



Discovery and Credentials

SL1 version 12.2.4

Table of Contents

Introduction to Discovery and Collection	7
Terminology	7
Who Should Read This Manual?	8
Requirements	8
Managing Credentials	9
What are Credentials?	10
Core Credential Types	10
Universal Credential Types	11
Viewing Information about Credentials	11
Defining Credentials	12
Defining a Basic/Snippet Credential	13
Defining a Database Credential	15
Monitoring Informix Databases	18
Defining an LDAP Credential	18
Defining a PowerShell Credential	21
Defining an SNMP Credential	24
Defining a SOAP/XML Credential	26
Defining an SSH/Key Credential	29
Defining an Aliyun Credential	31
Defining an AWS Credential	33
Defining an AWS Assume Role Credential	35
Defining an AWS EC2 Credential	38
Defining an AWS IAM Credential	41
Defining an Azure Credential	44
Defining a Citrix XenServer Credential	45
Defining an IBM Cloud Credential	47
Defining an S3 Backup Credential	48
Defining a VMware Credential	50
Testing a Credential	51
Using the Credential Tester Panel	52
Specifying Credentials For Discovery and Devices	52

Specifying Credentials for a Device/Dynamic Application Pair	53
Aligning One or More Organizations With a Credential	55
Editing a Credential	56
Duplicating a Credential	57
Deleting a Credential	57
Using Credentials in the Classic SL1 User Interface	58
Viewing Information about Credentials in the Classic SL1 User Interface	58
Filtering the List of Credentials in the Classic SL1 User Interface	59
Defining One or More SNMP Credentials for Initial Discovery in the Classic SL1 User Interface	60
Defining Credentials in the Classic SL1 User Interface	61
Defining an SNMP Credential in the Classic SL1 User Interface	61
Defining a Database Credential in the Classic SL1 User Interface	63
Defining a SOAP/XML Host Credential in the Classic SL1 User Interface	65
Defining an LDAP/AD Credential in the Classic SL1 User Interface	67
Defining a Basic/Snippet Credential in the Classic SL1 User Interface	70
Defining an SSH/Key Credential in the Classic SL1 User Interface	71
Defining a PowerShell Credential in the Classic SL1 User Interface	72
Specifying Credentials During Initial Classic Discovery	74
Defining the Primary and Secondary Credentials for a Single Device in the Classic SL1 User Interface	75
Defining the Credentials for a Specific Device/Dynamic Application Pair in the Classic SL1 User Interface	76
Specifying Credentials in a Device Template in the Classic SL1 User Interface	76
How SL1 Uses Credentials During Classic Discovery	77
Aligning One or More Organizations With a Credential in the Classic SL1 User Interface	77
Default Organizations Aligned with a Credential	79
Editing the Organizations Aligned with a Credential	79
Restricted Credentials in the Discovery Session Editor Page	80
Editing a Credential in the Classic SL1 User Interface	80
Deleting a Credential in the Classic SL1 User Interface	81
Managing Credential Tests	82
What are Credential Tests?	82
Default Credential Tests	83

AWS Credential Test	83
Azure Credential Test	83
Basic/Snippet Credential Test	83
Database Credential Test	84
PowerShell Credential Test	84
SNMP Credential Test	84
SOAP/XML Credential Test	85
SoftLayer Credential Test	85
SSH/Key Credential Test	85
VMware Credential Test	86
Viewing Information About Credential Tests	86
Filtering the List of Credential Tests	86
Testing a Credential	88
Testing a Credential During Guided or Unguided Discovery	88
Testing a Credential Using the Credential Tester Panel	88
Testing a Credential from the Credential Management Page in the Classic SL1 User Interface	89
Testing a Credential from the Credential Test Management Page in the Classic SL1 User Interface	90
Creating a Credential Test	91
Editing a Credential Test	92
Deleting Credential Tests	92
Available Step Functions	92
ping	92
nmap_udp	93
nmap_tcp	93
nslookup_forward	94
nslookup	94
dynapp_execute	95
snmp_getnext	95
ssh_request	96
db_query	96
curl	97
aws_connect	97

aws_service_scan	98
nmap_aws	99
nslookup_aws	99
ping_aws	100
Using External Credential Services	102
CyberArk	102
Installing the CyberArk Credential Agent on Every Collector	103
Reinstalling the CyberArk Credential Agent on a Rebuilt Data Collector	103
Defining a Custom CyberArk Delimiter	103
Starting the Credential Gateway Service on Every Data Collector	104
Troubleshooting Credential Gateway Service Issues	104
Setting Up Vault Tags in Your Credential	104
Example: Vault Tags	105
Overview of Discovery	107
What is Discovery?	107
What Happens During Discovery?	108
What Happens During Discovery when the SL1 Agent is Installed?	109
What is a Dynamic Application?	110
How Does SL1 Align Dynamic Applications During Discovery?	110
Before You Run Discovery	111
System Settings that Affect Discovery	112
Device Settings that Affect Auto-Discovery and Re-Discovery	116
How File Systems are Hidden During Discovery	117
Discovering Devices	119
Prerequisites for Discovering Devices	119
Adding Devices Using Guided Discovery	120
Adding Devices Using Unguided Discovery	123
Working with Discovery Sessions	129
Managing Classic Discovery Sessions	129
Viewing Information about a Classic Discovery Session	129
Running a Classic Discovery Session	132
Viewing Information about Classic Discovery	139



Creating a New Classic Discovery Session with the "Save As" Button	140
Editing a Classic Discovery Session	140
Scheduling a Classic Discovery Session	140
Viewing the Schedule Manager	141
Defining a Scheduled or Recurring Discovery Session	142
Enabling or Disabling One or More Scheduled Discovery Sessions	143
Deleting One or More Scheduled Discovery Sessions	144
Manually Re-Running Classic Discovery for a Dynamic Application	144
Manually Re-Running Classic Discovery for a Device	145
Collection	146
What is Collection?	146
Collection Processes	147
Duplication Protection During Discovery	149
Duplicate IP Addresses and Duplicate MAC Addresses During Discovery	149
Duplicate MAC Addresses for Component Devices	151
Managing MAC Vendor Records	152
Viewing the List of MAC Vendor Records	152
Filtering the List of MAC Vendor Records	153
Editing the Virtual Setting for MAC Vendor Records	153
Troubleshooting Discovery	155
Checking Network Security	155
Debugging the Discovery Processes	156
Checking Communication between Data Collectors and the Database Server	157

Introduction to Discovery and Collection

Overview

This chapter provides an overview of discovery, collection, and credentials 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:

Terminology	7
Who Should Read This Manual?	8
Requirements	8

Terminology

The following terms explain the key concepts used with discovery:

- **Discovery** is the tool that automatically discovers devices in your network. You supply the discovery tool with a range or list of IP addresses, and the discovery tool determines if a device exists at each IP address. For each device the discovery tool "discovers", the discovery tool can collect a list of open ports, DNS information, SSL certificates, a list of network interfaces, device classes to align with the device, 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 from the device, using each aligned Dynamic Application.

SL1 also uses discovery to update information about an already-discovered device and to add new information about an already-discovered device.

NOTE: Discovery collects a very specific set of information for each discovered device. Data that is not retrieved by discovery is retrieved by collection.

- **Collection** is the tool that retrieves policy-based information and Dynamic Application-based information from a device. After a device is discovered, you can define monitoring policies for that device in SL1. For example, if you define a policy to monitor a system process, the collection tool retrieves that information. For more information about collection processes, see the manual *Monitoring Device Infrastructure Health*.
- **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 from each discovered device. Dynamic Applications use credentials to retrieve SNMP information, database information, SOAP information, XML information, and LDAP and AD information. SL1 also 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. Another type of credential allows Dynamic Applications of type "Snippet" to use SSH to communicate with a remote device.

Who Should Read This Manual?

This manual is intended for users who are responsible for provisioning devices in SL1.

Requirements

To perform the troubleshooting steps in this manual, you must be allowed root-level access to SL1 appliances from a shell session.

Chapter



2

Managing Credentials

Overview

This chapter defines credentials and how they are used 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>What are Credentials?</i>	10
<i>Viewing Information about Credentials</i>	11
<i>Defining Credentials</i>	12
<i>Defining a Basic/Snippet Credential</i>	13
<i>Defining a Database Credential</i>	15
<i>Defining an LDAP Credential</i>	18
<i>Defining a PowerShell Credential</i>	21
<i>Defining an SNMP Credential</i>	24
<i>Defining a SOAP/XML Credential</i>	26
<i>Defining an SSH/Key Credential</i>	29
<i>Defining an Aliyun Credential</i>	31
<i>Defining an AWS Credential</i>	33
<i>Defining an AWS Assume Role Credential</i>	35
<i>Defining an AWS EC2 Credential</i>	38

<i>Defining an AWS IAM Credential</i>	41
<i>Defining an Azure Credential</i>	44
<i>Defining a Citrix XenServer Credential</i>	45
<i>Defining an IBM Cloud Credential</i>	47
<i>Defining an S3 Backup Credential</i>	48
<i>Defining a VMware Credential</i>	50
<i>Testing a Credential</i>	51
<i>Specifying Credentials For Discovery and Devices</i>	52
<i>Editing a Credential</i>	56
<i>Duplicating a Credential</i>	57
<i>Deleting a Credential</i>	57
<i>Using Credentials in the Classic SL1 User Interface</i>	58

What are Credentials?

Credentials are access profiles that allow SL1 to retrieve information from devices and from software applications on devices. Credentials typically include information such as a username and password, as well as any additional information required for accessing and monitoring devices. Dynamic Applications in SL1 use credentials to retrieve SNMP information, database information, SOAP information, XML information, XSLT information, and WMI information.

Core Credential Types

SL1 includes several **core credential** types that can be configured to access and monitor most device types:

- 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.
- **Basic/Snippet credentials** are 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.
- **Database Credentials** allow SL1 to access data on a database on a managed device. SL1 uses database credentials when collecting data for Database Dynamic Applications.
- **LDAP credentials** allow SL1 to communicate with an LDAP or Active Directory system. For details on integrating SL1 with LDAP or Active Directory, see the manual **Using Active Directory and LDAP**.
- **PowerShell credentials** allow 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.

- **SOAP/XML credentials** allow SL1 to access a web server on a managed device, and are used for SOAP, XML, XSLT, and snippet Dynamic Application types. With snippet Dynamic Applications, the snippet code must define the authentication protocol.
- **SSH credentials** allow Snippet-type Dynamic Applications in SL1 to use SSH to communicate with a remote device.

While these core credential types can be configured to monitor most device types, they use generic field labels that are not unique to each device type they might be used to access.

Universal Credential Types

SL1 also includes several **universal credential** types that are tailored to monitoring specific types of devices by using field names that correspond to the terminology used and the structures of data needed for those technologies. SL1 includes universal credentials for the following device types:

- Aliyun
- AWS, including credentials specific to Assume Role, EC2, and IAM
- Azure
- Citrix Xen
- IBM
- VMware

There are also several universal credential types that are used for SL1 configuration and administration, rather than to monitor specific device types. These include SL Service Connection and S3 Backup credential types.

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 example, suppose you want SL1 to discover a device that supports SNMP v2. To retrieve SNMP data from that device, SL1 must use a valid SNMP v2 read-only community string. So we would first go to the device and define the SNMP read-only community string. Then we would return to SL1 and create a credential in the SL1 system, using that community string. This new credential would allow discovery to retrieve SNMP data from the device.


Now suppose this same device also includes a MySQL database. Suppose you want SL1 to use a Dynamic Application to monitor that database. To retrieve data from the database, SL1 must use a valid username and password for that database. So we would first go to the device that hosts the MySQL database and create a database username and database password for SL1 to use. Then we would return to SL1 and create a credential in the SL1 system. The credential would include the database username and database password for the MySQL database. This credential would allow the Dynamic Application to retrieve data about the MySQL database.

Viewing Information about Credentials

The **Credentials** page (Manage > Credentials) allows you to view a list of all ScienceLogic credentials. From this page, you can also [create new credentials](#) and [edit](#), [duplicate](#), [test](#), or [delete](#) existing credentials.

For each credential, the **Credentials** page displays the following information:


- **ID.** Unique numeric ID, automatically assigned by SL1 to each credential.
- **Name.** Name of the credential.
- **Last Edit.** Date and time the credential was created or last edited.
- **Timeout (ms).** Time, in milliseconds, after which SL1 will stop trying to communicate with the external device or application.
- **Type.** Type of credential. Possible types are SNMP, Database, SOAP/XML, LDAP, Basic/Snippet, SSH/Key, and PowerShell.
- **Subtype.** Subtype of credential, for vendor-specific credentials. Possible subtypes are Aliyun, AWS, Azure, Citrix Xen, IBM, and VMware.

TIP: If you do not see one of these columns, click the gear icon () and then select *Column Preferences* to add or remove columns. You can also drag columns to different locations on the page or click on a column heading to sort the list by the values in that column. SL1 retains any changes you make to the columns that appear on the page and will automatically recall those changes the next time you visit the page.

TIP: You can filter the items on this inventory page by typing filter text or selecting filter options in one or more of the filters found above the columns on the page. For more information, see "Filtering Inventory Pages" in the *Introduction to SL1* manual.

TIP: You can adjust the size of the rows and the size of the row text on this inventory page. For more information, see the section on "Adjusting the Row Density" in the *Introduction to SL1* manual.

NOTE: [Editing a credential](#) and saving it overwrites the existing credential; it does not create a new credential. To create a new credential from an existing credential, see the section on [Duplicating a Credential](#).

NOTE: To automatically configure your SL1 Data Collector, Message Collector, or All-In-One Appliance to accept traps from monitored devices and communicate with those devices, click the **SNMPv3 Trap Configuration Reset** icon (). For more information, see the section on "Configuring SNMPv3 Traps" in the *Syslogs and Traps* manual.

Defining Credentials

To define a credential in SL1:

1. Collect the information you need to create each credential (usually username and password).
2. Go to the **Credentials** page (Manage > Credentials).
3. Click the **[Create New]** button and then select the type of credential you want to create from the drop-down list of options that appears. Your choices are:
 - - *Basic/Snippet Credential*
 - *Database Credential*
 - *LDAP Credential*
 - *PowerShell Credential*
 - *SNMP Credential*
 - *SOAP/XML Host Credential*
 - *SSH/Key Credential*
 - - *Aliyun Credential*
 - *AWS Credential*
 - *AWS Assume Role Credential*
 - *AWS EC2 Credential*
 - *AWS IAM Credential*
 - *Azure Credential*
 - *Citrix XenServer Credential*
 - *IBM Cloud Credential*
 - *S3 Backup Credential*
 - *VMware Credential*
4. The **Create Credential** modal page appears. In this page, you can define the new credential. The following sections explain how to create each type of credential.
5. Click the **[Save & Close]** button to save the new credential and close the window.

Defining a Basic/Snippet Credential

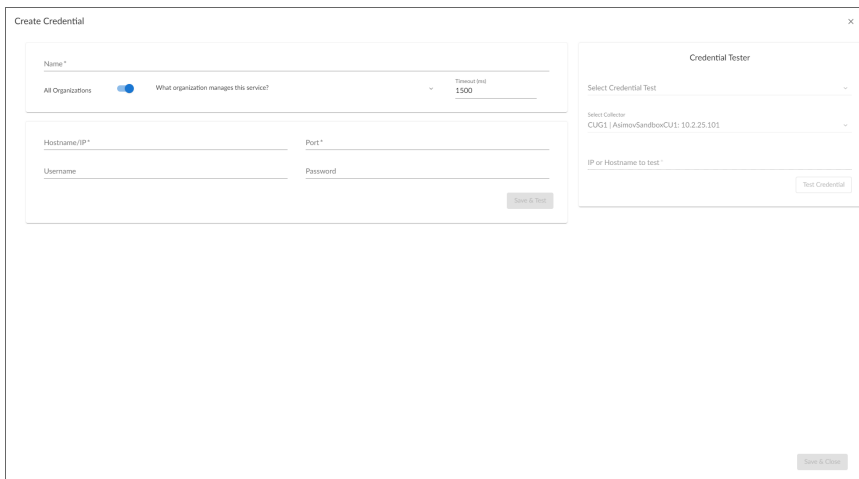
NOTE: Dynamic Applications of type "snippet" are not required to use only the Basic/Snippet Credential. In Dynamic Applications of type "snippet", the snippet code must define the authentication protocol. Therefore, Dynamic Applications of type "snippet" can use any type of credential.

Basic/Snippet credentials define standard authentication parameters, but are not tied to a specific authentication protocol. Basic/Snippet credentials are used in several places in SL1, including:

- **With Dynamic Applications of type "snippet"**. The snippet code must define the authentication protocol.
- **With Dynamic Applications of type "WMI"**. The authentication protocol is specific to WMI and is specified by SL1 when the Dynamic Application is executed. To access WMI information on a Windows server, ensure that the **Username** you specify is allowed access to the server and to the WMI namespace.
- **With Dynamic Applications of type "PowerShell"**. For information about configuring your environment for PowerShell collection, see the **Monitoring Windows Systems** manual.
- **When defining external backups**. The authentication protocol is defined in the **Backup Management** page (System > Settings > Backup).

