4	--	20	2	15	Enabled	Disabled	15	em7admin	2009-06-29 14:06:59
Data Collection: OS Service	service_collect.py	120	20	--	20	0	120	Enabled	Disabled	16	em7admin	2009-06-29 14:06:59
Data Collection: OS Service Check	service_check.py	5	0	--	30	2	15	Enabled	Disabled	17	em7admin	2009-06-29 14:06:59
Data Collection: RSS Event Feed	rss_collect.py	10	0	--	30	0	10	Enabled	Disabled	23	em7admin	2009-06-29 14:06:59
Data Collection: SNMP Detail	snmp_detail_collect.py	5	0	--	30	0	5	Enabled	Disabled	24	em7admin	2009-06-29 14:06:59
Data Collection: TCP Port Monitor	port_collect.py	5	0	--	30	0	5	Enabled	Disabled	20	em7admin	2009-06-29 14:06:59
Data Collection: Topology	topology_collect.py	60	12	--	30	0	60	Enabled	Disabled	25	em7admin	2009-06-29 14:06:59

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 SL1 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.

*number of tasks a process is responsible for completing/**Batch Factor** = 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 SL1 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 SL1.
- **Edited By**. Date and time the process settings were last edited.
- **Edit Date**. Date and time the process settings were last edited.

Searching and Filtering the List of ScienceLogic Processes

The **Process Manager** page includes 13 filters, in the top row in the list of processes. You can specify one or more parameters to filter the display of processes. Only processes that meet all the filter criteria will be displayed in the **Process Manager** page.

You can filter by one or more of the following parameters. The list of processes is dynamically updated as you select each filter.


- For each filter except **Edit Date**, you must enter text to match against. SL1 will search for processes that match the text, including partial matches. Text matches are not case-sensitive. You can use the following special characters in each filter:
 - , (comma). Specifies an "or" operation. For example:
"dell, micro" would match all values that contain the string "dell" OR the string "micro".
 - & (ampersand). Specifies an "and" operation. For example:
"dell & micro" would match all values that contain the string "dell" AND the string "micro".
 - ! (exclamation mark). Specifies a "not" operation. For example:
"!dell" would match all values that do not contain the string "dell".
- **Process Name**. You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Process Manager** page will display only processes that have a matching name.
- **Program File**. You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Process Manager** page will display only processes that have a matching program file.
- **Frequency**. You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Process Manager** page will display only processes that have a matching frequency number.
- **Runtime Offset**. You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Process Manager** page will display only processes that have a matching runtime offset.
- **Async Throttle**. You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Process Manager** page will display only processes that have a matching throttle number.
- **Batch Factor**. You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Process Manager** page will display only processes that have a matching batch factor.

- **Time Factor.** You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Process Manager** page will display only processes that have a matching time factor.
- **Run Length.** You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Process Manager** page will display only processes that have a matching run length.
- **State.** You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Process Manager** page will display only processes that have a matching state ("Enabled" or "Disabled").
- **Debug.** You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Process Manager** page will display only processes that have a matching debug state ("Enabled" or "Disabled").
- **ID.** You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Process Manager** page will display only processes that have a matching ScienceLogic process ID.
- **Edited By.** You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Process Manager** page will display only processes that have a matching "created by" or "edited by" value.
- **Edit Date.** You can select from a list of time periods. The **Process Manager** page will display only processes that have been created or edited within that time period:
 - *All.* Display all processes that match the other filters.
 - *Last Minute.* Display only processes that have been edited within the last minute.
 - *Last Hour.* Display only processes that have been edited within the last hour.
 - *Last Day.* Display only processes that have been edited within the last day.
 - *Last Week.* Display only processes that have been edited within the last week.
 - *Last Month.* Display only processes that have been edited within the last month.
 - *Last Year.* Display only processes that have been edited within the last year.

Editing the Parameters of a ScienceLogic Process

To view details about a specific process or edit the settings for a specific process:

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.

1. Go to the **Process Manager** page (System > Settings > Admin Processes).
2. In the **Process Manager** page, find the process you want to edit. Click its wrench icon ()

3. The **Process Editor** page appears and is populated with values from the selected process.

- **Process Name.** Name of the process. This field is read-only and cannot be changed.
- **Program File.** Name of the executable file associated with the process. This field is read-only and cannot be changed.
- **Operating State.** Current operational state of the process. Specifies whether the process is enabled and able to run. Select from the drop-down list. The choices are:
 - *Enabled.* Process can run.
 - *Disabled.* Process cannot run.
- **Debug Mode.** Enables or disables debugging information for a process. For more details on debugging a process, see the section [Debugging a Process](#).

NOTE: You cannot enable debug mode for the process *Message Collection: SNMP Trap*.

- **Frequency.** This field appears only for scheduled processes and asynchronous processes. Specifies the frequency with which SL1 launches the process. Select from the drop-down list. The choices are:
 - *Asynchronous.* For asynchronous processes, this is the only available option. You cannot edit the frequency.
 - *Scheduled.* For scheduled processes, you can edit the frequency. You can select from intervals ranging from 1 Minute to Daily.

NOTE: If a process is set to a frequency of *Asynchronous* or *Always*, this field cannot be changed. If a process is set to a time interval, this field cannot be changed to *Asynchronous* or *Always*.

- **Async Throttle.** This field appears only for asynchronous processes. This field indicates the number of jobs per process that can run simultaneously. This setting only applies to asynchronous processes.
- **Runtime Offset.** This field only appears for scheduled processes. This field allows SL1 to stagger the launch of a process. The value specified in this field specifies minutes after the default scheduled time for a process. For example, if a process has a **Frequency** of *5 Minutes* and the **Minute Offset** is set to "2", the process will execute at two minutes past the hour, seven minutes past the hour, 12 minutes past the hour, 17 minutes past the hour, etc. Choices range from 0–1439.

- **Batch Factor.** This field applies only to scheduled processes and determines how many multithreaded child processes are spawned on each execution of the process.

number of tasks a process is responsible for completing/**Batch Factor** = *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 SL1 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.** This field appears only for scheduled processes and asynchronous processes. This field determines how long a process can run before being killed.

- For scheduled processes, SL1 uses the formula **(Frequency * Time Factor) + Frequency**.

For example, suppose a process runs every 15 minutes. A factor of 2 means the process is allowed to run for 45 minutes. Factor of 0 means process is allowed to run for 15 minutes.

- For asynchronous processes, SL1 simply uses the value in this field as the number of minutes a process can run. This field does not appear for processes that are always running.

- **Appliance Types.** Specifies the appliance types where the process is allowed to run.

NOTE: All changes to the settings in the **Process Manager** page are logged in the **Audit Logs** page (System > Monitor > Audit Logs). The associated log entry will specify the user who altered a process, the process that was altered, and which settings for the process were changed.

4. If you make changes to one or more fields, click the **[Save]** button to save your changes.

Debugging a Process and Viewing Debug Logs


When you debug a process, you tell SL1 to use verbose logging for that process. You can then view SL1 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 and that you then immediately turn off debugging when you have completed troubleshooting. Don't leave the debug option enabled during normal operation of SL1. When you turn on debugging, SL1 will run significantly more slowly.

NOTE: You cannot enable debug mode for the process *Message Collection: SNMP Trap*.

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. The default password is **em7admin**.

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 **silو.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).

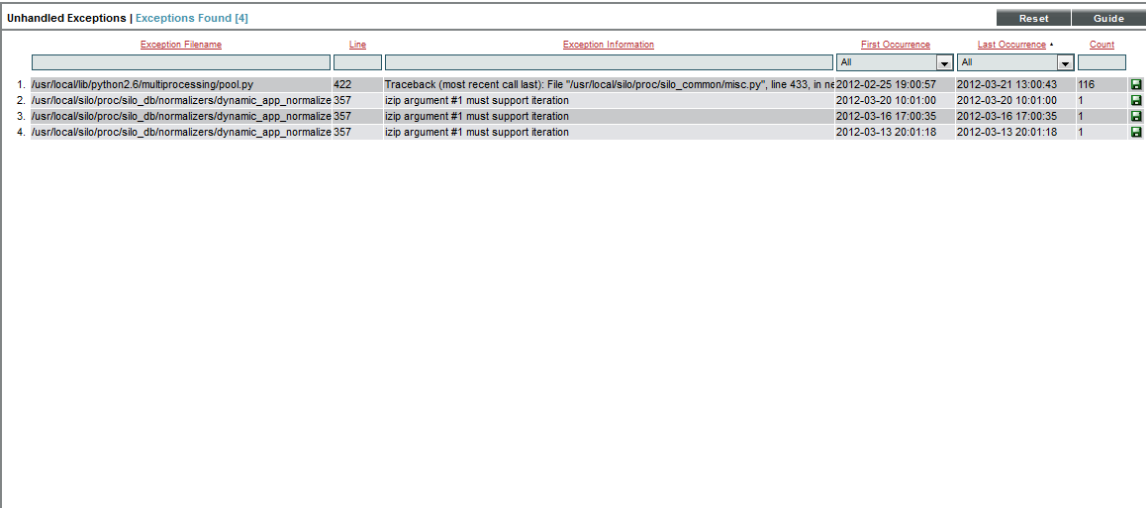
Viewing Information About Unhandled Exceptions

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, SL1 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 SL1 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	Line	Exception Information	First Occurrence	Last Occurrence	Count
/usr/local/lib/python2.6/multiprocessing/pool.py	422	Traceback (most recent call last): File "/usr/local/silo/proc/silo_common/misc.py", line 433, in nt	2012-02-25 19:00:57	2012-03-21 13:00:43	116
/usr/local/silo/proc/silo_db/normalizers/dynamic_app_normalize	357	izip argument #1 must support iteration	2012-03-20 10:01:00	2012-03-20 10:01:00	1
/usr/local/silo/proc/silo_db/normalizers/dynamic_app_normalize	357	izip argument #1 must support iteration	2012-03-16 17:00:35	2012-03-16 17:00:35	1
/usr/local/silo/proc/silo_db/normalizers/dynamic_app_normalize	357	izip argument #1 must support iteration	2012-03-13 20:01:18	2012-03-13 20:01:18	1

- **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.

Searching and Filtering the list of Unhandled Exceptions

The **Unhandled Exceptions** page includes six filters. You can filter the list of exceptions by one or multiple of the following parameters: exception filename, line number, exception descriptions, first occurrence, last occurrence, and count. Only exceptions that meet all the filter criteria will be displayed in the **Unhandled Exceptions** page.

You can filter by one or more of the following parameters. The list of devices is dynamically updated as you select each filter.


- For the first three filters, you must enter text to match against. SL1 will search for exceptions that match the text, including partial matches. Text matches are not case sensitive. You can use the following special characters in each filter:
 - `,` Specifies an "or" operation. For example:
 - "dell, micro" would match all values that contain the string "dell" OR the string "micro".
 - `!` Specifies a "not" operation. For example:
 - "!dell" would match all values that do not contain the string "dell".
- **Exception Filename**. You can enter text to match, including special characters, and the **Unhandled Exceptions** page will display only exceptions that have a matching filename.
- **Line**. You can enter text to match, including special characters, and the **Unhandled Exceptions** page will display only exceptions that have a matching line number.
- **Exception Information**. You can enter text to match, including special characters, and the **Unhandled Exceptions** page will display only exceptions that have a matching description.
- **First Occurrence**. Only those exceptions that match all the previously selected fields and have the specified first occurrence date will be displayed. The choices are:
 - *All*. Display exceptions with all first occurrence dates.
 - *Last Minute*. Display only exceptions that first occurred within the last minute.
 - *Last Hour*. Display only exceptions that first occurred within the last hour.
 - *Last Day*. Display only exceptions that first occurred within the last day.
 - *Last Week*. Display only exceptions that first occurred within the last week.
 - *Last Month*. Display only exceptions that first occurred within the last month.
 - *Last Year*. Display only exceptions that first occurred within the last year.
- **Last Occurrence**. Only those exceptions that match all the previously selected fields and have the specified last occurrence date will be displayed. The choices are:
 - *All*. Display exceptions with all last occurrence dates.
 - *Last Minute*. Display only exceptions that last occurred within the last minute.
 - *Last Hour*. Display only exceptions that last occurred within the last hour.
 - *Last Day*. Display only exceptions that last occurred within the last day.
 - *Last Week*. Display only exceptions that last occurred within the last week.

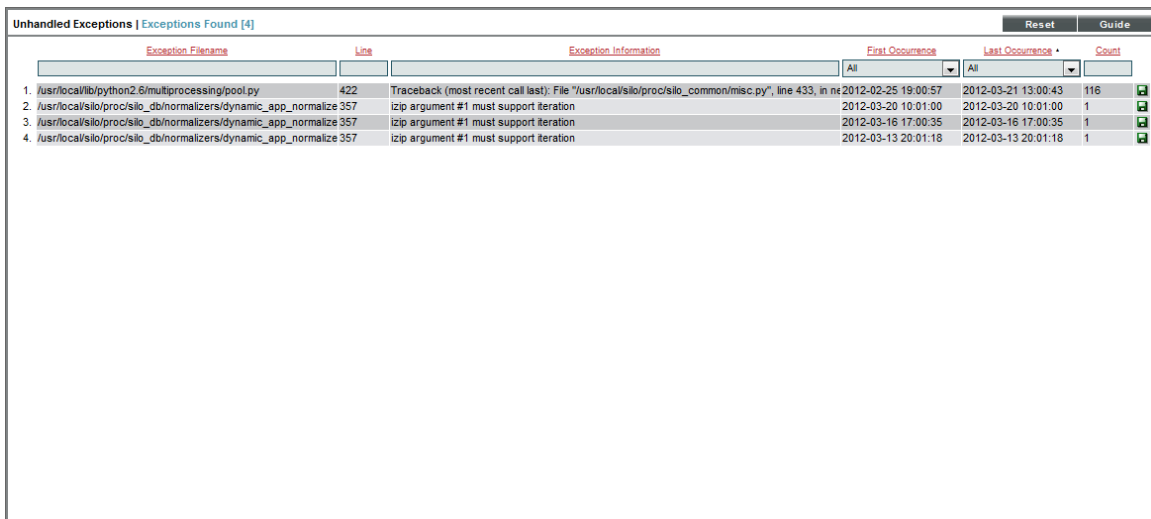
- *Last Month*. Display only exceptions that last occurred within the last month.
- *Last Year*. Display only exceptions that last occurred within the last year.
- **Count**. You can enter text to match, including special characters, and the **Unhandled Exceptions** page will display only exceptions that have a matching count number.

Saving the Unhandled Exception to the Local Computer

You can save the full text of the unhandled exception to a file on your local computer. You can then view the text in a text editor.

To save the full text of the unhandled exception to a file:

1. Go to the **Unhandled Exceptions** page (System > Monitor > Unhandled Exceptions).
2. In the **Unhandled Exceptions** page, find the exception you want to save to a file. Click its green diskette icon ().




Exception Filename	Line	Exception Information	First Occurrence	Last Occurrence	Count
			All	All	
1. /usr/local/lib/python2.6/multiprocessing/pool.py	422	Traceback (most recent call last): File "/usr/local/silo/proc/silo_common/misc.py", line 433, in ne	2012-02-25 19:00:57	2012-03-21 13:00:43	116
2. /usr/local/silo/proc/silo_db/normalizers/dynamic_app_normalize	357	zip argument #1 must support iteration	2012-03-20 10:01:00	2012-03-20 10:01:00	1
3. /usr/local/silo/proc/silo_db/normalizers/dynamic_app_normalize	357	zip argument #1 must support iteration	2012-03-16 17:00:35	2012-03-16 17:00:35	1
4. /usr/local/silo/proc/silo_db/normalizers/dynamic_app_normalize	357	zip argument #1 must support iteration	2012-03-13 20:01:18	2012-03-13 20:01:18	1


3. When prompted, you can either immediately view the text file with a text editor or save the file to your local computer for viewing later.

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).
- 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.

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:

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



2. 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.

