



Introduction to SL1

SL1 version 12.2.1 (Document revision 2)

Table of Contents

Overview of the SL1 User Interface	6
Prerequisites	8
Logging In and Out of the SL1 User Interface	8
The SL1 User Interface	8
The Classic User Interface	9
Using the Navigation Menus	10
The Setup and Config Page	10
Setting a Home Page in SL1	11
Filtering Inventory Pages	12
Filter Syntax	13
String and Numeric	13
String	14
Numeric	14
Examples	15
Editing the Settings for an Inventory Page	16
Adjusting and Sorting Columns	16
Creating a Multi Sort	17
Editing Column Preferences	17
Adjusting the Row Density	18
Exporting to a CSV File	18
Viewing the Device Summary Modal from an Inventory Page	19
Using Basic Search	20
Fields Used by an "ANY" Basic Search	22
Saving a Search	23
Performing an Advanced Search	24
Components of an Advanced Search	27
Fields	28
Operators	29
Values	30
Strings	31
Escape characters	31

Examples of Advanced Searches	32
Advanced Search Examples on the Events Page	33
Advanced Search Examples for Dynamic Component Mapping (DCM) Scenarios	33
Scenario 1: vCenters	33
Search 1a	34
Search 1b	35
Scenario 2: SQL Servers	36
Searches	37
Performing Bulk Actions	38
Customizing the SL1 User Interface	39
Getting Help and More Information	40
The About Page	41
Tips and Best Practices for Using the Product Documentation	42
General	43
Searching the Site	43
Links	43
Images	44
Guides from the "Classic" User Interface	44
Filtering the Items on a Classic Page	44
Tool Tips	46
Creating and Using Bookmarks	46
The Finder Tool	47
Searching for One or More Elements	47
Viewing a List of All Elements in SL1	48
Viewing the Results Table	48
The Toolbox	48
Overview of SL1 Features	50
SL1 Architecture	52
User Interface	52
Database Server	52
Data Collection	53
Message Collection	53

API	53
SL1 Extended Architecture	54
Compute	54
Load Balancer	54
Storage	54
Management	55
The SL1 Agent	55
Organizations and Users	56
Users	57
Credentials	58
Discovery	58
Device Administration	59
Virtual Device	59
Component Device	59
Business Services	60
Events	60
Automation	60
Anomaly Detection	61
Dynamic Applications	61
PowerPacks	62
Maps	62
Dashboards	63
Reports	64
Custom Reports	64
Embedded Reports	64
Ticketing	65
Asset Management	65
User Preferences	66
Account Preferences	67
Schedule Manager	71
Viewing the Schedule Manager	71
Defining a Scheduled or Recurring Calendar Item	73

Enabling or Disabling One or More Scheduled Calendar Items	75
Deleting One or More Scheduled Calendar Items	75
My Contact Information	76

Chapter

1

Overview of the SL1 User Interface



Overview

SL1 offers you the capabilities to monitor your hybrid cloud infrastructure, improve service visibility, and automate your IT workflows. This manual is intended for new users of SL1, and it covers features in both the new user interface that was released with version 8.12.0 of SL1 as well as the "classic" user interface of SL1 that was available before version 8.12.0.

Please note that the new user interface provides access to both new and "classic" user interface features.

For an overview of the SL1 user interface, view the video at <https://sciencelogic.com/product/resources/navigate-the-sl1-ui>.

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

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

This chapter covers the following topics:

Prerequisites	8
Logging In and Out of the SL1 User Interface	8
Using the Navigation Menus	10
The Setup and Config Page	10
Setting a Home Page in SL1	11
Filtering Inventory Pages	12
Using Basic Search	20
Performing an Advanced Search	24
Performing Bulk Actions	38

<i>Customizing the SL1 User Interface</i>	39
<i>Getting Help and More Information</i>	40
<i>Filtering the Items on a Classic Page</i>	44
<i>Tool Tips</i>	46
<i>Creating and Using Bookmarks</i>	46
<i>The Finder Tool</i>	47
<i>The Toolbox</i>	48

Prerequisites

This manual assumes that the initial installation and configuration (deployment) of SL1 version 8.12.0 or later has been completed. For details on the initial configuration of SL1, see the *Installation* manual.

Logging In and Out of the SL1 User Interface

This topic covers how to access the different user interfaces for SL1: the default SL1 user interface ("AP2") and the "classic" user interface.

NOTE: If you are upgrading from a version of SL1 prior to 11.1.0, the classic user interface might still be set as the default user interface. To change the default user interface (UI), go to the **Behavior Settings** page (System > Settings > Behavior) and select the **New UI Default** checkbox.

The SL1 User Interface

Starting with SL1 version 11.1.0, you no longer need to type **/ap2** at the end of the URL or IP address, such as **https://sl1.sciencelogic.com/ap2**, to get to the SL1 user interface. The **/ap2** is added to the IP address by default.

To log in to the SL1 user interface:

1. In a browser, type the URL or IP address for your SL1 system. The login page for SL1 appears.
2. Type the current user name and password you use with SL1 and click **Log In**. The SL1 login button displays a spinning status icon while you are being logged in.
3. If your company uses Single Sign-On (SSO) for authentication, you will be redirected to your company's SSO page, where you can log in to SL1 with your SSO credentials. When you log out, the logout screen redirects you to an SSO page instead of the typical login screen.
4. If you are logging in for the first time or you are using the default system password, you will be prompted to change your password. Type your username, your old password, and type your new password twice in the *New Password* and *Confirm Password* fields. Click **[Reset Password]**.
5. If your SL1 system is a Military Unique Deployment (MUD) system, a Disclaimer modal appears, and you will need to click **[OK]** after reading the text to use SL1.

6. For an Administrator user, the End User License Agreement (EULA) appears the first time the Administrator user logs in to SL1. The user must agree to the terms before using SL1.

TIP: If you can still see the horizontal NavBar or other features of the classic user interface in the SL1 user interface, click **[Refresh]** in your browser.

NOTE: If you navigate to the classic SL1 user interface by typing **/em7** in the URL, the top navigation might be missing. To address this issue, navigate back to the new user interface, log out, and then log in to the classic user interface.

To log out of the SL1 user interface:

1. Click your user name in the navigation bar at the top of any SL1 page.
2. Click **Log off**. You are logged out, and a logoff page appears with the option to log in again.

The Classic User Interface

To log in to the classic user interface for SL1:

1. In a browser, type the URL or IP address for your SL1 system, followed by **/em7**. The login page for the classic user interface appears.
2. Type the current user name and password you use with SL1 and click **Log In**.

TIP: You can log out of the classic user interface by typing **/em7/logout.em7** at the end of the URL or IP address.

3. If you are logging in for the first time or you are using the default system password, you will be prompted to change your password. Type your username, your old password, and type your new password twice in the *New Password* and *Confirm Password* fields. Click **[Reset Password]**.
4. To switch to the default SL1 user interface while you are using the classic user interface, type **/** at the end of the URL or IP address for your SL1 system and press **[Enter]**. The login page of the SL1 user interface appears.
5. To switch to the classic user interface while you are using the SL1 user interface, type **/em7** at the end of the URL or IP address for your SL1 system and press **[Enter]**. The login page for the classic user interface appears.
6. To make the classic user interface the default interface, go to the **Behavior Settings** page (System > Settings > Behavior) and unselect the **New UI Default** checkbox.

TIP: If you need to run both the classic and SL1 user interfaces simultaneously, you can open the classic and the SL1 version in two different browsers, or you can open private or incognito browser windows and then log in to the two user interfaces.

Using the Navigation Menus

The SL1 user interface includes content from both the "new" user interface (also called "ap2") and the "classic" user interface (also called "em7"). This unified user interface provides a more streamlined method of navigation that uses two menus: a **basic menu** and an **Advanced menu**.

The name of the SL1 page appears in the title bar of your browser, and the page title also appears in your browsing history to help you navigate through the pages you have visited.

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

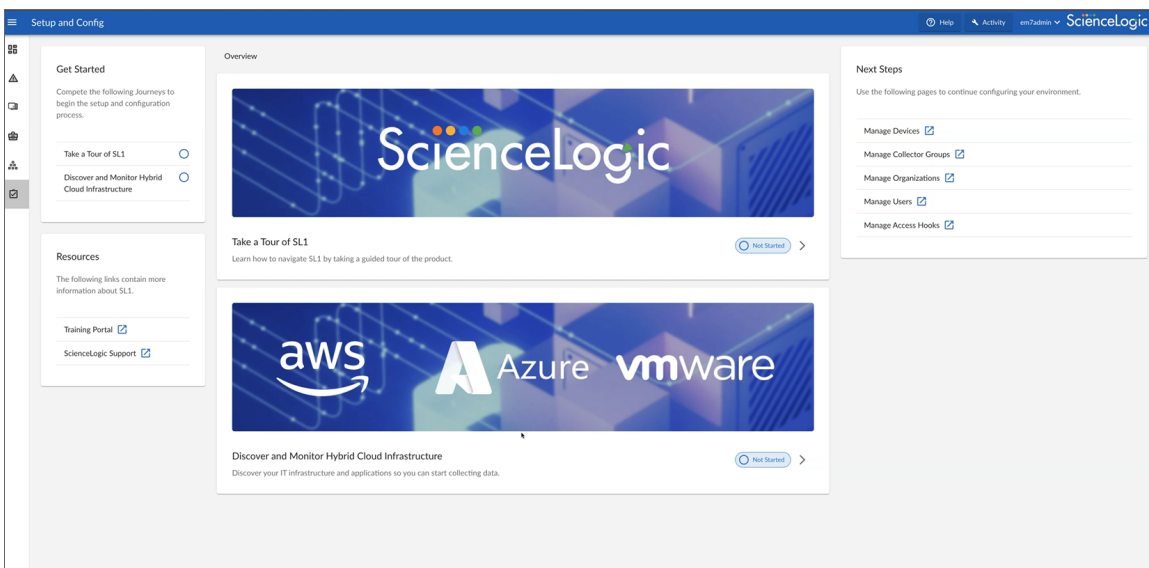
- To view a pop-out list of menu options for the basic menu, click the menu icon (☰) at the top left of any SL1 page. Use the up and down arrow buttons (^ v) to expand and contract the menu options.
- To view the **Advanced Menu** page, which contains links to *all* of the menu options, click the Advanced menu icon (⋮) at the bottom left of any SL1 page. Use **[Ctrl]+[F]** in your browser to quickly find a page.

The basic menu and the Advanced menu displays links for only the pages for which you have access.

The Setup and Config Page

The **Setup and Config** page (🔑) displays all information relevant to getting started in SL1 for administrator-level users. Included on this page are a number of **journeys**, intuitive self-service workflows that will guide you through the most common SL1 system tasks. Click the name of a workflow to get started.

This page also contains informational cards that provide you with the proper resources for SL1 setup and configuration.



The informational cards on this page include:

- **Get Started.** Displays a list of available user journeys and their journey status. Click the name of the journey to get started. The journeys include:
 - Take a Tour of SL1
 - Discover and Monitor Hybrid Cloud Infrastructure
- **Resources.** Hosts additional external resources to help you with setup and configuration; these links include:
 - Training Portal
 - ScienceLogic Support
- **Overview.** Provides links to the user journeys. These journeys include guided tours and interactive wizards that help you set up and refine your SL1 environment:
- **Next Steps.** Contains links to other pages in SL1 where you can continue working after completing some or all of a journey:
 - Manage Devices
 - Manage Collector Groups
 - Manage Organizations
 - Manage Users
 - Manage Access Hooks

Setting a Home Page in SL1

The **Dashboards** page is the default **home page** when you first launch SL1. Starting with version 12.2.0 of SL1, you can set any of the top-level pages that display in the left-hand navigation bar as your landing page, including **Events, Devices, Business Services, Maps, Machine Learning, and Setup and Config**. You can also set a specific dashboard from the **Dashboards** page as your home page. If you lose permission to the page or dashboard you set as your landing page, the **Events** page will be automatically set as your home page. You can click the **[Your Home]** tab in the left-hand navigation bar to be redirected to your home page.

NOTE: You cannot use any pages that display in the navigation bar below these pages as a landing page. This feature does not work with classic dashboards. For more information on what you can see and do on the default dashboards page, see the chapter on "Default Dashboards" in the **Dashboards** manual.

To set your home page:

1. In SL1, navigate to the page or the SL1 dashboard that you want to set as your home page.
2. Click your user name in the navigation bar at the top of the page and select *Set as Home Page*. That page or dashboard will be your home page the next time you launch SL1.


NOTE: If the page you are currently on cannot be used as a home page, the *Set as Home Page* option will be grayed out.

Filtering Inventory Pages

In SL1, an **inventory page** is any page (or tab) that contains a table with a list of items that you work with, like the **Events** page or the **Devices** page. At the top of most inventory pages in the SL1 user interface are a set of filters for the columns in the table on that page.

On an inventory page, you can start typing filter text or select filter options in one or more of these filters to narrow down the items in the list to just the items you want to view.

In the following image, the user is filtering the **Devices** inventory page by typing part of an IP address in the filter for the **IP Address** column:



The screenshot shows the SL1 interface for the 'Devices' page. At the top, there is a search bar and a table with columns: Device Name, IP Address, Category, Class, Organization, ID, Collection State, and State. The 'IP Address' filter field contains the text '10.2.23.4' and is highlighted with a red box. Below the table, four device entries are visible, all with IP addresses starting with '10.2.23.'. The interface includes a navigation menu on the left, a top navigation bar with 'Help', 'Activity', and 'Em7admin', and a 'ScienceLogic' logo.


Device Name	IP Address	Category	Class	Organization	ID	Collection State	State
SF-dist-mc1-discvm-1022346	10.2.23.46	System.EM7	ScienceLogic, Inc. EM7 Message Colle	System	1	active	Major
SF-dist-cug1-discvm-1022345	10.2.23.45	System.EM7	ScienceLogic, Inc. EM7 Data Collector	System	2	active	Major
SF-dist-ap-discvm-1022348	10.2.23.48	System.EM7	ScienceLogic, Inc. EM7 Admin Portal	System	13	active	Major
SF-dist-db-discvm-1022349	10.2.23.49	System.EM7	ScienceLogic, Inc. EM7 Database	System	14	active	Major

As you type text in the **Filter** field at the top of a column, SL1 starts to filter the list to include only those elements that include your search terms.

Some filters, like the **State** filter, let you click a drop-down arrow to select search options from a drop-down menu. Also, any filters related to date and time let you click a calendar icon (📅) to select a date or a date range, and a clock icon (🕒), where relevant, to select a time or a time range.

NOTE: If you were using a version of SL1 from before version 11.2.0, any list and column preferences you had configured will *not* be retained with the new features for inventory pages after you upgrade to 11.2.0 or later. This includes customizations such as which columns are visible, column widths, column order, and other customizations.

SL1 retains your search filters on pages in the SL1 user interface even after you refresh the browser or navigate to another page and then return to the page where you created your search.