To create a Basic/Snippet credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create Basic/Snippet Credential*. The **Create Credential** modal page appears:



3. Supply values in the following fields:

- **Name**. Name of the credential. Can be any combination of alphanumeric characters, up to 64 characters. This field is required.
- **All Organizations**. Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the **What organization manages this service?** drop-down field to align the credential with those specific organizations. This field is required.

NOTE: To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

- **Timeout (ms)**. Time, in milliseconds, after which SL1 will stop trying to communicate with the device from which you want to retrieve data.
- **Username**. Username for a user account on the device.
- **Password**. Password for a user account on the device.
- **Hostname/IP**. Hostname or IP address of the device from which you want to retrieve data. This field is required.
 - You can include the variable **%D** in this field. SL1 will replace the variable with the IP address of the current device (device that is currently using the credential).
 - You can include the variable **%N** in this field. SL1 will replace the variable with the hostname of the current device (device that is currently using the credential). If SL1 cannot determine the hostname, SL1 will replace the variable with the primary management IP address for the current device.
- **Port**. Port number associated with the data you want to retrieve. This field is required.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Defining a Database Credential

Database credentials allow SL1 to access data on a database on a managed device. SL1 uses database credentials when collecting data for Database Dynamic Applications.

To create a database credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create Database Credential*. The **Create Credential** modal page appears:

3. Supply values in the following fields:

- **Name**. Name of the credential. Can be any combination of alphanumeric characters, up to 64 characters. This field is required.
- **All Organizations**. Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the **What organization manages this service?** drop-down field to align the credential with those specific organizations. This field is required.

NOTE: To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

- **Timeout (ms)**. Time, in milliseconds, after which SL1 will stop trying to communicate with the database.
- **Database Type**. Type of database that will be accessed with the credential. Select from a list of databases supported by SL1. This field is required. Choices are:
 - MySQL
 - MS SQL Server
 - Oracle and *SQLNet
 - PostgreSQL
 - IBM DB2
 - Sybase ASE
 - Informix

NOTE: For information about monitoring Informix databases, see the [Monitoring Informix Databases](#) section.

- **Database Name.** Name of the database that will be accessed with the credential.
- **Database User.** Username associated with a valid account on the database.
- **Password.** Password associated with a valid account on the database.
- **Hostname/IP.** Hostname or IP address where the database resides. This field is required.
 - You can include the variable **%D** in this field. SL1 will replace the variable with the IP address of the current device (device that is currently using the credential).
 - You can include the variable **%N** in this field. SL1 will replace the variable with the hostname of the current device (device that is currently using the credential). If SL1 cannot determine the hostname, SL1 will replace the variable with the primary management IP address for the current device. This field is required.

NOTE: To use the localhost, in the **Hostname/IP** field, enter the IP address *127.0.0.1*. The credential will not work if you enter the string *localhost* in the **Hostname/IP** field.

- **Port.** Port number associated with the database you want to access with this credential. This field is required.
 - For **DB Type** of MySQL, the default value is 3306.
 - For **DB Type** of MS SQL Server, the default value is 1433.
 - For **DB Type** of Oracle and *SQLNet, the default value is 1521.
 - For **DB Type** of PostgreSQL, the default value is 5432.
 - For **DB Type** of IBM DB2, the default value is 523.
 - For **DB Type** of Sybase ASE, the default value is 4100.
 - For **DB Type** of Informix, see the 9088 section.

NOTE: SL1's Database Servers include a MySQL database running on port 7706. Data Collectors and Message Collectors include a MySQL database running on port 7707.

Oracle Settings

These fields should be completed only if you selected *Oracle & *SQLNet* in the **Database Type** field.

- **Oracle Connect Type.** Specifies the method SL1 should use to connect to the Oracle database. The choices are:

- Oracle System Identifier (SID)
- Oracle Real Application Clusters (SERVICE)
- Oracle Server Direct Connection (SERVER)

NOTE: In Oracle 11g, the "Oracle Server Direct Connection" option is deprecated. If you select this Oracle Connect Type for an Oracle 11g database, you must edit the file listener.ora and add the line "DEFAULT_SERVICE_LISTENER=<SID>", where <SID> is the SID value.

- **Oracle Database SID (if required).** Enter the value for the Oracle Connect Type (either Oracle SID, Oracle RAC, or Oracle Server) selected in the **Oracle Connect Type** field.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Monitoring Informix Databases

For SL1 to connect to an Informix database:

- The Informix database server must have a DRDA listener configured on a separate port than the current listener(s).
- The DRDA listener must be configured to share data with other listeners using a DBSERVERALIASES entry in the server's onconfig file.
- For servers that host multiple databases, multiple DRDA listeners are required with different port assignments.

For example Informix configuration files, please contact ScienceLogic Support.

Defining an LDAP Credential

LDAP or Active Directory credentials allow SL1 to access data on an LDAP server or an Active Directory server.

Authentication is the method by which SL1 determines if a user can access the SL1 system. For user accounts that are to be authenticated with LDAP or Active Directory, SL1 uses the LDAP or Active Directory credential to establish communication with the LDAP or Active Directory server. SL1 will then query the Active Directory or the LDAP server to determine if the username and password are legitimate and accurate.

Additionally, SL1 can automatically create accounts for one or more LDAP or Active Directory users. SL1 uses the LDAP or Active Directory credential to communicate with Active Directory or the LDAP server and:

- Determine if the username and password are legitimate and accurate.
- Gather information to populate fields in the user's automatically-created account.

For details on using Active Directory or LDAP for authentication, see the manual **Using Active Directory and LDAP**.

To create an LDAP credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create LDAP Credential*. The **Create Credential** modal page appears:

3. Supply values in the following fields:

- **Name**. Name of the credential. Can be any combination of alphanumeric characters, up to 64 characters. This field is required.
- **All Organizations**. Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the **What organization manages this service?** drop-down field to align the credential with those specific organizations. This field is required.

NOTE: To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

- **Timeout (ms)**. Time, in milliseconds, after which SL1 will stop trying to communicate with the directory server.
- **LDAP Type**. Specifies the type of LDAP running on the directory server. Choices are *LDAP* or *Active Directory*.
- **Hostname/IP**. Hostname or IP address of the LDAP or Active Directory server. This field is required.
- **Secure**. Specifies whether you are using LDAP over SSL.

- **Port.** Port number on the LDAP or Active Directory server to which SL1 will send requests. This field is required.
- **RDN (Bind DN / bind user).** Bind DN. The bind DN is a user on the LDAP or Active Directory server who is permitted to search the directory within the specified search base.
 - In many LDAP or AD configurations, each user has read-access to his/her own account. Therefore, you might find it most useful to include the %u variable in this field. When an LDAP or AD user logs in to SL1, SL1 stores the username in the %u variable. SL1 then uses the %u variable to build the bind DN, uses the bind DN to communicate with the LDAP or AD server, and then authenticates the current user.
 - An example entry in the RDN field might be:

```
uid=%u, ou=People, dc=sciencelogic, dc=com
```

This creates a DN using the current login name as the uid.

- You can also include the %d variable in this field. The %d variable represents the name of the LDAP domain, as specified in the **LDAP Domain** field.

NOTE: If you have configured SL1 to automatically create accounts when a user logs in with an LDAP/AD username, you must include the %u variable in the RDN field.

- **LDAP Domain.** If your LDAP or Active Directory configuration includes multiple domains, specify the domain components to bind to in this field. For example, you could specify:

```
dc=reston, dc=sciencelogic, dc=com.
```

This would bind to the sub-domain "reston", in the domain "sciencelogic", in the domain "com".

- **Bind Password.** Password that allows access to the LDAP or Active Directory server. In most cases, when you specify a bind password in a credential, you are creating a "write" credential (that is, a credential that allows SL1 to make changes to the LDAP or AD server).
- **User Search Base.** In this field, you specify the area in the directory where users to be authenticated reside, using RDN notation. For example, if you want to authenticate five users from the ou called "people", you could specify the RDN that includes that ou.

```
ou=People, dc=sciencelogic, dc=com.
```

This would allow SL1 to authenticate users in the ou called "people." In the **User Search Scope** field, you can specify whether SL1 should also authenticate all users in any ou underneath "people".

- **User Search Scope.** In this field, you specify whether SL1 should search only the directory specified in **User Search Base** or whether SL1 should search the directory specified in **User Search Base** and all its child branches. Choice are:

- *Subtree*. SL1 should search the directory specified in **User Search Base** and also search all its child branches.
- *One Level*. SL1 should search only the directory specified in **User Search Base**.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Defining a PowerShell Credential

Dynamic Applications can include PowerShell commands that collect data from Windows devices. If you want to use SL1's built-in transport agent (that is, run "agentless" on the Windows device), you can align a PowerShell credential with those Dynamic Applications.

NOTE: Consult the *Monitoring Windows* and *WMI and PowerShell Dynamic Application Development* manuals for detailed directions on configuring the Windows devices for agentless communication and on configuring a proxy server.

To define a PowerShell credential in SL1, you will need the following information:

- The username and password for a user on the Windows device.
- If the user is an Active Directory account, the hostname or IP address of the Active Directory server and the domain.
- Determine if an encrypted connection should be used.
- If you are using a Windows Management Proxy, the hostname or IP address of the proxy server.

To create a PowerShell credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create Powershell Credential*. The **Create Credential** modal page appears:

3. Supply values in the following fields:

- **Name.** Name of the credential. Can be any combination of alphanumeric characters, up to 64 characters. This field is required.
- **All Organizations.** Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the **What organization manages this service?** drop-down field to align the credential with those specific organizations. This field is required.

NOTE: To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

- **Timeout (ms).** Time, in milliseconds, after which SL1 will stop trying to communicate with the authenticating server. For collection to be successful, SL1 must connect to the authenticating server, execute the PowerShell command, and receive a response within the amount of time specified in this field.
- **Account Type.** Type of authentication for the username and password in this credential. Choices are:
 - *Active Directory.* On the Windows device, Active Directory will authenticate the username and password in this credential.
 - *Local.* Local security on the Windows device will authenticate the username and password in this credential.
- **Hostname/IP.** Hostname or IP address of the device from which you want to retrieve data. This field is required.
 - You can include the variable **%D** in this field. SL1 will replace the variable with the IP address of the device that is currently using the credential.

- You can include the variable **%N** in this field. SL1 will replace the variable with the hostname of the device that is currently using the credential. If SL1 cannot determine the hostname, SL1 will replace the variable with the primary, management IP address for the current device.
- You can include the prefix **HOST** or **WSMAN** before the variable **%D** in this field if the device you want to monitor uses a service principal name (for example, "HOST://%D" or "WSMAN://%D"). SL1 will use the WinRM service HOST or WSMAN instead of HTTP and replace the variable with the IP address of the device that is currently using the credential.
- **Username.** Type the username for an account on the Windows device to be monitored or on the proxy server. This field is required.

NOTE: The user should not include the domain name prefix in the username for Active Directory accounts. For example, use "em7admin" instead of "MSDOMAIN\em7admin".

- **Password.** Type the password for the account on the Windows device to be monitored or on the proxy server. This field is required.
- **Encrypted.** Select whether SL1 will communicate with the device using an encrypted HTTP or HTTPS connection:
 - Toggle on (blue) if SL1 will communicate with the device using an encrypted connection over HTTPS. If toggled on, when communicating with the Windows server, SL1 will use a local user account with authentication of type "Basic Auth". You must then use HTTPS and can use a Microsoft Certificate or a self signed certificate.

NOTE: In SL1 versions 11.3.0 and later, a newer Kerberos library is used that allows for message encryption over HTTP. This feature is on by default and may eliminate the need for you to configure an HTTPS certificate depending on your security requirements.

- Toggle off (gray) . The credential is encrypted over HTTP rather than HTTPS.
- **Port.** Type the port number used by the WinRM service on the Windows device. This field is required and is automatically populated with the default port based on the value you selected in the **Encrypted** field.
- **PowerShell Proxy Hostname/IP.** If you use a proxy server in front of the Windows devices you want to communicate with, type the fully-qualified domain name or the IP address of the proxy server in this field.
- **Active Directory Host/IP.** If you selected Active Directory in the **Account Type** field, type the hostname or IP address of the Active Directory server that will authenticate the credential.
- **Active Directory Domain.** If you selected Active Directory in the **Account Type** field, type the domain where the monitored Windows device resides.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Defining an SNMP Credential

SNMP credentials allow SL1 to access SNMP data on a managed device. SL1 uses SNMP credentials to perform discovery, run auto-discovery, and gather information from SNMP Dynamic Applications.

To create an SNMP credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create SNMP Credential*. The **Create Credential** modal page appears:

3. Supply values in the following fields:

- **Name.** Name of the credential. Can be any combination of alphanumeric characters, up to 64 characters. This is a required field.
- **All Organizations.** Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the **What organization manages this service?** drop-down field to align the credential with those specific organizations. This field is required.

NOTE: To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

- **Timeout (ms).** Time, in milliseconds, after which SL1 will stop trying to communicate with the device. The default value is 1500.
- **SNMP Version.** SNMP version. Choices are *SNMP V1*, *SNMP V2*, and *SNMP V3*. The default value is *SNMP V2*.
- **Port.** The port SL1 will use to communicate with the external device or application. The default value is 161. This field is required.
- **SNMP Retries.** Number of times SL1 will try to authenticate and communicate with the external device. The default value is 1.

SNMP V1/V2 Settings

If you selected *SNMP V1* or *SNMP V2* in the **SNMP Version** field, complete these fields. These fields are inactive if you selected *SNMP V3*.

- **SNMP Community (Read-Only).** The SNMP community string (password) required for read-only access of SNMP data on the remote device or application. For *SNMP V1* and *SNMP V2* credentials, you must supply a community string, either in this field or in the **SNMP Community (Read/Write)** field.
- **SNMP Community (Read/Write).** The SNMP community string (password) required for read and write access of SNMP data on the remote device or application. For *SNMP V1* and *SNMP V2* credentials, you must supply a community string, either in this field or in the **SNMP Community (Read Only)** field.

SNMP V3 Settings

If you selected *SNMP V3* in the **SNMP Version** field, complete these fields. These fields are inactive if you selected *SNMP V1* or *SNMP V2*.

- **Security Name.** Name for SNMP authentication. This field is required.
- **Security Passphrase.** Password to authenticate the credential. This value must contain at least 8 characters. This value is required if you use a **Security Level** that includes authentication.

In addition to alphanumeric characters, you **can** also use the following special characters in an SNMP V3 security passphrase: ? - _ = , . : # + % \$ [] { } & ! () | /

You **cannot** use the following special characters in an SNMP V3 security passphrase: " ' \

- **Authentication Protocol.** Select an authentication algorithm for the credential. This field is required. Choices are:
 - *MD5*. This is the default value.
 - *SHA*
 - *SHA-224*
 - *SHA-256*

- SHA-384
- SHA-512

NOTE: The SHA option is SHA-128.

- **Security Level.** Specifies the combination of security features for the credentials. This field is required. Choices are:
 - *No Authentication / No Encryption.*
 - *Authentication Only.* This is the default value.
 - *Authentication and Encryption.*
- **Engine ID.** The unique engine ID for the SNMP agent you want to communicate with. (SNMPv3 authentication and encryption keys are generated based on the associated passwords and the engine ID.) This field is optional.
- **Context.** A context is a mechanism within SNMPv3 (and AgentX) that allows you to use parallel versions of the same MIB objects. For example, one version of a MIB might be associated with SNMP Version 2 and another version of the same MIB might be associated with SNMP Version 3. For SNMP Version 3, specify the context name in this field. This field is optional.
- **Privacy Protocol.** The privacy service encryption and decryption algorithm. This field is required. Choices are:
 - *DES.* This is the default value.
 - *AES-128*
 - *AES-192*
 - *AES-256*
 - *AES-256-C.* This option is for discovering Cisco devices only.
- **Privacy Protocol Passphrase.** Privacy password for the credential. This field is optional.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Defining a SOAP/XML Credential

SOAP/XML credentials allow SL1 to access a web server on a managed device. SOAP/XML credentials are used in several places in SL1, including:

- **With Dynamic Applications of type "SOAP"**.
- **With Dynamic Applications of type "XML"**.
- **With Dynamic Applications of type "XSLT"**.
- **With Dynamic Applications of type "snippet"**. The snippet code must define the authentication protocol. Dynamic Applications of type "snippet" can use any type of credential.

To create a SOAP/XML credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create SOAP/XML Credential*. The **Create Credential** modal page appears:

3. Supply values in the following fields:

- **Name**. Name of the credential. Can be any combination of alphanumeric characters, up to 64 characters. This field is required.
- **All Organizations**. Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the **What organization manages this service?** drop-down field to align the credential with those specific organizations. This field is required.

NOTE: To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

- **Timeout (ms)**. Time, in milliseconds, after which SL1 will stop trying to communicate with the web service.
- **Content Encoding**. Tells the SOAP server or XML data-store how the content is encoded, so the SOAP server or XML data-store knows how to decode the message. Select the encoding that is appropriate for your request and response.