3. Click the **[Go]** button to execute the query.
4. The results from the query are displayed in the pane at the bottom of the page.
5. 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 SL1, 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 SL1 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 SL1 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 vi /etc/siteconfig/siloconf.siteconfig
```

4. This is the file where users can customize the silo.conf file. In step #7, you will execute a command that sends these changes to the system silo.conf file.
5. In the LOCAL section, add the following line:

```
rollups_disabled=ON
```

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

```
sudo /opt/em7/share/scripts/generate-silo-conf.py > silo.conf
```

8. You must 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 vi /etc/siteconfig/siloconf.siteconfig
```

4. This is the file where users can customize the silo.conf file. In step #7, you will execute a command that sends these changes to the system silo.conf file.
5. In the LOCAL section, add the following line:

```
rollups_disabled=OFF
```

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

```
sudo /opt/em7/share/scripts/generate-silo-conf.py > silo.conf
```

8. You must 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
```

9. 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, SL1 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 time stamp at this time or later will be normalized. SL1 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 time stamp at this time or earlier will be normalized. SL1 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 vi /etc/silo.conf
```

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 vi /etc/silo.conf
```

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 Log Settings

In rare cases, you may need to modify log levels or suppression of certain logs in SL1, usually at the request of ScienceLogic Customer Support. To do so, you will navigate to the **PHP Developer Logs** page (System > Tools > PHP Developer Logs). This section describes the options included on the PHP Developer Logs page.

NOTE: This page is only available in the classic user interface.

Setting UI Developer Log Levels

When configuring logging on an appliance, you must specify a log level. The log level controls the types of messages that are 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

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, you would sum one, two, and four to get a log level value of seven.

Setting UI/REST MySQL Query Log Levels

The UI/REST MySQL Query Log Levels settings let you specify the log level for the `mysql1.log`. This log file collects every PHP-based call to MySQL and includes general information about the query. Determine the granularity of data you want and select one or more checkboxes.

- **Error**
- **Warning**
- **Info (non-error)**

Configuring Advanced Log Settings

In the *Advanced Settings* section, you can configure the datetime format you want

- **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 datetime stamp.

Downloading Logs from the PHP Developer Logs page

To download the logs from the **PHP Developer Logs** page (System > Tools > PHP Developer Logs):

1. Go to the **PHP Developer Logs** page (System > Tools > PHP Developer Logs)
2. Under **Download Logs**, select a logfile to download.

Chapter

8

Changing Passwords



Overview

This chapter describes how to:

- Change every administrator password used in 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 SL1 user interface:

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

This chapter includes the following topics:

<i>Disabling phpMyAdmin</i>	208
<i>Changing the Password for the Default Account for the User Interface</i>	209
<i>Changing the Password for the Default Console User</i>	210
<i>Changing the Password for the Web Configuration Utility</i>	210
<i>Changing Database Passwords</i>	211
<i>Configuring a New Password on Collector Appliances</i>	211
<i>Editing Silo.Conf</i>	213

Updating the master.system_settings_licenses Table	214
Changing the MySQL Root Password on Database Appliances	214
Recovering the Root MySQL Password	215
Recovering the MySQL SNMP User Account on Collector Appliance	216

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 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. If you are using a distributed system, either go to the console of the Database Server or use SSH to access the Database Server. Open a shell session on the server. Log in as "root".
2. If you are using an All-In-One Appliance, either go to the console of the All-In-One Appliance or use SSH to access the All-In-One Appliance. Open a shell session on the server. Log in as "root".

NOTE: For details on enabling and using SSH, see the manual *System Administration*. For details and warnings about root access and instructions on how to make root access secure, see the manual *System Administration*.

3. Open a vi session to edit the file `/etc/siteconfig/firewalld-rich-rules.siteconfig`
4. Add the following lines:

```
rule service name="phpmyadmin" reject
rule port port="8008" protocol="tcp" reject
```

5. Save your changes and exit the file.
6. Tell SL1 to pick up the changes to firewalld. To do this, type the following at the command line:


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

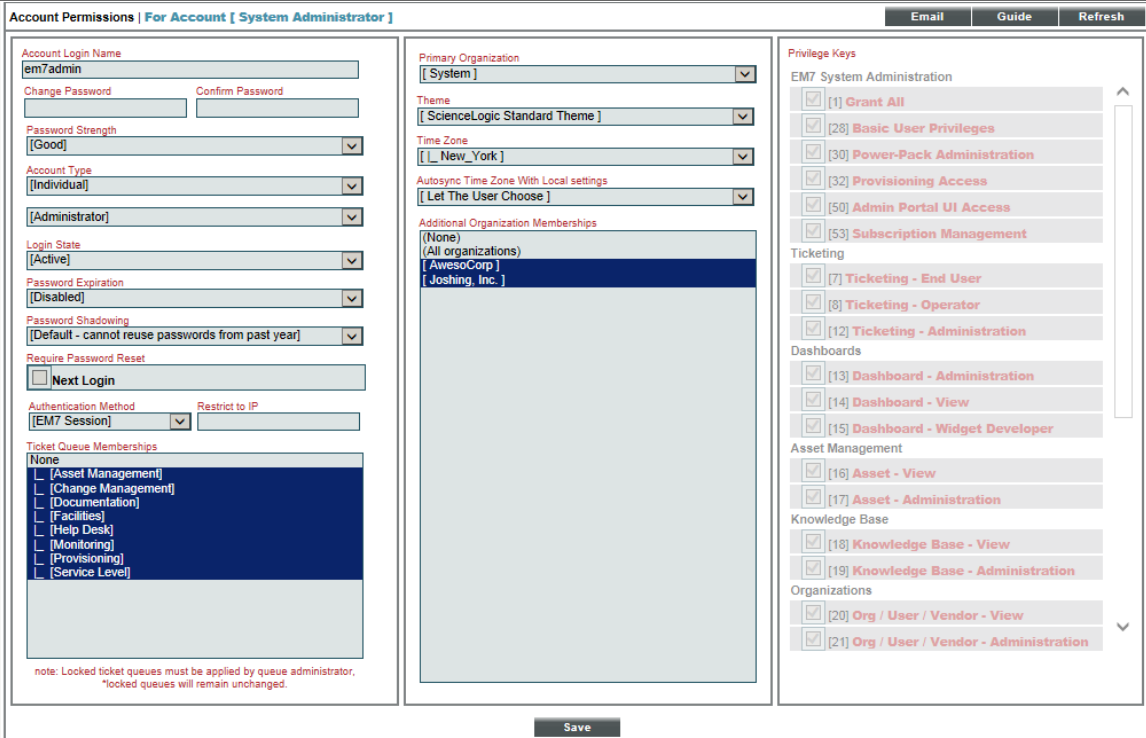
7. Restart the firewall services so that the phpMyAdmin service and port 8008 will no longer be allowed. To do this, type the following at the command line:

```
sudo systemctl restart firewalld
sudo systemctl restart nginx
```


Changing the Password for the Default Account for the User Interface

To change the password for the default em7admin user account, which can be used to access the user interface, perform the following steps:

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.



Account Permissions | For Account [System Administrator]

Account Login Name: em7admin

Change Password: [] Confirm Password: []

Password Strength: [Good]

Account Type: [Individual]

Login State: [Active]

Password Expiration: [Disabled]

Password Shadowing: [Default - cannot reuse passwords from past year]

Require Password Reset: Next Login

Authentication Method: [EM7 Session] Restrict to IP: []

Ticket Queue Memberships: [None]

Primary Organization: [System]

Theme: [ScienceLogic Standard Theme]

Time Zone: [New_York]

Autosync Time Zone With Local settings: [Let The User Choose]

Additional Organization Memberships: [None], [All organizations], [AwesoCorp], [Joshing, Inc.]

Privilege Keys: [11] Grant All, [28] Basic User Privileges, [30] Power-Pack Administration, [32] Provisioning Access, [50] Admin Portal UI Access, [53] Subscription Management

Ticketing: [7] Ticketing - End User, [8] Ticketing - Operator, [12] Ticketing - Administration

Dashboards: [13] Dashboard - Administration, [14] Dashboard - View, [15] Dashboard - Widget Developer

Asset Management: [16] Asset - View, [17] Asset - Administration

Knowledge Base: [18] Knowledge Base - View, [19] Knowledge Base - Administration

Organizations: [20] Org / User / Vendor - View, [21] Org / User / Vendor - Administration

Save

3. Enter the new password in the **Change Password** field.
4. Re-type the new password in the **Confirm Password** field.
5. Click the **[Save]** button. A pop-up window appears, asking you to confirm the change.
6. Click "OK" in the pop-up window. 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 SL1 appliance or use SSH to access the server.
2. Log in as user **em7admin** with the appropriate password. The default password is **em7admin**.
3. At the shell prompt, type the following:

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

Changing the Password for the Web Configuration Utility

You can change the password for the Web Configuration Utility.

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

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

To change the password for the Web Configuration Utility:

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

ScienceLogic™ Web Configuration Utility

Home Licensing Interfaces Device Settings PhoneHome Logout

Settings

Configure your appliance.

Web Configuration Username
em7admin

Web Config Password (change only) Confirm Web Config Password

Appliance Type
Database

Save

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

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

Changing Database Passwords

The following SL1 appliances include a database instance:

- All-In-One Appliances
- Database Servers
- Data Collectors
- Message Collectors

By default, SL1 appliances use the following user accounts to access appliance databases:

- **ap_user**. This user is used by the user interface to access the database on a Database Server or All-In-One Appliance. This user account exists only on the Administration Portal and does not exist by default on Data Collectors and Message Collectors. By default, this user has the user name **apuser** and the password **apuser**.
- **dbuser**. This user is used by ScienceLogic platform processes to access the database instance on all appliances. By default, this user has the user name **root**.

To change the password for the **ap_user** account, you must:

1. Configure a new password for the Administration Portal using the Web Configuration Utility for the Administration Portal.

To change the password for these **dbuser** account, you must:

1. Configure a new password in the database instance.
2. Configure SL1 to use the new password.

WARNING: Exercise caution when manipulating MySQL user accounts. Do not use these procedures unless you are confident and know how to undo changes, should something go wrong.

Configuring a New Password on Collector Appliances

Perform the following steps to change the password for a user on a 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 in CLI mode.
2. Log in as **em7admin** with the appropriate password.
3. Run the following command to launch the MySQL prompt:

```
silo_mysql mysql
```
4. From the MySQL prompt, change the root password by running the following SQL query:

```
UPDATE root SET password=PASSWORD('{new password}') WHERE User='root';
```
5. To effect the change immediately, run the following SQL query:

```
FLUSH PRIVILEGES;
```
6. 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 root -p
```
7. Edit the `/etc/silo.conf` file and change the **dbpasswd** variable to the new password. See [Editing Silo.Conf](#) for assistance.
8. From the SL1 interface, go to the **Appliances** page (System > Settings > Appliances) and retrieve the appliance ID for the Collector.

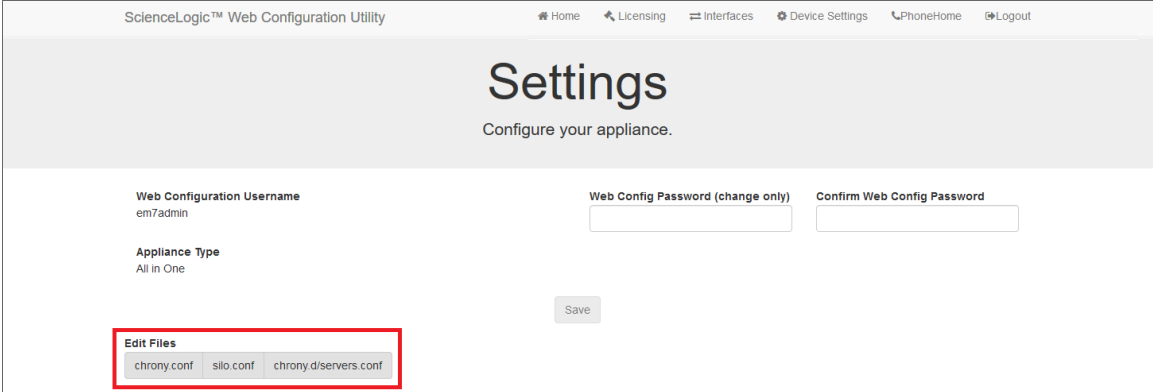


9. Go to the **DB Tool** page (System > Tools > DB Tool), and enter the following query to allow the Database Appliance to access the MySQL database of the Collector:

```
UPDATE master.system_settings_licenses SET db_pass='{new password}' WHERE id={appliance ID} LIMIT 1;
```
10. Confirm in the Collector Status page (System > Monitor > Collector Status) that the Collector is available.

Editing Silo.Conf

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



3. In the Edit Files section, click **silo.conf**. The Silo.conf Editor modal appears:



4. Edit the value assigned to **dbuser** and to **ap_user**. Assign the value you defined in the section [Configuring a New Password in the Database Instance](#).
5. To save your changes, click **Save** and then close the modal.

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).
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.

Changing the MySQL Root Password on Database Appliances

To change the MySQL root password on Database 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.

NOTE: If your Database Appliances are part of an HA cluster, place your HA cluster in maintenance mode using the steps found in the *High Availability and Disaster Recovery* manual.

3. Run the following command to launch the MySQL prompt:
`silo_mysql mysql`
4. From the MySQL prompt, change the root password by running the following SQL query:
`UPDATE user SET password=PASSWORD('{new password}') WHERE User='root';`
5. To effect this change immediately, run the following SQL query. Enter the new password when prompted.
`FLUSH PRIVILEGES;`

6. Edit the `/etc/silo.conf` file and change the **dbpasswd** variable to the new password in both the [LOCAL] and [CENTRAL] sections.
`mysql -u root -p`

NOTE: If you have clustered database appliances, be sure to update the `silo.conf` file for all cluster members.

7. If you have admin portals, update the **dbpasswd** variable in `silo.conf` on all admin portals.
8. If the Collectors' MySQL root user password is now different from the MySQL root user password on the Database Appliance, and the **db_pass** column in **master.system_settings_licenses** is "NULL", then the Database Appliance will attempt to use its own password to connect.

Change the **db_pass** column for the Collectors to their root MySQL user password using the instructions in [Updating the master.system_settings_licenses Table](#).

9. If you placed your HA cluster into maintenance mode to perform these steps, remember to return it to ready mode by setting `coro_config` to option 1. For more information, see the **High Availability and Disaster Recovery** manual.

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 by entering the following command:
`systemctl stop em7 mariadb`
4. Start the `mariadb` service with the "--skip-grant-tables" option using the following commands:
`systemctl set-environment MYSQLD_OPTS="--skip-grant-tables"`
`systemctl start mariadb`
5. Access the MySQL database by using the following command:
`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 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, using the following command and entering the new password when prompted:
`mysql -u root -p`
9. Restart the `em7` service using the following command:
`systemctl start em7`
10. Ensure that the password you set is also updated in the `/etc/silo.conf` `dbpasswd` variable.

Recovering the MySQL SNMP User Account on Collector Appliance

If you have removed the SNMP user account from the Collector's MySQL database in an attempt to harden your system, you must recover the account so that SL1 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 IP Address of an SL1 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 SL1 license, use the workflows in this section.