TIP: On most inventory pages, you can use the **Search** field at the top of the inventory page to perform a Basic Search. For more information, see [Using Basic Search](#). In addition, you can click the gear icon () and select *Advanced* to "translate" your basic search into an Advanced Search. For more information, see [Performing an Advanced Search](#).

Filter Syntax

When searching for a string in the **Filter** field on an inventory page, SL1 matches sub-strings 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 sub-string unless you use a special character.

NOTE: If you type invalid search terms or incorrect syntax in a filter on an updated inventory page, SL1 will highlight the headers in red. SL1 will ignore the incorrect input until you fix it.

You can use the following set of syntax and special characters to search and filter each column on an inventory page (except columns that display date and time):

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".

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

String

- **^ (caret)**. For strings only. Specifies "match the beginning". Matches any string that begins with the specified string.
For example: "^ sci" matches "scientific" and "scienceologic", but not "conscious".
- **\$ (dollar sign)**. For strings only. Specifies "match the ending". Matches any string that ends with the specified string.
For example: "ter\$" matches the string "renter" but not the string "terrific".

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" matches 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 matches 1, 2, 6, 345, etc. "-5" matches 5 and less. So matches 5, 3, 1, 0, etc.
- **> (greater than)**. Matches numeric values only. Specifies any value "greater than".
For example: ">7" matches all values greater than 7.
- **< (less than)**. Matches numeric values only. Specifies any value "less than". For example:
"<12" matches 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" matches 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" matches all values 12 and less.
- **= (equal)**. Matches numeric values only. For numeric values, allows you to match a negative value.
For example: "=-5" matches "-5" instead of being evaluated as the "half open range" as described above.

Examples

- "silo". Matches text that contains "silo".
- "!silo ". Matches text that does not contain "silo".
- "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 start with "silo".
- "!fer\$". Matches all values that do not end with "fer".
- "!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 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".
- "!vol&!silo". Matches text that does not contain "vol" AND also does not contain "silo". For example, "volume" matches, because it contains "vol" but not "silo".
- "!vol&02". Matches text that does not contain "vol" AND also contains "02". For example, "happy02" matches, 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".
- "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".
- "!^\$" matches all values that are not null.
- "!^" matches null values.
- "!\$" matches null values.
- "1-5" matches 1, 2, 3, 4, and 5.
- "1-" matches 1 and greater. So matches 1, 2, 6, 345, etc.
- "-5" matches 5 and less. So matches 5, 3, 1, 0, etc.

- " $=-5$ " matches "-5" instead of being evaluated as the "half open range" as described above.
- " <7 " matches 1,2,3,4,5,6
- " ≥ 3 " matches anything greater than or equal to 3

Editing the Settings for an Inventory Page

On an inventory page, you can perform the following actions:

- [Adjust and sort columns](#)
- [Create a Multi Sort](#)
- [Edit Column Preferences](#)
- [Adjust the row density](#)
- [Export to a CSV file](#)

TIP: To change how often SL1 updates this page, click **Refresh** and select a new timespan. You can also select *Refresh Now* to immediately update the page or select *Pause* to stop updating the page altogether.

Adjusting and Sorting Columns

You can adjust the columns on an inventory page in the following ways:

- To rearrange the columns in the list, click and drag the column name to a new location.
- To adjust the width of a column, click and drag the pipes (|) on either side of the column name.
- To quickly sort a column in ascending or descending order, click the up arrow (ascending) or down arrow (descending) that appears when you hover over the column name.

- To perform additional sorts and to hide or show columns, click the Menu icon (☰) next to the column name and select an option from the drop-down menu.

TIP: Hover over the column name to see the Menu icon (☰).

Your options include:

- *Multi Sort*. Creates multiple ascending or descending sorts, in order of importance. For more information, see [Creating a Multi Sort](#).
- *Unsort*. Clears the existing sort.
- *Sort by ASC*. Sorts the column in ascending order: smallest to largest for numbers, or alphabetically for strings.
- *Sort by DESC*. Sorts the column in descending order: largest to smallest for numbers, or reverse alphabetically for strings.
- *Hide*. Removes that column from the inventory page.
- *Show columns*. Opens the **Find column** menu, from which you can add hidden columns and hide columns as needed. You can also choose to show all columns, reset the columns to their default display settings, and clear any column filters on the inventory page. For more information, see [Editing Column Preferences](#).

NOTE: If you use the check boxes to select one or more items in a list, the relevant buttons appear at the top of the page, above the column names.

Creating a Multi Sort

On an inventory page, if you click the gear icon (⚙️) to open the **Grid Settings** menu and select *Multi Sort*, you can create multiple ascending or descending sorts, in order of importance:

- Click **Add Sort** to add one or more sorts.
- Click **[Clear Sort]** to clear all of the existing sorts. You can also click the delete icon (🗑️) next to a sort to remove that specific sort from the list.
- Click **[Apply Sort]** to run the sort and show the results on the list page.

The Multi Sort settings are cleared if you sort a column from the inventory page.

Editing Column Preferences

On an inventory page, if you click the gear icon (⚙️) to open the **Grid Settings** menu and select *Column Preferences*, the **Find column** menu appears. On the menu, you can update the display settings for that inventory page:

- Select a column from the menu to add a hidden column.
- De-select a column from the menu to hide that column.
- Click **Show All** to show all columns.
- Column widths, order, sort, refresh interval, and other column attributes will persist after you adjust them. You can click **Reset** to return the columns to their default display settings.

TIP: From the **Grid Settings** menu, click *Clear Filters* to remove any column filters on that page.

Adjusting the Row Density

On an inventory page, if you click the gear icon (⚙️) to open the **Grid Settings** menu and select *Density*, you can change the size of the rows and the size of the text in the rows on that page by selecting:

- *Compact*. Uses less padding between rows and a smaller font size to display more rows.
- *Comfortable*. Use more padding between rows and a larger font size to display fewer rows, but with more spacing between rows.


TIP: You can toggle the highlighting of the **Message** column by clicking the gear icon (⚙️) to open the **Grid Settings** menu and selecting *Severity Highlighting Off* or *Severity Highlighting On*.


Exporting to a CSV File

On an inventory page, if you click the gear icon (⚙️) to open the **Grid Settings** menu and select *Export (.csv)*, you can save the current list of items as a comma-separated list.

The list will include any column filters that you set up before clicking *Export*.

Viewing the Device Summary Modal from an Inventory Page

On the **Events** page and the **Devices** page, you can click the **Open** icon () next to an event or device to open a new **Device Summary** modal:




Device Summary ✕

Device: [rcarignan-aio-102821](#)

Message: [VM memory is not reserved or VM memory limit less than reservation.](#)

Tools

 Type to run an action on this device.

Vitals

100

50

0

06:00 06:30 07:00 07:30 08:00 08:30 09:00 09:30

Logs

Date/Time	Event Severity	Message
<input type="text" value="Filter"/>	<input type="text" value="Filter"/>	<input type="text" value="Filter"/>
Dec 1, 2023, 9:52 AM	● Critical	The SL1 license expires in -67 days (message repeats 1 time)
Dec 1, 2023, 9:50 AM	● Major	CPU runtime contention of 98.0% exceeds threshold (15%)

Total Rows: 10783

NOTE: On the **Events** page, the **Device Summary** modal displays only for events that are aligned with devices.


The detail window for that device contains the **Tools** pane, the **Vitals** graphs, and the **Logs** pane:

- The **Tools** pane enables you to run a set of diagnostic tools or user-initiated actions in the **Activity Center**, or to click on custom links that will open in a separate browser window. Click the search bar to search for tools, actions, or custom links that are available for the device.

- The **Vitals** pane displays graph data for the past four hours of CPU usage, memory usage, and latency for that device, where relevant. You can zoom in on a shorter time frame in the **Vitals** graph by clicking and dragging, and you can go back to the original time span by clicking the **[Reset zoom]** button.
- The **Logs** pane displays a list of events associated with that device.


TIP: To open the detail or Investigator page for an item, click the link for the item name at the top of the detail window.

Using Basic Search

At the top of most lists in the SL1 user interface, the **Search** field lets you look for specific elements in that list. The **Search** field contains a magnifying glass icon () next to the words "Type to search" or "Search".

As you type text in the **Search** field, SL1 filters the list to include only those elements that include your search terms. When searching, SL1 considers all relevant columns for the search, including those that are not currently displayed on the page.

SL1 retains your search criteria on pages in the SL1 user interface even after you refresh the browser or navigate to another page and then return to the search page using the **[Back]** button in the browser.


TIP: After you type search commands in the **Search** field, you can click the gear icon () and select *Advanced* to "translate" your basic search into an Advanced Search. For more information, see [Performing an Advanced Search](#).

To use the **Search** field:

1. Click the **Search** field and start typing search text. As you type, SL1 provides potential matching values in a drop-down menu and starts filtering the list with your search text.

TIP: For example, if you start searching for "database" by typing *data*, a drop-down list appears with a list of columns that might contain that word, and the list is filtered by items that have "data" in one of their fields.


2. If you select one of the suggested search criteria from the list, such as *message*, a **criteria button** displays under the **Search** field. You can click the criteria button and edit the search text under the button, if needed:



CAUTION: If you select the *ANY* option from the drop-down menu, the search looks through all relevant columns for matches to your search text. If you select the *ANY* option for **multiple** search terms, this can cause issues with search queries, as every term has to be searched against every column in the database. SL1 uses different search criteria for an "ANY" search based on the page you are currently on in SL1. ScienceLogic recommends that you do not use Basic Search in these situations, and instead use the [Advanced Search](#) feature. For more information, see [Fields Used by an "ANY" Basic Search](#).


NOTE: If you are unable to paste a copied string of text in the **Search** field, make sure that your copied text does not contain any hidden special characters.

3. To edit the search criteria for your current search even further, click the criteria button and click the *contains* field. You can choose from additional search operators in the drop-down that appears, such as *begins with*, *is null*, *equal to*, and *not equal to*.
4. You can add another set of search criteria to an existing search by typing additional text in the **Search** field, and then selecting additional fields from the drop-down list. The new search terms are added to a second criteria button:



TIP: The search criteria button under the **Search** field also displays the search operator and value for the search as hovertext.

- Alternately, you can click the menu icon (☰) to the right of the **Search** field to open a menu containing related search criteria. Select an element from the list of criteria and type additional search information in the criteria button that appears under the **Search** field:



TIP: If you select a date-related search criteria from the list of criteria, you can use a drop-down calendar to select a specific date and time.

- To remove a search criteria, click the close icon (X) on the criteria button under the **Search** field.
- To switch to an Advanced Search, click the gear icon (⚙️) to the right of the **Search** field and select *Advanced*. For more information, see [Advanced Search](#).
- To quickly clear a search, click the gear icon (⚙️) to the right of the **Search** field and select *Clear*.
- To save a search so you can use it again, see [Saving a Basic Search](#).

Fields Used by an "ANY" Basic Search

If you type search criteria into a **Search** field at the top of a page in SL1 and then select the *ANY* option from the drop-down menu, the search looks through *all* relevant columns for matches to your search text.

CAUTION: If you select the *ANY* option for **multiple** search terms, this can cause issues with search queries, as every term has to be searched against every column in the database. ScienceLogic recommends that you do not use Basic Search in these situations, and instead use the [Advanced Search](#) feature.

SL1 uses different search criteria for an "ANY" search, depending on the page you are currently on in SL1. For example, an "ANY" search on the **Dashboards** page searches for the search criteria in the Dashboard ID or Dashboard Name fields, while an "ANY" search on the **Devices** page searches for the search criteria in the Device Name, Hostname, Device Class, Collector Group, Organization, and additional fields. See the table below for details.

To see which fields SL1 is using for a search, click the gear icon (⚙️) to the right of the **Search** field and select *Advanced*. The syntax of the Basic Search is converted into an Advanced Search, which lists the relevant fields being searched by SL1.

The following table lists the fields that are used by an "ANY" search, based on the page you are on in SL1:

Page in SL1 user interface	Fields searched by SL1 for that page
Main Pages	

Page in SL1 user interface	Fields searched by SL1 for that page
Dashboards	Dashboard ID, Dashboard Name
Events	Event Message, Aligned Resource Name, Aligned Organization, Aligned Sub-entity Name, Device Name, Hostname, Device Class Logical Name, Device Class Description, Device Collector Group, Device Class, Device Group ID, Device Group Name, Device Class Category
Devices, Machine Learning, Device Services	Device Name, Hostname, Device Class Logical Name, Device Class Description, Device Class, Device Group, Device Class Category Name, Collector Group ID, Collector Group Name, Aligned Organization, Machine Learning Policy
Business Services and IT Services	Service Type, Service ID, Service Name, Service Label, Service Policy Name, Aligned Organization
Maps	Map Name, Map Description
Secondary Pages	
Agents	Agent Nickname, Agent Operating System, Hostname
Business Service Templates	Business Service Template Name, Business Service Template Descriptions
Credentials	Credential Name
Custom Attributes	Custom Attribute Label, Custom Attribute Name
Device Categories	Device Category Name
Device Classes	Device Class Name, Device Class Description, Virtual Type, Logical Name, Device Category Name
Discovery Sessions	Discovery Session Name, Aligned Collector Name, Aligned Organization
Event Policies	Event Policy Name
Subscription Usage (Current License Usage)	License Type, Device Name, Aligned Organization, Device Class Category Name

If a page from the SL1 user interface is not listed in the above table, then SL1 only uses the relevant Name field on that page for an "ANY" search.

Saving a Search

If you are creating a complicated search using Basic Search or Advanced Search, or if you have a search that you use on a regular basis, you can save that search criteria so you can quickly use it again later.

To save a search:

1. After you have created a Basic or Advanced Search, click the gear icon (⚙️) to the right of the **Search** field and select **Save**. A **Save Search** window appears.
2. In the **Search Name** field, type the name of your search and click **[Save]**. The search is added to the list of saved searches.

To use a saved search:

1. Click the gear icon (⚙️) to the right of the **Search** field and select *Saved Searches*. An **Apply Search** window appears.
2. From the **Select a search** drop-down, select the search you want to use and click **[Apply]**. That search is applied to the current list.

NOTE: By default, saved searches apply only to *your* user profile, and they are not shared with other users.

To share a saved search:

1. Create a Basic Search and change it to an Advanced Search by clicking the gear icon (⚙️) to the right of the **Search** field and selecting *Advanced*.
2. Copy the Advanced Search code from the **Search** field and paste it into the relevant documentation so you can share the search with other users.

Performing an Advanced Search

The SL1 user interface includes an Advanced Search option that lets you use customized search commands to search for data. The syntax for these Advanced Searches can be much more complex than a Basic Search, enabling you to find exactly what you need from a list of items.

Also, because the Basic Search only uses "AND" for multiple search criteria, you need to use an Advanced Search for an "OR" search using multiple search criteria, or if you want to create more complicated searches using Boolean Algebra.

At a minimum, an Advanced Search requires the following components, in the following order:

- A **field**. The general type of data for which you are searching, such as a device name or an event message.
- An **operator**. A word or symbol that specifies the relationship between the field and the value, such as equals or less than.
- A **value**. A specific aspect or version of the field, such as a name or an amount. If a value is a string, it should be surrounded by "quotation marks" or 'apostrophes'.