- **Method.** HTTP method to use to exchange credential data from the managed device. Choices are *GET* or *POST*.

NOTE: Typically, Dynamic Applications of type "XML" use GET methods. Dynamic Applications of type "SOAP" and of type "XSLT" use POST methods.

- **HTTP Version.** Version of HTTP to use. Choices are *1.0* or *1.1*.
- **URL.** Address of the SOAP server, HTML document, or XML document. This field is required and should be of the following format:

https://IP address:port/full path to desired SOAP, HTML, or XML document

NOTE: The port is stored if it is specified in the URL; otherwise, SL1 uses the default port values 80 for HTTP and 43 for HTTPS.

- You can include the variable **%D** in this field. SL1 will replace the variable with the IP address of the current device (device that is currently using the credential).

NOTE: For component devices, SL1 will replace %D with the IP address of the root device.

- You can include the variable **%N** in this field. SL1 will replace the variable with the hostname of the current device (device that is currently using the credential). If SL1 cannot determine the hostname, SL1 will replace the variable with the primary management IP address for the current device.

- **HTTP Auth User.** Username with which to log in to the web server.
- **HTTP Auth Password.** Password with which to access the web server.

Proxy Settings

If you use a proxy server in front of the SOAP server(s) or XML data-store(s) you want to communicate with, enter values in these fields. Otherwise, you can skip these fields.

- **Hostname/IP.** The host name or IP address of the proxy server.
- **Port.** Port on the proxy server to which you will connect.
- **User.** Username to use to access the proxy server.
- **Password.** Password to use to access the proxy server.

SOAP Options

These fields are optional. When a SOAP/XML credential is aligned with a SOAP or XSLT Dynamic Application, the requests defined in the Dynamic Application can use the values defined in these fields. To use a value defined in one of these fields, the request must include the substitution character associated with that value. For example, suppose a Dynamic Application request includes the XML tag <high_value=%1>. Suppose you specified "100" in the **Embed Value [%1]** field in the credential aligned with that Dynamic Application. The request will be sent with the XML tag <high_value=100>.

- **Embedded Password [%P]**. Specifies a password value to include in a request. The value defined in this field is substituted in to the %P substitution character. The value will be encrypted in the request, will be masked in the **Credential Editor**, and will be stored in an encrypted form in the database.
- **Embed Value [%1]**. The value defined in this field is substituted in to the %1 substitution character.
- **Embed Value [%2]**. The value defined in this field is substituted in to the %2 substitution character.
- **Embed Value [%3]**. The value defined in this field is substituted in to the %3 substitution character.
- **Embed Value [%4]**. The value defined in this field is substituted in to the %4 substitution character.

HTTP Headers

- If you require custom HTTP headers to communicate with the SOAP server, you can build the custom header here. To add a header, click the **[Add Header]** button

cURL Options

- You can include the cURL command and various options in your credential. The list of cURL options lists all the options you can include in your credential. To include a cURL option in the credential, click the **Add CURL Option** drop-down and then select it from the list. You can then supply arguments in the blank text field to the right of the option.
- For more information on cURL commands, see the cURL manpage at <http://curl.haxx.se/docs/manpage.html>.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Defining an SSH/Key Credential

Secure Shell (SSH) is a network protocol that enables users to securely access a command-line shell on a remote computer or server over an unsecured network. SSH provides strong encryption and authentication capabilities, making it an ideal method for securely administering commands or transferring data between a client and server.

To make SSH even more secure, you can use SSH keys instead of a simple password to log in to a server. SSH keys consist of two long strings of characters, called a public/private key pair, that are much less susceptible than passwords are to brute force attacks. The public key is placed on the server you want to access, while the private

key resides on the client. When you use SSH to log in to the server from the client, the key pair is used to authenticate the session.

In SL1, some Dynamic Applications of type "Snippet" use SSH to communicate with a remote device. To use these Dynamic Applications, you must define an SSH credential. This credential specifies the hostname or IP address of the system you want to monitor, the port number used to access that system, and the private key used for authentication.

NOTE: Consult the documentation associated with the PowerPack that contains the Dynamic Application of type "Snippet" to find detailed directions on configuring the remote device and generating a private key for SL1 to use.

To create an SSH/Key credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create SSH/Key Credential*. The **Create Credential** modal page appears:

3. Supply values in the following fields:
 - **Name.** Name of the credential. Can be any combination of alphanumeric characters, up to 64 characters. This field is required.
 - **All Organizations.** Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the **What organization manages this service?** drop-down field to align the credential with those specific organizations. This field is required.

NOTE: To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

- **Timeout (ms)**. Time, in milliseconds, after which SL1 will stop trying to communicate with the device from which you want to retrieve data.
- **Hostname/IP**. Hostname or IP address of the device from which you want to retrieve data. This field is required.
 - You can include the variable **%D** in this field. SL1 will replace the variable with the IP address of the current device (device that is currently using the credential).
 - You can include the variable **%N** in this field. SL1 will replace the variable with hostname of the current device (device that is currently using the credential). If SL1 cannot determine the hostname, SL1 will replace the variable with the primary, management IP address for the current device.
- **Port**. Port number associated with the data you want to retrieve. This field is required.

NOTE: The default TCP port for SSH servers is 22.

- **Username**. Username for an SSH or user account on the device to be monitored.
- **Password**. Password for an SSH user account on the device to be monitored.
- **Private Key (PEM Format)**. Enter the SSH private key that you want SL1 to use, in PEM format.

NOTE: The **Private Key (PEM Format)** field is only required in the current SL1 user interface. The **Private Key (PEM Format)** field is not required if you are using the classic SL1 user interface to define a credential.

NOTE: The private key can have a maximum of 64 characters per line. Therefore, you cannot use keys in the OpenSSH format, because that format uses 70 characters per line. When you attempt to save the credential, SL1 will validate that the private key entered is in the correct format. You will be able to save the credential only if the private key is correctly formatted.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Defining an Aliyun Credential

To configure SL1 to monitor Aliyun's Alibaba Cloud service, you must first create an Aliyun credential. This credential allows the Dynamic Applications in the *Alibaba Cloud: Aliyun* PowerPack to connect with the Aliyun service.

SL1 includes an Aliyun credential type that you can use to connect with the Aliyun service during guided discovery. This credential type uses field names and terminology that are specific to the Aliyun service.

NOTE: Alternatively, you could monitor Aliyun using a generic SOAP/XML credential that does not include Aliyun-specific fields. For more information, see the *Monitoring Alibaba Cloud* manual.

To create an Aliyun-specific credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create Aliyun Credential*. The **Create Credential** modal page appears:

3. Supply values in the following fields:

- **Name.** Name of the credential. Can be any combination of alphanumeric characters, up to 64 characters. This field is required.
- **All Organizations.** Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the *What organization manages this service?* drop-down field to align the credential with those specific organizations. This field is required.

NOTE: To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

- **Timeout (ms).** Time, in milliseconds, after which SL1 will stop trying to communicate with the device from which you want to retrieve data.
- **Account access key for Aliyun service (minimum 20 characters).** The account access key ID for the Aliyun service. This field is required.
- **Account access key password for Aliyun service (minimum 20 characters).** The account access key password for the Aliyun service. This field is required.

Proxy Settings

If you use a proxy server in front of the Aliyun services you want to communicate with, enter values in these fields. Otherwise, you can skip these fields.

- **Proxy Hostname/IP.** The host name or IP address of the proxy server.
- **Proxy Port.** Port on the proxy server to which you will connect.
- **Proxy User.** Username to use to access the proxy server.
- **Proxy Password.** Password to use to access the proxy server.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Defining an AWS Credential

To use the Amazon Web Services (AWS) Dynamic Applications, you must configure a credential that allows SL1 to connect to the AWS REST API.

SL1 includes an AWS credential type that you can use to connect with the AWS service during guided discovery. This credential type uses field names and terminology that are specific to the AWS service.

NOTE: Alternatively, you could monitor AWS using a generic SOAP/XML credential that does not include AWS-specific fields. For more information, see the [Monitoring Amazon Web Services](#) manual.

To define an AWS-specific credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create AWS Credential*. The **Create Credential** modal page appears:

3. Supply values in the following fields:

- **Name**. Name of the credential. Can be any combination of alphanumeric characters, up to 64 characters. This field is required.
- **All Organizations**. Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the **What organization manages this service?** drop-down field to align the credential with those specific organizations. This field is required.

NOTE: To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

- **Timeout (ms)**. Time, in milliseconds, after which SL1 will stop trying to communicate with the device from which you want to retrieve data.
- **AWS Access Key ID (minimum 20 characters)**. The Access Key ID for an account on the AWS device to be monitored. This field is required.
- **Cloud Type**. Type of cloud that will be accessed with the credential. Select from a list of AWS clouds supported by SL1. Choices are:
 - *Standard*. Select this option if you want to connect to a standard AWS account.
 - *GovCloud*. Select this option if you want to connect to an AWS GovCloud account.
 - *Beijing*. Select this option if you want to connect to AWS regions in China.
- **AWS Secret Access Key (minimum 20 characters)**. The Secret Access Key for an account on the AWS device to be monitored. This field is required.

Proxy Settings

If you use a proxy server in front of the AWS devices you want to communicate with, enter values in these fields. Otherwise, you can skip these fields.

- **Proxy Hostname/IP.** The host name or IP address of the proxy server.
- **Proxy Port.** Port on the proxy server to which you will connect.
- **Proxy User.** Username to use to access the proxy server.
- **Proxy Password.** Password to use to access the proxy server.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Defining an AWS Assume Role Credential

To use the Amazon Web Services (AWS) Dynamic Applications, you must configure a credential that allows SL1 to connect to the AWS REST API.

SL1 includes an AWS Assume Role credential type that you can use to connect with the AWS service during guided discovery using the Assume Role discovery method. The Assume Role discovery method provides an automated mechanism to discover all your AWS accounts within an organization using a single IAM key. This credential type uses field names and terminology that are specific to the AWS service.

NOTE: For more information about monitoring AWS using Assume Role, see the [Monitoring Amazon Web Services](#) manual.

To define an AWS Assume Role credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create AWS Assume Role Credential*. The **Create Credential** modal page appears:

3. Supply values in the following fields:

- **Name.** Type a unique name for the credential. Can be any combination of alphanumeric characters, up to 64 characters.
- **All Organizations.** Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the **Select the organizations the credential belongs to** drop-down field to align the credential with those specific organizations.

NOTE: To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

- **Timeout (ms).** Type the time, in milliseconds, after which SL1 will stop trying to communicate with the device from which you want to retrieve data.
- **AWS Access Key ID.** Type the Access Key ID for an account on the AWS device to be monitored.
- **AWS Secret Access Key.** Type the Secret Access Key for an account on the AWS device to be monitored.
- **Cloud Type.** Select the AWS cloud type that will be accessed with the credential. This field is required. Choices are:
 - *Standard.* Select this option if you want to connect to a standard AWS account.
 - *GovCloud.* Select this option if you want to connect to an AWS GovCloud account.
 - *Beijing.* Select this option if you want to connect to AWS regions in China.
- **Assume Role.** Type the AWS Role you created in each account. The default name is "ScienceLogic-Monitor".
- **Assume Role Session.** Optional. The default value is "SL1".

- **Organization Creation.** Auto-creates an SL1 organization for accounts using AssumeRole. You can type one of the following options:
 - *NAME.* The name of the organization will contain the name of the user.
 - *ID.* The name of the organization will contain the ID of the user.
 - *ID:NAME.* The name of the organization will contain both the ID and name of the user, in that order.
 - *NAME:ID.* The name of the organization will contain both the name and ID of the user, in that order.
- **Configuration.** Select the method used to control what AWS devices are discovered and monitored. Choices are:
 - *Default.* The default AWS discovery method.
 - *AwsConfig.* Select this option if your accounts have the AWS Config service enabled.
 - *AwsCloudwatch.* Select this option to discover only the AWS regions that are reporting CloudWatch metrics.
- **Regions.** Type the AWS regions that you want to discover. For example, entering "ap-southeast-2, us-east-2" will discover two regions. If left blank, all regions will be discovered. The default value is "ALL".
- **Filter by Tags.** To discover AWS devices and filter them by tags, type the tag operation, tag key, and tag value, in the following format: <operation> # <tag name> # <tag value>. For example, if you want to filter by Tag Name, you would type the following:

```
Tags:equals#Name#Example
```

Valid operations include:

- equals
- notEquals
- contains
- notContains

You can chain together multiple filters separating them by a comma. For example:

```
Tags:equals#Name#Example,contains#Owner#Someone
```

- **Proxy Hostname/IP.** Type the host name or IP address of the proxy server.
- **Proxy Port.** Type the port number on the proxy server to which you will connect.
- **Proxy User.** Type the username to use to access the proxy server.
- **Proxy Password.** Type the password to use to access the proxy server.

NOTE: If you use a proxy server in front of the AWS devices you want to communicate with, enter values in the proxy fields. Otherwise, you can skip these fields.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Defining an AWS EC2 Credential

To use the Amazon Web Services (AWS) Dynamic Applications, you must configure a credential that allows SL1 to connect to the AWS REST API.

SL1 includes an AWS EC2 credential type that you can use to connect with the AWS service during guided discovery when your Data Collectors are EC2 instances. This credential type uses field names and terminology that are specific to the AWS service.

NOTE: For more information about monitoring AWS accounts within an organization when your Data Collectors are EC2 instances, see the [Monitoring Amazon Web Services](#) manual.

To define an EC2 credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create AWS EC2 Credential*. The **Create Credential** modal page appears:

3. Supply values in the following fields:

- **Name.** Type a unique name for the credential. Can be any combination of alphanumeric characters, up to 64 characters.
- **All Organizations.** Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the **Select the organizations the credential belongs to** drop-down field to align the credential with those specific organizations.

NOTE: To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

- **Timeout (ms).** Type the time, in milliseconds, after which SL1 will stop trying to communicate with the device from which you want to retrieve data.
- **Cloud Type.** Select the AWS cloud type that will be accessed with the credential. This field is required. Choices are:
 - *Standard.* Select this option if you want to connect to a standard AWS account.
 - *GovCloud.* Select this option if you want to connect to an AWS GovCloud account.
 - *Beijing.* Select this option if you want to connect to AWS regions in China.
- **Organization Arn.** Type the Amazon Resource Name (ARN) for the Assume Role. This is the ARN of the role created in the master billing account.
- **Assume Role.** Type the AWS Role you created in each account. The default name is "ScienceLogic-Monitor".
- **Assume Role Session.** Optional. The default value is "SL1".

- **Organization Creation.** Auto-creates an SL1 organization for accounts using AssumeRole. You can type one of the following options:
 - *NAME.* The name of the organization will contain the name of the user.
 - *ID.* The name of the organization will contain the ID of the user.
 - *ID:NAME.* The name of the organization will contain both the ID and name of the user, in that order.
 - *NAME:ID.* The name of the organization will contain both the name and ID of the user, in that order.
- **Configuration.** Select the type of method used to control what AWS devices are discovered and monitored. Choices are:
 - *Default.* The default AWS discovery method.
 - *AwsConfig.* Select this option if your accounts have the AWS Config service enabled.
 - *AwsCloudwatch.* Select this option to discover only the AWS regions that are reporting CloudWatch metrics.
- **Regions.** Type the AWS regions that you want to discover. For example, entering "ap-southeast-2, us-east-2" will discover two regions. If left blank, all regions will be discovered. The default value is "ALL".
- **Filter by Tags.** To discover AWS devices and filter them by tags, type the tag operation, tag key, and tag value, in the following format: <operation> # <tag name> # <tag value>. For example, if you want to filter by Tag Name, you would type the following:

```
Tags:equals#Name#Example
```

Valid operations include:

- equals
- notEquals
- contains
- notContains

You can chain together multiple filters separating them by a comma. For example:

```
Tags:equals#Name#Example,contains#Owner#Someone
```

- **Proxy Hostname/IP.** Type the host name or IP address of the proxy server.
- **Proxy Port.** Type the port number on the proxy server to which you will connect.
- **Proxy User.** Type the username to use to access the proxy server.
- **Proxy Password.** Type the password to use to access the proxy server.

NOTE: If you use a proxy server in front of the AWS devices you want to communicate with, enter values in the proxy fields. Otherwise, you can skip these fields.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Defining an AWS IAM Credential

To use the Amazon Web Services (AWS) Dynamic Applications, you must configure a credential that allows SL1 to connect to the AWS REST API.

You can use IAM policies in AWS to restrict which regions and services SL1 will monitor. To do this, you can create another IAM policy and apply that along with the SL1 monitoring policy to the applicable user or role(s).

SL1 includes an AWS IAM credential type that you can use to connect with the AWS service during guided discovery using the IAM discovery method. This credential type uses field names and terminology that are specific to the AWS service.

NOTE: For more information about monitoring AWS using IAM permissions, see the [Monitoring Amazon Web Services](#) manual.

To define an AWS IAM credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create AWS IAM Credential*. The **Create Credential** modal page appears:

3. Supply values in the following fields:

- **Name.** Type a unique name for the credential. Can be any combination of alphanumeric characters, up to 64 characters.
- **All Organizations.** Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the **Select the organizations the credential belongs to** drop-down field to align the credential with those specific organizations.