Moving an SL1 appliance to a new network requires pre-planning. If your SL1 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 an SL1 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 SL1 user interface:

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

This chapter includes the following topics:

Changing the IP Address on an All-In-One Appliance	218
Step 1. Stop the EM7 Service	218
Step 2. Change the IP Address in the Configuration Files	219

Step 3. Change the IP Address in the /etc/hosts File	219
Step 4. Change the IP Address in the Network Interface Configuration File	219
Step 5. Update the IP Address in the MySQL Database	221
Step 6. Reboot the Appliance	221
Step 7. Change the Database Appliance IP Address in the Administration Portals	222
Changing the IP Address on a Database Server	224
Step 1. Stop the EM7 Service	224
Step 2. Change the IP Address in the Configuration Files	225
Step 3. Change the IP Address in the /etc/hosts File	225
Step 4. Change the IP Address in the Network Interface Configuration File	225
Step 5. Update the IP Address in the MySQL Database	227
Step 5a: For Database Servers Configured with PhoneHome	227
Step 5b For Clustered Database Appliances (using HA, DR, or HA+DR)	228
Step 6. Reboot the Appliance	230
Step 7. Change the Database Appliance IP Address in the Administration Portals, Data Collectors, and Message Collectors	230
Changing the IP Address on a Database Collector or Message Collector	233
Step 1. Stop the EM7 Service	233
Step 2. Change the IP Address in the Configuration Files	233
Step 3. Change the IP Address in the /etc/hosts File	233
Step 4. Change the IP Address in the Network Interface Configuration File	234
Step 5. Update the IP Address in the MySQL Database	235
Step 6. Reboot the Appliance	236
Step 7. Change the Database Appliance IP Address in the Administration Portals	236

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

```
sudo systemctl stop em7
```

Step 2. Change the IP Address in the Configuration Files

You must change the `ipaddress` value in the configuration files.

1. Either go to the console of the SL1 appliance or use SSH to access the server.
2. Enter the following at the command line:

```
sudo vi /etc/siteconfig/siloconf.siteconfig
```

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).
5. Do steps 2-4 to change the IP address in the `/etc/silo.conf` file.

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 SL1 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 All-In-One Appliance, 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 SL1 appliance or use SSH to access the server.
2. Login as user **em7admin** with the appropriate password.
3. Enter the following at the command line:

```
sudo ifconfig
```

4. Your output will look like this:

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

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

5. Examine the output, find the 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 SL1 appliance.
 - **IPADDR**=New IP address of the SL1 appliance.
 - **PREFIX**=netmask for the SL1 appliance.
 - **GATEWAY**=IP address of the network gateway that will be used by the SL1 appliance.
9. Save your changes and exit the file (:wq)
10. At the command line, enter the following:

```
sudo service network restart
```

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 SL1 is restarted, the new IP address is recognized as licensed.

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

```
silos_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 ~]$ silos_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.

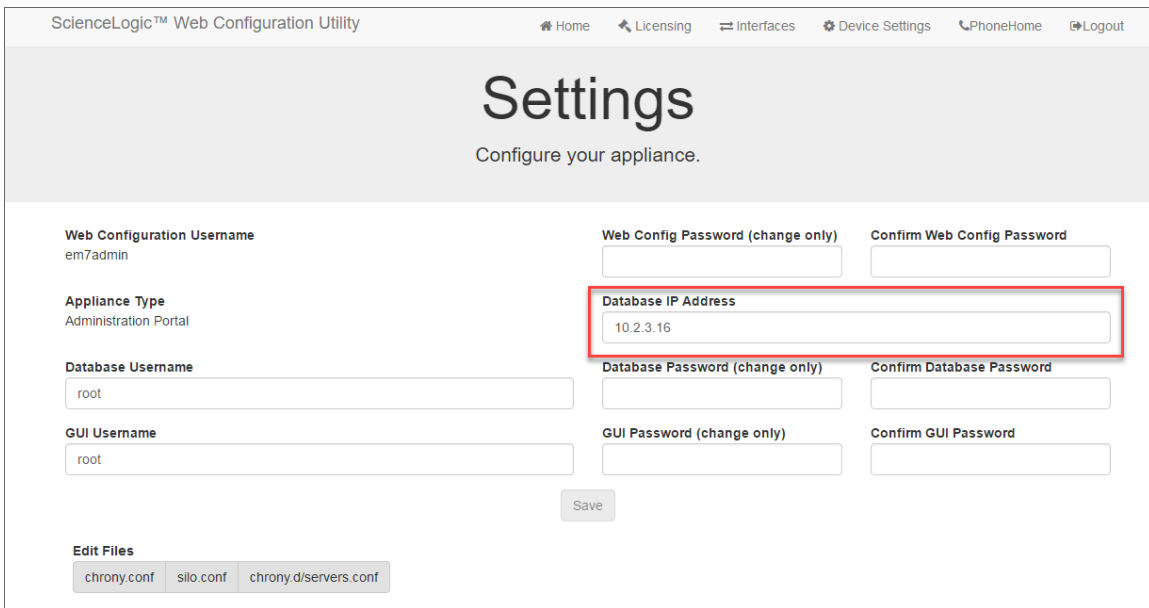
Step 6. Reboot the Appliance

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. SL1 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, SL1 will keep the current license for the appliance.

Step 7. Change the Database Appliance IP Address in the Administration Portals

If the All-In-One Appliance uses one or more Administration Portals, you must update the configuration of the Administration Portals to use the new IP address. . Open the Web Configuration Utility ([https://\[ip address of admin portal\]:7700](https://[ip address of admin portal]:7700)) to make the change in Device Settings, as highlighted below.



The screenshot shows the ScienceLogic™ Web Configuration Utility interface. The page title is "Settings" with the subtitle "Configure your appliance." The navigation bar includes Home, Licensing, Interfaces, Device Settings, PhoneHome, and Logout. The main content area contains several configuration sections:

- Web Configuration Username:** em7admin
- Web Config Password (change only):** [Empty text box]
- Confirm Web Config Password:** [Empty text box]
- Appliance Type:** Administration Portal
- Database IP Address:** 10.2.3.16 (This field is highlighted with a red box in the original image)
- Database Username:** root
- Database Password (change only):** [Empty text box]
- Confirm Database Password:** [Empty text box]
- GUI Username:** root
- GUI Password (change only):** [Empty text box]
- Confirm GUI Password:** [Empty text box]

At the bottom, there is a "Save" button and an "Edit Files" section with buttons for "chrony.conf", "silo.conf", and "chrony.d/servers.conf".


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 SL 1 . 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.

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



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



- 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.
 - For an All-In-One Appliance with multiple Administration Portals, enter the IP address for the All-In-One Appliance.
- **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.

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, or 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.

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 Configuration Files

You must change the `ipaddress` value in the configuration files.

1. Either go to the console of the SL1 appliance or use SSH to access the server.
2. Enter the following at the command line:

```
sudo vi /etc/siteconfig/siloconf.siteconfig
```

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`).
5. Do steps 2-4 to change the IP address in the `/etc/silo.conf` file.

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 SL1 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, SL1 recognizes the new IP address as licensed.

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

```
silos_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 ~]$ silos_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

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:

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.

phonehome status

- 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
```

- Note the Device ID for the Database Server. .
- 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
```

- 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.

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

NOTE: 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 the Database Server has 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 SL1 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
  }
}
```

For the Database Server with the new IP address, edit the value for `ring0_addr` to match the new IP address. Save the file.

6. 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 SL1 configured for High Availability, DRBD uses the private IPs to synchronize high-availability data.
- In SL1 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 SL1 configured for Disaster Recovery, DRBD uses the virtual IP and the public IP addresses to synchronize data.

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;
  }
}
```

Replace instances of the old IP address with the new IP address and save the file.

Shut down the Database Server .

7. 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 SL1 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. SL1 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, SL1 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 SL1 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 SL1 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 SL 1 . The address of the Web Configuration Utility is in the following format:


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

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


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



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



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.

Changing the IP Address on a Database Collector or Message Collector

To change the primary IP address of a Data Collector or Message Collector:

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 SL1 appliance or use SSH to access the server.
2. Log in 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 Configuration Files

You must change the `ipaddress` value in the configuration files.

1. Either go to the console of the SL1 appliance or use SSH to access the server.
2. Enter the following at the command line:

```
sudo vi /etc/siteconfig/siloconf.siteconfig
```

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`).
5. Do steps 2-4 to change the IP address in the `/etc/silo.conf` file.

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 SL1 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 Data Collector or Message Collector, 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. 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 SL1 appliance or use SSH to access the server.
2. Log in as user **em7admin** with the appropriate password.
3. Enter the following at the command line:

```
sudo ifconfig
```

4. Note that your output will look like this:¹

```
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 SL1 appliance.
- **IPADDR**=New IP address of the SL1 appliance.
- **PREFIX**=netmask for the SL1 appliance.
- **GATEWAY**=IP address of the network gateway that will be used by the SL1 appliance.

9. Save your changes and exit the file (:wq)

10. At the command line, enter the following:

```
sudo service network restart
```

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 SL1 is restarted, the new IP address is recognized as licensed.

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

```
silos_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 ~]$ silos_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 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. SL1 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, SL1 will keep the current license for the appliance.

Step 7. Change the Database Appliance IP Address in the Administration Portals

If the Data Collector or Message Collector uses one or more Administration Portals, you must update the configuration of the Administration Portals to use the new IP address. Open the Web Configuration Utility ([https://\[ip address of admin portal\]:7700](https://[ip address of admin portal]:7700)) to make the change in Device Settings, as highlighted below.

The screenshot shows the ScienceLogic™ Web Configuration Utility interface. The page title is "Settings" with the subtitle "Configure your appliance." The navigation bar includes Home, Licensing, Interfaces, Device Settings, PhoneHome, and Logout. The main content area contains several configuration sections:

- Web Configuration Username:** em7admin
- Web Config Password (change only):** [Empty field]
- Confirm Web Config Password:** [Empty field]
- Appliance Type:** Administration Portal
- Database IP Address:** 10.2.3.16 (This field is highlighted with a red box in the original image)
- Database Username:** root
- Database Password (change only):** [Empty field]
- Confirm Database Password:** [Empty field]
- GUI Username:** root
- GUI Password (change only):** [Empty field]
- Confirm GUI Password:** [Empty field]

A "Save" button is located below the password fields. At the bottom, there is an "Edit Files" section with buttons for "chrony.conf", "silo.conf", and "chrony.d/servers.conf".


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 SL 1 . The address of the Web Configuration Utility is in the following format:


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

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


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



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



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.
 - For a Data Collector or Message Collector with multiple Administration Portals, enter the IP address for the Data Collector or Message Collector.
- **Database Username.** Username for the database account that the Administration Portal will use to communicate with the Data Collector or Message Collector.
- **Accept the default values in all other fields.**

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

Chapter



10

Changing Name Servers on an SL1 Appliance

Overview

This chapter describes how to change domain name servers (DNS) on an SL1 appliance.

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

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

This chapter includes the following topics:

Changing Name Servers on an SL1 Appliance	240
Installing Additional RPMs on an SL1 Appliance	242

Changing Name Servers on an SL1 Appliance

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.

To change the DNS settings:

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 appliance is rebooted.

To add one or more domain name servers to the interface configuration file:

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

Installing Additional RPMs on an SL1 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 SL1 system.

Chapter



11

Backup Management

Overview

SL1 allows you to define three types of backups for your system: Configuration Backup, Full Backup, and Disaster Recovery Backup. A configuration backup stores a copy of the core database tables, while both full and disaster recovery backups backs up everything in your ScienceLogic database. This chapter describes how to define and restore from different backup types.

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

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

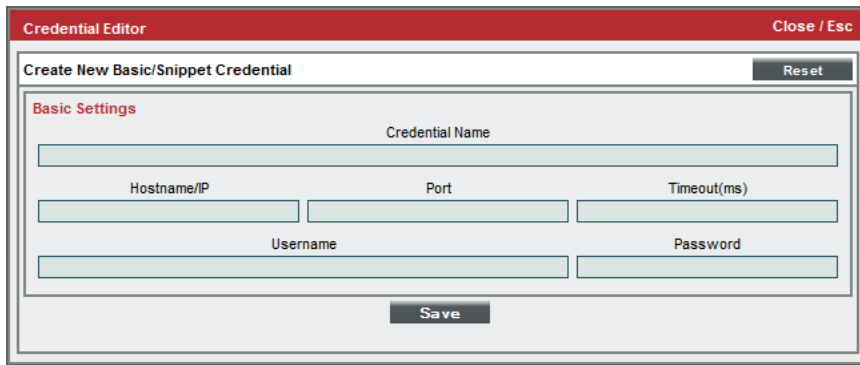
This chapter includes the following topics:

Creating a Backup Credential

To configure a backup, you must create a **Basic/Snippet** Credential that allows SL1 to write to the external systems where you will store the backups. To create a backup credential:

1. Go to the **Credential Management** page (System > Manage > Credentials).
2. In the **Credential Management** page, click the **[Actions]** button in the upper right corner of the page. Select **Create Basic/Snippet Credential**.

3. The **Credential Editor** page appears, where you can define values in the following fields:



- **Credential Name**. Name of the credential. Can be any combination of alphanumeric characters.
 - **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 the **[Save]** button.

Configuration Backups

A configuration backup stores a local copy of the core database tables that are required to restore a SL1 system, and optionally transfers the copy to an external system. A configuration backup entails:

- Configuration backup includes scope and policy information, but **not performance data, data collected using configuration Dynamic Applications, events, or logs**. The following databases are backed up during a configuration backup:
 - **master**. Includes system-level settings for SL1, 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_biz**. Includes asset information, dashboards, distribution lists, document templates, IT Service policy information, knowledge base information, organization information, product SKU information, RSS feeds, ticketing information, and user preferences.

CAUTION: Due to security vulnerabilities, ScienceLogic recommends that customers who installed SL1 prior to 8.9.2 disable the Knowledge Base. For details, see the release notes for version 8.9.2 of SL1.

- **master_custom.** Includes GUI customizations, dashboard widget definitions, PowerPack files, and custom attributes for the Integration Server.
- **master_dns.** Includes DNS information.
- **master_dev.** Includes information associated with device records, excluding performance data, data collected using configuration Dynamic Applications, events, or logs.
- **master_events.** The configuration backup includes only the event_suppressions database table from this database. This table stores event suppression settings.
- **master_filestore.** Includes information about files, PowerPacks, and notes.

NOTE: In SL1 versions prior to 8.14, configuration backups contained RPM update and package files. In version 8.14 and higher, configuration backups no longer contain these files, making the backup process complete faster and the backup file smaller.

- **master_platform.** Includes information about ScienceLogic appliances.
- **master_reports.** Includes custom report definitions.
- **scheduler.** Includes all instances of scheduled items: reports, discovery sessions, etc.

NOTE: You can configure the staging and remote directories used for backups in the master.system_settings_backup database, to ensure that the backup can be placed on a directory that has enough disk space. For remote directories, the current Unix time will be appended to the directory name to ensure the directory name is unique each time the backup runs.

- SL1 automatically launches this backup at the interval you specify.
- SL1 saves local copies of the last seven days of backups and stores the first backup of the month for the current month and the three previous months.
- A configuration backup also contains all the files and folders specified in /etc/backup.conf.
- During configuration backup, the ScienceLogic database remains online.

Defining a Configuration Backup

A configuration backup stores a local copy of the core database tables that are required to restore a SL1 system. A configuration backup includes scope and policy information, but **not performance data or logs**. Configuration backups save copies of the files and folders specified in /etc/backup.conf.

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:

The screenshot shows a configuration pane titled "Configuration Backup". It contains the following fields:

- Every**: Input field containing "1".
- Interval**: Dropdown menu showing "[Disabled]".
- Start Time / Date**: Input field containing "2011-01-01 00:00".
- Timezone**: Dropdown menu showing "[_ UTC]".
- Configuration Credentials**: Dropdown menu showing "[**restricted credential**]".
- Configuration Protocol**: Dropdown menu showing "[NFS-Remote]".
- Configuration Subdirectory**: Empty input field.

- **Every**. Specifies the frequency of the backup.
- **Interval**. You must specify how frequently SL1 should automatically execute a full backup. Your choices are:
 - *Disabled*. Full backups are disabled.
 - *Daily*. SL1 will execute full backups every day.
 - *Weekly*. SL1 will execute full backups once a week.
 - *Monthly*. SL1 will execute full backups once a month.
- **Start Time / Date**. If you enabled configuration backups, you must specify the daily start time. This is the time at which SL1 will automatically execute configuration backups every day. You must also specify the date on which the daily backups will begin. Use the drop-down lists to select the date and time.