TIP: As you type your Advanced Search, a red icon (❗) or a green icon (✅) appears at the end of the text field to show that your search is incorrectly or correctly formatted.

TIP: To view a list of all possible search commands in an Advanced Search, press **[Ctrl] + [Space]**.

The Advanced Search fields and values vary based on the page you are on in the SL1 user interface. For more information about fields, operators, and values, see [Components of an Advanced Search](#).

Below are simple examples of Advanced Search syntax:

```
message contains 'risk is high'
```



```
organization has (company contains 'system')
```

```
attribute has (id = year and value = 2021)
```

```
name contains 'web_tier' and deviceClass has (description contains  
'AppDynamics')
```

```
deviceClass has (description contains 'em7 admin portal')
```


For additional examples, including examples with more complex syntax, see [Examples of Advanced Searches](#).

TIP: You can type search commands in the **Search** field for a Basic Search, and then click the gear icon (⚙️) and select *Advanced* to "translate" your basic search into an Advanced Search. You cannot go from an Advanced Search back to a Basic Search, however, without losing your search criteria.


To create an Advanced Search:

1. Click the gear icon (⚙️) to the right of the **Search** field and select *Advanced*. The search type changes from Basic to Advanced (note the change in font style).


2. Start typing a field name for your search. As you type, SL1 displays a list of available fields in a drop-down menu:



3. Select or type a field name.
4. To view a list of all possible search commands at any point in an Advanced Search, press **[Ctrl+Space]**. For example, the following operator options appear if you press **[Ctrl+Space]** after typing "message" and inserting a space:




5. Select or type an operator name. If you are typing, SL1 provides a list of available options.



TIP: As you type your search command, a red icon (❗) appears at the end of the text field if your command is incorrectly formatted or incomplete. Click the red icon to view additional details.

6. Type a value to complete your search, and type additional search commands as needed. When your search is complete and formatted correctly, a green icon (✅) appears at the end of the text field:



7. Click the **[Search]** button. The results of your search appear.

NOTE: Even if you have correct search syntax, SL1 will save your search query only *after* you click **[Search]**. For example, some pages, such as a Device Service search, might not show all of the search results until you click **[Search]**.

8. To clear a search, click the gear icon (⚙️) to the right of the **Search** field and select *Clear*.
9. You can save an Advanced Search to use later. By default, saved searches apply only to *your* user profile, and they are not shared with other users. For more information, see [Saving a Search](#).

Components of an Advanced Search

At the minimum, an Advanced Search requires the following components, in the following order:

- A **field**. The general type of data for which you are searching, such as a device name or event message.
- An **operator**. A word or symbol that specifies the relationship between the field and the value, such as equals or less than.
- A **value**. A specific aspect or version of the field, such as a name or an amount. You must use either "quotation marks" or 'apostrophes' for your search strings, and strings are not case-sensitive.

You can also include the operators "and" or "or" to your search command. Basic Search in SL1 uses only "AND" searches, unless you specify "Any" in your Basic Search.

NOTE: When SL1 evaluates an Advanced Search command, it evaluates the "OR" expressions first, followed by the "AND" filters.

For example, the following search command looks for events that have a status of *Critical* *and* contain a message with the word "error":

```
status = critical and message contains 'error'
```

The following search command looks for devices with a name of "device-name" or messages containing the word "error":

```
name = "device-name" or message contains "error"
```


You can use parentheses () to group expressions and to ensure that the expressions are evaluated in the correct order. The following search command looks for either devices with a name of "device-name" and a status of *Critical*, or devices with a name of "device-name" and a status of *Major*:

```
(name = "device-name" and status = critical) or (name = "device-name" and status = Major)
```


TIP: Searches in SL1 are *not* case-sensitive, so you can use any combination of upper-case and lower-case letters.

Fields

For most searches, you start your search command with a field name. When you start typing in an Advanced Search field, SL1 provides a list of potential fields in a drop-down menu that you can select for your search command:



The list of potential fields depends upon the page you are currently on in SL1. The example above is from the **Advanced Search** field on the **Devices** page. If you typed the same letters in the **Advanced Search** field on the **Events** page, the drop-down menu would look like this:



The following table lists some of the more common fields, along with how to use them and examples of search commands that use those fields:

Field name	Purpose	Example
<code>alignedResourceName</code>	Search for the name of a device aligned with an event.	<code>alignedResourceName contains "lab"</code>
<code>asset</code>	Search for an asset aligned with a device.	<code>asset has (assetTag contains 1)</code>
<code>attribute</code>	Search for devices based on custom attributes. In the example, the custom attribute is "year" and the value is "2021".	<code>attribute has (id = year and value = 2021)</code>
<code>dateCreated</code>	Search for the date and time a device was created.	<code>dateCreated isNotNull</code>

Field name	Purpose	Example
deviceClass	Search for devices belonging to a device class.	deviceClass has (class contains 'Cisco')
deviceGroup	Search for devices belonging to a device group.	deviceGroup has (name contains "Network")
hostname	Search for a device hostname	device has (hostname = "srv")
id	Search for the unique numeric ID assigned by SL1.	id contains "10"
isAcknowledged	Search for events that have or have not been acknowledged.	isAcknowledged = true
message	Search for details about an event message.	message contains "problem"
name	Search for the name of the device.	name = "server"
organization	Search for the organization to which the device is assigned.	organization has (company = "System") organization has (company doesNotContain 'ABC Systems')
severity	Search for the severity of an event; severities range from 0 to 4, from Healthy to Critical.	severity in 3,4 Searches for all Major and Critical events.
state	Search for the state of a device; states range from 0 to 4: Healthy, Notice, Minor, Major, and Critical.	state in 0,1,2 Searches for all devices with a state of Healthy, Notice, and Minor.
suppressGroup	Hide data related to the specified group.	suppressGroup = sciencelogic

Operators

For most searches, you follow a field with an operator. The operator establishes a relationship between the field and the value that comes after the operator.

TIP: The list of available operators changes based on the page where you are making your search.

The following table lists some of the more common operators, along with how to use them and examples of search commands that use those operators:

Operator name	Purpose	Example
and	Include two or more search criteria before producing search results	deviceClass has (description contains 'rds instance') and name contains 'wordpress'
or	Include at least one of multiple	name = "server" or message

Operator name	Purpose	Example
	search criteria.	contains "error"
=, ==, eq	The field and the value are equal.	name == 'ECS 23' and ip doesNotContain '.'
<>, !=, neq	The field and the search value are <i>not</i> equal.	field != abc
<, lt	The field is less than the search value.	state < 2
>, gt	The field is greater than the search value.	severity > 3
<=, lte	The field is less than or equal to the search value.	state lte 2
>=, gte	The field is greater than or equal to the search value.	severity gte 3
contains	The field includes the specified string.	deviceClass has (description contains 'em7 admin portal')
doesNotContain	The field does not include the specified string.	name contains 'SL1 Classic' or (description contains 'PowerFlow' and description doesNotContain 'Not Deployed') or (description contains 'Extended Architecture' and description doesNotContain 'Not Deployed')
has	The field contains a specific value. The value following "has" must be enclosed in parentheses.	deviceClass has (deviceCategory has (name contains 'Server')) and deviceClass has (description doesNotContain 'vcenter')
in	The field must be part of a specific set of values.	severity in 2,3,4
not	Opposite values; this operator precedes the field name.	not field = abc
isNull	The field is empty.	extTicketRef isNull
isNotNull	The field is not empty.	counter isNotNull

Values

The value you type at the end of a search command depends on the field name and the operator you use. For most searches, you can type the value instead of picking it from the drop-down menu that lists possible search options.

In the following example, the first search value is a string (red text) and the second search value is a numeric value (blue text):

```
Q name contains 'np' and ip beginsWith 192.168
```

You must use either "quotation marks" or 'apostrophes' for search strings, and strings are not case-sensitive.

Strings

You can create a search command that searches for a specific set of words in a string. You must use either "quotation marks" or 'apostrophes' for your search strings, and strings are not case-sensitive.

The following table lists some of the more common string operators, along with how to use them and examples of search commands that use those string operators:

String operator name	Purpose	Example
<code>beginsWith</code>	Search for strings beginning with a specified value	<code>message beginsWith "Host Resource"</code>
<code>endsWith</code>	Search for strings ending with a specified value	<code>message endsWith 'shutdown'</code>
<code>contains</code>	Search for strings containing a specified value	<code>message contains "problem"</code>
<code>doesNotBeginWith</code>	Search for strings that do not begin with a specified value	<code>message doesNotBeginWith "front"</code>
<code>doesNotEndWith</code>	Search for strings that do not end with a specified value	<code>message doesNotEndWith 'warning'</code>
<code>doesNotContain</code>	Search for strings that do not contain a specified value	<code>message doesNotContain "codec"</code>

Escape characters

In double-quoted strings (strings surrounded by quotation marks), you can include quotation marks in the search by *escaping* the quotation marks. To escape those characters, add a backslash before each quotation mark, such as `\`.

For example:

```
"Error in \"process x\""
```

In single-quoted strings, you can include the single-quote character by escaping it with a backslash, such as `\'`.

For example:

```
'Eric\'s Laptop'
```

```
'Error in "process x"'
```

TIP: You do *not* need to add quotes around strings in your search commands. However, if your string contains only numbers, you might want to add quotes around it to ensure that SL1 interprets it as a string.

If you do not include quotes around strings in your search commands, you must escape the following characters with a backslash:

- all empty spaces or white spaces
- comma
- end parenthesis

Examples:

```
Eric's\ Laptop
```

```
Error\ in\ "process\ x"
```

```
devices\ \(system\,\ server\)
```

Other than the escape characters mentioned above, you can escape any character. You must escape the backslash character if you want to use it in a string, such as `\\`.

The normal whitespace escape sequences can be used: `\t` (tab), `\n` (new line), `\b` (backspace), `\r` (carriage return), and `\f` (form feed).

You can also use four-digit Unicode hex escape codes in the form `\uXXXX`.

Examples of Advanced Searches

Because the search commands differ for each page in SL1, this section contains a set of search examples based on context.

TIP: To view a list of all possible search commands at any point in an Advanced Search, press **[Ctrl+Space]**.

Advanced Search Examples on the Events Page

The following table contains a set of sample Advanced Searches for the **Events** page:

Purpose of Advanced Search on Events Page	Advanced Search Syntax
Search for events on devices by Device ID of 1, 2, or 3.	<code>device has (id in 1,2,3)</code>
Search for all events that contains the word "error"	<code>message contains "error"</code>
Search for all events on devices with a Device Category Name of "xtremio".	<code>device has (deviceClass has (deviceCategory has (name contains 'xtremio')))</code>

Advanced Search Examples for Dynamic Component Mapping (DCM) Scenarios

Dynamic Component Mapping (DCM) allows SL1 to collect data from a single management system, such as a VMware ESX server, and then use that data to create multiple device records for the entities managed by that single management system. For example, the managed entities for a VMware ESX server would be the Guest VMs hosted by that ESX server.

The following table contains a set of sample Advanced Searches for DCM devices on the **Devices** page:

Purpose of Advanced Search	Advanced Search Syntax
Search for all host devices from a specific vCenter and add those devices to a Device Service	<code>componentRoot has (name contains "VCSA") and (componentParent has (deviceClass has (description contains 'Host Server')) or componentParent has (deviceClass has (description contains 'ESX')))</code>
Search for SQL servers than run as the master.	<code>deviceClass has (description contains 'SQL Instance') and componentRoot has (attribute has (id == 'DB_Role'and value contains 'Master'))</code>

The following sections contain more details about these two Advanced Searches and how they work.

Scenario 1: vCenters

An SL1 system has multiple vCenters, each of which has multiple host servers. You want to include in a Device Service all Virtual Machines (VMs) from just one of the vCenters. This scenario includes the following elements:

- The vCenter has a host name of **VCSA hayward-hq.loc**
- One of the hosts has a Device Class of **Host Server**, and it contains four VMs
- One of the hosts has a Device Class of **ESX**, and it contains 16 VMs

Device Name	IP Address	Device Category	Device Class Sub-class	DID	Organization	Current State	Collection Group	Collection State
VCSA hayward-hq loc	172.22.101.116	VMware	VMware vCenter Server Appliance	5181	Hayward HHQ	Healthy	HHQ	Active
Datacenter	--	Infrastructure	VMware Datacenter	6588	Hayward HHQ	Healthy	HHQ	Active
Datstores	--	Infrastructure	VMware Folder	6590	Hayward HHQ	Healthy	HHQ	Active
Hosts	--	Infrastructure	VMware Folder	6589	Hayward HHQ	Healthy	HHQ	Active
172.22.101.12	--	Host	VMware Host Server	6592	Hayward HHQ	Healthy	HHQ	Active
HHQ-Collector-Beta3	--	Guest	VMware Virtual Machine	7932	Hayward HHQ	Healthy	HHQ	Active
lin-hhq-09t1	172.22.101.76	Servers	Linux CentOS release 6.7 (Final)	322	Hayward HHQ	Healthy	HHQ	Active
Restorepoint Appliance	172.22.101.89	Servers	Restorepoint Appliance	4865	Hayward HHQ	Major	HHQ	Active
simulators	172.21.101.22	Servers	Linux CentOS	331	Hayward HHQ	Healthy	HHQ	Active
vm2 hayward-hq loc	172.22.101.10	VMware	VMware ESX(i)	7	Hayward HHQ	Major	HHQ	Active
GNS3 VM	--	Guest	VMware Virtual Machine	6606	Hayward HHQ	Healthy	HHQ	Active
HHQ-Collector-2	--	Guest	VMware Virtual Machine	6605	Hayward HHQ	Healthy	HHQ	Active
hhq-message-collector	172.22.101.75	EM7	ScienceLogic, Inc. EM7 Message Collector	323	Hayward HHQ	Healthy	HHQ	Active
Lin-HHQ-Mail01	--	Guest	VMware Virtual Machine	7943	Hayward HHQ	Healthy	HHQ	Active
nessus	172.21.101.8	Servers	Linux CentOS	11	Hayward HHQ	Healthy	HHQ	Active
pfsense hayward-hq loc	172.22.101.1	Firewall	Fraunhofer FOKUS pfsense	5	Hayward HHQ	Healthy	HHQ	Active
test-delete	--	Guest	VMware Virtual Machine	7976	Hayward HHQ	Major	HHQ	Unavailable
VCSA	--	Guest	VMware Virtual Machine	6609	Hayward HHQ	Healthy	HHQ	Active
Win-HHQ-AD1 hayward-hq loc	172.22.101.120	Servers	Microsoft Windows Server 2012 R2 Domain Co	50	Hayward HHQ	Healthy	HHQ	Active

Search 1a

componentRoot has (name contains "VCSA") and deviceClass has (description contains "Virtual Machine")

This search works, but it only returns seven devices instead of the 16 devices you might have expected:

Query for the right set of devices.

componentRoot has (name contains "VCSA") and deviceClass has (description contains "Virtual Machine")

Preview: 7 Devices

NAME	STATE	IP ADDRESS	CATEGORY	CLASS	SUB-CLASS
GNS3 VM	Healthy	--	Virtual.Guest	VMware	Virtual Machine
HHQ-Collector-2	Healthy	--	Virtual.Guest	VMware	Virtual Machine
HHQ-Collector-Beta3	Major	--	Virtual.Guest	VMware	Virtual Machine
Lin-HHQ-Mail01	Healthy	--	Virtual.Guest	VMware	Virtual Machine
test-delete	Major	--	Virtual.Guest	VMware	Virtual Machine
VCSA	Healthy	--	Virtual.Guest	VMware	Virtual Machine
Win-HHQ-SRV2	Healthy	--	Virtual.Guest	VMware	Virtual Machine

How the query works:

componentRoot has (name contains "VCSA")

This part of the query returns a list of all devices that have a root with a name that includes **VCSA**.

```
and deviceClass has (description contains "Virtual Machine")
```

This part of the query filters the results of the first part of the query to isolate the devices with Device Class of **Virtual Machine**

NOTE: The reason why this Advanced Search does not result in the 16 VMs that we can see are hosted by both hosts is that some of the VMs have been *merged*. As a result, their Device Class has been reset to that of the operating system they are running.