NOTE: To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

- **Timeout (ms).** Type the time, in milliseconds, after which SL1 will stop trying to communicate with the device from which you want to retrieve data.
- **AWS Access Key ID.** Type the Access Key ID for an account on the AWS device to be monitored.
- **AWS Secret Access Key.** Type the Secret Access Key for an account on the AWS device to be monitored.
- **Cloud Type.** Select the AWS cloud type that will be accessed with the credential. This field is required. Choices are:
 - *Standard.* Select this option if you want to connect to a standard AWS account.
 - *GovCloud.* Select this option if you want to connect to an AWS GovCloud account.
 - *Beijing.* Select this option if you want to connect to AWS regions in China.
- **Configuration.** Select the method used to control what AWS devices are discovered and monitored. Choices are:

- *Default*. The default AWS discovery method.
 - *AwsConfig*. Select this option if your accounts have the AWS Config service enabled.
 - *AwsCloudwatch*. Select this option to discover only the AWS regions that are reporting CloudWatch metrics.
- **Regions**. Type the AWS regions that you want to discover. For example, entering "ap-southeast-2, us-east-2" will discover two regions. If left blank, all regions will be discovered. The default value is "ALL".
 - **Filter by Tags**. To discover AWS devices and filter them by tags, type the tag operation, tag key, and tag value, in the following format: <operation>#<tag name>#<tag value>. For example, if you want to filter by Tag Name, you would type the following:

```
Tags:equals#Name#Example
```

Valid operations include:

- equals
- notEquals
- contains
- notContains

You can chain together multiple filters separating them by a comma. For example:

```
Tags:equals#Name#Example,contains#Owner#Someone
```

- **Proxy Hostname/IP**. Type the host name or IP address of the proxy server.
- **Proxy Port**. Type the port number on the proxy server to which you will connect.
- **Proxy User**. Type the username to use to access the proxy server.
- **Proxy Password**. Type the password to use to access the proxy server.

NOTE: If you use a proxy server in front of the AWS devices you want to communicate with, enter values in the proxy fields. Otherwise, you can skip these fields.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Defining an Azure Credential

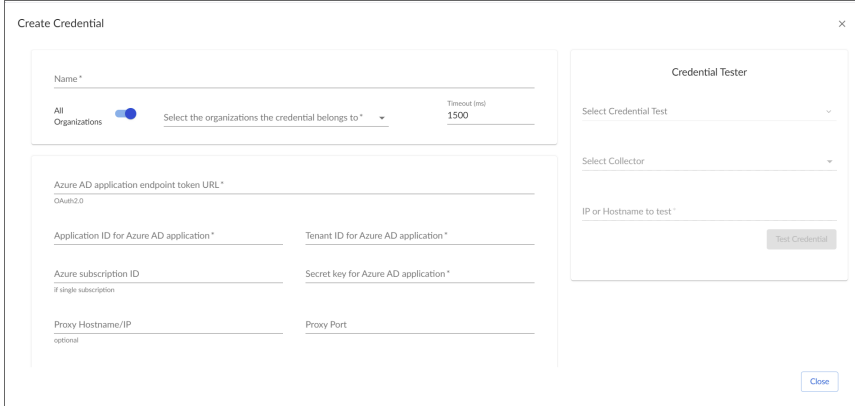
To configure SL1 to monitor Microsoft Azure, you must first create an Azure credential. This credential allows the Dynamic Applications in the *Microsoft: AzurePowerPack* to connect with the Azure Active Directory Application.

SL1 includes an Azure credential type that you can use to connect with the Azure service during guided discovery. This credential type uses field names and terminology that are specific to the Azure service.

NOTE: Alternatively, you could monitor Azure using a generic SOAP/XML credential that does not include Azure-specific fields. For more information, see the **Monitoring Microsoft Azure** manual.

To define an Azure-specific credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create Azure Credential*. The **Create Credential** modal page appears:



The screenshot shows the 'Create Credential' modal window. It has a title bar with 'Create Credential' and a close button. The main area is divided into two columns. The left column contains several input fields: 'Name*' (required), 'All Organizations' (toggle, currently on), 'Select the organizations the credential belongs to*' (dropdown), 'Timeout (ms)' (set to 1500), 'Azure AD application endpoint token URL*' (with 'OAuth2.0' below it), 'Application ID for Azure AD application*' and 'Tenant ID for Azure AD application*' (both required), 'Azure subscription ID' (with 'if single subscription' below it) and 'Secret key for Azure AD application*' (both required), and 'Proxy Hostname/IP' (optional) and 'Proxy Port' (optional). The right column is titled 'Credential Tester' and contains 'Select Credential Test' (dropdown), 'Select Collector' (dropdown), and 'IP or Hostname to test*' (input field). A 'Test Credential' button is located below the 'IP or Hostname to test*' field. A 'Close' button is at the bottom right of the modal.

3. Supply values in the following fields:
 - **Name.** Name of the credential. Can be any combination of alphanumeric characters, up to 64 characters. This field is required.
 - **All Organizations.** Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the **What organization manages this service?** drop-down field to align the credential with those specific organizations. This field is required.

NOTE: To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

- **Timeout (ms).** Time, in milliseconds, after which SL1 will stop trying to communicate with the device from which you want to retrieve data.
- **Azure AD application endpoint token URL (OAuth2.0).** The AD application endpoint token URL for the Azure Active Directory application. This field is required.
- **Application ID for Azure AD application.** The Application ID for the Azure Active Directory application. This field is required.
- **Tenant ID for Azure AD application.** The Tenant ID for the Azure Active Directory application. This field is required.
- **Azure subscription ID (if single subscription).** The subscription ID for the Azure Active Directory application. This field is required only if you are monitoring a single Azure subscription.
- **Secret key for Azure AD application.** The secret key for the Azure Active Directory application. This field is required.

Proxy Settings

If you use a proxy server in front of the Azure Active Directory applications you want to communicate with, enter values in these fields. Otherwise, you can skip these fields.

- **Proxy Hostname/IP.** The host name or IP address of the proxy server.
- **Proxy Port.** Port on the proxy server to which you will connect.
- **Proxy User.** Username to use to access the proxy server.
- **Proxy Password.** Password to use to access the proxy server.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Defining a Citrix XenServer Credential

To use the Dynamic Applications in the *Citrix: Xen PowerPack*, you must first define a credential in SL1 that enables SL1 to communicate with your XenCenter system and XenServer devices.

SL1 includes a Citrix Xen credential type that you can use to connect with your XenServer devices during guided discovery. This credential type uses field names and terminology that are specific to XenServer.

NOTE: Alternatively, you could monitor Citrix XenServer using a generic Basic/Snippet credential that does not include XenServer-specific fields. For more information, see the **Monitoring Citrix XenCenter** manual.

To define a XenServer-specific credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create Citrix Xen Credential*. The **Create Credential** modal page appears:

3. Supply values in the following fields:

- **Name**. Name of the credential. Can be any combination of alphanumeric characters, up to 64 characters. This field is required.
- **All Organizations**. Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the **What organization manages this service?** drop-down field to align the credential with those specific organizations. This field is required.

NOTE: To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

- **Timeout (ms)**. Time, in milliseconds, after which SL1 will stop trying to communicate with the device from which you want to retrieve data.
- **XenServer username**. The username for an account on the XenServer device to be monitored.
- **XenServer password**. The password for an account on the XenServer device to be monitored.
- **Hostname/IP of the Xen server**. The Hostname or IP address of the XenServer device from which you want to retrieve data. This field is required.
- **Port**. The port number associated with the data you want to receive. This field is required.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Defining an IBM Cloud Credential

To configure SL1 to monitor an IBM Cloud/SoftLayer account, you must create an IBM credential. This credential allows the Dynamic Applications in the *SoftLayer: Cloud PowerPack* to communicate with your IBM Cloud/SoftLayer account.

SL1 includes an IBM credential type that you can use to connect with your IBM Cloud/SoftLayer service during guided discovery. This credential type uses field names and terminology that are specific to IBM Cloud/SoftLayer.

NOTE: Alternatively, you could monitor IBM Cloud/SoftLayer using a generic Basic/Snippet credential that does not include IBM-specific fields. For more information, see the *Monitoring SoftLayer* manual.

To define an IBM-specific credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create IBM Credential*. The **Create Credential** modal page appears:

3. Supply values in the following fields:

- **Name**. Name of the credential. Can be any combination of alphanumeric characters, up to 64 characters. This field is required.
- **All Organizations**. Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the *What organization manages this service?* drop-down field to align the credential with those specific organizations. This field is required.

NOTE: To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

- **Timeout (ms).** Time, in milliseconds, after which SL1 will stop trying to communicate with the device from which you want to retrieve data.
- **IBM/SoftLayer account username.** The IBM/SoftLayer username for an account on the IBM device to be monitored.
- **IBM/SoftLayer account API key.** The IBM/SoftLayer API key for an account on the IBM device to be monitored.
- **Hostname/IP.** The Hostname or IP address of the IBM device from which you want to retrieve data. The default value is %D. SL1 will replace the variable with the IP address of the device that is currently using the credential. This field is required.
- **Port.** The port number associated with the data you want to receive. This field is required.

NOTE: The default TCP port for IBM devices is 80.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Defining an S3 Backup Credential

You can use an S3 storage service to store configuration backups for SL1. To do so, you will need to create a credential that enables SL1 to connect to the S3 service. SL1 includes an **S3 Backup** credential type, which uses field names and terminology specific to S3 services, that you can use to connect with your S3 service.

NOTE: SL1 supports the use of Amazon Web Services (AWS) or MinIO for S3 backup storage.

NOTE: For more information about configuration backups, see the **System Administration** manual.

To define an S3 backup credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create S3 Backup Credential*. The **Create Credential** modal page appears:

3. Supply values in the following fields:

- **Name.** Type a unique name for the credential. Can be any combination of alphanumeric characters, up to 64 characters.
- **All Organizations.** Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the **What organization manages this service?** drop-down field to align the credential with those specific organizations.

NOTE: To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

- **Timeout (ms).** Type the time, in milliseconds, after which SL1 will stop trying to communicate with the S3 storage service.
- **Provider.** Select the S3 storage provider you want to use to store the backup. Choices are *Amazon Web Services (AWS) S3* and *Minio Object Storage*.
- **Access Key ID.** Type the Access Key ID for the S3 account on which you want to store the backup.
- **Secret Access Key.** Type the Secret Access Key for the S3 account on which you want to store the backup.
- **Endpoint.** Type the URL of the S3 endpoint. The endpoint URL should not include the bucket name.
- **Region.** Select the region of the S3 endpoint.
- **Bucket.** Type the name of the S3 bucket on which you want to store the backup.
- **Encryption Password.** Type the encryption password for the backup file.
- **Encryption Salt.** Type the encryption salt used to safeguard the backup file encryption password.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Defining a VMware Credential

SL1 includes a VMware credential type that you can use to connect with the VMware web service during guided discovery. This credential type uses field names and terminology that are specific to VMware vSphere.

NOTE: Alternatively, you could monitor VMware using a generic SOAP/XML credential that does not include VMware-specific fields. For more information, see the [Monitoring VMware Systems](#) manual.

To define a VMware-specific credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **[Create New]** button and then select *Create VMware Credential*. The **Create Credential** modal page appears:

The screenshot shows the 'Edit Credential' modal page. It has a title bar with a close button. The main content is split into two columns. The left column has a 'Name *' field with the value 'VMware web service'. Below it is a toggle for 'All Organizations' and a dropdown for 'Select the organizations the credential belongs to *' with 'System' selected. To the right is a 'Timeout (ms)' field with '1500'. Below these are fields for 'VMware web service username *' (sample), 'VMware web service password *' (masked), and 'URL *' (https://43:%D). A blue 'Save & Test' button is at the bottom right of this column. The right column is titled 'Credential Tester' and has a dropdown for 'Select Credential Test', a dropdown for 'Select Collector' with the value 'CUG | sl1aio1: 10.128.68.26', and a field for 'IP or Hostname to test *'. A 'Test Credential' button is at the bottom right of this column. A 'Save & Close' button is at the bottom right of the modal.

3. Supply values in the following fields:

- **Name.** Name of the credential. Can be any combination of alphanumeric characters, up to 64 characters. This field is required.
- **All Organizations.** Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the **What organization manages this service?** drop-down field to align the credential with those specific organizations. This field is required.

NOTE: To learn more about aligning credentials and organizations, see [Aligning Organizations With a Credential](#).

- **Timeout (ms).** Time, in milliseconds, after which SL1 will stop trying to communicate with the device from which you want to retrieve data.
- **VMware web service username.** The VMware username for the VMware web service account. This field is required.
- **VMware web service password.** The VMware password for the VMware web service account. This field is required.
- **URL.** The URL of the VMware web service that you want to monitor. This field is required.
 - You can include the variable **%D** in this field. SL1 will replace the variable with the IP address of the device that is currently using the credential. For component devices, SL1 will replace **%D** with the IP address of the root device.
 - You can include the variable **%N** in this field. SL1 will replace the variable with the hostname of the device that is currently using the credential. If SL1 cannot determine the hostname, SL1 will replace the variable with the primary management IP address for the current device.

TIP: Click the **URL** field for a set of options for formatting the URL: `http(s)://Host:Port/Path | %D = Aligned Device IP | %N = Aligned Device Name`. For example: `https://%D:433/<sample path>`.

4. Click **[Save & Close]**.

NOTE: If you would like to test your credential using the Credential Tester panel, click **[Save & Test]**. For detailed instructions on using the Credential Tester panel, see the [Using the Credential Tester Panel](#) section.

Testing a Credential

You can test a credential using a predefined credential test. For more information about creating and managing credential tests, see the chapter on [Managing Credential Tests](#).

Using the Credential Tester Panel

When defining or editing a credential in SL1, you can test the credential using the **Credential Tester** panel.

To test a credential using the **Credential Tester** panel:

1. From the **Credentials** page (Manage > Credentials) or from the **Credential Selection** page during **guided** or **unguided** discovery, click **[Create New]** to create a new credential or click the **Actions** icon (⋮) of a credential that you want to test and then select *Edit/Test*.
2. While **defining** or **editing** the credential, supply values in the required fields. Required fields may vary depending on the type of credential you create.
3. Click the **[Save & Test]** button. This activates the **Credential Tester** fields.
4. In the **Credential Tester** panel, supply values in the following fields:
 - **Select Credential Test.** Select a credential test to run. This drop-down list includes the **ScienceLogic Default Credential Tests**, credential tests included in any PowerPacks that have been optionally installed on your system, and credential tests that users have created on your system.
 - **Select Collector.** Select the All-In-One Appliance or Data Collector that will run the test.
 - **IP or Hostname to test.** Type a hostname or IP address that will be used during the test. For example, if you are testing an SNMP credential, the hostname/IP address you supply will be used to perform a test SNMP request.
5. Click the **[Test Credential]** button to run the credential test. The **Testing Credential** window appears.

The **Testing Credential** window displays a log entry for each step in the credential test. The steps performed are different for each credential test. The log entry for each step includes the following information:

- **Step.** The name of the step.
- **Description.** A description of the action performed during the step.
- **Log Message.** The result of the step for this execution of the credential test.
- **Status.** Whether the result of this step indicates the credential and/or the network environment is configured correctly (Passed) or incorrectly (Failed).
- **Step Tip.** Mouse over the question mark icon (?) to display the tip text. The tip text recommends what to do to change the credential and/or the network environment if the step has a status of "Failed".

Specifying Credentials For Discovery and Devices

Discovery is the process by which SL1 discovers what types of hardware and applications exist on the network and then retrieves data from the discovered hardware and applications.

Before running discovery, you must:

- Determine the SNMP credentials for the devices and applications in your network. Define correlating credentials in SL1, to allow discovery to retrieve as much information as possible.
- If you want SL1 to immediately start collecting data from devices using Dynamic Applications, you should also define any additional credentials required for those Dynamic Applications. For example, if you want SL1 to immediately start monitoring all MySQL databases in your network, you should define credentials that allow SL1 to communicate with each MySQL database in your network. During discovery, SL1 will determine which devices can be monitored with a Dynamic Application for MySQL. After discovery, SL1 will use the database credential to collect data about each MySQL database in your network.

Use the previous sections to define credentials for your network.

When you run discovery, you must specify one or more of these credentials to use. The more credentials you align with a discovery session, the more access SL1 will have to devices and their data during discovery.


- For information about specifying credentials during guided discovery, see the section on [Adding Devices Using Universal or Guided Discovery](#).
- For information about specifying credentials during unguided discovery, see the section on [Adding Devices Using Unguided Discovery](#).
- For information about specifying credentials during initial classic discovery, see the section on [Specifying Credentials During Initial Classic Discovery](#).