- **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.** Optional. If you want to store the configuration backup on a remote NFS mount or a remote SMB mount, you must select a credential in this field. The credential must be of type **Basic/Snippet** and specify the hostname or IP, user name, and password.
- **Configuration Protocol.** Optional. If you specified a configuration credential in the **Configuration Credentials** field, you must select the type of external system where the backup will be stored.
 - **NFS-Remote.** When you select this option, SL1 stores the full backup on an NFS mount. You specify the NFS mount with the **Configuration Credentials** field and the **Configuration Subdirectory** field. **The system creates the backup file directly on the external NFS mount**, rather than creating the backup file on the appliance, transferring the backup file to the NFS mount, and then removing the backup file from the appliance.

NOTE: 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.

- **SMB-Remote.** When you select this option, SL1 stores the full backup on an SMB mount. You specify the SMB mount with the **Configuration Credentials** field and the **Configuration Subdirectory** field. **The system creates the backup file directly on the external SMB mount.**
- **Configuration Subdirectory.** Optional. If you specified NFS mount or SMB mount by selecting a credential in the **Configuration Credentials** field, you can specify a directory on the NFS mount or SMB mount in which you would like to store the configuration backup. When entering the subdirectory path, omit the leading slash ("/").

NOTE: SL1 always maintains local copies of configuration backups from the last seven days and the first day of the month for the current month and the previous three months. Even if you choose to store configuration backups on a remote NFS mount or a remote SMB mount, SL1 still maintains local copies of configuration backups.

3. Click the **[Save]** button to save your settings. SL1 will execute the configuration backup every day, starting on the date you specified in the **Start Time / Date** field, at the time you specified in the **Start Time / Date** field.
4. To run the backup immediately, click the **[Backup Now]** button under **Configuration Backup**. SL1 will immediately run the backup, and will still run the backup every day at the time you specified in the **Start Time** field.

Restoring a Configuration Backup

If your database has been backed up using a configuration backup, in the event of data corruption or other failure, you will need to restore your system using the configuration backup. The backup file contains one .sql file for each database that was included in the backup. To restore a database using the backup file:

1. Either go to the console of the Database Server or use SSH to access the server.
2. Login as user **em7admin** and sudo to the root account.

```
sudo -s
```

3. Perform the following commands to uncompress the backup file:

```
mkdir /data.local/db/restore_tmp  
cd /data.local/db/restore_tmp  
pigz -dc <full path and file name for backup.tgz> | tar xv
```

4. Navigate to the directory that contains the .sql files from the backup:

```
cd data/backup/staging
```

Where: *unix timestamp* is appended to each remote directory to ensure that naming is unique.

5. The directory will contain one .sql file for each database included in the backup. To restore a database, execute the following command using the username of a user that has administrative privileges in MySQL (by default, the user is **root** and the password is **em7admin**):

```
silo_mysql <name_of_database> -u <username> -p<password> < <name_of_ database>.sql
```

NOTE: Do not include a space between "-p" and the password.

For example, to restore the database "master" as the user "root" with the default password of "em7admin", perform the following command:

```
silo_mysql master -u root -pem7admin < master.sql
```

6. Re-license the Database Server using the standard licensing procedure.
7. To restore all the databases that are included in the backup file, repeat steps 5 and 6 for each .sql file.

Full Backup

Running a full backup creates a complete backup of the ScienceLogic database. This type of backup is recommended for SL1 systems in small-to-medium enterprises. Note the following about full backups:

- A full backup comprises all configuration data, performance data, and log data.
- Full backup is disabled by default. You can configure SL1 to automatically launch this backup at a frequency and time you specify.
- During full backup, the ScienceLogic database remains online.
- The full backup is stored on a remote NFS mount or a remote SMB mount.

NOTE: *ScienceLogic does not recommend Full Backups for large SL1 systems.* For large SL1 systems, ScienceLogic recommends you use the **DR Backup** option or a SAN with snapshot technology to backup and restore data.

NOTE: You can configure the staging and remote directories used for backups in the master.system_settings_backup database, to ensure that the backup can be placed on a directory that has enough disk space. For remote directories, the current unix time will be appended to the directory name to ensure the directory name is unique each time the backup runs.

NOTE: If your SL1 System uses AWS RDS (remote database), the **Full Backup** option is disabled.

Defining a Full Backup

Full backup includes all configuration data, performance data, and log data. SL1 automatically launches this backup and 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:

The screenshot shows a configuration pane titled "Full Backup" with the following fields and values:

- Every**: 1
- Interval**: [Disabled]
- Start Time / Date**: 2011-01-01 00:00
- Timezone**: [|_ UTC]
- Full Credentials**: [**restricted credential**]
- Full Protocol**: [NFS-Remote]
- Full Subdirectory**: (empty)
- Custom innobackupex Options**: (empty)

- **Every**. Specifies the frequency of the backup.
- **Interval**. You must specify how frequently SL1 should automatically execute a full backup. Your choices are:
 - *Disabled*. Full backups are disabled.
 - *Day*. SL1 will execute full backups every day.
 - *Week*. SL1 will execute full backups once a week.
 - *Month*. SL1 will execute full backups once a month.
- **Start Time / Date**. Specify the time of day at which SL1 should automatically execute full backups. You must also specify the date on which the backups will begin. If you selected *Weekly* or *Monthly* in the **Backup Frequency** field, the date you select will determine which day of the week or month the backups will be scheduled. Use the drop-down lists to select the date and time.
- **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 Backup Credentials.** You must store full backups on an external system (not the Database Server or All-In-One Appliance). You can specify that full backups be stored on a remote NFS mount or a remote SMB mount. To specify where to store full backups, you must select a credential in this field. The credential must be of type **Basic/Snippet** and specify the hostname or IP, username, and password.
- **Full Protocol.** Specify the type of external system where the full backup will be stored. Choices are:
 - **NFS-Remote.** When you select this option, SL1 stores the full backup on an NFS mount. You specify the NFS mount with the **Configuration Credentials** field and the **Configuration Subdirectory** field. **The system creates the backup file directly on the external NFS mount**, rather than creating the backup file on the appliance, transferring the backup file to the NFS mount, and then removing the backup file from the appliance.

NOTE: 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 guide.

- **SMB-Remote.** When you select this option, SL1 stores the full backup on an SMB mount. You specify the SMB mount with the **Configuration Credentials** field and the **Configuration Subdirectory** field. **The system creates the backup file directly on the external SMB mount.**
 - **Full Subdirectory.** Specify a directory on the remote NFS mount or the remote SMB mount in which you would like to store the full backup. When entering the subdirectory path, omit the leading slash ("/").
 - **Custom innobackupex Options.** Specify one or more custom backup options. For details on these options, see http://www.percona.com/doc/percona-xtrabackup/2.1/innobackupex/innobackupex_option_reference.htm.
3. Select the **[Save]** button to save your settings. SL1 will execute the full backup at the frequency and time you specified in the **Backup Frequency** and **Start Time** fields.
 4. To run the backup immediately, click the **[Backup Now]** button under **Full Backup**. SL1 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.

Restoring a Full Backup

To restore a SL1 system using a full backup file, perform the following steps:

NOTE: These steps assume that the Database Server has not been previously configured.

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. 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 log in to that the Database Server.
3. Log in as user **em7admin** and then sudo to the root account:

```
sudo -s
```

4. Execute the following commands:

WARNING: Executing this command will stop the database. SL1 will not be operational until you complete the restore procedure.

```
systemctl stop em7
systemctl stop mariadb
rm -rf /data/db/*
```

5. Execute the following two commands, substituting in the full path name of your backup file:

```
cd /data/db
```

```
pigz -dc <full_path_and_name_of_backup_file.gz> | xbstream -x -C .
```

6. Execute the following command:

```
sudo more /data/db/backup-my.cnf
```

7. Locate the line that looks like the following. Copy or write down the exact text that appears:

```
innodb_data_file_path=ibdata1:354M;ibdata2:500M:autoextend:max:8925M
```

8. Execute the following command to edit the /etc/my.cnf.d/silo_mysql.cnf file:

```
sudo vi /etc/my.cnf.d/silo_mysql.cnf
```

9. Locate the line that looks like the following. Change the line so that it is identical to the line you noted in step 7:

```
innodb_data_file_path=ibdata1:354M;ibdata2:500M:autoextend:max:8925M
```

10. Save the file and exit the vi editor. To do this, enter **:wq**.

11. Execute the following command to build the updated configuration file:

```
sudo /opt/em7/share/scripts/generate-my-conf.py -o /etc/my.cnf.d/silo_mysql.cnf
```

12. Execute the following commands:

NOTE: Depending on the size of the backup, the `innobackupex` command might take a long time to complete.

```
cd /data/db
innobackupex --apply-log /data/db 2>&1 | tee /data/tmp/restore.log
chown -R mysql:mysql /data/db/*
systemctl start mariadb
systemctl start em7
```

11. Re-license the Database Server using the standard licensing procedure.

Additional Configuration for Solaris NFS Mounts

To use the *NFS-remote* backup protocol with an NFS mount 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.sciencelogic.local -d "ScienceLogic Backup Share" /export/home/backup
```

- In `/etc/defaults/nfs` on the Solaris system, include the line `"NFMAPID_DOMAIN=<domain of Database Server or All-In-One Appliance>".` For example:

```
NFMAPID_DOMAIN=ScienceLogic.local
```

You can test this configuration by mounting the NFS file system from the console of your SL1 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 owned by the `root` user.

Defining a DR Backup

For SL1 systems configured for disaster recovery, DR Backup temporarily stops replication, performs a full backup of the disaster-recovery database, and then re-enables replication and performs a partial resync from the primary.

DR backup includes all configuration data, performance data, and log data. During DR backup, the primary ScienceLogic database remains online.

NOTE: The *DR Backup* fields appear only for systems configured for Disaster Recovery. DR Backup is not available for the two-node DRBD-HA cluster.

NOTE: If your SL1 System uses AWS RDS (remote database), the *DR Backup* option is disabled.

To define and schedule a DR backup:

1. Go to the **Backup Management** page (System > Settings > Backup).

DR Backup

Every Interval

Start Time / Date

Timezone

DR Credentials

DR Protocol

DR Subdirectory

2. In the **DR Backup** pane, provide values in the following fields:

- **Every**. Specifies the frequency of the backup.
- **Interval**. You must specify how frequently SL1 should automatically execute a full backup. Your choices are:
 - *Disabled*. Full backups are disabled.
 - *Daily*. SL1 will execute full backups every day.
 - *Weekly*. SL1 will execute full backups once a week.
 - *Monthly*. SL1 will execute full backups once a month.
- **Start Time/Date**. Specify the time of day at which SL1 should automatically execute DR backups. You must also specify the date on which the backups will begin. If you selected *Weekly* or *Monthly* in the **Backup Frequency** field, the date you select will determine which day of the week or month the backups will be scheduled. Use the drop-down lists to select the date and time.
- **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**. You must store DR backups on an external system (not the Database Server or All-In-One Appliance). You can specify that DR backups be stored on an NFS mount or SMB mount. To specify where to store full backups, you must select a credential in this field. The credential must be of type **Basic/Snippet** and specify the hostname or IP, username, and password to access the external system. For more information on credentials, see the **Credential Management** page.

NOTE: Your organization membership(s) might affect the list of credentials you can see in the **Full Backup Credentials** field. For details, see the **Discovery and Credentials** manual.

- **DR Protocol**. Specifies where SL1 should store the full backups. Choices are:
 - *NFS-Remote*. When you select this option, SL1 stores the full backup on an NFS mount. You specify the NFS mount with the **Full Backup Credentials** field and the **Full Subdirectory** field. **The system creates the backup file directly on the external NFS mount.**

NOTE: 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, SL1 stores the full backup on an SMB mount. You specify the SMB mount with the **Full Backup Credentials** field and the **Full Subdirectory** field. **The system creates the backup file directly on the external SMB mount.**

NOTE: ScienceLogic strongly recommends that you use the **DR Protocol** option *NFS-Remote* or *SMB-Remote* when performing a DR backup; the other **DR Protocol** options cause the resync step to take much longer.

- **DR Subdirectory.** Specify a directory on the NFS mount or SMB mount in which you would like to store the DR backup.
3. Click the **[Save]** button to save your settings. SL1 will execute the DR backup at the frequency and time you specified in the **Backup Frequency** and **Start Time** fields.
 4. To run the backup immediately, click the **[Backup Now]** button under **DR Backup**. SL1 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.

Restoring a DR Backup

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. SL1 will not be operational until you complete the restore procedure.

```
systemctl stop em7
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 SL1 and the database:

```
systemctl start em7
systemctl start mariadb
```

Chapter

12

Subscription Licenses

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
- Download subscription license usage data for manual upload to the ScienceLogic billing server
- Upload a receipt from the ScienceLogic billing server

If your SL1 system is configured to communicate with the ScienceLogic billing server, usage data will be sent automatically from your SL1 system to the ScienceLogic billing server once a day. After the ScienceLogic billing server receives the usage data, SL1 will automatically mark the license usage file as delivered.

If your SL1 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 SL1 user interface:

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

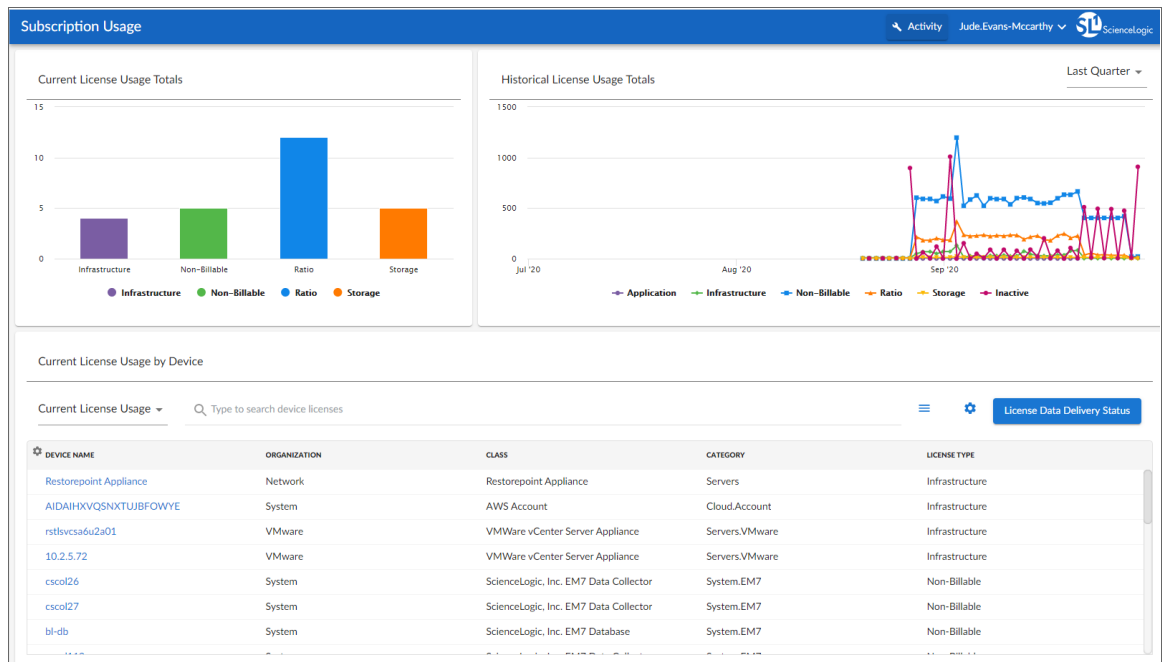
This chapter includes the following topics:

Viewing the Subscription License Usage	259
Viewing Delivery Status	259
Manually Uploading License Usage to ScienceLogic	260
Downloading the Daily License Usage File	261
Manually Uploading the Daily License Usage File to ScienceLogic	262

Viewing the Subscription License Usage

If you have a subscription license, you can view a dashboard on license usage for your SL1 system. To view the dashboard containing the license usage for your SL1 system:

1. Go to the **System Usage** page (Manage > Subscription Usage). The **Subscription Usage** page appears.



2. The dashboard displays three reports:
 - Current License Usage Totals
 - Historical License Usage Totals
 - Current License Usage by Device or Current License Counts by Device Class, depending on your selection

Device categories depend upon your specific license agreement with ScienceLogic. The types of license categories that appear in the dashboard reports will reflect your license agreement with ScienceLogic.