Search 1b

```
componentRoot has (name contains "VCSA") and (componentParent has (deviceClass has (description contains 'Host Server')) or componentParent has (deviceClass has (description contains 'ESX')))
```

This search successfully returns all 16 VMs:

The screenshot shows a search interface with a query input field containing the following query: `componentRoot has (name contains "VCSA") and (componentParent has (deviceClass has (description contains "Host Server")) or componentParent has (deviceClass has (description contains "ESX")))`. Below the query, there is a button labeled "Preview: 16 Devices". Below the button is a table with the following columns: NAME, STATE, IP ADDRESS, CATEGORY, CLASS, and SUB-CLASS. The table contains three rows of data:

NAME	STATE	IP ADDRESS	CATEGORY	CLASS	SUB-CLASS
GNS3 VM	Healthy	-	Virtual.Guest	VMware	Virtual Machine
HHQ-Collector-2	Healthy	-	Virtual.Guest	VMware	Virtual Machine
HHQ-Collector-Beta3	Healthy	-	Virtual.Guest	VMware	Virtual Machine

How the query works:

```
componentRoot has (name contains "VCSA")
```

This part of the query returns a list of all devices that have a root with a name that includes **VCSA**.

```
and (componentParent has (deviceClass has (description contains 'Host Server')))
```

This part of the query filters the results of the first part of the query to isolate the devices that have a parent of the Device Class **Host Server**.

```
or componentParent has (deviceClass has (description contains 'ESX'))
```

This part of the query then does a second search of the results of the first query, looking for devices with parents that have a Device Class of **ESX**.

Scenario 2: SQL Servers


A DB Cluster has two SQL Servers that run as the master and two SQL servers that run as the slave:

Device Components Devices Found [4]						
	Device Name	IP Address	Device Category	Device Class Sub-class	DID	Organization
1. -	db-lab4-ha.phx3.llnw.net	10.12.217.4	Pingable	Ping ICMP	3649	LLNW
1. +	MySQL Server	--	Software	Oracle MySQL Server	3653	LLNW
2. -	db-lab5-ha.phx3.llnw.net	10.12.217.7	Pingable	Ping ICMP	3663	LLNW
1. +	MySQL Server	--	Software	Oracle MySQL Server	3664	LLNW
3. -	db-llnw34-ha.phx3.llnw.net	10.12.61.56	Pingable	Ping ICMP	4668	LLNW
1. +	MySQL Server	--	Software	Oracle MySQL Server	4670	LLNW
4. -	db-llnw35-ha.phx7.llnw.net	10.14.205.42	Pingable	Ping ICMP	4672	LLNW
1. +	MySQL Server	--	Software	Oracle MySQL Server	4676	LLNW

Four times a year the master and slave is swapped between the two pairs. As this is a planned, infrequent, and manual activity, you can manually swap a custom attribute on the four hosting devices to designate if a device is acting as master or slave:

The screenshot shows the 'Attributes' configuration for a device. The 'DB_Role' attribute is currently set to 'Master'. The interface includes a dropdown menu for the attribute value, a 'String' type indicator, and an 'Extended' checkbox.

Label	Value Type	Value	Attribute Type
DB_Role	String	Master	Extended

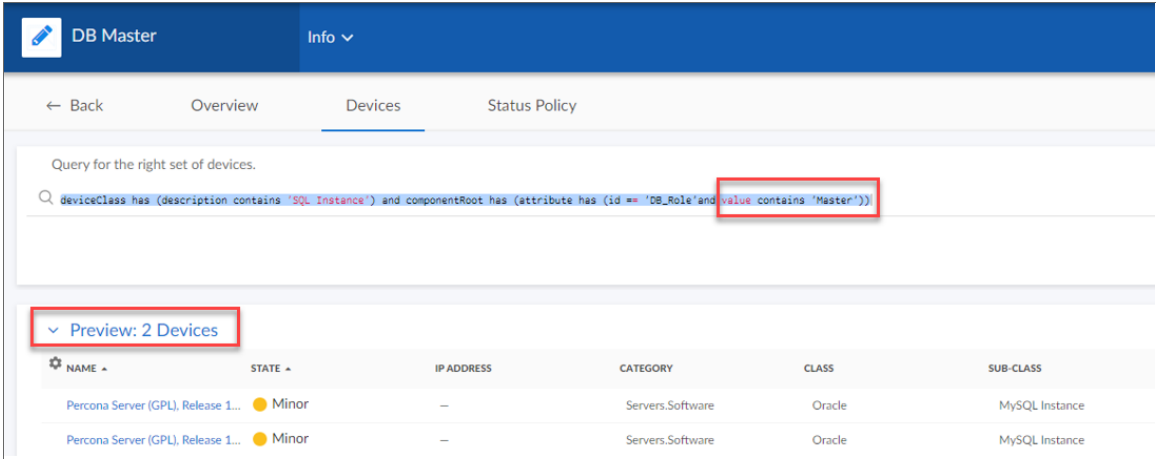


NOTE: For production or higher frequency and automated swapping, making the switch of the custom attributes could be embedded into a switching script and use the API or GQL interfaces to change the value of the custom attributes.

Searches

```
deviceClass has (description contains 'SQL Instance') and componentRoot has (attribute has (id == 'DB_Role'and value contains 'Master'))
```

This search returns the two SQL instances that are running on the two SQL hosts designated as the *master* pair:



```
deviceClass has (description contains 'SQL Instance') and componentRoot has (attribute has (id == 'DB_Role'and value contains 'Master'))
```

This search returns the two SQL instances that are running on the two SQL hosts designated as the *slave* pair:

DB Slave Info

← Back Overview **Devices** Status Policy

Query for the right set of devices.

deviceClass has (description contains 'SQL Instance') and componentRoot has (attribute has (id == 'DB_Role' and value contains 'Slave'))

Preview: 2 Devices

NAME	STATE	IP ADDRESS	CATEGORY	CLASS	SUB-CLASS
Percona Server (GPL), Release 2...	Minor	–	Servers.Software	Oracle	MySQL Instance
Percona Server (GPL), Release 7...	Minor	–	Servers.Software	Oracle	MySQL Instance

How the query works:

```
deviceClass has (description contains 'SQL Instance')
```

This part of the query returns a list of all devices that have a Device Class of **SQL Instance**.

```
and componentRoot has (attribute has (id == 'DB_Role' and value contains 'Slave'))
```

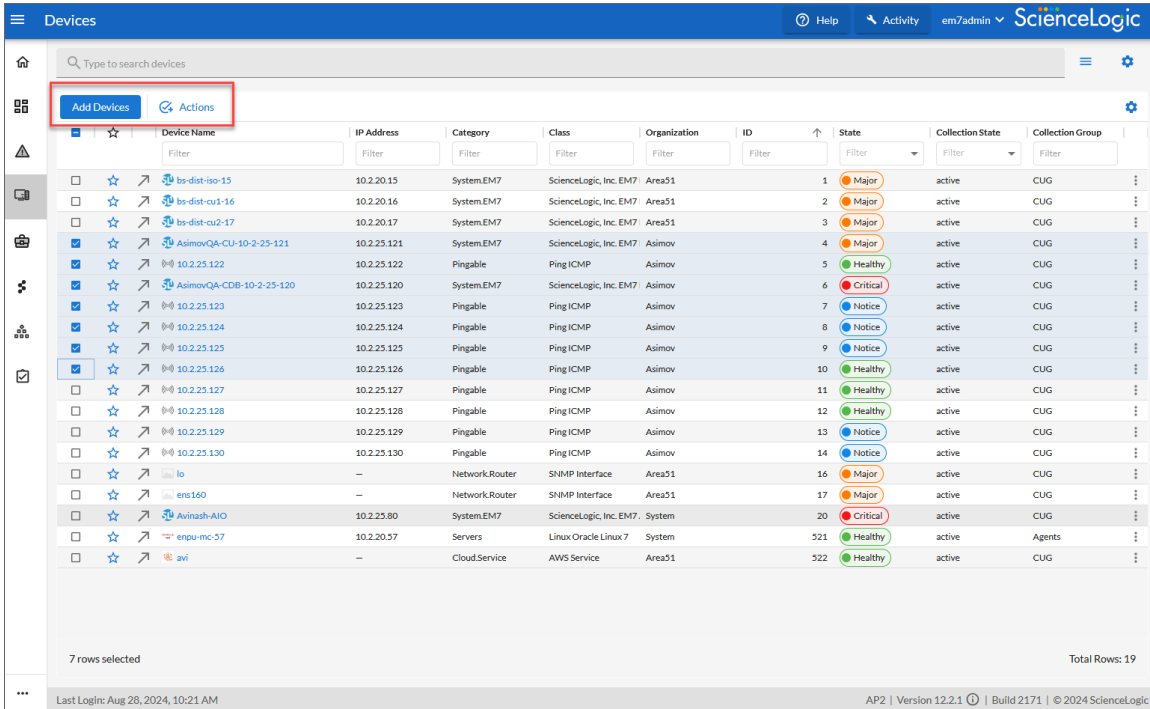
This second part of the query then filters the results of the first query to isolate the devices that have a root device with a custom attribute of **DB_Role** set to *Slave*.

Performing Bulk Actions

If an inventory page in SL1 displays a list of items, and that page contains a checkbox () to the left of each item in the list, you can select two or more items to perform bulk actions on all of the selected items at the same time.

When you use the checkbox to select one or more items in a list, a set of relevant buttons appear at the top of the page. If there are many bulk actions available, the **[Actions]** button appears, and you can search or click the down arrow to select an action. Select an action or click the relevant button to run that action on all of the selected items.

For example, on the **Devices** page, you can select two, twenty, or all of the devices on the page, and then you can click the **[Actions]** button and select from a number of actions to run on all of the selected devices.



To select all of the items on a page, click the checkbox at the top of the list. To clear all of the selected items, click the checkbox at the top again.

TIP: To select a series of adjacent rows, you can hold down the **[Shift]** key as you click the first item and then click the last item in the series.

TIP: Pages that contain lists use "infinite scrolling", where the list continues to populate as you scroll toward the bottom of the list. The scrolling stops when you reach the end of the list.

Customizing the SL1 User Interface

A **theme** is a graphic template this is applied to the user interface. SL1 includes one System Default theme, but you can completely customize the look and feel of your SL1 system by creating new themes.

For example, you could create a theme that replaces the SL1 logo with your company's logo and updates the colors used in the user interface to match those used in your company's branding. You can also choose between a light theme or a dark theme for the user interface.

For more information about customizing the user interface using themes, see the manual *Customizing the User Experience*.


Getting Help and More Information

For product documentation about any page in SL1, click the **[Help]** button at the top right of any page in SL1.

The **Help** menu appears in a new pane on the right side of the SL1 window:

The screenshot shows the SL1 Events page interface. At the top, there's a navigation bar with 'Events', 'Help', 'Activity', and 'em7admin' (ScienceLogic). Below the navigation bar, a summary shows 'Total Events: 21' with filters for Critical (5), Major (11), Minor (0), Notice (5), and Healthy (0). A search bar is present with the text 'Type to search events'. The main area is a table of events with columns for Organization, Severity, Name, Message, Last Det., Age, Ticket ID, Count, Event Type, Automated Actions, Event ID, Masked Events, and Event Source. The table lists various events such as 'ens160 SHMP DCM Interface Inbound Utiliz', 'io SHMP DCM Interface Inbound Utiliz', and 'AsimovQA-CDB-1 CPU runtime contention of 52% esse'. On the right side, the 'Help Menu' is open, titled 'The Events Page'. It contains text explaining how to monitor network health, what events are, and what alerts are. It also lists actions users can take: 'View detailed information about an event on the Event Investigator page', 'View summary information about the device aligned with an event in the Device Summary modal', 'Search for an event and save a search for an event', 'Acknowledge or clear an event', and 'View or edit the Event Note'. At the bottom of the Help menu, there is a link to 'View Product Documentation'. The footer of the page shows 'AP2 | Version 12.2.1 | Build 2171 | © 2024 ScienceLogic'.

If you click the **View Product Documentation** link at the bottom of the **Help** menu, a product documentation topic specific to the current page appears in a new browser window:



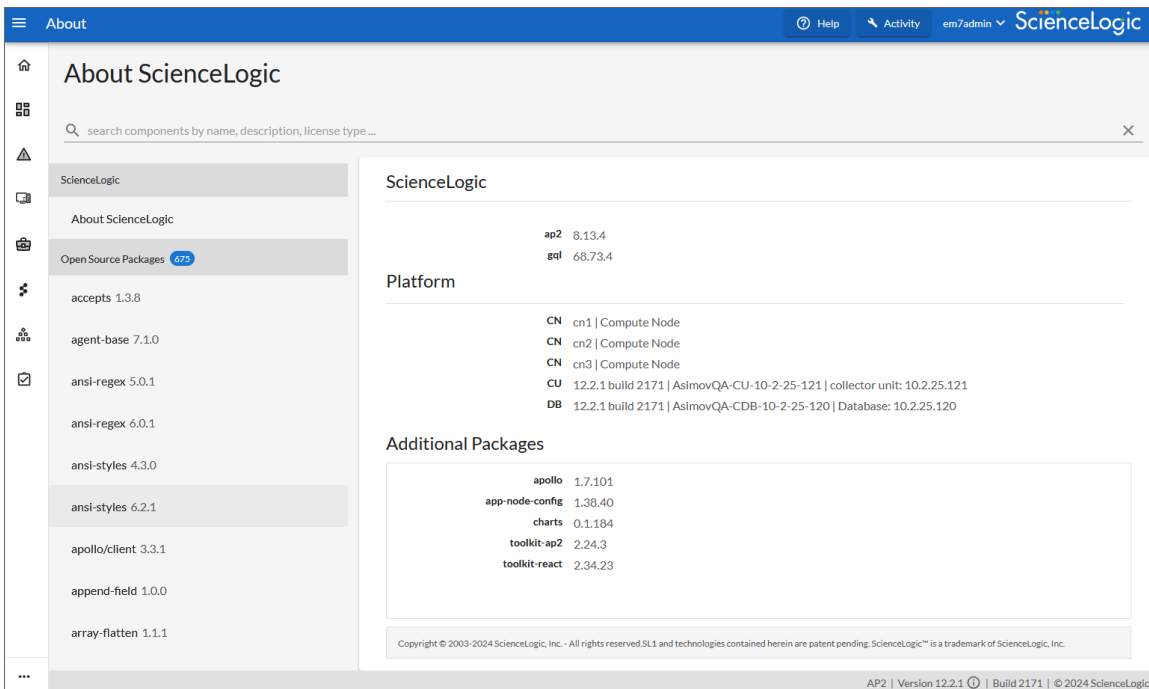
The product documentation includes a **Search** field at the top of the page that you can use to find additional topics related to the SL1 user interface. To find a specific topic that is longer than one word, enclose all of the key words in parentheses, such as "business services".