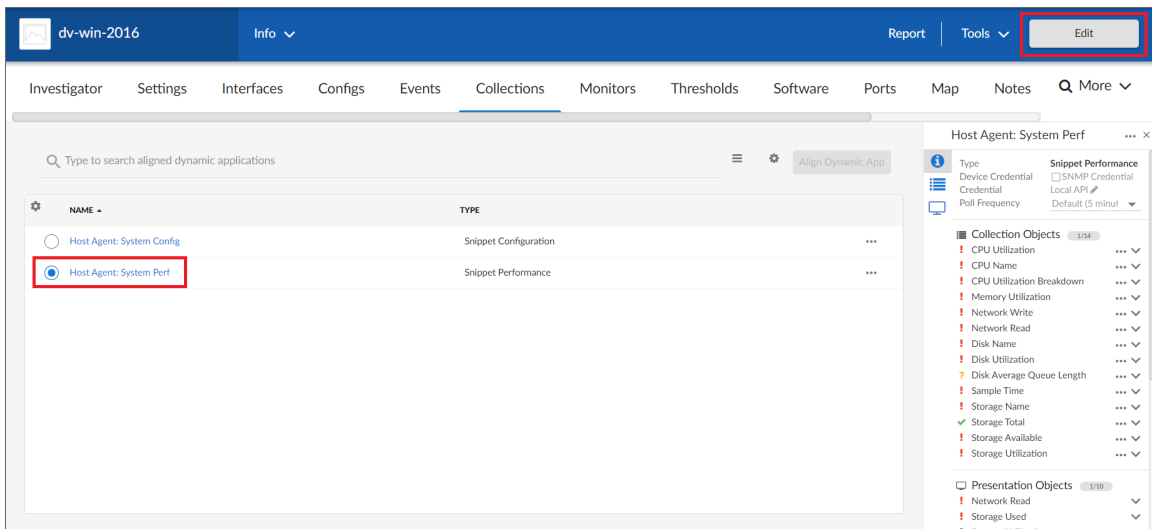
Specifying Credentials for a Device/Dynamic Application Pair

After a device has been discovered by SL1 and one or more Dynamic Applications have been aligned with the device, you can manually assign the credential to use for each Dynamic Application.

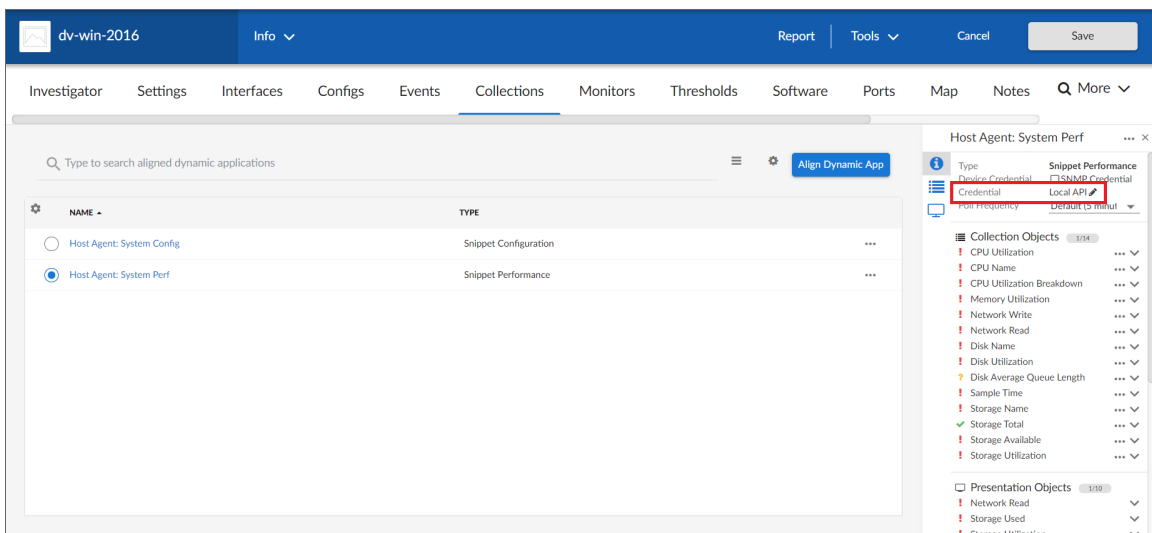
The manually assigned credential will be used by SL1 only for this specific Dynamic Application associated with this specific device. For all other devices, SL1 will use the default credential associated with each device, or will use the credential defined on the **Collections** tab for each device.

To manually associate a credential with a Dynamic Application aligned to a device:

1. Go to the **Devices** page by clicking the **Devices** icon (.
2. Find the device for which you want to define a credential. Click its hyperlink in the **Device Name** column.
3. On the **Device Investigator**, click the **[Collections]** tab.
4. Find the Dynamic Application for which you want to define a credential. Click its radio button and then click **[Edit]**.



5. In the information pane for the Dynamic Application, click the **Edit** icon () next to the **Credential** field.



6. From the **Choose Credential** modal page, select the credential that you want to align to the Dynamic Application, and then click the **[Select]** button.

NOTE: Your organization membership(s) might affect the list of credentials you can see on the **Choose Credential** modal page. To learn more about credentials and organizations, see the section [Aligning Organizations With a Credential](#).

NOTE: If this Dynamic Application has already been aligned with a credential to which you do not have access, the **Credential** field will display the value *Restricted Credential*. If you align the Dynamic Application with a different credential, you will not be able to re-align the device with the *Restricted Credential*.

7. The selected Dynamic Application will now use the manually selected credential when collecting data from this device. You should see your change reflected in the **Credential** field in the information pane for the Dynamic Application on the **[Collections]** tab.

For more information about the **[Collections]** tab, see the chapter about **Using the Devices Page** in the **Device Management** manual.

Aligning One or More Organizations With a Credential

To support multi-tenancy, SL1 allows you to align each credential with one, multiple or all organizations in SL1. You can also align a credential with no organizations.

When you align an organization with a credential, you control who can view details about the credential, who can view the name of the credential, and who can apply the credential in SL1.

NOTE: When you align an organization with a credential, you are restricting only the users who can view and assign the credential. You are not restricting the devices and actions that can be associated with the credential. For example, you can align a credential only with the organization "Operations" but assign the credential to a device in the "Finance" organization.

By default, newly created credentials are aligned to all organizations. However, when defining a new credential, you can opt to align the credential with one or more specific organizations rather than all credentials. To do so, toggle off the **All Organizations** field on the **Create Credential** modal and then select one or more specific organizations from the **What organization manages this service?** drop-down field.

Credentials that are aligned with an organization have the following behavior:

- For each credential that is aligned with an organization, only administrators and users who are members of the aligned organization can view the credential on the **Credentials** page.
- When aligning credentials to devices or Dynamic Applications, non-administrator users can view and align only those credentials that are assigned to organizations common to both the user and the device's collector group, plus those credentials that are assigned to all organizations or otherwise required for that collector group.
- Users can change the organization setting of only those credentials that are not currently aligned to any devices.

- In any field or column that displays the name of the credential, users who are not members of the aligned organization will not see the credential name. Instead, these users will see either a dash character (-) or the text "Restricted Credential".
- In any list from which users can select a credential, users who are not members of the aligned organization will not see the credential as an entry in the list.
- In any page where the credential has already been assigned, users who are not members of the aligned organization will see only the name "Restricted Credential".
- In any page where the credential has already been assigned, users who are not members of the aligned organization can save the page and maintain the credential. The credential will still appear to that user as "Restricted Credential".
- In any page where the credential has already been assigned, users who are not members of the aligned organization can change the credential to a credential aligned with their organization(s). However, those users cannot change the credential again and re-assign the "Restricted Credential". The entry for "Restricted Credential" is removed from the list of possible credentials.
- If you attempt to run a discovery session where the devices, collector group, and credentials do not all belong to the same organization, you will receive an error and will not be able to save or execute the discovery session.

To understand the behavior of a credential aligned with an organization, consider the following example:

- Suppose you have a user account of type "Administrator". Suppose you create an SNMP credential called "ops_cred". Suppose you align that credential with the organization "Operations".
- In the **Credentials** page, only administrators and users who are members of the organization "Operations" will be able to see the credential "ops_cred" in SL1.
- In SL1, in any field or column that displays the name of the credential, users who are not members of the organization "Operations" will not see the "ops_cred" name displayed. Instead, these users will see either a dash character (-) or the text "Restricted Credential".
- In SL1, in any list from which users can select a credential, users who are not members of the organization "Operations" will not see the "ops_cred" credential as an entry in the field.
- In SL1, in any page where the credential "ops_cred" has already been assigned, users who are not members of the organization "Operations" will see only the name "Restricted Credential".
- In SL1, in any page where the credential "ops_cred" has already been assigned, users who are not members of the organization "Operations" can save the page and maintain the "ops_cred" credential. The credential will still appear to that user as "Restricted Credential".
- In SL1, in any page where the credential "ops_cred" has already been assigned, users who are not members of the organization "Operations" can change the credential to a credential aligned with their organization. However, that user cannot change the credential again and re-assign the "Restricted Credential". The entry for "Restricted Credential" is removed from the list of possible credentials.

Editing a Credential

The **Credentials** page allows you to edit credentials from SL1.

To edit a credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **Actions** icon (☰) of the credential you want to edit and then select *Edit/Test*. The **Edit Credential** page appears.
3. After editing the fields in the **Edit Credential** modal page, click the **[Save & Close]** button.

NOTE: Editing a credential and saving it overwrites the existing credential; it does not create a new credential. To create a new credential from an existing credential, see the section on [Duplicating a Credential](#).

NOTE: When editing a credential, the current password displays as a masked string of characters. If you make any changes to this field, SL1 completely removes the previous credential password and replaces it with what you type; therefore, you must either completely replace the password or leave it unchanged. If you accidentally type anything in the password field but do not actually want to change the password, you should close the **Edit Credential** modal page without saving your changes to avoid overwriting the current password.

Duplicating a Credential

The **Credentials** page allows you to create a duplicate of an existing credential. When you do so, SL1 copies all of the original credential's values into the new credential. You can then edit the new credential to make changes as needed while still retaining the values that you want to keep from the original credential.

To duplicate a credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **Actions** icon (☰) of the credential you want to duplicate and then select *Duplicate*.
3. A copy of the credential appears on the **Credentials** page.
4. To edit the copied credential, click its **Actions** icon (☰) and then select *Edit/Test*. The **Edit Credential** modal appears.
5. After editing the fields in the **Edit Credential** modal, click the **[Save & Close]** button.

Deleting a Credential

The **Credentials** page allows you to delete one or more credentials from SL1.

NOTE: SL1 displays an error message if you attempt to delete a credential that is currently assigned to one or more devices.

To delete a credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Click the **Actions** icon (☰) of the credential you want to delete and then select *Delete*.
3. The credential is deleted.

Using Credentials in the Classic SL1 User Interface

This section describes how to view, define, and manage credentials using the **Credential Management** page in the classic SL1 user interface.

Viewing Information about Credentials in the Classic SL1 User Interface

The **Credential Management** page allows you to view a list of all ScienceLogic credentials. From this page, you can also create new credentials and editing existing credentials.

TIP: To sort the list of credentials, click on a column heading. The list will be sorted by the column value, in ascending order. To sort by descending order, click the column heading again. The **Last Edited** column sorts by descending order on the first click; to sort by ascending order, click the column heading again.

For each credential, the page displays:

- **Profile Name.** Name of the credential.
- **Organization.** If you have an account of type "User" and are a member of only one ScienceLogic organization, this field will not appear in the **Credential Management** page. The **Credential Management** page will display only credentials that are aligned with your organization. For all other users, this column specifies the organization(s) aligned with the credential. Possible values are all orgs, multiple orgs, a single organization name, or none. For details, see the section [Aligning One or More Organizations with a Credential](#).
- **RO Use.** Specifies the number of devices that SL1 can retrieve read-only information from, using the credential.
- **RW Use.** Specifies the number of devices that SL1 can both read from and write to, using the credential.
- **DA Use.** Specifies the number of Dynamic Applications aligned with this credential.
- **Type.** Type of credential. Possible types are SNMP, Database, SOAP/XML, LDAP/AD, Basic/Snippet, SSH/Key, and PowerShell.
- **Credential User.** Username associated with the credential.
- **Host.** Hostname or IP address that SL1 will use the credential to communicate with.
- **Port.** Port used by the credential to communicate with the external device or application.
- **Timeout.** Time, in milliseconds, after which SL1 will stop trying to communicate with the external device or application.

- **ID**. Unique numeric ID, automatically assigned by SL1 to each credential.
- **Last Edited**. Date and time the credential was created or last edited.
- **Edited By**. ScienceLogic user who created or last edited the credential.

Filtering the List of Credentials in the Classic SL1 User Interface

To filter the list of credentials in the **Credential Management** page, use the search fields at the top of each column. The search fields are find-as-you-type filters; as you type, the page is filtered to match the text in the search field, including partial matches. Text matches are not case-sensitive. Additionally, you can use the following special characters in each filter:

- , (comma). Specifies an "or" operation. For example:
 "dell, micro" would match all values that contain the string "dell" OR the string "micro".
- & (ampersand). Specifies an "and" operation. For example:
 "dell & micro" would match all values that contain the string "dell" AND the string "micro".
- ! (exclamation mark). Specifies a "not" operation. For example:
 "!dell" would match all values that do not contain the string "dell".
- ^ (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.

Defining One or More SNMP Credentials for Initial Discovery in the Classic SL1 User Interface

Before running discovery, you must first define credentials for the devices and applications in the network to be managed. You must either define or note the credentials on the device to be managed, and then you must define matching credentials in SL1.

To create credentials for initial discovery, you must first:

1. Determine which devices or IP ranges you want to discover.
2. Determine which of these devices support SNMP.
3. Determine the SNMP community string or SNMP credentials for each device that supports SNMP.

NOTE: If you do not know which devices in your network support SNMP, consult your system administrator. In some cases, you might also need to consult your system administrator about enabling SNMP, and defining SNMP community strings or SNMP credentials on these devices.

4. In SL1, define one or more SNMP credentials to use during discovery. These credentials should match those SNMP community strings and SNMP credentials that already exist in your network.
5. Initially, discovery uses only SNMP credentials. However, when SL1 collects data specified in Dynamic Applications, SL1 can use other types of credentials.

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.

Defining Credentials in the Classic SL1 User Interface

To define a credential in the classic SL1 user interface:

1. Collect the information you need to create each credential (usually username and password).
2. Go to the **Credential Management** page (System > Manage > Credentials).
3. In the **Credential Management** page, click the **[Create]** menu. Select the type of credential you want to create. Your choices are:
 - [SNMP Credential](#)
 - [Database Credential](#)
 - [SOAP/XML Host Credential](#)
 - [LDAP/AD Credential](#)
 - [Basic/Snippet Credential](#)
 - [SSH/Key Credential](#)
 - [PowerShell Credential](#)
4. The **Credential Editor** modal page appears. In this page, you can define the new credential. The following sections explain how to create each type of credential.
5. Click the **[Save]** button to save the new credential.

Defining an SNMP Credential in the Classic SL1 User Interface

SNMP Credentials allow SL1 to access SNMP data on a managed device. SL1 uses SNMP credentials to perform discovery, run auto-discovery, and gather information from SNMP Dynamic Applications.

To create an SNMP credential:

1. Go to the **Credential Management** page (System > Manage > Credentials).
2. Click the **[Actions]** button and select *Create SNMP Credential*. The **Credential Editor** page appears.
3. Supply values in the following fields:
 - **Profile Name**. Name of the credential. Can be any combination of alphanumeric characters. This field is required.
 - **SNMP Version**. SNMP version. Choices are *SNMP V1*, *SNMP V2*, and *SNMP V3*. The default value is *SNMP V2*.
 - **Port**. The port SL1 will use to communicate with the external device or application. The default value is *161*. This field is required.
 - **Timeout (ms)**. Time, in milliseconds, after which SL1 will stop trying to communicate with the SNMP device. The default value is *1500*.
 - **Retries**. Number of times SL1 will try to authenticate and communicate with the external device. The default value is *1*.

SNMP V1/V2 Settings

These fields appear if you selected *SNMP V1* or *SNMP V2* in the **SNMP Version** field. The fields are inactive if you selected *SNMP V3*.

- **SNMP Community (Read-Only)**. The SNMP community string (password) required for read-only access of SNMP data on the remote device or application. For *SNMP V1* and *SNMP V2* credentials, you must supply a community string, either in this field or in the **SNMP Community (Read/Write)** field.
- **SNMP Community (Read/Write)**. The SNMP community string (password) required for read and write access of SNMP data on the remote device or application. For *SNMP V1* and *SNMP V2* credentials, you must supply a community string, either in this field or in the **SNMP Community (Read Only)** field.

SNMP V3 Settings

These fields appear if you selected *SNMP V3* in the **SNMP Version** field. These fields are inactive if you selected *SNMP V1* or *SNMP V2*.

- **Security Name**. Name for SNMP authentication. This field is required.
- **Security Passphrase**. Password to authenticate the credential. This value must contain at least 8 characters. This value is required if you use a **Security Level** that includes authentication.
- **Authentication Protocol**. Select an authentication algorithm for the credential. This field is required. Choices are:
 - *MD5*. This is the default value.
 - *SHA*
 - *SHA-224*
 - *SHA-256*
 - *SHA-384*
 - *SHA-512*

NOTE: The *SHA* option is *SHA-128*.