Viewing Delivery Status

The **License Data Delivery Status** page displays the status of one or more daily license usage files. To view the **License Data Delivery Status** page:

1. Go to the **Subscription Usage** page (Manage > Subscription Usage).
2. Click the **[License Data Delivery Status]** button.
4. 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 SL1 system by the automated maintenance process.
 - **Summary Size.** Size of the daily license usage file.

	Summary Date	Delivery Status	Summary Size (kB)	
1.	2020-10-04 00:00:00	0	10948.4	<input type="checkbox"/>
2.	2020-10-03 00:00:00	0	10948.4	<input type="checkbox"/>
3.	2020-10-02 00:00:00	0	10948.4	<input type="checkbox"/>
4.	2020-10-01 00:00:00	0	10948.4	<input type="checkbox"/>
5.	2020-09-30 00:00:00	0	10948.4	<input type="checkbox"/>
6.	2020-09-29 00:00:00	0	10948.4	<input type="checkbox"/>
7.	2020-09-28 00:00:00	0	10948.4	<input type="checkbox"/>
8.	2020-09-27 00:00:00	0	10948.4	<input type="checkbox"/>
9.	2020-09-26 00:00:00	0	10948.4	<input type="checkbox"/>
10.	2020-09-25 00:00:00	0	10948.4	<input type="checkbox"/>
11.	2020-09-24 00:00:00	0	10948.4	<input type="checkbox"/>

Manually Uploading License Usage to ScienceLogic

If your SL1 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, SL1 will automatically mark the license usage file as delivered.

If your SL1 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 SL1 system, allowing SL1 to mark the license usage file as delivered.
- License usage files will not be deleted from your system until they are delivered.

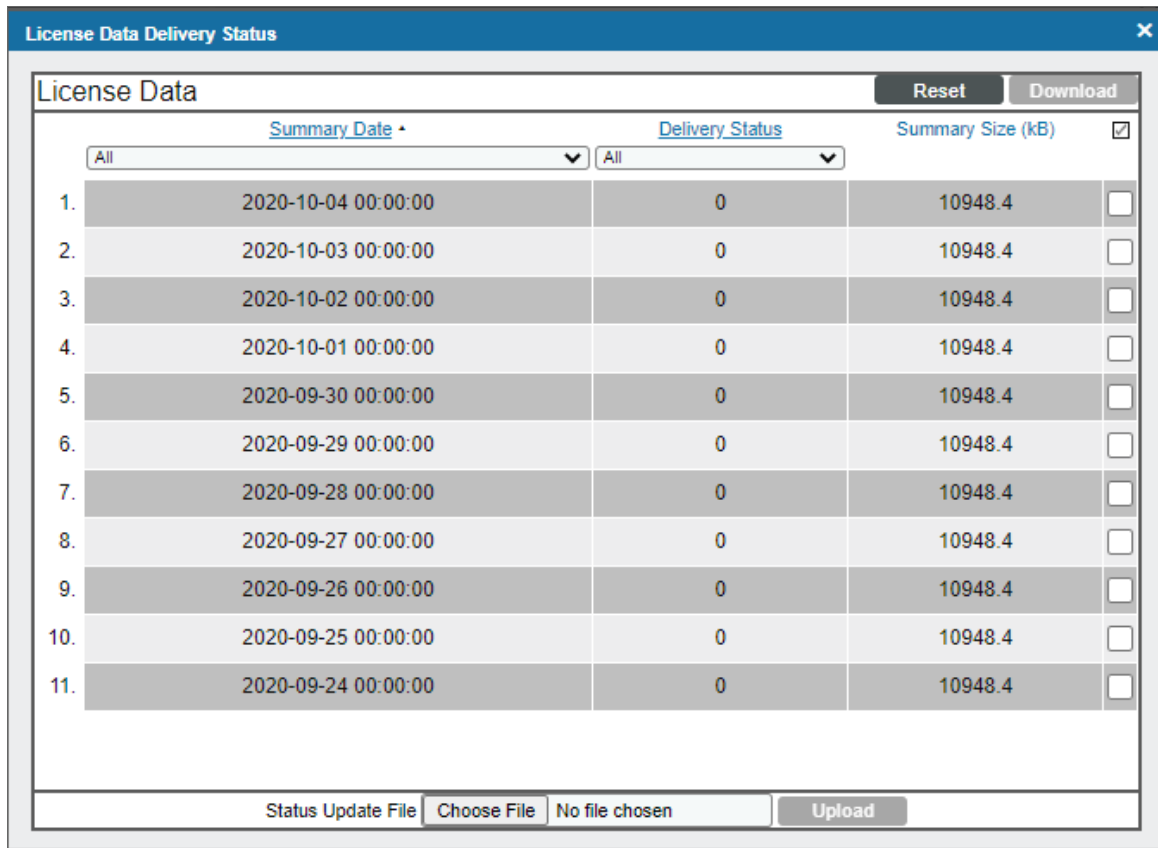
Downloading the Daily License Usage File

If your SL1 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 [**License Data Delivery Status**].

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 50MB, 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. To do this:

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.

The screenshot shows a web interface for 'Subscription Data'. At the top, it says 'Subscription Data | For [System Administrator] | Organization: System'. Below this, there are three main sections: 1. 'Subscription Data Update' with a 'License Data File' field containing 'Choose File' and 'No file chosen' buttons, and a 'Get Update' button highlighted in red. 2. 'Subscription Data Receipt Status' with 'From' and 'To' date pickers set to '07/11/2015' and '10/11/2015' respectively, and a 'Get Status' button. 3. 'Device Count by License Type' which is currently empty and displays '- No Data -'.

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 SL1 system.

Uploading the ScienceLogic Receipt

After uploading the daily license usage file to the ScienceLogic Billing server, 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" file to your SL1 system to inform your SL1 system that the upload was successful and that the SL1 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 **[License Data Delivery Status]**.
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 SL1 system.

License Data				Reset	Download
	Summary Date	Delivery Status	Summary Size (kB)	<input checked="" type="checkbox"/>	
1.	2020-10-04 00:00:00	0	10948.4	<input type="checkbox"/>	
2.	2020-10-03 00:00:00	0	10948.4	<input type="checkbox"/>	
3.	2020-10-02 00:00:00	0	10948.4	<input type="checkbox"/>	
4.	2020-10-01 00:00:00	0	10948.4	<input type="checkbox"/>	
5.	2020-09-30 00:00:00	0	10948.4	<input type="checkbox"/>	
6.	2020-09-29 00:00:00	0	10948.4	<input type="checkbox"/>	
7.	2020-09-28 00:00:00	0	10948.4	<input type="checkbox"/>	
8.	2020-09-27 00:00:00	0	10948.4	<input type="checkbox"/>	
9.	2020-09-26 00:00:00	0	10948.4	<input type="checkbox"/>	
10.	2020-09-25 00:00:00	0	10948.4	<input type="checkbox"/>	
11.	2020-09-24 00:00:00	0	10948.4	<input type="checkbox"/>	

Status Update File Choose File No file chosen Upload

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**. 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 twelve months.
 - **Subscriber Device Usage Data**. 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 six 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 SL1 system. Default value is twelve months.

- **Subscriber System Usage Data**. 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 six 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 six 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. SL1 will not prune data that has not yet been delivered to the ScienceLogic Licensing and Billing server.

Chapter

13

CAC Authentication

Overview

SL1 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.

Each CAC contains a client-side security certificate from the DoD certificate authority (DoD Root CA). This client-side certificate allows the CAC to authenticate with web servers that include the server-side security certificate from the DoD certificate authority. Web servers with the server-side security certificate are deemed secure for DoD use.

SL1 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. You can install server-side certificates on the user interface appliances and then authenticate access to those web servers with a CAC or a client-side certificate associated with a user's web browser.

When authentication of the client-side certificate against the server-side certificate is successful, SL1 displays the ScienceLogic login page to the end user. The client-side certificate does not authenticate the username and password. The client-side certificate authenticates only that the user is permitted access to the user interface, and in the case of a DoD issued server-side certificate, that the user interface appliance (Administration Portal, All-In-One Appliance, or the Database Server) is deemed secure for DoD use. If CAC authentication is enabled, SL1 verifies the client-side certificate every time a page is accessed in the user interface; if a client-side certificate is unavailable or invalid, the user is immediately logged out of the user interface.

NOTE: Currently, SL1 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 SL1 user interface:

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

This chapter includes the following topics:

<i>Prerequisites</i>	267
<i>Importing an SSL Certificate</i>	268
<i>Updating the ScienceLogic Configuration File</i>	269
<i>Defining the Client Certificate</i>	269
<i>Testing the Configuration</i>	271
<i>Prerequisites</i>	267
<i>Importing an SSL Certificate</i>	268
<i>Updating the ScienceLogic Configuration File</i>	269
<i>Defining the Client Certificate</i>	269
<i>Testing the Configuration</i>	271

Prerequisites

To use client certificate authentication with SL1, you must first perform the following tasks:

1. Users must have either:
 - Valid CACs with valid client-side certificates already loaded onto the cards.
 - Valid client-side certificates installed in their web browser.
2. 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. For Mozilla Firefox, users can install <http://www.forge.mil/Resources-Firefox.html>. For Internet Explorer, users must purchase and install commercial software for reading security certificates.
3. In the **Behavior Settings** page (System > Settings > Behavior), you must enable the following field
 - **Force Secure HTTPS**. Only when the Administration Portal, All-In-One Appliance, or the Database Server uses HTTPS will the appliance request a security certificate from the CAC or client web browser.

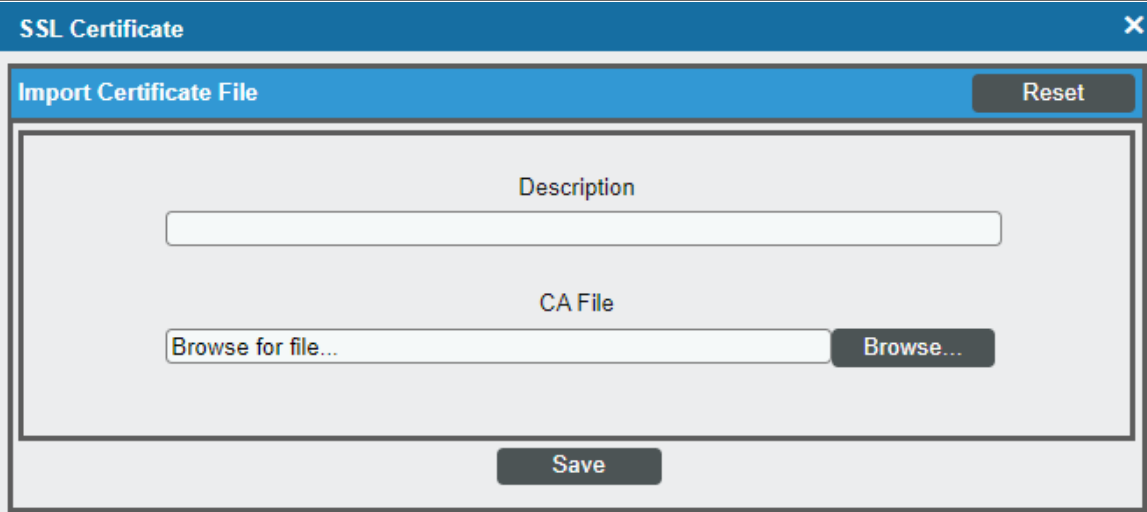
4. In the **SSL Certificates** page (System > Settings > Authentication > Admin SSL Certificates), you must install the server-side certificate on the Administration Portal, All-In-One Appliance, or the Database Server. For CAC authentication, the server-side certificate is issued by the DoD. To learn more about importing a certificate, see the section [Importing an SSL Certificate](#).
5. You can customize the user name that is displayed in SL1 after CAC authentication. You can edit the ScienceLogic configuration file to customize the displayed user name.
6. In the **Client Certificate & CAC Authentication** page (System > Settings > Authentication > CAC Client Cert Auth), you must configure the server-side certificate and test it against your client-side certificate.

Importing an SSL Certificate

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 SL1, 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.

To import an SSL certificate for CAC authentication:

1. Go to the **SSL Certificates** page (System > Settings > Authentication > SSL Certificates).
2. In the **SSL Certificates** page, click the **[Actions]** menu. Select **Import DoD Root CA Certificate**. The **Import Certificate** modal page appears.



The screenshot shows a modal window titled "SSL Certificate" with a close button (X) in the top right corner. Below the title bar is a sub-header "Import Certificate File" and a "Reset" button. The main area contains a "Description" label above a text input field, and a "CA File" label above a "Browse for file..." text input field with a "Browse..." button. At the bottom center is a "Save" button.

3. In the **Import Certificate** modal page, 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.

Updating the ScienceLogic Configuration File

By default, the certificate configuration file (`em7_certificate.conf`) is configured to display a CAC user's common name (CN) as a user name in SL1 after CAC authentication. If this is your preference, then you do not need to update the configuration file and can skip this section.

However, if you prefer that SL1 display only the user name portion of the CAC user's CN, then you can edit the certificate configuration file to parse out the user name from the certificate CN.

To do so:

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 `em7_certificate.conf` with a text editor like `vi`:

```
vi em7_certificate.conf
```