TIP: Click the **[Feedback]** button on the right side of the product documentation window to send comments directly to the ScienceLogic Documentation team, such as reporting typos, inaccuracies, questions, or other comments about that specific page in the Product Documentation.

NOTE: As of version 8.12.2 of SL1, ScienceLogic no longer updates the content that appears when you click the **[Guide]** button in the classic user interface. All help content is maintained in the **Help** menu and the online product documentation, which is located at <https://docs.sciencelogic.com>.

The About Page

For more information about the components used by SL1, click your user name in the navigation bar at the top of any SL1 page and select *About*. The **About ScienceLogic** page appears:



The current version number of the SL1 user interface displays in the **ap2** row at the top of the list of components in the right-hand pane, along with the current version of GraphQL used by SL1. The **Platform** section lists SL1 version information for the various components in SL1, including Application Server, Collector Unit, and Database details, and All-In-One configurations (AO) where relevant.

In the left-hand pane, click any of the components in the **Open Source Components** pane to view licensing information about those components, along with links to relevant websites where relevant.

To search for a specific open-source component, type the name of that component in the **Search** field at the top of the page. The list of components is filtered by your search terms.

In addition, every page in SL1 contains a footer that displays the SL1 version, the build number, and the copyright date on the right side of the footer. You can click the SL1 version text to go to the **About** page:

AP2 | Version 12.2.1 ⓘ | Build 2171 | © 2024 ScienceLogic

If you select the **Display Previous Login In Footer** option on the **Behavior Settings** page (System > Settings > Behavior), the footer will also display the date and time of the previous login or the last failed login (if applicable) on the left side of the footer:

Last Login: Aug 28, 2024, 10:21 AM

Tips and Best Practices for Using the Product Documentation

Use the following tips and best practices when using the ScienceLogic Help menu and product documentation:

General

- To ensure that you are always using the most recent version of the product documentation, check the URL for this site to make sure you are using <https://docs.sciencelogic.com/latest/>.
- The documentation at this site contains all manuals for the various features of SL1 and SL1 PowerFlow. You can click the *Download manual as PDF* link at the top of each chapter to download the PDF version of the manual.
- The documentation at this site is the same documentation that appears when you click the **[Help]** button at the top of any SL1 page. When you click **[Help]**, the **Help Menu** appears:
 - The content in the **Help Menu** is based on your current page in SL1, and it is a condensed version of what appears at this site.
 - If you click the **View Product Documentation** link at the bottom of the **Help Menu**, a product documentation topic specific to the current page appears in a new browser window:
- If you use the **Version** drop-down to navigate to a previous version of the product release notes or product documentation, the **Version** drop-down on the older version might contain links to archived, end-of-life versions of the site that display a 404 error.
- To hide the Table of Contents that displays at the left of the Help window, click the **[Hide Table of Contents]** button (☰). To show the Contents again, click the **[Show Table of Contents]** button (☰).

Searching the Site

- To search for a specific item in the product documentation, type the relevant information in the **Search** bar at the top of the page. After you select a document from the search results page, type the same search into your browser's Search (**Ctrl+F**).
- If you are using a multiple-word search for a specific term, enclose those words in quotation marks ("").
 - For example: "**business services**".
 - If you do not use quotes around a term, the online docs assume you want to search for any of those terms, such as **business OR services**, as opposed to a search for **business AND services** with quotes.
- Use the following command to search the product documentation via the SL1 API:
https://docs.sciencelogic.com/latest/#search-<your-query-here>. For example:
https://docs.sciencelogic.com/latest/#search-business services.
- After you click a link to a page from the search results, your search term or terms display as one or more highlighted colors. To get rid of the highlights, click the **Remove Highlights** button (✖) on the toolbar at top right of the Help content pane.

Links

- If you want to share or save a link to specific location in the documentation, click the link icon (🔗) to the right of the heading for a topic or sub-topic. That topic moves to the top of the reading pane, and you can copy the URL in the **Address** bar of the browser to share or save that location.

- If a link to an external site does not open properly in a new browser window or tab, right-click the link in the product release notes and select the option to open the link in a new tab or window.

Images

- An image in the product documentation might appear slightly blurry until you click that image to expand it in a pop-up window.
- If you clicked to expand an image in a pop-up window, you will need to click the expanded image in its pop-up window to return it to its regular size before you can click another image or link in the product documentation.
- You can also right-click an image and select *Open Image in New Tab* from the pop-up menu to make the image display at full size in the new browser tab.

Guides from the "Classic" User Interface

- As of version 8.12.2 of SL1, ScienceLogic is no longer updating the help content that appears when you click the **[Guide]** button in the user interface. All help content will be updated and maintained here at <https://docs.sciencelogic.com>.

Filtering the Items on a Classic Page

Many pages in the classic user interface for SL1 that display data in tabular format include a "filter-while-you-type" filter above each table column. For example, the following filters appear on the **Device Manager** page (Devices > Device Manager or Registry > Devices > Device Manager in the classic user interface).

You can filter the list of items on a page by entering values in one or more filters. The list of items is dynamically updated as you enter values. 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. There are two general types of filter:

- Text filters, where you enter text to match against. SL1 will search for the items that contain the text you entered, including partial matches. For example, if you enter "server" in the **Device Name** filter in the **Device Manager** page, the list of devices will be filtered to include only the devices that include "server" in the device name. Most text filters support the following special characters:
 - , (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.

- `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 it would match 1, 2, 6, 345, etc.
 - "`-5`" matches 5 and less, so it 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.

- Drop-down list filters, where you select a value from a list of pre-defined values. SL1 will search the items that match the value you selected. For example, if you select ">=Notice" in the **Current State** filter in the **Device Manager** page, the list of devices will be filtered to include only the devices that have a current state of "Notice" or above.

If you select multiple filters, the list of values will be filtered to include only items that meet all the filter criteria.

NOTE: Not all pages include "filter-while-you-type" filters and not all filters support all special characters.

Tool Tips

Some pages in SL1 include question-mark icons next to one or more fields. These question-mark icons are called Tool Tips. When you move your mouse over a Tool Tip, SL1 displays a brief description of the corresponding field.

Creating and Using Bookmarks

The **Administer Bookmarks** page (Misc > Bookmarks) allows you to create bookmarks in the classic user interface and quickly navigate to a selected page or to a selected result.

You can bookmark:

- The current page (classic user interface only)
- The results of a search, such as a search for all devices where the name begins with "cisco" (classic user interface only)
- A filtered list, such as a list of all tickets in the "Engineering" queue (classic user interface only)
- An editor page, with drop-down lists already selected. For example, you could bookmark the **Dynamic Applications Create New Application** page, with the Application Type of "XML Performance" already selected. The selected drop-down determines the other fields that appear in the subsequent pages.

To create a new bookmark:

1. In SL1, navigate to the page and/or results you want to bookmark (classic user interface only).
2. Use the shortcut keys Ctrl + Alt + B to display the **Administer Bookmarks** page.
3. In the **Administer Bookmarks** page (Misc > Bookmarks), supply a value in the **Save Current Page** field. SL1 automatically supplies the page name in this field. You can edit the field to include more descriptive text.
4. Click the **[Save]** button to save the new bookmark. The new bookmark appears at the bottom of the **Administer Bookmarks** page.

To access an existing **bookmark**, click its **star** icon (★). SL1 will display the bookmarked page.

To delete a **bookmark**, click its **bomb** icon (💣).

The Finder Tool

The **Finder** tool (Misc > Finder) allows you to easily find one or multiple elements in SL1. This prevents you from having to navigate through multiple pages to find the element you are interested in.

The **Finder** tool allows you to search for one or more of the following types of elements:

- Organizations
- Devices
- Assets
- IP networks
- Interfaces
- Vendors
- User Accounts
- Virtual Interfaces
- Device Groups

NOTE: To access the **Finder** page, accounts of type "user" must be granted one or more access keys that includes the following access hook: Finder. Accounts of type "user" will then be able to view the **Finder** tool.

- To search organizations, the user must be granted the access hook Org:View.
- To search for device, the user must be granted the access hook Dev:View.
- To search asset records, the user must be granted the access hook Asset:View.
- To search IPv4 Networks, the user must be granted the access hook Networks:IPv4:View.
- To search interfaces, the user must be granted the access hook networks:Interfaces:View.
- To search vendor records, the user must be granted the access hook Vendor:View.
- To search user accounts, the user must be granted the access hook User:View.
- To search virtual interfaces, the user must be granted the access hook Networks:Interfaces:View.

Searching for One or More Elements

To use the **Finder** tool to search for one or more elements:

1. From any page in SL1, press Ctrl + Alt + F. The **Finder** tool appears.
2. In the **Finder** modal page, select the checkboxes for the element types you want to search. Choices are:

- Organizations
 - Devices
 - Assets
 - IP networks
 - Interfaces
 - Vendors
 - User Accounts
 - Virtual Interfaces
3. In the **Search** field, type the whole string of text or a partial string of text for which you want to search.
 4. You can include the asterisk (*) wildcard in the **Search** field.
 - If placed at the beginning of a string, the asterisk says "match on this string, preceded by any number of any characters".
 - If placed at the end of a string, the asterisk says "match on this string, followed by any number of any characters."
 5. Click the **[Search]** button.
 6. The elements that match the search requirements are displayed at the bottom of the **Finder** page.

Viewing a List of All Elements in SL1



To use the **Finder** page to view a list of all elements in SL1 :

1. From any page in SL1 , enter Ctrl + Alt + F.
2. On the **Finder** page, select all the checkboxes.
3. In the **Search** field, type "%" (percent).
4. Click the **[Search]** button. A list of all elements in SL1 is displayed at the bottom of the **Finder** page.

Viewing the Results Table

After clicking the **[Search]** button, SL1 will display a list of elements that match the selected element type and search string.

From this list of results, you can:



- **View an element's properties.** To do this, click on its icon (the icon will vary, depending upon the element).
- **View reports for some elements.** To do this, click on the bar graph icon (.
- **Create a ticket about an element.** To do this, click on the life ring icon (.

The Toolbox

The Toolbox appears in the upper right of the window in the classic user interface of SL1 .

The Toolbox displays the following:

- **Logged in.** Displays the name of the current user. When you hover over the username, the IP address for the current session appears.
- **Finder.** If you enter a string in this field and click the **[Go]** button, SL1 uses the **Finder** tool to search all possible element types. The results are returned in the **Finder** tool page.
- **Toolbox.** Clicking the **[Toolbox]** button displays a set of links to commonly accessed pages and the shortcut keys for those pages. The **[Toolbox]** button also displays links for *License Information* and *SL1 Version Information*.
 - The *License Information* link leads to a pop-up page with a list of all third-party licenses included in SL1.
 - The *SL1 Information* link leads to a pop-up page with information about the current version of SL1.

You can also access the following Toolbox options from the menu bar on the left () and from the Advanced menu (), under the **Misc** heading:

- *Bookmarks* (Misc > Bookmarks). Displays the Administer Bookmarks page for saved and saving bookmarked SL1 pages.
- *Finder* (Misc > Finder). Displays the Finder search tool for searching for specific SL1 features.
- *Guides* (Misc > Guides). Displays the Guide Browser, which you can search to find more information about the classic user interface. The Guides are no longer being updated. For the most recent Help information, click the **[Help]** button at the top right of any page in SL1.
- *Clear SL1 Cache* (Misc > Clear SL1 Cache). Removes cached items from SL1.
- *Regular Expression Tester* (Misc > Regular Expression Tester). Lets you type regular expressions and search text into fields and then test them without impacting anything in SL1.
- *SL1 License Info* (Misc > SL1 License Info). Displays a page of licensed software used by SL1.

Chapter



2

Overview of SL1 Features

Overview

This chapter provides an overview of the features and terminology 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 of the menu options, click the Advanced menu icon ().

This chapter covers the following topics:

<i>SL1 Architecture</i>	52
<i>SL1 Extended Architecture</i>	54
<i>The SL1 Agent</i>	55
<i>Organizations and Users</i>	56
<i>Credentials</i>	58
<i>Discovery</i>	58
<i>Device Administration</i>	59
<i>Business Services</i>	60
<i>Events</i>	60
<i>Automation</i>	60
<i>Anomaly Detection</i>	61
<i>Dynamic Applications</i>	61
<i>PowerPacks</i>	62
<i>Maps</i>	62

<i>Dashboards</i>	63
<i>Reports</i>	64
<i>Ticketing</i>	65
<i>Asset Management</i>	65

SL1 Architecture

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

User Interface

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

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

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

Database Server

The node or appliance that provides the database function is responsible for:

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

The following appliances can perform these database functions:

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

Data Collection

Data Collectors are the SL1 nodes or appliances that retrieve data from monitored devices. In a distributed system, nodes or appliances that perform the data collection function also perform some pre-processing of collected data and execute automation actions.

The following appliances can perform the collection function:

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

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

Message Collection

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

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

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

API

SL1 provides a REST-based API that external systems can use to configure SL1 and access collected data. For example, the API can be used to automate the provisioning of devices in SL1.

The API is available through the **Administration Portal**, the **All-In-One Appliance**, and the **Database Server**.

For more information about SL1 architecture, see the **Architecture** manual.

SL1 Extended Architecture

SL1 Extended Architecture includes additional types of SL1 nodes or appliances. The following SL1 features require the SL1 Extended Architecture:

- **Expanded Agent Capabilities.** You can configure the SL1 Agent to communicate with SL1 via a dedicated Message Collector. However, this configuration limits the capabilities of the SL1 Agent. If you configure the SL1 Agent to communicate with SL1 via a Compute Cluster, you expand the capabilities of the SL1 Agent to include features like extensible collection and application monitoring.
- **Data Pipelines.** Data pipelines transport and transform data. Data transformations include enrichment with metadata, data rollup, and pattern-matching for alerting and automation. The Data Pipelines provide an alternative to the existing methods of data transport (data pull, config push, streamer, and communication via encrypted SQL) in SL1. Data pipelines introduce message queues and communicate using encrypted web services.
- **Publisher.** Publisher enables the egress of data from SL1. Publisher can provide data for long-term storage or provide input to other applications that perform analysis or reporting.
- **Scale-out storage of performance data.** Extended Architecture includes a non-SQL database (Scylla) for scalable storage of performance data.
- **Anomaly Detection and future AI/ML developments.** Anomaly detection is a technique that uses machine learning to identify unusual patterns that do not conform to expected behavior. SL1 does this by collecting data for a particular metric over a period of time, learning the patterns of that particular device metric, and then choosing the best possible algorithm to analyze that data. Anomalies are detected when the actual collected data value falls outside the boundaries of the expected value range.