- **Security Level**. Specifies the combination of security features for the credentials. This field is required. Choices are:
 - *No Authentication / No Encryption*.
 - *Authentication Only*. This is the default value.
 - *Authentication and Encryption*.

- **SNMP v3 Engine ID.** The unique engine ID for the SNMP agent you want to communicate with. (SNMPv3 authentication and encryption keys are generated based on the associated passwords and the engine ID.) This field is optional.
- **Context Name.** A context is a mechanism within SNMPv3 (and AgentX) that allows you to use parallel versions of the same MIB objects. For example, one version of a MIB might be associated with SNMP Version 2 and another version of the same MIB might be associated with SNMP Version 3. For SNMP Version 3, specify the context name in this field. This field is optional.
- **Privacy Protocol.** The privacy service encryption and decryption algorithm. This field is required. Choices are:
 - *DES.* This is the default value.
 - *AES-128*
 - *AES-192*
 - *AES-256*
 - *AES-256-C.* This option is for discovering Cisco devices only.
- **Privacy Protocol Passphrase.** Privacy password for the credential. This field is optional.

4. Click the **[Save]** button to save the new SNMP credential.
5. Repeat steps 1-4 for each SNMP-enabled device in your network that you want to monitor with SL1.

NOTE: When you define an SNMP Credential, SL1 automatically aligns the credential with all organizations of which you are a member.

Defining a Database Credential in the Classic SL1 User Interface

Database Credentials allow SL1 to access data on a database on a managed device. SL1 uses database credentials when collecting data for Database Dynamic Applications.

To define a database credential:

1. Collect the information you need to create each credential (usually username and password).
2. Go to the **Credential Management** page (System > Manage > Credentials).
3. In the **Credential Management** page, click the **[Actions]** menu. Select **Create Database Credential**.
4. The **Credential Editor** modal page appears. In this page, you can define the new database credential. To define the new credential, supply values in the following fields:

Basic Settings

- **Profile Name.** Name of the credential. Can be any combination of alphanumeric characters. This field is required.
- **DB Type.** Type of database that will be accessed with the credential. Select from a list of databases supported by SL1. This field is required.

NOTE: For information about monitoring Informix databases, see the [Monitoring Informix Databases](#) section.

- **DB Name.** Name of the database that will be accessed with the credential.
- **DB User.** Username associated with a valid account on the database.
- **Password.** Password associated with a valid account on the database.
- **Hostname/IP.** Hostname or IP address where the database resides. This field is required.
 - You can include the variable **%D** in this field. SL1 will replace the variable with the IP address of the current device (device that is currently using the credential).
 - You can include the variable **%N** in this field. SL1 will replace the variable with the hostname of the current device (device that is currently using the credential). If SL1 cannot determine the hostname, SL1 will replace the variable with the primary management IP address for the current device.

NOTE: To use the localhost, in the **Hostname/IP** field, enter the IP address *127.0.0.1*. The credential will not work if you enter the string *localhost* in the **Hostname/IP** field.

- **Port.** Port number associated with the database you want to access with this credential. This field is required.
 - For **DB Type** of MySQL, the default value is 3306.
 - For **DB Type** of MS SQL Server, the default value is 1433.
 - For **DB Type** of Oracle and SQLNet, the default value is 1521.
 - For **DB Type** of PostgreSQL, the default value is 5432.
 - For **DB Type** of IBM DB2, the default value is 523.
 - For **DB Type** of Sybase ASE, the default value is 4100.
 - For **DB Type** of Informix, see the 9088 section.
 - For **DB Type** of Ingress, the default value is 1572.

NOTE: SL1's Database Servers include a MySQL database running on port 7706. Data Collectors and Message Collectors include a MySQL database running on port 7707.

Oracle Settings

These fields appear if you selected Oracle & *SQLNet in the **DB Type** field. Otherwise, these fields are grayed out.

- **Oracle Connect Type.** Specifies the method SL1 should use to connect to the Oracle database. The choices are:
 - Oracle System Identifier (SID)
 - Oracle Real Application Clusters (SERVICE)
 - Oracle Server Direct Connection (SERVER)

NOTE: In Oracle 11g, the "Oracle Server Direct Connection" option is deprecated. If you select this Oracle Connect Type for an Oracle 11g database, you must edit the file listener.ora and add the line "DEFAULT_SERVICE_LISTENER=<SID>", where <SID> is the SID value.

- **SID (if required).** Enter the value for the Oracle Connect Type (either Oracle SID, Oracle RAC, or Oracle Server) selected in the **Oracle Connect Type** field.

5. Click the **[Save]** button to save the new database credential.
6. Repeat steps 1-5 for each database credential in your network.

NOTE: When you define a Database Credential, the credential will automatically be aligned with the organization(s) you are a member of. To learn more about credentials and organizations, see [Aligning One or More Organizations with a Credential](#).

Defining a SOAP/XML Host Credential in the Classic SL1 User Interface

SOAP/XML credentials allow SL1 to access a web server on a managed device. SOAP/XML credentials are used in several places in SL1, including:

- **With Dynamic Applications of type "SOAP".**
- **With Dynamic Applications of type "XML".**
- **With Dynamic Applications of type "XSLT".**
- **With Dynamic Applications of type "snippet".** The snippet code must define the authentication protocol. Dynamic Applications of type "snippet" can use any type of credential.

To define a SOAP/XML credential:

1. Collect the information you need to create each credential (usually username and password).
2. Go to the **Credential Management** page (System > Manage > Credentials).
3. In the **Credential Management** page, click the **[Actions]** menu. Select **Create SOAP/XML Host Credential**.
4. The **Credential Editor** modal page appears. In this page, you can define the new SOAP/XML credential. To define the new credential, supply values in the following fields:

Basic Settings

- **Profile Name.** Name of the credential. Can be any combination of alphanumeric characters. This field is required.
- **Content Encoding.** Tells the SOAP server or XML data-store how the content is encoded, so the SOAP server or XML data-store knows how to decode the message. Select the encoding that is appropriate for your request and response.
- **Method.** HTTP method to use to exchange credential data from the managed device. Choices are *GET* or *POST*.

NOTE: Typically, Dynamic Applications of type "XML" use GET methods. Dynamic Applications of type "SOAP" and of type "XSLT" use POST methods.

- **HTTP Version.** Version of HTTP to use. Choices are *1.0* or *1.1*.
- **URL.** Address of the SOAP server, HTML document, or XML document. This field is required and should be of the format:

https://IP address:port/full path to desired SOAP, HTML, or XML document

- You can include the variable **%D** in this field. SL1 will replace the variable with the IP address of the current device (device that is currently using the credential).

NOTE: For component devices, SL1 will replace %D with the IP address of the root device.

- You can include the variable **%N** in this field. SL1 will replace the variable with the hostname of the current device (device that is currently using the credential). If SL1 cannot determine the hostname, SL1 will replace the variable with the primary management IP address for the current device.
- **HTTP Auth User.** Username with which to log in to the web server.
- **HTTP Auth Password.** Password with which to access the web server.
- **Timeout (seconds).** Time, in seconds, after which SL1 will stop trying to communicate with the web server.

Proxy Settings

This pane displays optional fields. If you use a proxy server in front of the SOAP server(s) or XML data-store (s) you want to communicate with, enter values in these fields.

- **Hostname/IP.** The host name or IP address of the proxy server.
- **Port.** Port on the proxy server to which you will connect.
- **User.** Username to use to access the proxy server.
- **Password.** Password to use to access the proxy server.

cURL Options

- You can include the cURL command and various options in your credential. The list of cURL options lists all the options you can include in your credential. To include a cURL option in the credential, select it and then select the right-arrow icon. You can then supply arguments in the field to the left of the option.
- For more information on cURL commands, see the cURL manpage at <http://curl.haxx.se/docs/manpage.html>.

SOAP Options

These fields are optional. When a SOAP/XML credential is aligned with a SOAP or XSLT Dynamic Application, the requests defined in the Dynamic Application can use the values defined in these fields. To use a value defined in one of these fields, the request must include the substitution character associated with that value. For example, suppose a Dynamic Application request includes the XML tag <high_value=%1>. Suppose you specified "100" in the **Embed Value [%1]** field in the credential aligned with that Dynamic Application. The request will be sent with the XML tag <high_value=100>.

- **Embedded Password [%P]**. Specifies a password value to include in a request. The value defined in this field is substituted in to the %P substitution character. The value will be encrypted in the request, will be masked in the **Credential Editor**, and will be stored in an encrypted form in the database.
- **Embed Value %1**. The value defined in this field is substituted in to the %1 substitution character.
- **Embed Value %2**. The value defined in this field is substituted in to the %2 substitution character.
- **Embed Value %3**. The value defined in this field is substituted in to the %3 substitution character.
- **Embed Value %4**. The value defined in this field is substituted in to the %4 substitution character.

HTTP Headers

- If you require custom HTTP headers to communicate with the SOAP server, you can build the custom header here.

5. Click the **[Save]** button to save the new SOAP/XML credential.
6. Repeat steps 1-5 for each SOAP/XML credential in your network.

NOTE: When you define a SOAP/XML Credential, the credential will automatically be aligned with the organization(s) you are a member of. To learn more about credentials and organizations, see [Aligning One or More Organizations with a Credential](#).

Defining an LDAP/AD Credential in the Classic SL1 User Interface

LDAP or Active Directory credentials allow SL1 to access data on an LDAP server or an Active Directory server.

Authentication is the method by which SL1 determines if a user can access the SL1 system. For user accounts that are to be authenticated with LDAP or Active Directory, SL1 uses the LDAP or Active Directory credential to establish communication with the LDAP or Active Directory server. SL1 will then query the Active Directory or the LDAP server to determine if the username and password are legitimate and accurate.

Additionally, SL1 can automatically create accounts for one or more LDAP or Active Directory users. SL1 uses the LDAP or Active Directory credential to communicate with Active Directory or the LDAP server and:

- Determine if the username and password are legitimate and accurate.
- Gather information to populate fields in the user's automatically-created account.

For details on using Active Directory or LDAP for authentication, see the manual **Using Active Directory and LDAP**.

To define an LDAP/AD credential:

1. Collect the information you need to create each credential (usually username and password).
2. Go to the **Credential Management** page (System > Manage > Credentials).
3. In the **Credential Management** page, click the **[Actions]** menu. Select **Create LDAP/AD Credential**.
4. The **Credential Editor** modal page appears. In this page, you can define the new LDAP/AD Credential. To define the new credential, supply values in the following fields:

Basic Settings

- **Profile Name**. Name of the credential. Can be any combination of alphanumeric characters. This field is required.
- **LDAP Type**. Specifies the "flavor" or LDAP running on the directory server. Choices are LDAP or Active Directory.
- **Hostname/IP**. Hostname or IP address of the LDAP or Active Directory server. This field is required.
- **Port**. Port number on the LDAP or Active Directory server to which SL1 will send requests. This field is required.
- **Secure**. Specifies whether you are using LDAP over SSL.
- **RDN (Bind DN / bind user)**. Bind DN. The bind DN is a user on the LDAP or Active Directory server who is permitted to search the directory within the specified search base.
 - In many LDAP or AD configurations, each user has read-access to his/her own account. Therefore, you might find it most useful to include the **%u** variable in this field. When an LDAP or AD user logs in to SL1, SL1 stores the username in the **%u** variable. SL1 then uses the **%u** variable to build the bind DN, uses the bind DN to communicate with the LDAP or AD server, and then authenticates the current user.
 - An example entry in the **RDN** field might be:

```
uid=%u, ou=People, dc=sciencelogic, dc=com
```

This creates a DN using the current login name as the uid.

- You can also include the **%d** variable in this field. The **%d** variable represents the name of the LDAP domain, as specified in the **LDAP Domain** field.

NOTE: If you have configured SL1 to automatically create accounts when a user logs in with an LDAP/AD username, you must include the **%u** variable in the **RDN** field.

- **LDAP Domain.** If your LDAP or Active Directory configuration includes multiple domains, specify the domain components to bind to in this field. For example, you could specify:

```
dc=reston, dc=sciencelogic, dc=com.
```

This would bind to the sub-domain "reston", in the domain "sciencelogic", in the domain "com".

- **Bind Password.** Password that allows access to the LDAP or Active Directory server. In most cases, when you specify a bind password in a credential, you are creating a "write" credential (that is, a credential that allows SL1 to make changes to the LDAP or AD server).
- **User Search Base.** In this field, you specify the area in the directory where users to be authenticated reside, using RDN notation. For example, if you want to authenticate five users from the ou called "people", you could specify the RDN that includes that ou.

```
ou=People, dc=sciencelogic, dc=com.
```

This would allow SL1 to authenticate users in the ou called "people." In the **User Search Scope** field, you can specify whether SL1 should also authenticate all users in any ou underneath "people".

- **User Search Scope.** In this field, you specify whether SL1 should search only the directory specified in **User Search Base** or whether SL1 should search the directory specified in **User Search Base** and all its child branches. Choice are:
 - *Subtree.* SL1 should search the directory specified in **User Search Base** and also search all its child branches.
 - *One Level.* SL1 should search only the directory specified in **User Search Base**.

5. Click the **[Save]** button to save the new LDAP/AD credential.
6. Repeat steps 1-5 for each LDAP/AD credential in your network.

NOTE: When you define an LDAP/AD Credential, the credential will automatically be aligned with the organization(s) you are a member of. To learn more about credentials and organizations, see the section [Aligning One or More Organizations with a Credential](#).

Defining a Basic/Snippet Credential in the Classic SL1 User Interface

NOTE: Dynamic Applications of type "snippet" are not required to use only the Basic/Snippet Credential. In Dynamic Applications of type "snippet", the snippet code must define the authentication protocol. Therefore, Dynamic Applications of type "snippet" can use any type of credential.

Basic/Snippet credentials define standard authentication parameters, but are not tied to a specific authentication protocol. Basic/Snippet credentials are used in several places in SL1, including:

- **With Dynamic Applications of type "snippet"**. The snippet code must define the authentication protocol.
- **With Dynamic Applications of type "WMI"**. The authentication protocol is specific to WMI and is specified by SL1 when the Dynamic Application is executed. To access WMI information on a Windows server, ensure that the **Username** you specify is allowed access to the server and to the WMI namespace.
- **With Dynamic Applications of type "PowerShell"**. For information about configuring your environment for PowerShell collection, see the *Monitoring Windows Systems* manual.
- **When defining external backups**. The authentication protocol is defined in the **Backup Management** page (System > Settings > Backup).

To define a Basic/Snippet credential:

1. Collect the information you need to create each credential (usually username and password).
2. Go to the **Credential Management** page (System > Manage > Credentials).
3. In the **Credential Management** page, click the **[Actions]** menu. Select **Create Basic/Snippet Credential**.
4. The **Credential Editor** modal page appears. In this page, you can define the new Basic/Snippet credential. To define the new credential, supply values in the following fields:
 - **Credential Name**. Name of the credential. Can be any combination of alphanumeric characters. This field is required.
 - **Hostname/IP**. Hostname or IP address of the device from which you want to retrieve data. This field is required.
 - You can include the variable **%D** in this field. SL1 will replace the variable with the IP address of the current device (device that is currently using the credential).
 - You can include the variable **%N** in this field. SL1 will replace the variable with the hostname of the current device (device that is currently using the credential). If SL1 cannot determine the hostname, SL1 will replace the variable with the primary management IP address for the current device.
 - **Port**. Port number associated with the data you want to retrieve. This field is required.
 - **Timeout (ms)**. Time, in milliseconds, after which SL1 will stop trying to communicate with the authenticating server.

- **Username.** Username for a user account on the device.
 - **Password.** Password for a user account on the device.
5. Click the **[Save]** button to save the new Basic/Snippet credential.
 6. Repeat steps 1-5 for each Basic/Snippet credential in your network.

NOTE: When you define a Basic/Snippet credential, the credential will automatically be aligned with the organization(s) you are a member of. To learn more about credentials and organizations, see the section [Aligning One or More Organizations with a Credential](#).

Defining an SSH/Key Credential in the Classic SL1 User Interface

Secure Shell (SSH) is a network protocol that enables users to securely access a command-line shell on a remote computer or server over an unsecured network. SSH provides strong encryption and authentication capabilities, making it an ideal method for securely administering commands or transferring data between a client and server.

To make SSH even more secure, you can use SSH keys instead of a simple password to log into a server. SSH keys consist of two long strings of characters, called a public/private key pair, that are much less susceptible than passwords are to brute force attacks. The public key is placed on the server you want to access, while the private key resides on the client. When you use SSH to log into the server from the client, the key pair is used to authenticate the session.

In SL1, some Dynamic Applications of type "Snippet" use SSH to communicate with a remote device. To use these Dynamic Applications, you must define an SSH credential. This credential specifies the hostname or IP address of the system you want to monitor, the port number used to access that system, and the private key used for authentication.

NOTE: Consult the documentation associated with the PowerPack that contains the Dynamic Application of type "Snippet" to find detailed directions on configuring the remote device and generating a private key for SL1 to use.