4. Edit the file to look like this:

```
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

When you define a CAC or client-side certificate on a web browser, you are actually selecting a server-side certificate on the SL1 appliance and testing the client-side certificate (on your browser or your CAC) against the certificate on the appliance.

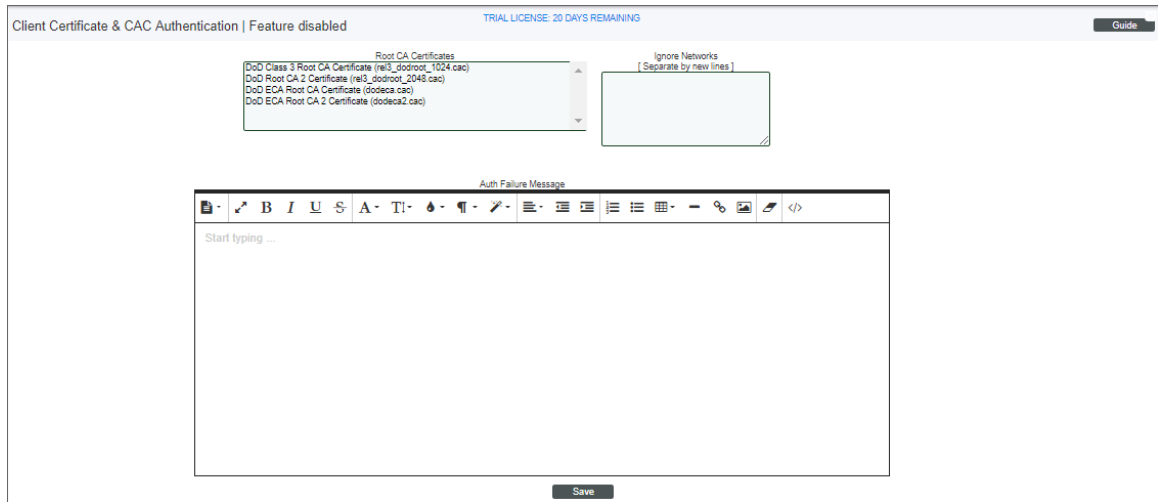
You can also define some custom settings for client-side certificate authentication. You can define error messages that are displayed to the end user when authentication fails. You can also define IP addresses for which the user interface will not perform authentication.

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).



3. Supply a value in each of the following fields:

- **Root CA Certificates.** Select from a list of certificates installed on the **SSL Certificates** page (System > Settings > Authentication > Admin SSL Certificates). Your client-side certificate will be authenticated against the selected server-side certificate.
- **Certificate User Field.** Specifies which field in the certificate the platform will use to find the username. The choices are:
 - *Common Name*
 - *MS UPN*
- **Auth Failure Message.** Enter text for the error message that appears to users if authentication fails.
- **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, SL1 will not require certificate authentication from that client.
 - 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. SL1 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.


5. Click the **Test** link to test the configuration against your current client-side certificate.

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. Access the user interface with your CAC or a browser with the your client-side certificate installed.
2. Go to the **Client Certificate & CAC Authentication** page (System > Settings > Authentication > CAC Client Cert Auth).



3. Define the authentication settings.
4. 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.

5. Click the **TEST** link. SL1 will attempt to authenticate your client-side certificate against the selected server-side certificate.

6. If the test authentication is successful, SL1 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.

7. A new field, **Client Cert / CAC Auth**, appears with a default value of *Allowed*. You can select one of the following values for this field:
 - *Allowed*. This is the default value. Users with CAC and a corresponding account in the platform are automatically logged in to the platform when the user enters the URL of the Administration Portal or the All-In-One. If a CAC user does not have an account defined in the platform, the login screen is displayed. When a non-CAC user enters the URL of the Administration Portal or the All-In-One, he/she will see the message defined in the **Auth Failure Message** field (defined in the previous section).
 - *Locked*. Users with CAC and an aligned account in the platform are automatically logged in to the platform when the user enters the URL of the Administration Portal or the All-In-One. If a CAC user does not have an account defined in the platform, or a user does not have a CAC, when that user enters the URL of the Administration Portal or the All-In-One, he/she will see the message defined in the **Auth Failure Message** field (defined in the [certificate authentication settings](#)).

NOTE: ScienceLogic recommends that you set this field to *Locked* unless your implementation specifically requires one of the other options.

- *Disabled*. When a user enters the URL of the Administration Portal or the All-In-One Appliance, the platform displays the login screen.
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.

Chapter

14

Installing an SSL Certificate

Overview

SSL is an acronym for Secure Sockets Layer. SSL is a protocol for securely transmitting data via the internet. SSL uses a private key to encrypt data to be transferred over the Internet connection. Usually, URLs that include "HTTPS" are using SSL for security.



To implement SSL, an SSL certificate resides on the web server and is used to encrypt the data and to identify the website. The SSL 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 an SSL certificate:

- You can purchase 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 from a certificate authority.

SL1 includes a self-signed certificate from ScienceLogic. Self-signed certificates can trigger a warning message in some browsers. For this reasons, some customers might prefer to purchase an SSL certificate from a certificate authority and install the certificate on one or more servers.

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

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

This chapter includes the following topics:

Certificates for ScienceLogic Servers	274
Requesting a Commercial SSL Certificate	274

Creating Your Own Certificate	275
Installing the Certificate on an SL1 Appliance	277

Certificates for ScienceLogic Servers

Each SL1 appliance includes a self-signed certificate from ScienceLogic.

Each SL1 appliance uses the Nginx web server and OpenSSL.

If you want to use commercial SSL certificates with SL1, you must purchase certificates for the following SL1 appliances:

- For each Administration Portal, Database Server, or All-In-One Appliance you must purchase **two** certificates, one for the standard user interface and one for the Configuration Utility.
- For each Data Collector, you must purchase one certificate, for use with the Configuration Utility.
- For each Message Collector, you must purchase one certificate, for use with the Configuration Utility.

Requesting a Commercial SSL Certificate

To purchase a commercial SSL 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 SL1 appliance.

1. Either go to the console of the SL1 appliance or use SSH to access the server. Open a shell session on the SL1 appliance. Log in as "em7admin".
2. First, you must 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*.

NOTE: Make sure the file is **not** named **silossl.key**. This is the name of the pre-existing ScienceLogic, self-signed certificate file.

3. You will be prompted to enter a pass phrase for the key.
4. Best practice is to make a backup copy of the key file and the pass phrase and store both in a secure location.

5. You must remove the pass phrase from the key before generating a Certificate Signing Request (CSR). To do this, enter the following at the shell prompt, inserting the keyname you used where indicated:

```
sudo openssl rsa -in [keyname].key -out [keyname].key.insecure
```

6. Next, you must create a Certificate Signing Request (CSR) for the private key you created in the previous steps. To do this, enter the following 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 *adminport.key* and name the CSR for that key *adminport.csr*.

NOTE: Make sure the keyname is **not** *silossl.key*. This is the name of the pre-existing ScienceLogic, self-signed certificate file.

7. You will be prompted to enter the Common Name. Enter the fully qualified domain name of the server where the certificate will be used and SSL and https will be run.

For example, if the SL1 appliance is accessed at <https://company.adminportal.com>, you would enter "company.adminportal.com" as the Common Name.

8. You can now send the .csr file to a Certificate Authority. The Certificate Authority will provide details on how to send the .csr file. The Certificate Authority will send you a .crt file. The .crt file is the public key that matches your private key for the SL1 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 ways to create your own SSL certificate:

- If your organization is a root Certificate Authority (for example, some departments of the US government), you can create your own private key and public key for each ScienceLogic server.
- If your security requirements allow a self-signed certificate, you can create your own private key and public key for each SL1 appliance.

Remember to create key pairs for all for each SL1 appliance in your SL1 system and also remember to create two key pairs for each Administration Portal in your SL1 system.

If your organization is a Certificate Authority, see your organization's internal documentation on creating a certificate for Nginx.

If you want to create a self-signed certificate, perform the following:

1. Either go to the console of the SL1 appliance or use SSH to access the server. Open a shell session on the SL1 appliance. Log in as "em7admin".
2. First, you must 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*.

NOTE: Make sure the file is **not** named *silossl.key*. This is the name of the pre-existing ScienceLogic, self-signed certificate file.

3. You will be prompted to enter a pass phrase for the key. Best practice is to make a backup copy of the key file and the pass phrase and store both in a secure location.
5. Next, you must 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 SL1 appliance .
- *[keyname].crt* is the public key (certificate) for the SL1 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 SL1 appliance.

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.

6. Remove the private key pass phrase. To do this, enter the following at the shell prompt:

```
sudo openssl rsa -in [keyname].key -out [keyname].key.insecure
```

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 an SL1 Appliance

ScienceLogic does not provide support for third party certificates. Be advised that installing a new SSL certificate can affect the operation of SSL 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 SL1 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 SSL certificate on a SL1 appliance, perform the following:

1. Purchase a certificate from a certificate authority.
2. Copy the certificate files (*.key and all *.crt files) to a server that can access the SL1 appliance via SFTP.

NOTE: Make sure the files are **not** named **silssl.crt** and **silssl.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 SL1 appliance in the /etc/nginx directory.
4. Either go to the console of the SL1 appliance or use SSH to access the server. Open a shell session on the SL1 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].cert [bundle name].cert > [combined certificate name].cert
```

Use the combined .cert 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.cert and silossl.key and replacing with the names of the new .key and .certfiles:

```
ssl_certificate /etc/nginx/[name of .cert 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/em7ngx_web_ui.conf
- /etc/nginx/conf.d/em7ngx_em7proxy_web_ui.conf
- Edit the following lines, removing references to silossl.pem and silossl.key and replacing with the names of the new key files:


```
ssl_certificate /etc/nginx/[name of .cert 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 SSL 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 System > Settings > Appliances.
- In the **Appliance Manager** page, select the toolbox icon () for each server. Notice that the URL for the Configuration Utility includes https.



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 SL1 should communicate with a user store.

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

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

This chapter includes the following topics:

Authentication Profiles	280
<i>Viewing the List of Authentication Profiles</i>	280
<i>Filtering the List of Authentication Profiles</i>	281
<i>The "default" Authentication Profile</i>	282
<i>Creating an Authentication Profile</i>	283
<i>Editing an Authentication Profile</i>	286
<i>Deleting One or More Authentication Profiles</i>	286
Authentication Resources	286
<i>Viewing the List of Authentication Resources</i>	287
<i>Filtering the List of Authentication Resources</i>	288
<i>The "EM7 Internal" Resource</i>	289

The Legacy Authentication Resources	289
Creating an LDAP/AD Authentication Resource	290
Creating an SSO Authentication Resource	296
Editing an Authentication Resource	301
Deleting an Authentication Resource	301

Authentication Profiles

Authentication profiles are policies that align user accounts with one or more types of authentication:

- **Alignment by pattern matching.** SL1 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, SL1 will automatically use the matching profile to perform user authentication.
- **Credential Source.** Specifies from where SL1 should extract the user name and password or certificate to be authenticated. These credentials are passed to SL1 via HTTP. SL1 then passes the credentials to each authentication resource specified in the authentication profile. The authentication resources communicate with user stores that can authenticate the credentials.
- **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. 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](#).

Viewing the List of Authentication Profiles

To view a list of all authentication profiles in SL1:

1. Go to the **Authentication Profiles** page (System > Settings > Authentication > Profiles).

	Profile Name *	ID	Hostname Pattern	Priority Order	Edited By	Last Edited
1.	em7admin	2	*	0	em7admin	2016-02-16 21:52:52
2.	default	1	--	--	em7admin	--

2. The following information is displayed about each authentication profile:

- **Profile Name.** Name of the authentication profile.
- **ID.** Unique numeric ID, automatically assigned by SL1 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, SL1 applies the authentication profile to the user for the current session.
- **Priority Order.** If your SL1 System includes multiple authentication profiles, SL1 evaluates the authentication profiles in priority order, ascending. This column displays the priority order value for the authentication profiles.
- **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.

Filtering the List of Authentication Profiles

You can filter the list of authentication profiles on the **Authentication Profiles** page by one or more of the following parameters: **Profile Name**, **ID**, **Hostname Pattern**, **Priority Order**, **Edited By**, and **Last Edited**. The list of authentication profiles is dynamically updated as you select each filter. For each filter except **Last Edited**, you must enter text to match against. SL1 will search for authentication profiles that match the text, including partial matches. Text matches are not case-sensitive. You can use the following special characters in each filter except **Last Edited**:

- , (comma). Specifies an "or" operation. For example:
"dell, micro" would match all values that contain the string "dell" OR the string "micro".
- & (ampersand). Specifies an "and" operation. For example:
"dell & micro" would match all values that contain the string "dell" AND the string "micro".
- ! (exclamation mark). Specifies a "not" operation. For example:
"!dell" would match all values that do not contain the string "dell".
- ^ (caret mark). Specifies "starts with". For example:
"^ micro" would match all strings that start with "micro", like "microsoft".
"^" will include all rows that have a value in the column.
"!^" will include all rows that have no value in the column.

- \$ (dollar sign). Specifies "ends with". For example:
 "\$ware" would match all strings that end with "ware", like "VMware".
 "\$" will include all rows that have a value in the column.
 "!\$" will include all rows that have no value in the column.

By default, the cursor is placed in the first Filter-While-You-Type field. You can use the <Tab> key or your mouse to move your cursor through the fields.

Only authentication profiles that meet all the following filter criteria will be displayed in the **Authentication Profiles** page:

- **Profile Name**. Name of the authentication profile. You can enter text to match, including special characters, and the **Authentication Profiles** page will display only authentication profiles that have a matching name.
- **ID**. Unique numeric ID, automatically assigned by SL1 to each authentication profile. You can enter text to match, including special characters, and the **Authentication Profiles** page will display only authentication profiles that have a matching ID.
- **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, SL1 applies the authentication profile to the user for the current session. You can enter text to match, including special characters, and the **Authentication Profiles** page will display only authentication profiles that have a matching hostname pattern.
- **Priority Order**. You can enter text to match, including special characters, and the **Authentication Profiles** page will display only authentication profiles that have a matching priority number.
- **Edited By**. The user who created or last edited the authentication profile. You can enter text to match, including special characters, and the **Authentication Profiles** page will display only authentication profiles that have been created or edited by a matching username.
- **Last Edited**. Date and time the authentication profiles was created or last edited. You can select from a list of time periods. The **Authentication Profiles** page will display only authentication profiles that have been created or edited within that time period.

The "default" Authentication Profile

SL1 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, SL1 applies the *default* authentication profile.
- For users running version 7.7 or earlier of SL1 who apply one or more patches to upgrade to version 7.8, the **default** profile allows ScienceLogic authentication to perform as it did prior to version 7.8.
 - On patched systems, the *default* profile is included in the patch.
 - On patched systems, the *default* profile is pre-configured to allow ScienceLogic administrators to log in via the ScienceLogic login page and the authentication resource *EM7 Internal*.

- On patched systems, the *default* profile is pre-configured to allow credentials via *CAC/Client Certificate*, *HTTP Auth*, or the *EM7 Login Page*.
- On patched systems, the *default* profile is pre-configured to use all legacy authentication resources: *SSO (legacy)*, *LDAP/AD (legacy)*, and *EM7 Internal*.

NOTE: Administrators can edit the default profile and use the new, non-legacy authentication resources but are not required to do so.


- For users who installed version 7.8 or later of SL1 using an ISO, initially the *default* profile is pre-configured to allow ScienceLogic administrators to log in via *CAC/Client Certificate*, *HTTP Auth*, or the *EM7 Login Page* and the authentication resource *EM7 Internal*. This allows administrators to log in and perform initial configuration on the SL1 system.
 - On ISO systems, the *default* profile is included in the patch.
 - On ISO systems, the *default* profile is pre-configured to allow credentials via *CAC/Client Certificate*, *HTTP Auth*, or the *EM7 Login Page*.
 - On ISO systems, the *default* profile is pre-configured to use only the authentication resource *EM7 Internal*.

NOTE: After initial configuration, administrators can edit the **default** profile as best fits their organization.

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 page appears.

The screenshot shows a modal window titled "Create Authentication Profile". Inside, there's a sub-header "Authentication Profile Editor | Creating New Authentication Profile" with a "Reset" button. The form has several input fields: "Name", "Priority Order", "Pattern Type" (a dropdown menu currently showing "Wildcard"), and "AP Hostname Pattern". Below these are two main sections. The first is "Available Credential Sources" with a list containing "CAC/Client Cert", "EM7 Login Page", and "HTTP Auth". The second is "Available Authentication Resources" with a list containing "asdfasdf", "asdfasdfasdf", and "EM7 Internal". Between the available and aligned lists for both sections are "Add" (») and "Remove" («) buttons. On the right side of each aligned list, there are "Up" (↑) and "Down" (↓) arrow buttons. At the bottom center of the modal is a "Save" button.

3. Enter values in the following fields:

- **Name**. Name of the authentication profile.
- **Priority Order**. If your SL1 System includes multiple authentication profiles, SL1 evaluates the authentication profiles in ascending priority order. SL1 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 SL1 will evaluate the value in the **AP Hostname Pattern** field. Choices are:
 - *Wildcard*. SL1 will perform a text match, with wildcard characters (asterisks).
 - *Regex*. SL1 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, SL1 applies the authentication profile to the user for the current session.
 - For example, if you specify "*" (asterisk), any IP address or URL will match. SL1 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", SL1 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", SL1 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.
- **Available Credential Sources.** This field tells SL1 how to retrieve the user's credentials from the HTTP request to SL1. 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:
 - *CAC/Client Cert.* SL1 will retrieve a certificate from the HTTP request.
 - *EM7 Login Page.* SL1 will retrieve a user name and password from the ScienceLogic login page fields.
 - *HTTP Auth.* SL1 will retrieve a user name and password from the HTTP request.


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 SL1.