SL1 Extended Architecture includes the following additional SL1 nodes or appliances:

Compute

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

Load Balancer

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

Storage

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

Management

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

The SL1 Agent

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

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

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

NOTE: The following minimum agent versions are required for SL1 12.1.1: **Windows** version 131; **Linux** version 174; **AIX** version 180; and **Solaris** version 180.

In the **SL1 Extended Architecture** (which includes Compute Nodes, Storage Nodes, and a Management Node), the Gen 3 agent collects the following data:

- **Device Availability.** SL1 can determine the availability state of a device (available or unavailable) and generate trended availability graphs based on uptime data collected by the agent.
- **Logs.** The SL1 agent can be configured to push logs to SL1 that match specific criteria from a log file or the Windows Event Log. You can view logs collected by the SL1 agent on the Logs pane of the **Device Investigator** page. The same logs also appear on the **[Logs]** tab in the **Device Properties** and **Device Summary** pages for that device. You can define event policies that specify how logs collected by an agent will trigger events.
- **Host Performance Metrics.** Using Dynamic Applications, SL1 translates data provided by an SL1 agent to trend the following metrics:
 - Overall CPU Utilization
 - CPU Utilization Breakdown
 - Disk Average Queue Length
 - Disk IO Utilization
 - Memory Utilization
 - Network Bytes Read
 - Network Bytes Written
 - Storage Available

- Storage Total
- Storage Utilization
- Swap Utilization

You can view these metrics on the **Device Investigator** page and the **[Performance]** tab of the **Device Summary** panel for a specific device.

- **Host Configuration.** Using a Dynamic Application, SL1 collects the following configuration data based on data provided by the SL1 Agent:
 - The number and speed of the installed CPUs
 - The amount of installed memory
 - The overall and per-disk storage size
 - The total swap capacity (SL1 Extended Architecture only)

You can view the collected configuration data on the **[Configs]** tab of the **Device Investigator** page and the **Device Summary** panel.

- **Network Interface.** The SL1 agent collects a list of the network interfaces running on the device. You can view the list of interfaces on the **[Interfaces]** tab of the **Device Investigator** page and the **Device Summary** page. This list includes attributes such as the interface MAC address, IP address, position, and speed as well as inbound and outbound utilization, number of errors, and discard and usage percentage.
- **File System.** The SL1 agent collects data about the of configuration of the file systems found within a device, such as name, size and, type as well as utilization data such as free space, size, and usage percentage. You can view the file system data on the **[Hardware?]** tab of the **Device Investigator** page and the **Device Summary** page.
- **System Processes.** The SL1 agent collects a list of all processes running on the device, such as name, process ID (PID), and state. You can view the list of processes on the **[Processes]** tab of the **Device Investigator** page and the **[Processes]** tab of the **Device Summary** page. Monitoring policies can be configured to trend and alert on process availability, process CPU usage, and process memory usage.
- **Windows Services.** The SL1 agent collects a list of all Windows services enabled on the device. This list includes attributes such as the service name and run state. You can view the list of Windows Services on the **[Services]** tab of the **Device Investigator** page and the **Device Summary** page.
- **Installed Software.** The SL1 agent collects a list of the software running on the device. This list includes attributes such as software name, version, and installation date. You can view the list of software on the **[Software]** tab of the **Device Investigator** page and the **Device Summary** page.

Organizations and Users

All policies, events, tickets, users, and other elements in SL1 are associated with an organization. An **organization** is a group for managing elements and user accounts.

The basic characteristics of an organization are:

- A unique name (required).
- Users who are members of the organization.
- Elements (for example, devices) associated with the organization.

Organizations can be defined by geographic areas, departments, types of devices, or any structure that works best for your needs. For example, for a business with multiple locations, an administrator might create organizations named Boston, New York, and DC. Another administrator might create organizations named for departments, like Finance, Sales/Marketing, and Engineering.

Users

In SL1, there are two broad types of user accounts:

- **Administrators.** By default, users of type "administrator" are granted all permissions available in SL1. Administrators can access all tabs and pages, and perform all actions and tasks on all entities, regardless of organization.
- **Users.** Accounts of type "user" are assigned key privileges. Key privileges are customizable by the administrator and grant users access to pages and tabs and permit users to view information and perform tasks in SL1. These key privileges are defined by the SL1 system administrator from the **Access Keys** page (System > Manage > Access Keys).

An account of type "user" can be granted the privileges that allow him/her to create or modify other users' accounts. However, for accounts of type "user", certain restrictions apply:

- An account of type "user" cannot create or modify an account of type "administrator".
- An account of type "user" cannot change his/her own account to type "administrator" or change another user's account to type "administrator".
- An account of type "user" cannot add additional Access Keys to his/her own account.
- An account of type "user" cannot grant or remove Access Keys to other accounts that he/she has not also been granted.

Regardless of access keys, **accounts of type "user" can access only pages and actions associated with their organization.** For example:

- Suppose your organization includes three regional offices. Suppose you define three organizations: Northeast, Headquarters, and West Coast.
- Suppose each organization includes the hardware located at the corresponding office.
- Now suppose the account "JohnDoe" is of type "user" and is a member of the organization "West Coast". User JohnDoe would be able to view and act upon only devices that are included in the organization "West Coast". User JohnDoe would not be able to view or act upon the hardware at the other offices.
- **SL1 allows you to assign each user a primary organization and optional additional organizations.**
- Now suppose that user "JohnDoe" needs to view the status of a device at headquarters. If you add "Headquarters" as a secondary organization in JohnDoe's account information, that user will now be able to view and act upon all the devices in the "Headquarters" organization.

For more information about Organizations and Users, see the **Organizations and Users** manual.

Credentials

Credentials are access profiles (usually username, password, and any additional information required for access) that allow SL1 to retrieve information from devices and from software applications on devices. Discovery uses SNMP credentials to retrieve SNMP information during initial discovery and nightly auto-discovery. If SL1 can connect to a device with an SNMP credential, SL1 deems that device "manageable" in SL1.

Dynamic Applications use credentials to retrieve SNMP information, database information, SOAP information, XML information, XSLT information, and WMI information. Proxied Web Services use SOAP/XML Host credentials to pass authentication information to external web services.

SL1 includes a type of credential called "Basic/Snippet" that is not bound to a specific authentication protocol. You can use this type of credential for Dynamic Applications of type "WMI", of type "snippet", and when defining system backups. "Basic/Snippet" credentials can also be used for monitoring Windows devices using PowerShell.

SL1 includes a type of credential that allows Dynamic Applications of type "Snippet" to use SSH to communicate with a remote device. To use these Dynamic Applications, you must define an SSH credential.

SL1 includes a type of credential that allows Dynamic Applications to retrieve data from Windows devices. If you align a Dynamic Application for PowerShell with a PowerShell credential, SL1 assumes that you want to use its built-in agentless transport to communicate with Windows devices.

NOTE: If necessary, a single device can use multiple credentials. If more than one agent or application is running on the device, each agent or application can be associated with its own credential. During discovery, SL1 will use the appropriate credential for each agent.

For more information about Credentials, see the *Discovery and Credentials* manual.

Discovery

Discovery is the tool that automatically finds all the hardware-based devices, hardware components, and software applications in your network. You must provide the discovery tool with a range or list of IP addresses and/or a list of fully-qualified domain names (hostnames), and the discovery tool determines if a device, hardware component, or software application exists at each IP address. For each device, hardware component, or software application the discovery tool "discovers", the discovery tool can collect a list of open ports, DNS information, SSL certificates, list of network interfaces, device classes to align with the device, topology information, and basic SNMP information about the device.

The Discovery tool also determines which (if any) Dynamic Applications to align with the device. If the discovery tool finds Dynamic Applications to align with the device, the discovery tool triggers collection for each aligned Dynamic Application.

For more information about Discovery, see the *Discovery and Credentials* manual.

Device Administration

As part of monitoring your network, SL1 collects data using common networking protocols. Most collected data is associated with a **device** in SL1. A device in SL1 is a record that can represent:

- Physical network hardware, for example, servers, switches, routers, printers, etc.
- A component of a larger system, for example, a data store in a hypervisor system, a blade server, etc.
- Any other entity about which you want to collect data, but want or need to associate that data with a container that does not correspond directly to a physical device or a component. For example, you might configure a device record that represents a web site or a cloud service.

SL1 allows you to monitor and manage hardware and applications within your network. SL1 provides a network-wide view through a "single pane of glass." This means that you can monitor status, create policies, define thresholds, and receive notifications, all through a single, browser-based application.

Virtual Device

A **virtual device** is a container for collected data. A virtual device can be used when you want to:

- Monitor a device or application that doesn't support TCP/IP, SNMP, or both. The device's data can be pushed to SL1 via another method (for example, email) and stored in a virtual device.
- Monitor multiple SNMP agents on a single device. In such a case, one of the SNMP agents (for example, a hardware agent) can be associated with the device and another SNMP agent (for example, an agent that monitors a software application) can be associated with a virtual device.
- Isolate and monitor specific parameters separately from their originating device. For example, you might want to monitor a database and keep its data separate from the hardware data you are collecting from the host device.

Component Device

SL1 uses Dynamic Applications to retrieve data from a management device and discover each entity managed by that management device. SL1 then uses that retrieved data to create a device for each managed entity. In some cases, the managed entities are nested.

- In SL1 a managed entity is called a **component device**. A component device is an entity that runs under the control of a physical management device.
- In SL1, the **root device** is the physical device that manages one or more component devices.
- In SL1, a **parent device** is a device that has associated entities modeled as component devices. A parent device can be either a root device or another component device.

For more information about Devices, see the **Device Management** manual.

Business Services

A **business service** includes one or more technical services that provide value to internal or external customers. Some examples of business services include verifying Internet access or website hosting, online banking, remote backups, and remote storage. Usually a business service includes an associated Service Level Agreement (SLA) that specifies the terms of the service.

Create the following types of services on the **Business Services** page, in the following order:

1. **Device Service**. Monitors a set of related devices, such as all devices from a specific region.
2. **IT Service**. Monitors a service that IT provides to your organization. An IT service is made up of one or more device services.
3. **Business Service**. Monitors a service your organization provides to your customers. A business service is made up of one or more IT services.

For more information about Business Services, see the **Business Services** manual.

Events

One of the quickest ways to monitor the health of your network is to look at events. You can view events on the **Events** page in SL1.

Events are messages that are triggered when a specific condition is met. For example, an event can signal if a server has gone down, if a device is exceeding CPU or disk-space thresholds, or if communication with a device has failed. Alternately, an event can simply display the status of a managed element.

SL1 generates log messages from incoming trap and syslog data, and also when SL1 executes user-defined policies. SL1 then uses these log messages to generate events. SL1 examines each log message and compares it to each event definition. If a log message matches an event's definition, SL1 generates an event instance and displays the event on the **Events** page.

Each event includes a description of the problem, where the problem occurred (device, network hardware, software, policy violation), a pre-defined severity, the time of first occurrence, the time of most recent occurrence, and the age of the event.

SL1 includes pre-defined events for the most commonly encountered conditions in the most common environments. You can also create custom events for your specific environment or edit the pre-defined events to better fit your specific environment.

For more information about Events, see the **Events** manual.

Automation

SL1 includes automation features that allow you to specify actions you want SL1 to execute automatically when specific event conditions are met. Automation in SL1 is divided into two parts:

- An **automation policy** defines the event conditions that can trigger an automatic action.
- An **action policy** defines an action that can be triggered by an automation policy. An action policy can perform one of the following tasks:
 - Send an email message to a pre-defined list of users and/or external contacts.
 - Send an SNMP trap from SL1 to an external device.
 - Create a new ticket (using ticket templates defined in the **Ticket Templates** page [Registry > Ticketing > Templates]).
 - Update an existing ticket. An action policy can change the status and/or severity of an existing ticket and/or add a note to an existing ticket. For this action policy to trigger successfully, a ticket must be associated with the event that triggered the action.
 - Write an SNMP value to an existing SNMP object on an external device.
 - Query a database.
 - Run a custom python script, called a snippet.
 - Send an SNS Message to a Topic ARN (Amazon Resource Name). All subscribers to the Topic ARN will receive the message.

For more information about Automation, see the **Run Book Automation** manual.

Anomaly Detection

Anomaly detection is a technique that uses machine learning to identify unusual patterns that do not conform to expected behavior. SL1 does this by collecting data for a particular metric over a period of time, learning the patterns of that particular device metric, and then choosing the best possible algorithm to analyze that data.

SL1 uses the resulting combination of collected data and the auto-selected algorithm to build a model that is unique to that specific device and metric. That model is then used to anticipate the expected behavior for that device metric. Anomalies are detected when the actual collected data value falls outside the boundaries of the expected value range. SL1 then continuously refines the model as it collects more data.

TIP: Anomalies do not necessarily represent problems or events to be concerned about; rather, they represent unexpected behavior that you might want to investigate.

For more information about Anomaly Detection, see the **Anomaly Detection** manual.

Dynamic Applications

Dynamic Applications are the customizable policies that tell SL1 what data to collect from devices and applications. For example, suppose you want to monitor a MySQL database running on a device in your network. Suppose you want to know how many insert operations are performed on the MySQL database. You can create

or edit a Dynamic Application that monitors inserts. Every five minutes (for example), SL1 could check the number of insert operations performed on the MySQL database. SL1 can use the retrieved data to trigger events and/or to create performance reports.

SL1 includes Dynamic Applications for the most common hardware and software. You can customize these default Dynamic Applications to suit your environment. You can also create custom Dynamic Applications.

Dynamic Applications in SL1 support a variety of protocols to ensure that SL1 can always communicate with the devices and applications in your network and retrieve information from them. Dynamic Applications can use the following protocols to communicate with devices:

- SNMP
- SQL
- XML
- SOAP
- XSLT (uses SOAP and XSLT to convert XML data to a new format)
- WMI (Windows Management Instrumentation), including WMI and WBEM
- Windows PowerShell
- Custom Python applications (called "snippets") for proprietary or more complex data retrieval

For more information about Automation, see the **Run Book Automation** manual.

PowerPacks

A **PowerPack** is an exportable and importable package of one or more Dynamic Applications, device classes, device templates, event policies, custom reports, dashboard widgets, dashboards, run book policies, run book actions, ticket templates, credentials, proxy XML transformations, themes, device categories, device dashboards, and/or IT service policies.

You can use PowerPacks to share customized content among SL1 systems and to download customized content from ScienceLogic.