To define an SSH/Key credential:

1. Collect the information you need to create each credential (usually username and password).
2. Go to the **Credential Management** page (System > Manage > Credentials).
3. In the **Credential Management** page, click the **[Actions]** menu. Select **Create SSH/Key Credential**.
4. The **Credential Editor** modal page appears. In this page, you can define the new SSH/Key credential. To define the new credential, supply values in the following fields:
 - **Credential Name.** Name of the credential. Can be any combination of alphanumeric characters. This field is required.
 - **Hostname/IP.** Hostname or IP address of the device from which you want to retrieve data. This field is required.

- You can include the variable **%D** in this field. SL1 will replace the variable with the IP address of the current device (device that is currently using the credential).
- You can include the variable **%N** in this field. SL1 will replace the variable with hostname of the current device (device that is currently using the credential). If SL1 cannot determine the hostname, SL1 will replace the variable with the primary, management IP address for the current device.
- **Port**. Port number associated with the data you want to retrieve. This field is required.

NOTE: The default TCP port for SSH servers is 22.

- **Timeout (ms)**. Time, in milliseconds, after which SL1 will stop trying to communicate with the authenticating server.
- **Username**. Username for an SSH or user account on the device to be monitored.
- **Password**. Password for an SSH user account on the device to be monitored.
- **Private Key (PEM Format)**. Enter the SSH private key that you want SL1 to use, in PEM format.

NOTE: The private key can have a maximum of 64 characters per line. Therefore, you cannot use keys in the OpenSSH format, because that format uses 70 characters per line. When you attempt to save the credential, SL1 will validate that the private key entered is in the correct format. You will be able to save the credential only if the private key is correctly formatted.

5. Click the **[Save]** button to save the new SSH/Key credential.
6. Repeat steps 1-5 for each SSH/Key credential in your network.

NOTE: When you define a SSH/Key credential, the credential will automatically be aligned with the organization(s) you are a member of. To learn more about credentials and organizations, see the section [Aligning One or More Organizations with a Credential](#).

Defining a PowerShell Credential in the Classic SL1 User Interface

Dynamic Applications can include PowerShell commands that collect data from Windows devices. If you want to use SL1's built-in transport agent (that is, run "agentless" on the Windows device), you can align a PowerShell credential with those Dynamic Applications.

NOTE: Consult the *Monitoring Windows* and *WMI and PowerShell Dynamic Application Development* manuals for detailed directions on configuring the Windows devices for agentless communication and on configuring a proxy server.

To define a PowerShell credential in SL1:

1. Collect the information you need to create the credential:
 - The username and password for a user on the Windows device.
 - If the user is an Active Directory account, the hostname or IP address of the Active Directory server and the domain.
 - Determine if an encrypted connection should be used.
 - If you are using a Windows Management Proxy, the hostname or IP address of the proxy server.
2. Go to the **Credential Management** page (System > Manage > Credentials).
3. In the **Credential Management** page, click the **[Actions]** menu. Select **Create PowerShell Credential**.
4. The **Credential Editor** page appears, where you can define the following fields:
 - **Profile Name**. Name of the credential. Can be any combination of alphanumeric characters. This field is required.
 - **Hostname/IP**. Hostname or IP address of the device from which you want to retrieve data. This field is required.
 - You can include the variable **%D** in this field. SL1 will replace the variable with the IP address of the device that is currently using the credential.
 - You can include the variable **%N** in this field. SL1 will replace the variable with the hostname of the device that is currently using the credential. If SL1 cannot determine the hostname, SL1 will replace the variable with the primary, management IP address for the current device.
 - You can include the prefix **HOST** or **WSMAN** before the variable **%D** in this field if the device you want to monitor uses a service principal name (for example, "HOST://%D" or "WSMAN://%D"). SL1 will use the WinRM service HOST or WSMAN instead of HTTP and replace the variable with the IP address of the device that is currently using the credential.
 - **Username**. Type the username for an account on the Windows device to be monitored or on the proxy server. This field is required.

NOTE: The user should not include the domain name prefix in the username for Active Directory accounts. For example, use "em7admin" instead of "MSDOMAIN\em7admin".

- **Encrypted**. Select whether SL1 will communicate with the device using an encrypted connection. Choices are:
 - **yes**. When communicating with the Windows server, SL1 will use a local user account with authentication of type "Basic Auth". You must then use HTTPS and can use a Microsoft Certificate or a self-signed certificate.
 - **no**. When communicating with the Windows server, SL1 will not encrypt the connection.
- **Port**. Type the port number used by the WinRM service on the Windows device. This field is automatically populated with the default port based on the value you selected in the **Encrypted** field. This field is required.

- **Account Type.** Type of authentication for the username and password in this credential. Choices are:
 - *Active Directory.* On the Windows device, Active Directory will authenticate the username and password in this credential.
 - *Local.* Local security on the Windows device will authenticate the username and password in this credential.
- **Timeout (ms).** Type the time, in milliseconds, after which SL1 will stop trying to collect data from the authenticating server. For collection to be successful, SL1 must connect to the authenticating server, execute the PowerShell command, and receive a response within the amount of time specified in this field.
- **Password.** Type the password for the account on the Windows device to be monitored or on the proxy server. This field is required.
- **PowerShell Proxy Hostname/IP.** If you use a proxy server in front of the Windows devices you want to communicate with, type the fully-qualified domain name or the IP address of the proxy server in this field.
- **Active Directory Hostname/IP.** If you selected Active Directory in the **Account Type** field, type the hostname or IP address of the Active Directory server that will authenticate the credential.
- **Domain.** If you selected Active Directory in the **Account Type** field, type the domain where the monitored Windows device resides.

5. To save the credential, click the **[Save]** button. To clear the values you set, click the **[Reset]** button.

Specifying Credentials During Initial Classic Discovery

Discovery is the process by which SL1 discovers what types of hardware and applications exist on the network and then retrieves data from the discovered hardware and applications.

Before running discovery, you must:

- Determine the SNMP credentials for the devices and applications in your network. Define correlating credentials in SL1, to allow discovery to retrieve as much information as possible.
- If you want SL1 to immediately start collecting data from devices using Dynamic Applications, you should also define any additional credentials required for those Dynamic Applications. For example, if you want SL1 to immediately start monitoring all MySQL databases in your network, you should define credentials that allow SL1 to communicate with each MySQL database in your network. During discovery, SL1 will determine which devices can be monitored with a Dynamic Application for MySQL. After discovery, SL1 will use the database credential to collect data about each MySQL database in your network.

Use the previous sections to define credentials for your network.

When you run discovery, you must specify one or more of these credentials to use. The more credentials you align with a discovery session, the more access SL1 will have to devices and their data during discovery.

To specify credentials during initial discovery:

1. Go to the **Discovery Control Panel** page (System > Manage > Classic Discovery).
2. In the **Discovery Control Panel** page, click the **[Create]** button.

3. In the **Discovery Session Editor** page, supply values in each field.
4. In the **SNMP Credentials** field and in the **Other Credentials** field, you can select one or more credentials to use during discovery. In these fields, you should see a list of all credentials in SL1.
5. When trying to communicate with discovered hardware and applications, SL1 will look at the list of selected credentials and use the appropriate credential to get permission to access data on the external system.


NOTE: During discovery, SL1 tries each SNMP credential specified in the discovery session on each discovered device, to determine if SL1 can collect SNMP details from the device. Later in the discovery session, during alignment of Dynamic Applications, discovery again tries each SNMP credential specified in the discovery session. If one of the SNMP credentials times out three times **without any response**, discovery will stop trying to use that SNMP credential to align SNMP Dynamic Applications. Note that "no response" means that a device did not respond at all. Note that if a device reports that "no OID was found" or "the end of the OID tree was reached", these are considered a legitimate response and would not cause SL1 to abandon the credential.

Defining the Primary and Secondary Credentials for a Single Device in the Classic SL1 User Interface

You can define multiple credentials for a single device. This allows SL1 to align multiple agents and applications for a single device. For example, SL1 might use an SNMP credential to discover hardware information about a device and a database credential to retrieve information about the database on the same device.

To define primary and secondary credentials for a single device:

NOTE: When defining primary and secondary credentials for a device, you will see only the credentials aligned to organizations you are a member of. If a primary or secondary credential has already been defined on the device, and is aligned to an organization you are not a member of, the credential will be restricted.

1. Go to the **Device Manager** page (Devices > Device Manager).
2. In the **Device Manager** page, find the device you want to edit. Click its wrench icon (). The **Device Properties** page is displayed.
3. In the **Device Properties** page, you can select two SNMP credentials in the fields **SNMP Read** and **SNMP Write**.
4. To select a second credential (either of type SNMP or of another type), click the **[Actions]** menu. Select **Secondary Credentials**.
5. The **Secondary Credentials** modal page appears. In the **Secondary Credentials** modal page, you can select one or more credentials to associate with the device. To add a secondary credential to a device, highlight an entry in the list of credentials. To select multiple credentials, hold down the **<CTRL>** key and select the entries by left-clicking with your mouse.
6. During discovery (either nightly, manual, or associated with device policies), SL1 will first try the primary

credentials for the device and then will try the secondary credentials.


7. Click the **[Save]** button to save the change to the device.

Defining the Credentials for a Specific Device/Dynamic Application Pair in the Classic SL1 User Interface

After a device has been discovered by SL1 and one or more Dynamic Applications have been aligned with the device, you can manually assign the credential to use for each Dynamic Application.

The manually assigned credential will be used by SL1 only for this specific Dynamic Application associated with this specific device. For all other devices, SL1 will use the default credential associated with each device, or will use the credential defined in the **Dynamic Application Collections** page for each device.

To manually associate a credential with a Dynamic Application aligned to a device:

1. Go to the **Device Manager** page (Devices > Device Manager).
2. In the **Device Manager** page, find the device for which you want to define a credential. Click its wrench icon ().
3. In the Device Administration panel, click the **[Collections]** tab.
4. In the **Dynamic Application Collections** page, find the Dynamic Application you want to define a credential for. Select its checkbox (). To apply a credential to multiple Dynamic Applications, select the checkbox for each Dynamic Application.
5. From the **Select Action** drop-down list, select the credential from the list of all credentials that you are allowed to use, and then click the **[Go]** button.

NOTE: Your organization membership(s) might affect the list of credentials you can see in the **Select Action** drop-down list.

NOTE: If this Dynamic Application has already been aligned with a credential to which you do not have access, the **Credential** column will display the value *Restricted Credential*. If you align the Dynamic Application with a different credential, you will not be able to re-align the device with the *Restricted Credential*.

6. The selected Dynamic Applications will now use the manually selected credential when collecting data from this device. You should see your change reflected in the **Credential** column in the **Dynamic Application Collections** page.

Specifying Credentials in a Device Template in the Classic SL1 User Interface

You can specify the primary SNMP credentials in a device template. Then, when you use the device template to create a new device or when you apply the device template to a device group, the primary credentials are automatically applied to each device and appear in the **Device Properties** page, in the **SNMP Read** field.

If you include a device template during discovery or re-discovery, SL1 will discover devices first and then apply the device template to each discovered device. During discovery, SL1 automatically assigns a default SNMP credential to each discovered device (that is not a **pingable** device) and then applies the device template.

CAUTION: If you include a primary SNMP credential in a device template and then apply that device template during discovery, you might overwrite the default SNMP credential assigned by SL1. In some cases, this could prevent SL1 from communicating further with the device.

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

How SL1 Uses Credentials During Classic Discovery

During initial discovery, nightly discovery, discovery associated with device policies, and any manually triggered discovery, SL1 uses credentials in the following order:

1. For devices that have not yet been discovered, SL1 uses the credentials supplied in the **Discovery Session Editor** page to collect both SNMP data and Dynamic Application data.
2. For devices that have already been discovered at least once, SL1 uses the **SNMP credentials** specified in the **Device Properties** page.
3. For devices that have already been discovered at least once, SL1 uses **the secondary credentials** defined in the **Device Properties** page if the primary credentials don't work.
4. For devices that have already been discovered at least once, SL1 will use the **credentials defined in the Dynamic Application Collections page** for specific Dynamic Applications.

Aligning One or More Organizations With a Credential in the Classic SL1 User Interface

To support multi-tenancy, SL1 allows you to align each credential with one, multiple or all organizations in SL1. You can also align a credential with no organizations.

When you align an organization with a credential, you control who can view details about the credential, who can view the name of the credential, and who can apply the credential in SL1.

NOTE: When you align an organization with a credential, you are restricting only the users who can view and assign the credential. You are not restricting the devices and actions that can be associated with the credential. For example, you can align a credential only with the organization "Operations" but assign the credential to a device in the "Finance" organization.

If you have an account of type "User" and are a member of only one organization, the **Organization** column will not appear in the **Credential Management** page. The **Credential Management** page will display only credentials that are aligned with your organization.

Credentials that are aligned with an organization have the following behavior:

- For each credential that is aligned with an organization, only administrators and users who are members of the aligned organization will be able to see the credential in the **Credential Management** page.
- When aligning credentials to devices or Dynamic Applications, non-administrator users can view and align only those credentials that are assigned to organizations common to both the user and the device's collector group, plus those credentials that are assigned to all organizations or otherwise required for that collector group.
- Users can change the organization setting of only those credentials that are not currently aligned to any devices.
- In SL1, in any field or column that displays the name of the credential, users who are not members of the aligned organization will not see the credential name. Instead, these users will see either a dash character (-) or the text "Restricted Credential".
- In SL1, in any list from which users can select a credential, users who are not members of the aligned organization will not see the credential as an entry in the list.
- In SL1, in any page where the credential has already been assigned, users who are not members of the aligned organization will see only the name "Restricted Credential".
- In SL1, in any page where the credential has already been assigned, users who are not members of the aligned organization can save the page and maintain the credential. The credential will still appear to that user as "Restricted Credential".
- In SL1, in any page where the credential has already been assigned, users who are not members of the aligned organization can change the credential to a credential aligned with their organization(s). However, those users cannot change the credential again and re-assign the "Restricted Credential". The entry for "Restricted Credential" is removed from the list of possible credentials.

To understand the behavior of a credential aligned with an organization, consider the following example:

- Suppose you have a user account of type "Administrator". Suppose you create an SNMP credential called "ops_cred". Suppose you align that credential with the organization "Operations".
- In the **Credential Management** page, only administrators and users who are members of the organization "Operations" will be able to see the credential "ops_cred" in SL1.
- In SL1, in any field or column that displays the name of the credential (for example, in the **SNMP Credential** column in the **Device Manager** page), users who are not members of the organization "Operations" will not see the "ops_cred" name displayed. Instead, these users will see either a dash character (-) or the text "Restricted Credential".
- In SL1, in any list from which users can select a credential (for example, in the **SNMP Read** field, in the **Device Properties** page), users who are not members of the organization "Operations" will not see the "ops_cred" credential as an entry in the field.
- In SL1, in any page where the credential "ops_cred" has already been assigned, users who are not members of the organization "Operations" will see only the name "Restricted Credential".
- In SL1, in any page where the credential "ops_cred" has already been assigned (for example, in the **SNMP Read** field, in the **Device Properties** page), users who are not members of the organization "Operations" can save the page and maintain the "ops_cred" credential. The credential will still appear to that user as "Restricted Credential".

- In SL1, in any page where the credential "ops_cred" has already been assigned, (for example, in the **SNMP Read** field, in the **Device Properties** page), users who are not members of the organization "Operations" can change the credential to a credential aligned with their organization. However, that user cannot change the credential again and re-assign the "Restricted Credential". The entry for "Restricted Credential" is removed from the list of possible credentials.

Default Organizations Aligned with a Credential

When you create a new credential, SL1 automatically aligns the credential with all your organizations. For example:


Account Type	Organization(s) Aligned with the ScienceLogic Account	Default Organizations Aligned with Credential
Administrator	All	All
User	Primary Organization = NOC Additional Organization Memberships = All organizations	All
User	Primary Organization = NOC Additional Organization Memberships = Sales	NOC, Sales
User	Primary Organization = NOC Additional Organization Memberships = None	NOC

After you save the credential, you can edit the organization(s) aligned with the credential.

Editing the Organizations Aligned with a Credential

After a credential has been created, you change the default organizations aligned with a credential.

To edit the organization alignment on a credential:

1. Go to the **Credential Management** page (System > Manage > Credentials).
2. In the **Credential Management** page, find the credential for which you want to edit the organization. In the **Organization** column, click its org icon (.
3. The **Align Organizations** modal page appears. In this page, provide values in the following fields:
 - **Credential Availability**. Specifies whether you want to align all organizations with the credentials or manually select one, multiple, or no organizations to align with the credential. Choices are:
 - *Aligned Organizations Only*. Selecting this option will make the *Aligned Organizations* pane available. You can select one or multiple organizations to be aligned with the credential.

- *System (All Organizations)*. This option is only available if you are a system administrator or a member of all organizations in SL1. All organizations will be aligned with the credential. If another organization is created, it will be aligned to the credential, by default.

NOTE: The **Credential Availability** field appears only for users who are Administrators and users who are members of all organizations.

- **Aligned Organizations**. Displays a list of all organizations to which you belong. Select one, multiple, or no organizations to align with the credential.
 - To select a single organization, highlight it and left-click.
 - To unselect a single organization, highlight it and left-click.
 - To select multiple organizations, hold down the CTRL key and select the entries by left-clicking.
 - To unselect multiple organizations, hold down the CTRL key and select the entries by left-clicking.