- **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 SL1 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 SL1. 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

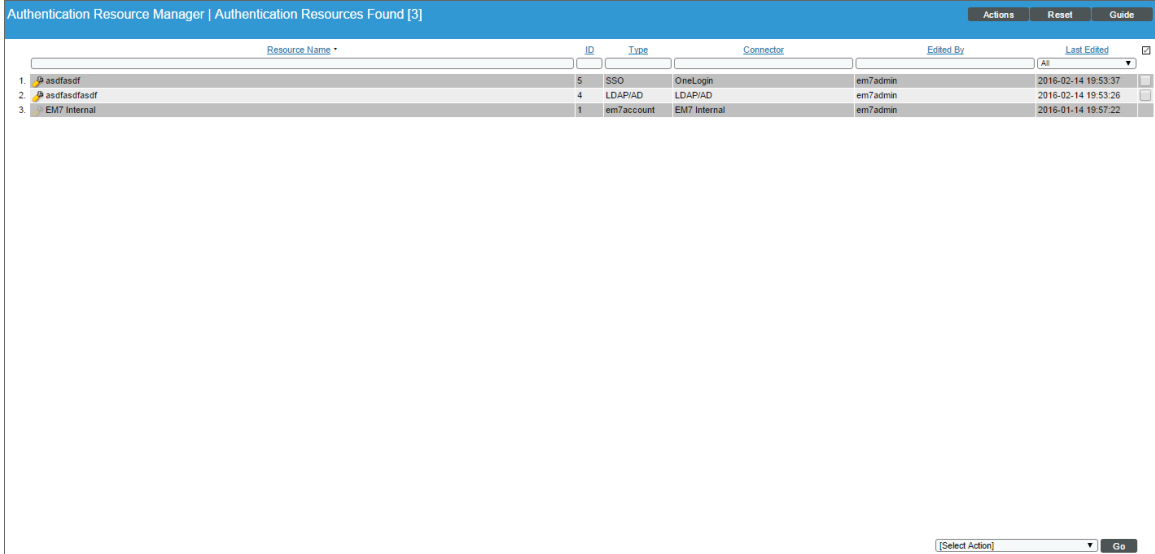
An authentication resource is a configuration policy that describes how SL1 should communicate with a user store. An authentication resource specifies the connector to use to communicate with the user store, the credential to use to connect to the user store (if applicable), and the URLs to examine during 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 SL1 System.

To view the list of authentication resources :

1. Go to the **Authentication Resource Manager** page (System > Settings > Authentication > Resources).



The screenshot shows the 'Authentication Resource Manager' interface. At the top, it says 'Authentication Resources Found [3]'. Below this is a table with columns: Resource Name, ID, Type, Connector, Edited By, and Last Edited. There are three rows of data. At the bottom right, there is a '[Select Action]' dropdown and a 'Go' button.

	Resource Name *	ID	Type	Connector	Edited By	Last Edited	
1.	asofasof	5	SSO	OneLogin	em7admin	2016-02-14 19:53:37	
2.	asofasof	4	LDAP/AD	LDAP/AD	em7admin	2016-02-14 19:53:26	
3.	EM7 Internal	1	em7account	EM7 Internal	em7admin	2016-01-14 19:57:22	

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.
- **ID.** Unique numeric ID, automatically assigned by SL1 to each authentication resource.
- **Type.** Specifies the user store that is associated with the resource. Possible types are:
 - *EM7 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:
 - *EM7 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 SL1 . SL1 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 SL1 . SL1 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.

Filtering the List of Authentication Resources

You can filter the list of authentication resources on the **Authentication Resource Manager** page by one or more of the following parameters: **Resource Name**, **ID**, **Type**, **Connector**, **Edited By**, and **Last Edited**. The list of authentication resources is dynamically updated as you select each filter. For each filter except **Last Edited**, you must enter text to match against. SL1 will search for authentication resources that match the text, including partial matches. Text matches are not case-sensitive. You can use the following special characters in each filter except **Last Edited**:

- , (comma). Specifies an "or" operation. For example:
 - "dell, micro" would match all values that contain the string "dell" OR the string "micro".
- & (ampersand). Specifies an "and" operation. For example:
 - "dell & micro" would match all values that contain the string "dell" AND the string "micro".
- ! (exclamation mark). Specifies a "not" operation. For example:
 - "!dell" would match all values that do not contain the string "dell".
- ^ (caret mark). Specifies "starts with". For example:
 - "^ micro" would match all strings that start with "micro", like "microsoft".
 - "^" will include all rows that have a value in the column.
 - "!^" will include all rows that have no value in the column.

- \$ (dollar sign). Specifies "ends with". For example:
 "\$ware" would match all strings that end with "ware", like "VMware".
 "\$" will include all rows that have a value in the column.
 "!\$" will include all rows that have no value in the column.

By default, the cursor is placed in the first Filter-While-You-Type field. You can use the <Tab> key or your mouse to move your cursor through the fields.

Only authentication resources that meet all the following filter criteria will be displayed in the **Authentication Resource Manager** page:

- **Resource Name**. Name of the authentication resource. You can enter text to match, including special characters, and the **Authentication Resource Manager** page will display only authentication resources that have a matching name.
- **ID**. Unique numeric ID, automatically assigned by SL1 to each authentication resource. You can enter text to match, including special characters, and the **Authentication Resource Manager** page will display only authentication resources that have a matching ID.
- **Type**. Specifies the user store that is associated with the resource. You can enter text to match, including special characters, and the **Authentication Resource Manager** page will display only authentication resources that have a matching type.
- **Connector**. The specific software that allows communication between the authentication resource and the user store. You can enter text to match, including special characters, and the **Authentication Resource Manager** page will display only authentication resources that have a matching connector.
- **Last Edited**. Date and time the authentication resources was created or last edited. You can select from a list of time periods. The **Authentication Resource Manager** page will display only authentication resources that have been created or edited within that time period.
- **Edited By**. ScienceLogic user who created or last edited the authentication resource. You can enter text to match, including special characters, and the **Authentication Resource Manager** page will display only authentication resources that have been created or edited by a matching username.

The "EM7 Internal" Resource

The *EM7 Internal* resource allows you to access the user store in the ScienceLogic database.

- By default, each SL1 System, whether upgraded to version 7.8 or built from a 7.8 ISO, includes the *EM7 Internal* authentication resource.
- You cannot create an *EM7 Internal* authentication resource.
- You cannot edit or delete the *EM7 Internal* authentication resource included with your SL1 System.
- Each SL1 System can include only one the *EM7 Internal* authentication resource.

The Legacy Authentication Resources

SL1 includes two "legacy" authentication resources:

- *LDAP/AD* with connector *LDAP/AD - Legacy*
- *SSO (legacy)* with connector *SimpleSAML - Legacy*

These legacy authentication resources allow patched systems (systems that upgraded to version 7.8 of SL1) to continue using the same authentication as used prior to upgrading to version 7.8.

- Legacy authentication resources are available only on systems that have upgraded from a previous version of SL1.
- You cannot create a new authentication resource using the legacy connectors.
- If you edit and save changes to the *LDAP/AD* authentication resource, SL1 updates the connector from *LDAP/AD - Legacy* to the non-legacy connector *LDAP/AD*.

Creating an LDAP/AD Authentication Resource

The **LDAP/AD Auth Resource Editor** modal 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 SL1:

- You can configure SL1 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 SL1.
- You can use LDAP or Active Directory to authenticate one or more ScienceLogic users when they log in to SL1.

To create an LDAP/AD authentication resource:

1. Go to the **Authentication Resource Manager** page (System > Settings > Authentication > Resources).

- Click the **[Actions]** menu and then select *Create LDAP/AD Resource*. The **LDAP/AD Auth Resource Editor** modal page appears.

- Enter values in the following fields:

Basic Settings

- **Name**. Name of the LDAP/AD authentication resource.
- **Read Credential**. Credential that allows SL1 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 SL1 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*.

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.

- **User Name Suffix**. Optional field. Because SL1 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 SL1 with the user name **bishopbrennan**.
- SL1 will log that user in as **bishopbrennan@ad.local**.

NOTE: A best practice to avoid collisions is to use email addresses as user names.

- **Search Filter.** Specifies where to find the user's account information in LDAP or Active Directory. You must tell SL1 where to find the LDAP or AD attribute that maps to the user's account name in SL1.

For example, an LDAP user might use his/her uid value to log in to SL1. 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 SL1 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 SL1 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 to EM7 on login.** If an LDAP or AD administrator makes changes to an LDAP or AD account, SL1 will automatically retrieve those updates and apply them to the user's account in SL1 (in the **Account Properties** page) the next time the user logs in to SL1. (For more information about user account properties, see the **Organizations & Users** manual.)
- **Sync EM7 values to directory on save.** If a ScienceLogic administrator made changes to the ScienceLogic account, SL1 will automatically write those changes to the user's account in LDAP or Active Directory.

NOTE: The **Sync EM7 values to directory on save** option requires a write credential.

Attribute Mapping

If you have configured SL1 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.)

SL1 automatically populates as many of these fields as possible. You can edit or delete the default values provided by SL1. For example, SL1 automatically inserts the value of the LDAP/AD attribute "sn" (surname) into the **Last Name** field in each user's **Account Properties** page.

NOTE: SL1 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, SL1 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, SL1 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, SL1 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, SL1 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, SL1 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, SL1 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, SL1 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, SL1 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, SL1 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 SL1 should automatically create ScienceLogic accounts for each LDAP or Active Directory user in the search base (which is specified in the credential), whether SL1 should simply use LDAP or Active Directory to authenticate one or more users, or whether SL1 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 SL1 can log in to SL1. However, if one or more users' **Account Permissions** page specifies *LDAP /Active Directory* in the **Authentication Method** field, SL1 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 SL1 using the LDAP or AD attribute specified in the **Search Filter** field, SL1 will automatically create an account for that user. SL1 will use **one user policy** (specified in the **Policy** field) to create all imported LDAP or AD user accounts. SL1 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 SL1 using the LDAP or AD attribute specified in the **Search Filter** field, SL1 will automatically create an account for that user. SL1 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". SL1 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). SL1 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 SL1 and for users who already have an account in SL1.
 - *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 SL1 using the Sales User Policy. User accounts from the NOC department would be imported into SL1 using the NOC User Policy.
- To define additional **Value** and **Policy** fields, click on the green 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 map 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 SL1 :

- You can configure SL1 to automatically create user accounts for existing SSO users and then always use SSO to authenticate those users when they log in to SL1 .
- You can use SSO to authenticate one or more ScienceLogic users when they log in to SL1 .

To create an SSO authentication resource:

1. Go to the **Authentication Resource Manager** page (System > Settings > Authentication > Resources).

- Click the **[Actions]** menu and then select *Create SSO Resource*. The **SSO Auth Resource Editor** page appears.

- Enter values in the following fields:

Basic Settings

- **Name**. Name of the SSO authentication resource.
- **IdP Entity ID**. Globally unique name for the identity provider or service provider, in the format of an absolute URL.
- **IdP Cert Fingerprint**. The SHA1 certificate fingerprint, provided by the identity provider or service provider.
- **User Name Suffix**. Optional field. 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:
 - Suppose we entered **@ad.local** in this field.
 - Suppose the next LDAP/AD user logs in to SL1 with the user name **bishopbrennan**.
 - SL1 will log in that user as **bishopbrennan@ad.local**.

NOTE: A best practice to avoid collisions is to use email addresses as user names.

- **IdP SSO URL**. The URL to which SL1 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 SL1, enter the URL to which SL1 will post the logout request to the IdP. If you leave this field blank, a user can log out of SL1 without automatically logging out of the IdP.
- **Sync directory values to EM7 on login.** If an SSO administrator makes changes to an SSO account, SL1 will automatically retrieve those updates and apply them to the user's account in SL1 (in the **Account Properties** page) the next time the user logs in to SL1. (For more information about user account properties, see the **Organizations & Users** manual.)

Attribute Mapping

If you have configured SL1 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.)

SL1 automatically populates as many of these fields as possible. You can edit or delete the default values provided by SL1. For example, SL1 automatically inserts the value of the SAML attribute "sn" (surname) into the **Last Name** field in each user's **Account Properties** page.

NOTE: SL1 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, SL1 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, SL1 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, SL1 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, SL1 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, SL1 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, SL1 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, SL1 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, SL1 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, SL1 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 SL1 should automatically create ScienceLogic accounts for each SSO user, whether SL1 should simply use SSO to authenticate one or more users, or whether SL1 will refuse to authenticate specific users. Choices are:
 - *Do not authenticate new users.* Only those users who have an account already created in SL1 can log in to SL1, which will authenticate those users with SSO using the settings specified in this page.
 - *Static policy alignment.* If an SSO user tries to access SL1, SL1 will automatically create an account for that user. SL1 will use **one user policy** (specified in the **Policy** field) to create the imported SSO user accounts for this authentication resource. SL1 will also use the settings specified in this page when creating the account.
 - *Dynamic policy alignment.* If an SSO users tries to access SL1, SL1 will automatically create an account for that user. SL1 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". SL1 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). SL1 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 SL1 and for users who already have an account in SL1.
 - *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 SL1 using the *Sales User Policy*.
 - User accounts from the *NOC* department would be imported into SL1 using the *NOC User Policy*.
- To define additional **Value** and **Policy** fields, click on the green 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 SL1. 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 *EM7 Internal* authentication resource.

Chapter

16

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 SL1 system. When you create or edit an entry in the **Host File Entry Manager** page, SL1 automatically sends an update to every Data Collector in the specified Collector Group.

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

- The SL1 system does not reside in the end-customer's domain
- The SL1 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 SL1 system monitors

You can create host file entries for each device managed by the SL1 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 SL1 user interface:

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

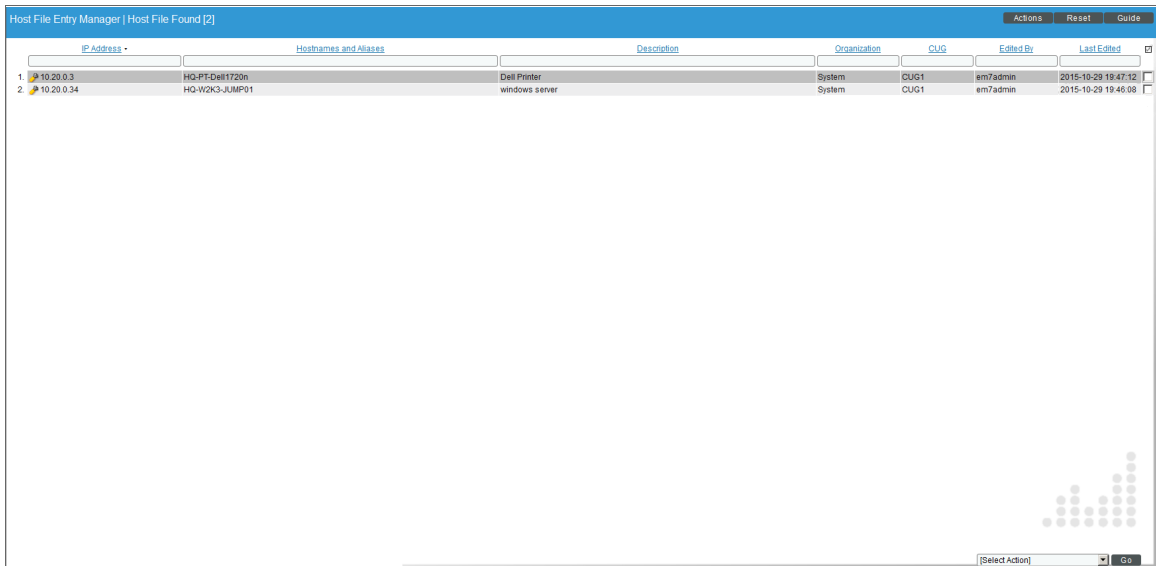
This chapter includes the following topics:

Viewing the List of Host Entries	304
Creating a New Host Entry	305
Editing a Host Entry	306
Using an Existing Host File Entry to Create a New Host File Entry (Save As)	308
Deleting One or More Host Entries	309

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).