You can create a PowerPack on a SL1 system to export one or more Dynamic Applications, device classes, device templates, event policies, custom reports, dashboard widgets, dashboards, run book policies, run book actions, ticket templates, credentials, proxy XML transformations, themes, device categories, device dashboards, and/or IT service policies. You can then import that PowerPack on another SL1 system to install the Dynamic Applications, device classes, device templates, event policies, custom reports, dashboard widgets, dashboards, run book policies, run book actions, ticket templates, credentials, proxy XML Transformation, themes, device categories, device dashboards, and/or IT service policies.


For more information about PowerPacks, see the **PowerPacks** manual.

Maps

A **map** is a visual representation of the various devices and related elements, also called **nodes**, in your environment that have been discovered by SL1. A map displays the important details about the nodes, their hierarchy, and the relationships associated with those nodes.

Maps can display business services, component maps (DCM, DCM+R), CDP topology, LLDP topology, Layer-2 topology, Layer-3 topology, and Virtual Infrastructure (VMware and virtual machines).

You can also create your own maps with your most important devices, and add images, text, and shapes to customize your maps.

To view a map, go to the **Maps** page () and click the name of the map from the **Maps** page.

A map includes the following graphical elements:

- **Nodes.** Shapes that represent Devices, Topology Elements, and Business Services defined in SL1. The shape of a node represents its type, and the color of its outline specifies the current state of the node.
- **Links.** Lines with or without arrows that represent the relationships and hierarchies between nodes. All device relationships are displayed as child and parent relationships. If the nodes on a map contain arrows, then the arrows represent the direction of the relationship, pointing from the child node to its parent node. If a node does not contain an arrow, then the relationship is bi-directional, or *undirected*.

For more information about Maps, see the **Maps** manual.


Dashboards

A **dashboard** is a page that displays one or more graphical reports, called **widgets**. These widgets appear in their own pane, and display charts, tables, and text. Access to dashboards is based on your login credentials, so you can view only dashboard data for which you have access. Also, some dashboards might be private instead of public.

To define a widget, you first select from a list of pre-defined widget definitions, and then customize what will be displayed by the selected widget by supplying values in the option fields provided by that widget.

NOTE: If an animated line appears under a widget name, the widget is in the process of updating its data. When the line disappears, the widget is done updating.

TIP: If an item name displays as a hyperlink in a dashboard, you can click that link to go to the relevant detail or Investigator page for that item. You can click dashboard links to the Investigator pages for devices, events, and services.

To navigate to the **Dashboards** page, click the Dashboards icon (). You can also access "classic" dashboards from the **Classic Dashboards** page (Dashboards > Classic Dashboards):

For more information about Dashboards, see the **Dashboards** manual.

Reports

Custom Reports

A **custom report** in SL1 provides you with a collection of data from one or more tables in the SL1 database. This information is populated and generated in different user-defined formats. You can select from default custom reports provided by ScienceLogic, edit these default reports, or create your own reports. You can also schedule reports, view a list of archived reports, and email reports to other users.

Custom reports include **Quick Reports**, which are custom report templates in SL1. You can access Quick Reports on the **Reports** page, in the **Run Report** category (Reports > Run Report).

A report includes three components:

- An input form where you select the options and data you want to include in the report.
- An **.ods** output template that specifies the format of the generated report.
- **Gluecode**, the code that specifies how to handle your input, which data to retrieve, and any processing that needs to be performed on the data.

SL1 includes predefined reports, with defined forms, output templates, and the gluecode. These predefined reports can be modified, and you can create your own custom reports.

Embedded Reports

Several pages in SL1 allow you to generate a report that contains the information displayed in the page. Reports that are specific to a page are called **embedded reports**. The embedded reports cover the following elements:

- Devices
- Device Interfaces
- System Processes
- Windows Services
- Hardware Components
- Installed Software
- Organizations
- User Accounts
- Access Keys
- Tickets
- Asset Records
- Product Subscriptions
- Vendors

If a feature includes embedded reports, the SL1 page for that feature will include a **[Reports]** button. Also, the section in the documentation that covers that feature will include a description of the embedded reports for that feature.

For more information about Reports, see the **Reports** manual.

Ticketing

A **ticket** is a request for work. This request can be in response to a problem that needs to be fixed, for routine maintenance, or for any type of work you require. Tickets are assigned a severity based on the severity of the issue that needs to be fixed or worked on. For example, a server going down might require a critical ticket, whereas a routine maintenance issue might require only a minor ticket. These severities range from healthy to notice, minor, major, and critical.

A ticket can be created manually, or created based on an event. If a ticket is created based on a selected event, most of the ticket fields are populated automatically by SL1. The SL1 can also automatically create a ticket, using Run Book Automation and user-defined parameters.

In SL1 you can view a list of active tickets, create new tickets, edit one or more existing tickets, and generate reports for one or more tickets, among other features.

For more information about Ticketing, see the **Ticketing** manual.

Asset Management

An asset is a piece of equipment owned by an organization. An asset record is a collection of information about that asset. In SL1, asset records are usually created for hardware devices, with some of the information populated automatically from collected data. Users can also manually enter information into an asset record.

In SL1, asset records can contain information about:

- The name, make, and model of a device.
- The serial number of a device.
- Function and status of a device.
- Networking information, like host ID, IP address, or DNS server for the device.
- Physical location of the device.
- Description of the network interface.
- Vendor information for the device, including PO or check number, warranty policy, and service policy.
- Hardware information like the amount of memory, CPU, and BIOS or EPROM version.
- Description of each hardware component (if applicable).
- Description of installed software (if applicable).

When possible, SL1 can automatically populate fields in each asset record. SL1 also allows users to create their own tabs and form fields in addition to the ones provided by default.

For more information about Assets, see the **Asset Management** manual.

Chapter



3

User Preferences

Overview

This chapter describes the different tools in SL1 that allow you to manage a restricted set of properties for your own account.

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

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

This chapter covers the following topics:

<i>Account Preferences</i>	67
<i>Schedule Manager</i>	71
<i>My Contact Information</i>	76

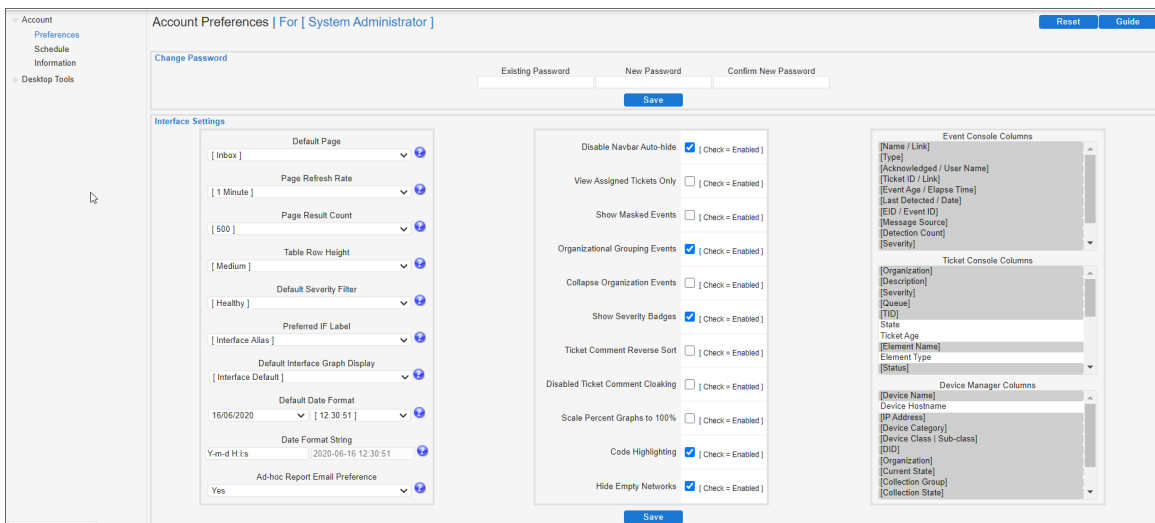
Account Preferences

The **Account Preferences** page allows you to change your SL1 password and customize some of the behavior and appearance of SL1. The customizations that you choose will appear each time you log in to SL1. They will not affect how SL1 appears to other users.

NOTE: To access the **Account Preferences** page, accounts of type "user" must be granted one or more access keys that includes the following access hook: MyPreferences. Accounts of type "user" will then be able to view and edit the settings in the **Account Preferences** page.

To access the **Account Preferences** page:

1. Go to Preferences > Account > Preferences.
2. In the **Account Preferences** page, you can edit one or more of the following fields:



3. The **Change Password** pane allows you to change your password. The following fields appear:
 - **Existing Password.** Your current password. This value must be at least four characters in length and can be up to 64 characters in length.
 - **New Password.** The new password. This value must be at least four characters in length and can be up to 64 characters in length.
 - **Confirm Password.** The new password again. This value must be at least four characters in length and can be up to 64 characters in length.
 - **[Save].** Click this button to save changes in the **Change Password** pane.
4. The **Interface Settings** pane allows a you to define the appearance and behavior of some pages. The **Interface Settings** pane contains the following fields:

- **Default Page.** Displays a drop-down list of pages. The selected page will automatically appear when you log in. If you select *None*, the default page (the **[Views]** tab) will appear when you log in.
- **Page Refresh Rate.** Specifies how often Event, Ticket, and Views pages in SL1 will be refreshed. The possible choices are from 15 seconds to 60 minutes.
- **Page Result Count.** Specifies the number of results to be displayed on each page. The choices are 50 to 500.
- **Table Row Height.** Affects the row height of all pages that display a table in the main content pane. You can also change this setting in the **Event Console Preferences** page, the **Ticket Console Preferences** page, and the system **Account Preferences** page. Changing the setting for row height in this page, the **Event Console Preferences** page, the **Ticket Console Preferences** page, or the system **Account Preferences** page affects the row height in all pages that display a table in the main content pane. Choices are:
 - *Small.* Sets row height to 17 px and font size to 11 px.
 - *Medium.* Sets row height to 27 px and font size to 12 px.
 - *Large.* Sets row height to 35 px and font size to 13 px.
- **Default Severity Filter.** When a severity is selected, you will see only events of the selected severity and greater in the **Event Console** page.
 - *Healthy.* Will display all events, including events with a severity of Healthy.
 - *Notice.* Will display all events with a severity of Notice, Major, Minor, and Critical.
 - *Minor.* Will display all events with a severity of Minor, Major, and Critical.
 - *Major.* Will display all events with a severity of Major and Critical.
 - *Critical.* Will display all events with a severity of Critical
- **Preferred IF Label.** Specifies how interfaces will be labeled in all pages and reports that reference network interfaces.
 - *Interface Alias.* Easy-to-remember, human-readable name for the network interface.
 - *Interface Name.* The name of the network interface.
- **Default Interface Graph Display.** Specifies the default unit of measure for the Hourly Interface Usage graph in the **Device Summary** page. Choices are:
 - *Interface Default.* The Hourly Interface Usage graph displays the amount traffic in the unit of measure specified in the **Measurement** field in the **Interface Properties** page for the interface.
 - *% Utilization.* The Hourly Interface Usage graph displays utilization in percent.
- **Default Date Format.** Specifies the default date format that will be used throughout SL1. You can select from a list of possible formats.

- **Date Format String.** Specifies a user-defined date format that will be used throughout SL1. If defined, this date format overrides the default date format. Any date variables supported by the PHP date function can be used.
- **Ad-hoc Report Email Preferences.** Specifies whether the user will receive an email after an ad-hoc report (Reports > Create Report > Report Jobs > lightning bolt) is generated.

5. In the **Checkboxes** pane, you can configure features that are toggled on and off.

- **Disable NavBar Auto-hide.** If you select this checkbox, the NavBar pane persists after you click a link. This option is selected by default.
- **View Assigned Tickets Only.** If you select this checkbox, by default, only tickets assigned to you are displayed in the **Ticket Console** page.
- **Show Masked Events.** If you select this checkbox, all events that have been grouped together under a single event description will be displayed in the **Event Console** page. The default behavior of SL1 is to roll up related events under a single description.
- **Organizational Grouping Events.** If you select this checkbox, events will be grouped by organization in the **Event Console** page. The filter-while-you-type fields and the advanced filter tool will appear for each organization grouping and will act only on the events in that organization grouping. You will not be able to apply a single filter to events in multiple organizations.
- **Collapse Organization Events.** If you select this checkbox, all organizations with assigned events will be displayed but will be contracted; the **Event Console** page will display only a list of contracted organizations, which can be expanded by clicking on the plus sign (+). The default behavior of SL1 is to expand each organization and display the list of events for each organization.
- **Show Severity Badges.** If you select this checkbox:
 - The value in the **Severity** column will be displayed as a color-coded badge in the **Event Console** page and the **Ticket Console** page.
 - The value in the **Current State** column will be displayed as a color-coded badge in the **Device Manager** page.

If you do not select the **Show Severity Badges** checkbox:

- In the **Event Console** page, the value in the **Event Message** column and the value in the **Severity** column will be painted with the severity color.
- In the **Ticket Console** page, the value in the **Description** column and the **Severity** column will be painted with the severity color.
- In the **Device Manager** page, the value in the **Device Name** column and the value in the **Current State** column will be painted with the severity color.
- **Ticket Comment Reverse Sort.** In the Notes section of a ticket, sort notes by newest first. If you do not select this checkbox, the user interface displays ticket notes from oldest to newest, with oldest displayed first.
- **Disabled Ticket Comment Cloaking.** When you add comments to a ticket, by default the comments are viewable by all (not cloaked).

- **Scale Percent Graphs to 100%.** Graphs that display percentage on the y-axis will display from 0% to 100%, regardless of the highest actual value. Default behavior is to display from 0% to highest actual value.
- **Code Highlighting.** If selected, enables syntax highlighting in areas of SL1 that display HTML, PHP, Python, and SQL code. If selected, syntax highlighting appears in:
 - The **Snippet Editor & Registry** page for Dynamic Applications of type "snippet" (System > Manage > Applications > create/edit > Snippets).
 - The **Dashboard Widget Editor** page (System > Customize > Dashboards > Widgets > create/edit).
 - The **Database Tool** page (System > Tools > DB Tool).

NOTE: The **Database Tool** page is available only in versions of SL1 prior to 12.2.1 and displays only for users that have sufficient permissions to access the page.

- The **Action Policy Editor** page for actions of type "Snippet" and "SQL Query" (Registry > Run Book > Actions > create/edit).
 - The **Report Template Editor** page (Reports > Management > Report Manager > create/edit).
- **Hide Empty Networks.** If you select this checkbox, the **IPv4 Networks** page hides networks that do not include any devices or interfaces.
6. In the **Event Console Columns** pane and the **Ticket Console Columns** pane, you specify the columns that will be displayed by default in the **Event Console** page and the **Ticket Console** page.
- **Event Console Columns.** In this list, you can select the default columns to be displayed in the **Event Console** page.

NOTE: From the **Event Console** page, you can also go to the **Event Console Preferences** modal page and edit the list of columns to be displayed in the **Event Console** page. When you edit the list of columns in the **Event Console Preferences** page, the selected list of columns in the **Account Preferences** page is automatically updated. When you edit the list of columns in the **Account Preferences** page, the selected list of columns in the **Event Console Preferences** page is updated.