The screenshot shows the 'Host File Entry Manager' interface with a table of host entries. The table has columns for IP Address, Hostnames and Aliases, Description, Organization, CUG, Edited By, and Last Edited. Two entries are visible:

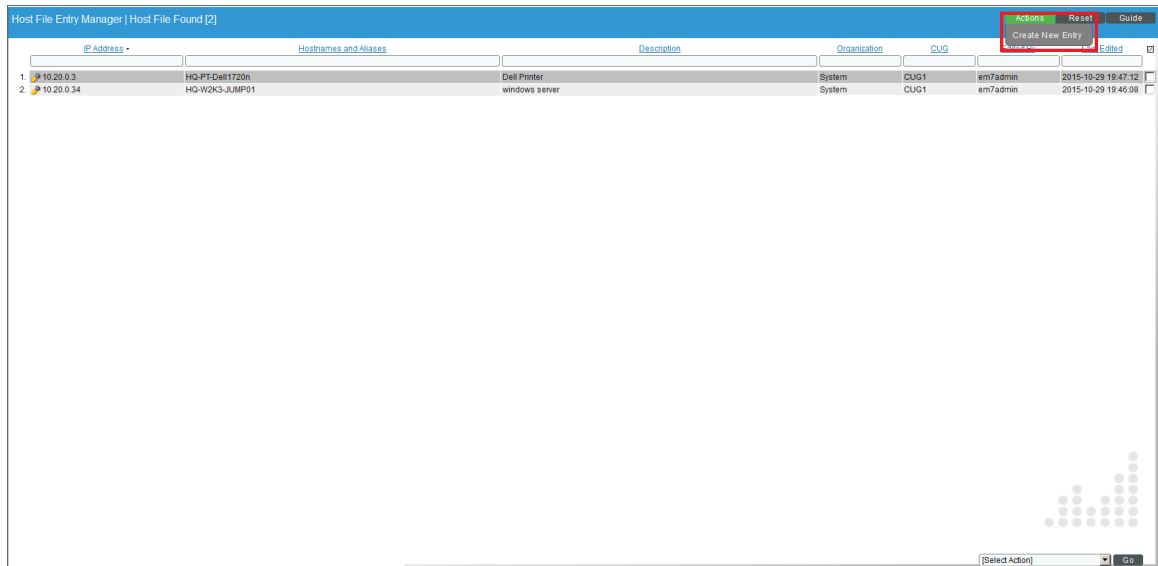
	IP Address	Hostnames and Aliases	Description	Organization	CUG	Edited By	Last Edited
1	10.20.0.2	HQ-RT-DellT720n	Dell Printer	System	CUG1	em7admin	2015-10-29 19:47:12
2	10.20.0.34	HQ-W2K3-JUMP01	windows server	System	CUG1	em7admin	2015-10-29 19:46:08

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 SL1 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.

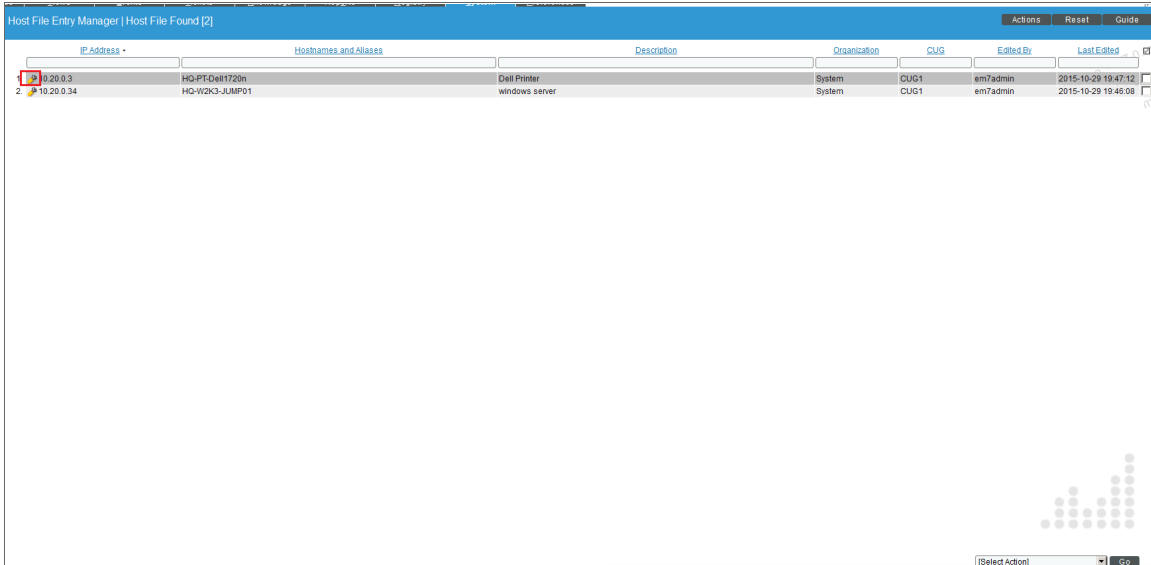
- **Hostnames and Aliases.** The hostname to align with the specified IP address. You can also include a space-delimited list of aliases for the host name.
- **Description.** Description of the host entry. This field is not written to the host file. This field is for administrators to use when managing 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:

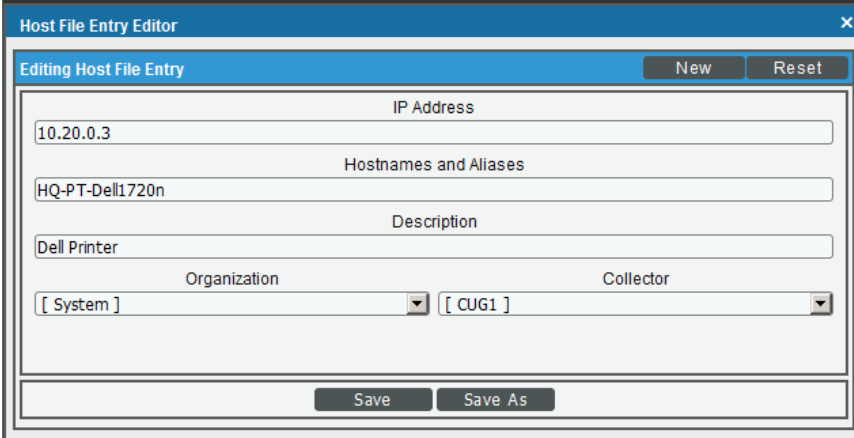
1. Go to the **Host File Entry Manager** page (System > Customize > Host Files).



The screenshot shows the 'Host File Entry Manager' interface with a table of host entries. The table has columns for IP Address, Hostnames and Aliases, Description, Organization, CUG, Edited By, and Last Edited. Two entries are visible: one for a Dell Printer and one for a windows server. The first entry is highlighted with a red box.

IP Address	Hostnames and Aliases	Description	Organization	CUG	Edited By	Last Edited
19.20.0.3	HQ-PT-Dell1720n	Dell Printer	System	CUG1	em7admin	2015-10-29 19:47:12
10.20.0.34	HQ-W2K3-JUMP01	windows server	System	CUG1	em7admin	2015-10-29 19:46:08

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.



The screenshot shows a modal window titled "Host File Entry Editor" with a close button (X) in the top right corner. The main title bar of the modal is "Editing Host File Entry" and contains "New" and "Reset" buttons. The form fields are as follows:

- IP Address:** 10.20.0.3
- Hostnames and Aliases:** HQ-PT-Dell1720n
- Description:** Dell Printer
- Organization:** [System]
- Collector:** [CUG1]

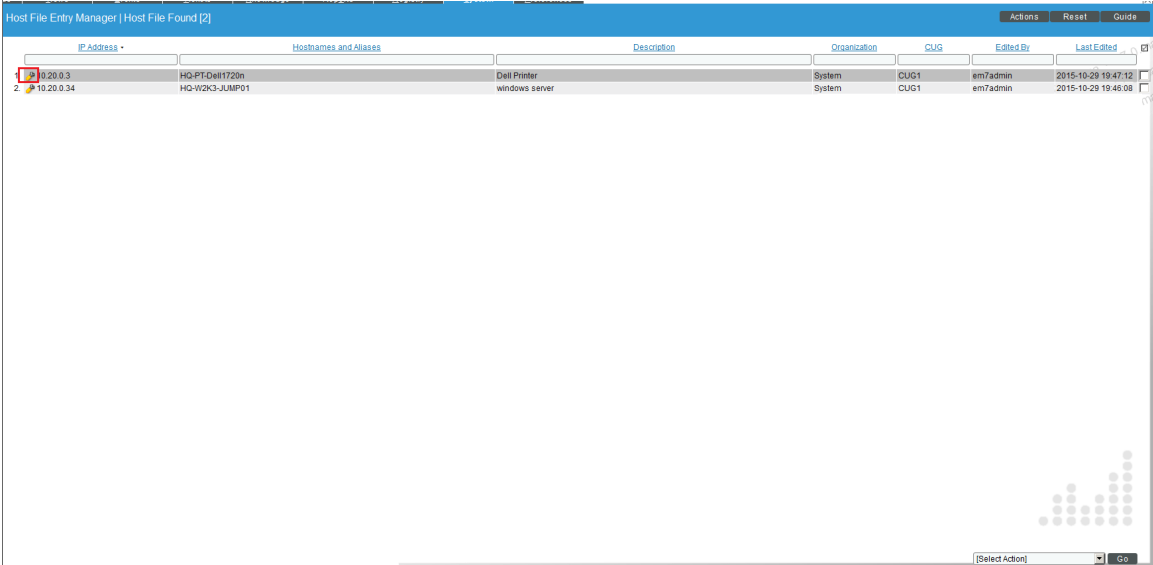
At the bottom of the modal are "Save" and "Save As" buttons.

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:

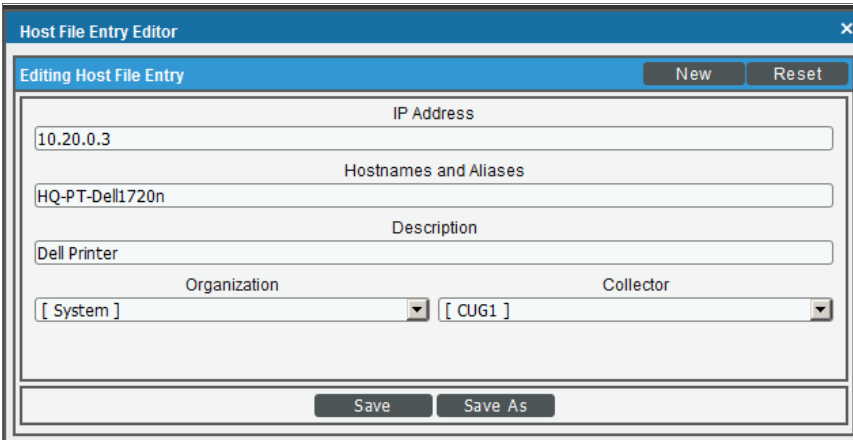
1. Go to the **Host File Entry Manager** page (System > Customize > Host Files).



The screenshot shows the 'Host File Entry Manager' interface. At the top, there are tabs for 'Actions', 'Reset', and 'Guide'. Below the tabs is a table with the following columns: IP Address, Hostnames and Aliases, Description, Organization, CUG, Edited By, and Last Edited. The table contains two entries:

IP Address	Hostnames and Aliases	Description	Organization	CUG	Edited By	Last Edited
10.20.0.3	HQ-PT-Dell1720n	Dell Printer	System	CUG1	em7admin	2015-10-29 19:47:12
10.20.0.34	HQ-W2K3-JUMP01	windows server	System	CUG1	em7admin	2015-10-29 19:48:08

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.



The screenshot shows the 'Host File Entry Editor' modal page. It has a title bar with 'Host File Entry Editor' and a close button. Below the title bar is a sub-header 'Editing Host File Entry' with 'New' and 'Reset' buttons. The form contains the following fields:

- IP Address: 10.20.0.3
- Hostnames and Aliases: HQ-PT-Dell1720n
- Description: Dell Printer
- Organization: [System] (dropdown)
- Collector: [CUG1] (dropdown)

At the bottom of the form are 'Save' and 'Save As' buttons.

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.

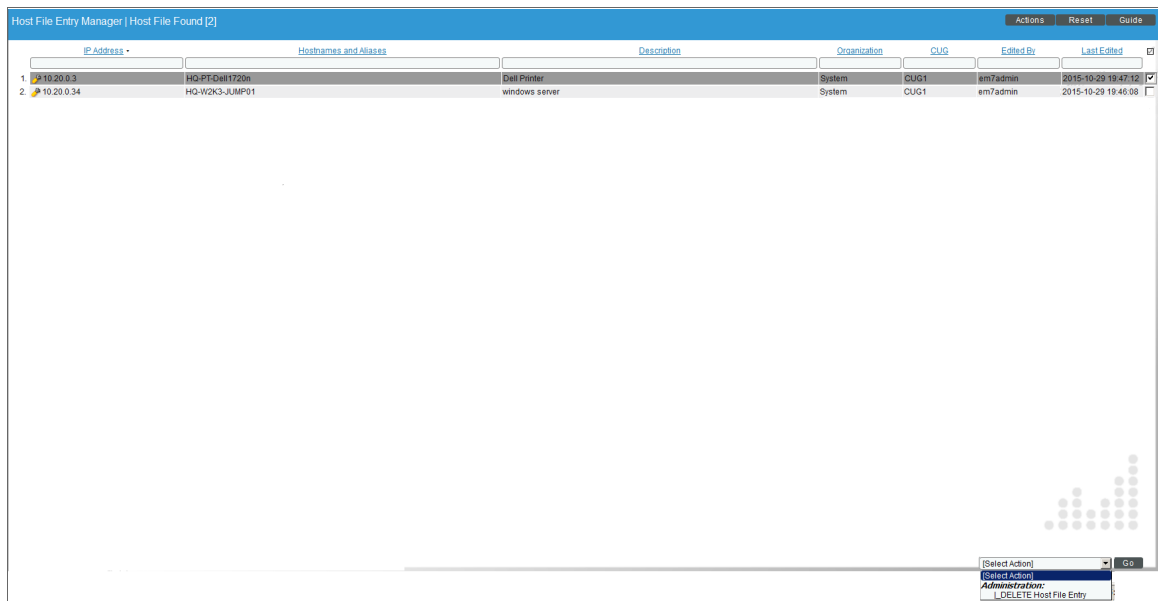
- **Hostnames and Aliases.** The host name to align with the specified IP address. You can also include a space-delimited list of aliases for the host name.
- **Description.** Description of the host entry. This field is not written to the host file. This field is for administrators to use when managing 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 - 2020, ScienceLogic, Inc.

All rights reserved.

LIMITATION OF LIABILITY AND GENERAL DISCLAIMER

ALL INFORMATION AVAILABLE IN THIS GUIDE IS PROVIDED "AS IS," WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED. SCIENCELOGIC™ AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT.

Although ScienceLogic™ has attempted to provide accurate information on this Site, information on this Site may contain inadvertent technical inaccuracies or typographical errors, and ScienceLogic™ assumes no responsibility for the accuracy of the information. Information may be changed or updated without notice. ScienceLogic™ may also make improvements and / or changes in the products or services described in this Site at any time without notice.

Copyrights and Trademarks

ScienceLogic, the ScienceLogic logo, and EM7 are trademarks of ScienceLogic, Inc. in the United States, other countries, or both.

Below is a list of trademarks and service marks that should be credited to ScienceLogic, Inc. The ® and ™ symbols reflect the trademark registration status in the U.S. Patent and Trademark Office and may not be appropriate for materials to be distributed outside the United States.

- ScienceLogic™
- EM7™ and em7™
- Simplify IT™
- Dynamic Application™
- Relational Infrastructure Management™

The absence of a product or service name, slogan or logo from this list does not constitute a waiver of ScienceLogic's trademark or other intellectual property rights concerning that name, slogan, or logo.

Please note that laws concerning use of trademarks or product names vary by country. Always consult a local attorney for additional guidance.

Other

If any provision of this agreement shall be unlawful, void, or for any reason unenforceable, then that provision shall be deemed severable from this agreement and shall not affect the validity and enforceability of any remaining provisions. This is the entire agreement between the parties relating to the matters contained herein.

In the U.S. and other jurisdictions, trademark owners have a duty to police the use of their marks. Therefore, if you become aware of any improper use of ScienceLogic Trademarks, including infringement or counterfeiting by third parties, report them to Science Logic's legal department immediately. Report as much detail as possible about the misuse, including the name of the party, contact information, and copies or photographs of the potential misuse to: legal@sciencelogic.com



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

International: +1-703-354-1010