- **Ticket Manager Columns.** In this list, you can select the default columns to be displayed in the **Ticket Console** page.

NOTE: From the **Ticket Console** page, you can also go to the **Ticket Console Preferences** modal page and edit the list of columns to be displayed in the **Ticket Console** page. When you edit the list of columns in the **Ticket Console Preferences** page, the selected list of columns in the **Account Preferences** page is automatically updated. When you edit the list of columns in the **Account Preferences** page, the selected list of columns in the **Ticket Console Preferences** page is updated.

Schedule Manager

The **Schedule Manager** page allows you to enter one-time or recurring appointments, meetings, and vacation leave. You can use the **Schedule Manager** page to specify the following:

- Your normal work schedule (for example, in the office on Monday – Friday, but out of the office on Saturday and Sunday)
- Vacation time
- Recurring meetings and appointments (for example, a weekly status meeting that occurs every Tuesday)
- One-time meetings and appointments (for example, a doctor's appointment)



The schedule information you enter in the **Schedule Manager** page is then visible to other users when they click the **View User Account Details** icon () for your account and then click the **[Schedule]** tab.

NOTE: To access the **Schedule Manager** page, accounts of type "user" must be granted one or more access keys that includes the following access hook: ACT_MY_SCHEDULE_PAGE. Accounts of type "user" will then be able to view and edit the settings in the **Schedule Manager** page.

NOTE: You can also view and manage all scheduled processes from the **Schedule Manager** page (Registry > Schedules > Schedule Manager). For more information, see the **System Administration** manual.


Viewing the Schedule Manager

The **Schedule Manager** page (Preferences > Account > Schedule) displays the following information about each scheduled or recurring calendar item:

Schedule Manager Schedules Found [2]														Create	Reset
Schedule Summary	Schedule Description	Event ID	sch_id	Context	Timezone	Start Time	Duration	Recurrence Interval	End Date	Last Run	Owner	Organization	Visibility	Enabled	
1.  Team Meeting	Team Meeting	179	0	Users	America/New_York	2017-01-20 15:00:00	30 minu	Every 1 Week	2017-12-29 13:00:00	--	banderton	System	Organiz	Yes	
2.  Vacation	Vacation	180	0	Users	America/New_York	2017-01-21 00:00:00	10080 n	--	--	--	banderton	System	World	Yes	

(Select Action)

- **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".

To edit a scheduled or recurring calendar item, click its wrench icon () and update the calendar item as needed on the **Schedule Editor** modal page.

Defining a Scheduled or Recurring Calendar Item

You can add a scheduled or recurring meeting, appointment, vacation, or other calendar item for the user in SL1 from the **Schedule Manager** page.

To define a scheduled or recurring calendar item:

1. Go to the **Schedule Manager** page (Preferences > Account > Schedule).
2. Click **[Create]**. The **Schedule Editor** modal page appears.
3. On the **Schedule Editor** modal page, make entries in the following fields:

The screenshot shows the 'Schedule Editor | Creating New schedule' modal. It is divided into two main sections: 'Basic Settings' and 'Time Settings'.
Basic Settings: Includes fields for 'Schedule Name', 'Users', 'Schedule Type', 'Visibility' (set to '[World]'), 'Organization' (set to '[System]'), 'Owner' (set to '[banderton]'), and a 'Description' field.
Time Settings: Includes 'Start Time' and 'End Time' (both with YYYY-MM-DD HH:MM:SS format), 'Time Zone' (set to '[America/New_]'), and an 'All Day' checkbox. Below this are 'Recurrence' (set to 'By Interval'), 'Interval' (set to '0'), and 'Interval' unit (set to 'Hours'). At the bottom, there are 'Recur Until' (set to '[Specified Date]') and 'Last Recurrence' (with YYYY-MM-DD HH:MM:SS format) fields. A 'Save' button is at the bottom center.

Basic Settings

- **Schedule Name.** Type a name for the scheduled process.
- **Schedule Type.** Indicates the scheduled process type (such as Tickets, Reports, or Devices).
- **Visibility.** Select the visibility for the scheduled process. You can select one of the following:
 - *Private.* The scheduled process is visible only to the owner selected in the **Owner** field.
 - *Organization.* The scheduled process is visible only to the organization selected in the **Organization** field.
 - *World.* The scheduled process is visible to all users.
- **Organization.** Select the organization to which you want to assign the scheduled process.
- **Owner.** Select the owner of the scheduled process. The default value is the username of the user who created the scheduled process.
- **Preserve Schedule.** Select this checkbox to exclude this schedule from being pruned after expiration.
- **Description.** Type a description of the scheduled process.

Time Settings

- **Start Time.** Click in the field and select the date and time you want the scheduled process to start.
- **End Time.** Click in the field and select the date and time you want the scheduled process to end.
- **Time Zone.** Select the region or time zone for the scheduled start time.

NOTE: If you want SL1 to automatically adjust for daylight savings time (if applicable), then you must select a named region (such as *America/New York*) in the **Time Zone** field. If you select a specific time zone (such as *EST*) or a specific time offset (such as *GMT-5*), then SL1 will not automatically adjust for daylight savings time. In addition, if you select a specific time zone, such as *EST*, that does not exist during daylight savings time observance, your schedules will be saved and execute at unexpected times.

- **All Day.** Select this checkbox if the scheduled process occurs all day rather than during a specific period of time. If you do so, the **End Time** field becomes disabled.
- **Recurrence.** Select whether you want the scheduled process to occur once or on a recurring basis. You can select one of the following:
 - *None.* The scheduled process occurs only once.
 - *By Interval.* The scheduled process recurs at a specific interval.
 - *Every Xth day of the Week.* The scheduled process occurs at a monthly interval based on a day of the week. The day of the week displayed in this option matched the day selected in the **Start Time** field. For example, if you set the **Start Time** to Thursday, August 5th and that day is the first Thursday of the month, then the recurrence option will be *Every 1st Thursday*, and the scheduled process will occur monthly on the first Thursday of the month.

If you select *By Interval*, the following additional fields appear:

- **Interval.** In the first field, enter a number representing the frequency of the scheduled process, then select the time interval in the second field. Choices are *Minutes, Hours, Days, Weeks, or Months*. For example:
 - If you specify "6 Hours", then the scheduled process recurs every six hours from the time listed in the **Start Time** field.
 - If you specify "10 Days", then the scheduled process recurs every 10 days from the date listed in the **Start Time** field.
 - If you specify "2 Weeks", then the scheduled process recurs every two weeks, on the same day of the week as the **Start Time**.
 - If you specify "3 Months" the ticket recurs every three months, on the same day of the month as the **Start Time**.

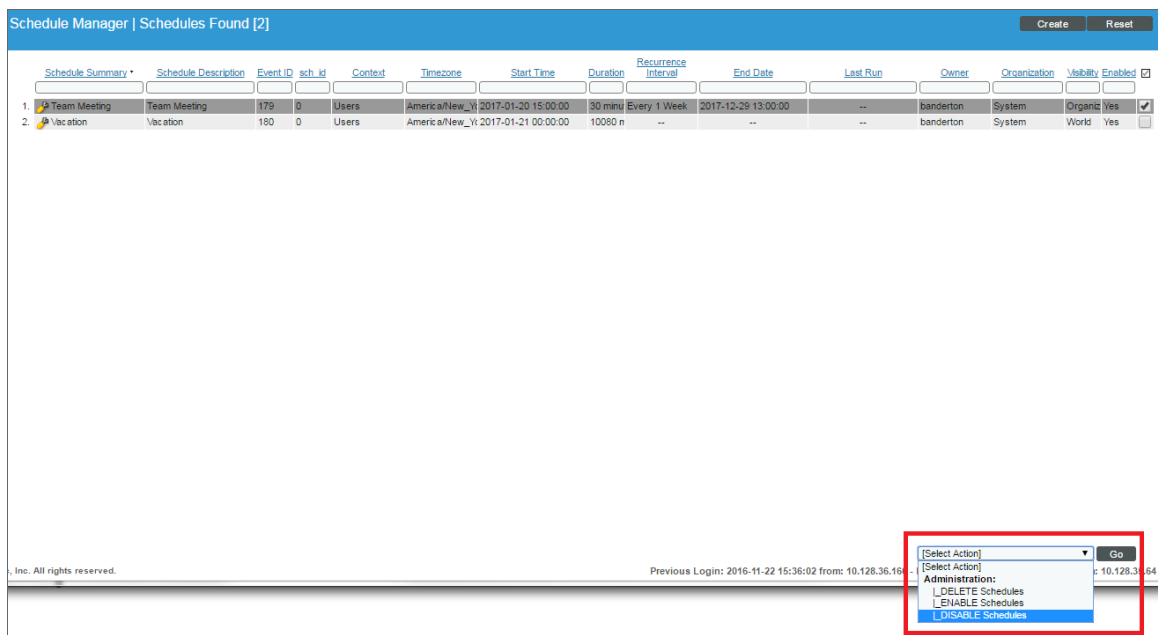
- **Recur Until.** Specifies when the scheduled process stops recurring. You can select one of the following:
 - *No Limit.* The scheduled process recurs indefinitely until it is disabled.
 - *Specified Date.* The scheduled process recurs until a specific date and time. If you select *Specified Date*, you must select a date and time in the **Last Recurrence** field.
- **Last Recurrence.** Click in the field and select the date and time you want the scheduled process to stop recurring.

4. Click **[Save]**.

Enabling or Disabling One or More Scheduled Calendar Items

You can enable or disable one or more scheduled or recurring calendar items from the **Schedule Manager** page (Preferences > Account > Schedule). To do this:

1. Go to the **Schedule Manager** page (Preferences > Account > Schedule).

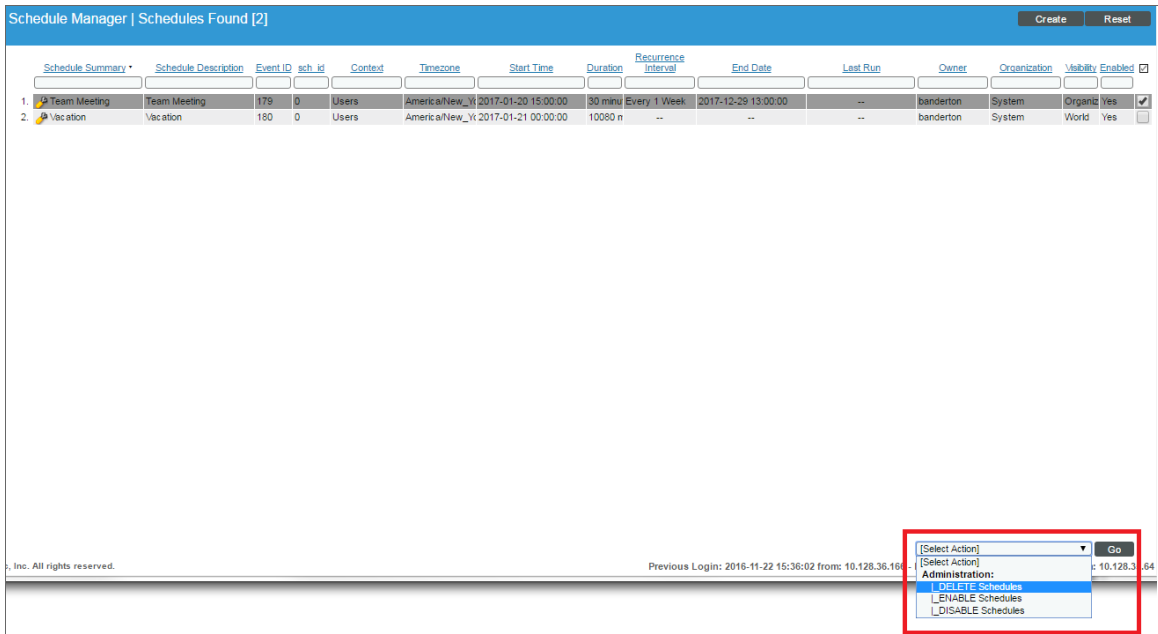


2. Select the checkbox icon for each scheduled process you want to enable or disable.
3. Click the **Select Action** menu and choose *Enable Schedules* or *Disable Schedules*.
4. Click the **[Go]** button.

Deleting One or More Scheduled Calendar Items

You can delete one or more scheduled or recurring calendar items from the **Schedule Manager** page (Preferences > Account > Schedule). To do this:

1. Go to the **Schedule Manager** page (Preferences > Account > Schedule).



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.

My Contact Information

The **My Contact Information** page allows you to define or edit your contact information. The contact information you enter in this page is then visible to other users when they click the **View User Account Details** icon (👤) associated with your account anywhere in SL1.


By default, the following fields will be automatically populated with values from your parent organization (if the values have been defined for your parent organization):

- Fax
- Address
- City
- State
- Postal Code
- Country
- Toll Free
- Time Zone

NOTE: To access the **My Contact Information** page, accounts of type "user" must be granted one or more access keys that includes the following access hook: MyInfo. Accounts of type "user" will then be able to view and edit the settings in the **My Contact Information** page.

To define or edit your contact information:

1. Go to Preferences > Account > Information.
2. In the **My Contact Information** page, position your cursor in the field you want to define or edit.



3. You can edit one, multiple, or all of the following fields:

- **First Name.** Your first name.
- **Last Name.** Your last name.
- **Title.** Your title. This field can be up to 32 characters in length.
- **Department.** Your department. This field can be up to 36 characters in length.
- **Phone.** Your phone number at work. This field can be up to 24 characters in length.
- **Fax.** Your fax number. This field can be up to 24 characters in length.
- **Mobile.** Your cell phone number. This field can be up to 24 characters in length.
- **Pager.** Any other phone numbers for contacting you. This field can be up to 24 characters in length.
- **Primary Email.** Your primary email address. This field can be up to 64 characters in length.
- **Secondary Email.** Your mobile email address, for contacting you via cell-phone or mobile device. This field can be up to 64 characters in length.
- **Alternate Email.** Additional email address for contacting you. This field can be up to 64 characters in length.
- **Pass Phrase.** Questions that verify your account if you forget your password. SL1 does not use this field.
- **Answer.** This field contains the answer to the question selected in the **Pass Phrase** field. This field can be up to 64 characters in length.
- **Street Address.** Your street address at work. This field can be up to 64 characters in length.

- **Suite/Building.** Suite/Building for your location at work. This field can be up to 64 characters in length.
- **City.** City where you work. This field can be up to 64 characters in length.
- **State.** State where you work. You can select from a list of states.
- **Postal Code.** Postal code where you work. This field can be up to 12 characters in length.
- **Country.** Country where you work. You can select from a list of countries.
- **Toll Free.** Toll-free phone number for your work address. This field can be up to 24 characters in length.
- **Time Zone.** Time zone associated with your work address. Select from a list of time zones.

4. Click the **[Save]** button to save your changes.

© 2003 - 2024, 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. For more information, see <https://sciencelogic.com/company/legal>.

ScienceLogic

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

International: +1-703-354-1010