NOTE: Only users who are Administrators and users who are members of all organizations can unselect all organizations in the **Aligned Organizations** list.

4. To save the new organization alignment, click the **[Save]** button.

Restricted Credentials in the Discovery Session Editor Page


The **Discovery Session Editor** page allows you to select multiple credentials to align with a discovery session.

In the **SNMP Credentials** field and the **Other Credentials** field, the **Discovery Session Editor** page might include credentials that have been aligned with one or more organizations. If one of these credentials has been previously selected, users who are not members of the aligned organization(s) will see "Restricted Credential" appear in the **SNMP Credentials** field or the **Other Credentials** field.

If multiple aligned credentials have been previously selected for a discovery session, users who are not members of the aligned organization(s) will see only a single "Restricted Credential" entry appear in the **SNMP Credentials** field or the **Other Credentials** field. This single entry of Restricted Credential represents all restricted credentials. If a user who is not a member of the aligned organization(s) removes the entry Restricted Credential from the discovery session, all restricted credentials are removed. That user cannot change the credential again and re-assign Restricted Credential. The entry for Restricted Credential is removed from the list of possible credentials.

Editing a Credential in the Classic SL1 User Interface

The **Credential Management** page allows you to edit credentials from SL1. To do so:

1. Go to the **Credential Management** page (System > Manage > Credentials).
2. In the **Credential Management** page, click the wrench icon () for the credential you want to edit.
3. The **Credential Editor** modal appears.

4. After editing the fields in the **Credential Editor**, click the **[Save]** button. If you want to save your changes as a new credential, click the **[Save As]** button.

Deleting a Credential in the Classic SL1 User Interface

The **Credential Management** page allows you to delete one or more credentials from SL1. To do so:

NOTE: You cannot delete a credential until all aligned devices, Dynamic Applications, backup settings, LDAP/AD settings, and discovery sessions that use the credential are aligned with another credential.



1. Go to the **Credential Management** page (System > Manage > Credentials).
2. In the **Credential Management** page, select the checkbox for each credential you want to delete.
3. Go to the **Select Actions** menu (in the lower right). Select *DELETE Credential Policy*. Click the **[Go]** button.
4. The selected credentials will be deleted.

Managing Credential Tests

Overview

This chapter describes how to create, run, and manage Credential Tests 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:

What are Credential Tests?	82
Default Credential Tests	83
Viewing Information About Credential Tests	86
Testing a Credential	88
Creating a Credential Test	91
Editing a Credential Test	92
Deleting Credential Tests	92
Available Step Functions	92

What are Credential Tests?

Credential Tests define a series of steps that SL1 can execute on-demand to validate whether a credential works as expected. This chapter describes how to manage existing credential tests and create new credential tests. For information about executing a credential test, see the [Credentials section](#).

A number of commonly user Credential Tests, such as AWS, Azure, and PowerShell, are included in SL1 by default.

You can also include Credential Tests in PowerPacks. For information about including a credential test in a PowerPack, see the **PowerPacks** manual.

Default Credential Tests

This section describes the credential tests included in the default SL1 installation.

AWS Credential Test

The AWS Credential Test can only be used to test a SOAP/XML credential or a SOAP/XML with AWS subtype credential for monitoring AWS using the Dynamic Applications in the "Amazon Web Services" PowerPack. The AWS Credential Test performs the following steps:

- **Test Reachability.** Performs an ICMP ping request to the URL for the EC2 service in the region specified in the credential. If a region is not specified in the credential, the us-east-1 region is used.
- **Test Port Availability.** Performs an NMAP request to TCP port 443 on the URL for the EC2 service in the region specified in the credential. If a region is not specified in the credential, the us-east-1 region is used.
- **Test Name Resolution.** Performs an nslookup request on the URL for the EC2 service in the region specified in the credential. If a region is not specified in the credential, the us-east-1 region is used.
- **Make connection to AWS account.** Attempts to connect to the AWS service using the account specified in the credential.
- **Scan AWS services.** Verifies that the account specified in the credential has access to the services.

Azure Credential Test

The Azure Credential Test can be used to test a SOAP/XML credential for monitoring Microsoft Azure using the Dynamic Applications in the "Microsoft: Azure" PowerPack. The Azure Credential Test performs the following steps:

- **Test Port Availability.** Performs an NMAP request to TCP port 443 on management.azure.com.
- **Test Name Resolution.** Performs an nslookup request on management.azure.com.
- **Make connection to Azure account.** Attempts to connect to the Azure service using the account specified in the credential.
- **Make Azure Active Directory Request.** Verifies that the account specified in the credential has the permissions required to discover the Azure account.

Basic/Snippet Credential Test

The Basic/Snippet Credential Test can be used to test a Basic/Snippet credential for connectivity. The Basic/Snippet Credential Test performs the following steps:

- **Test Reachability.** Performs an ICMP ping request to the host specified in the credential.
- **Test Port Availability.** Performs an NMAP request to the TCP port specified in the credential on the host specified in the credential.
- **Test Name Resolution.** Performs an nslookup request on the host specified in the credential.

Database Credential Test

The Database Credential Test can be used to test a Database credential for connectivity. The Database Credential Test performs the following steps:

- **Test Reachability.** Performs an ICMP ping request to the host specified in the credential.
- **Test Port Availability.** Performs an NMAP request to the TCP port specified in the credential on the host specified in the credential.
- **Test Name Resolution.** Performs an nslookup request on the host specified in the credential.
- **Make DB Connection.** Attempts to make a database connection using the credential and executes the query "SELECT 1;".
- **Verify Table Existence.** Attempts to make a database connection using the credential and executes the query "SELECT * FROM master.system_settings_core;".

PowerShell Credential Test

The PowerShell Credential Test can be used to test a PowerShell credential for connectivity. The PowerShell Credential Test performs the following steps:

- **Test Reachability.** Performs an ICMP ping request to the host specified in the credential.
- **Test Port Availability.** Performs an NMAP request to the TCP port specified in the credential on the host specified in the credential.
- **Test Name Resolution.** Performs an nslookup request on the host specified in the credential.
- **Test Kerberos.** If the credential does not specify local authentication, attempts to acquire a kerberos ticket using the credential.
- **Test WinRM Connection.** Attempts a WinRM connection using the credential.
- **Execute PowerShell Cmdlet.** Attempts to execute the 'Get-WmiObject Win32_Process | Select Name' PowerShell Cmdlet using the credential.

SNMP Credential Test

The SNMP Credential Test can be used to test an SNMP credential for connectivity. The SNMP Credential Test performs the following steps:

- **Test Reachability.** Performs an ICMP ping request to the host specified in the credential.
- **Test Port Availability.** Performs an NMAP request to the UDP port specified in the credential on the host specified in the credential.
- **Test SNMP Availability.** Attempts an SNMP getnext request to .1.3.6.1 using the credential.

SOAP/XML Credential Test

The SOAP/XML Credential Test can be used to test a SOAP/XML credential for connectivity. The SOAP/XML Credential Test performs the following steps:

- **Test Reachability.** Performs an ICMP ping request to the host specified in the credential.
- **Test Port Availability.** Performs an NMAP request to the TCP port specified in the credential on the host specified in the credential.
- **Test Name Resolution.** Performs an nslookup request on the host specified in the credential.
- **Make cURL Request.** Attempts to make a cURL request connection using the credential.
- **Verify Content.** Attempts to make a cURL request connection using the credential and verifies whether "discovery_session" appears in the response.

SoftLayer Credential Test

The SoftLayer Credential Test can be used to test a SOAP/XML credential for monitoring SoftLayer using the Dynamic Applications in the *SoftLayer: Cloud* PowerPack. The SoftLayer Credential Test performs the following steps:

- **Test Reachability.** Performs an ICMP ping request to api.softlayer.com.
- **Test Port Availability.** Performs an NMAP request to TCP port 443 on api.softlayer.com.
- **Test Name Resolution.** Performs an nslookup request on api.softlayer.com.
- **Make connection to SoftLayer account.** Attempts to connect to the Softlayer Account endpoint using the account specified in the credential.
- **Query SoftLayer Resource.** Performs a getDatacenters request to the Softlayer Location endpoint using the account specified in the credential.

SSH/Key Credential Test

The SSH/Key Credential Test can be used to test a SSH/Key credential for connectivity. The SSH/Key Credential Test performs the following steps:

- **Test Reachability.** Performs an ICMP ping request to the host specified in the credential.
- **Test Port Availability.** Performs an NMAP request to the TCP port specified in the credential on the host specified in the credential. If no port is specified in the credential, port 22 is used.
- **Test Name Resolution.** Performs an nslookup request on the host specified in the credential.
- **Make SSH Connection.** Attempts to make an SSH connection using the credential.
- **Execute Command via SSH.** Attempts to make an SSH connection using the credential and executes the command "ping localhost -c1".

VMware Credential Test

The VMware Credential Test can be used to test a SOAP/XML credential for monitoring VMware using the Dynamic Applications in the "VMware: vSphere Base Pack" PowerPack. The VMware Credential Test performs the following steps:

- **Test Reachability.** Attempts to reach the vCenter server using ICMP.
- **Attempt VMware Connection.** Attempts to connect to the VMware service using the account specified in the credential.

Viewing Information About Credential Tests

The **Credential Test Management** page (System > Customize > Credential Tests) allows you to view a list of all credential tests. From this page, you can also create, edit, run, and delete credential tests.

TIP: To sort the list of credential tests, click on a column heading. The list will be sorted by the column value, in ascending order. To sort by descending order, click the column heading again. The **Last Edited** column sorts by descending order on the first click; to sort by ascending order, click the column heading again.

For each credential test, the page displays:

- **Test Name.** Name of the credential test.
- **Type.** The type of credential that can be tested using this credential test. Possible types are SNMP, Database, SOAP/XML, LDAP/AD, Basic/Snippet, SSH/Key, and Powershell.
- **PowerPack.** The PowerPack that contains the credential test.
- **ID.** Unique numeric ID, automatically assigned by SLI to each credential test.
- **Last Edited.** Date and time the credential test was created or last edited.
- **Edited By.** The username of the user who created or last edited the credential test.

Filtering the List of Credential Tests

To filter the list of credentials in the **Credential Test Management** page (System > Customize > Credential Tests), use the search fields at the top of each column. The search fields are find-as-you-type filters; as you type, the page is filtered to match the text in the search field, including partial matches. Text matches are not case-sensitive. Additionally, you can use the following special characters in each filter:

- **,** (comma). Specifies an "or" operation. For example:
"dell, micro" would match all values that contain the string "dell" OR the string "micro".
- **&** (ampersand). Specifies an "and" operation. For example:
"dell & micro" would match all values that contain the string "dell" AND the string "micro".

- ! (exclamation mark). Specifies a "not" operation. For example:
"!dell" would match all values that do not contain the string "dell".
- ^ (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.

Testing a Credential

This section describes the following methods for testing a credential in SL1:

- [During guided or unguided discovery](#)
- [From the Credential Tester panel when creating or editing a credential](#)
- [From the Credential Management page in the classic SL1 user interface](#)
- [From the Credential Test Management page in the classic SL1 user interface](#)

Testing a Credential During Guided or Unguided Discovery

You can test a credential from the **Credentials** page during guided or unguided discovery. To do so:

1. Complete steps 1-6 from the section on [Adding Devices Using Universal or Guided Discovery](#), or steps 1-5 from the section on [Adding Devices Using Unguided Discovery](#).
2. Complete the steps from the next section, [Testing a Credential Using the Credential Tester Panel](#).

Testing a Credential Using the Credential Tester Panel

When defining or editing a credential in SL1, you can test the credential using the **Credential Tester** panel.

To test a credential using the **Credential Tester** panel:

1. From the **Credentials** page (Manage > Credentials) or from the **Credential Selection** page during [guided](#) or [unguided](#) discovery, click **[Create New]** to create a new credential or click the **Actions** icon (☰) of a credential that you want to test and then select *Edit/Test*.
2. While [defining](#) or [editing](#) the credential, supply values in the required fields. Required fields may vary depending on the type of credential you create.
3. Click the **[Save & Test]** button. This activates the **Credential Tester** fields.
4. In the **Credential Tester** panel, supply values in the following fields:
 - **Select Credential Test.** Select a credential test to run. This drop-down list includes the [ScienceLogic Default Credential Tests](#), credential tests included in any PowerPacks that have been optionally installed on your system, and credential tests that users have created on your system.
 - **Select Collector.** Select the All-In-One Appliance or Data Collector that will run the test.
 - **IP or Hostname to test.** Type a hostname or IP address that will be used during the test. For example, if you are testing an SNMP credential, the hostname/IP address you supply will be used to perform a test SNMP request.

5. Click the **[Test Credential]** button to run the credential test. The **Testing Credential** window appears.

The **Testing Credential** window displays a log entry for each step in the credential test. The steps performed are different for each credential test. The log entry for each step includes the following information:

- **Step.** The name of the step.
- **Description.** A description of the action performed during the step.
- **Log Message.** The result of the step for this execution of the credential test.
- **Status.** Whether the result of this step indicates the credential and/or the network environment is configured correctly (Passed) or incorrectly (Failed).
- **Step Tip.** Mouse over the question mark icon (?) to display the tip text. The tip text recommends what to do to change the credential and/or the network environment if the step has a status of "Failed".

Testing a Credential from the Credential Management Page in the Classic SL1 User Interface

You can test a credential from the **Credential Management** page in the classic SL1 user interface using a predefined credential test.

To run a credential test from the **Credential Management** page:

1. Go to the **Credential Management** page (System > Manage > Credentials).
2. Click the **[Actions]** menu, and then select **Test Credential**. The **Credential Tester** modal page appears:
3. Supply values in the following fields:
 - **Test Type.** Select a credential test to run. This list includes the [ScienceLogic Default Credential Tests](#), credential tests included in any PowerPacks that have been optionally installed on your system, and credential tests that users have created on your system.
 - **Credential.** Select the credential to test. This drop-down list includes only credentials that you have access to that can be tested using the selected credential test.
 - **Hostname/IP.** Enter a hostname or IP address that will be used during the test. For example, if you are testing an SNMP credential, the hostname/IP address you supply will be used to perform a test SNMP request.
 - **Collector.** Select the All-In-One Appliance or Data Collector that will run the test.
4. Click the **[Run Test]** button to run the credential test. The **Test Credential** window appears.

The **Test Credential** window displays a log entry for each step in the credential test. The steps performed are different for each credential test. The log entry for each step includes the following information:

- **Step.** The name of the step.
- **Description.** A description of the action performed during the step.
- **Log Message.** The result of the step for this execution of the credential test.
- **Status.** Whether the result of this step indicates the credential and/or the network environment is configured correctly (Passed) or incorrectly (Failed).

- **Step Tip.** Mouse over the question mark icon (❓) to display the tip text. The tip text recommends what to do to change the credential and/or the network environment if the step has a status of "Failed".
5. Optionally, you can click the **[Execute Discovery Session]** button to run a discovery session using the **Credential**, **Hostname/IP**, and **Collector** you selected in the **Credential Tester** modal page.

Testing a Credential from the Credential Test Management Page in the Classic SL1 User Interface

The **Credential Test Management** page in the classic SL1 user interface allows you to run a credential test to validate that a credential works as expected. To do so:


1. Go to the **Credential Test Management** page (System > Customize > Credential Tests).
2. Find the credential test that you want to run and click its lightning bolt icon (⚡). The **Credential Tester** modal page appears:
3. Supply values in the following fields:
 - **Test Type.** This field is pre-populated with the credential test you selected.
 - **Credential.** Select the credential to test. This drop-down list includes only credentials that you have access to that can be tested using the selected credential test.
 - **Hostname/IP.** Enter a hostname or IP address that will be used during the test. For example, if you are testing an SNMP credential, the hostname/IP address you supply will be used to perform a test SNMP request.
 - **Collector.** Select which All-In-One Appliance or Data Collector will run the test from the drop-down list.
4. Click the **[Run Test]** button to run the credential test. The **Test Credential** window appears.

The **Test Credential** window displays a log entry for each step in the credential test. The steps performed are different for each credential test. The log entry for each step includes the following information:

- **Step.** The name of the step.
 - **Description.** A description of the action performed during the step.
 - **Log Message.** The result of the step for this execution of the credential test.
 - **Status.** Whether the result of this step indicates the credential and/or the network environment is configured correctly (Passed) or incorrectly (Failed).
 - **Step Tip.** Mouse over the question mark icon (❓) to display the tip text. The tip text recommends what to do to change the credential and/or the network environment if the step has a status of "Failed".
5. Optionally, you can click the **[Execute Discovery Session]** button to run a discovery session using the **Credential**, **Hostname/IP**, and **Collector** you selected in the **Credential Tester** modal page.

Creating a Credential Test

The **Credential Test Management** page allows you to create a new credential test. To do so:

1. Go to the **Credential Test Management** page (System > Customize > Credential Tests).
2. Click the **[Create]** button. The **Add Credential Test** modal page appears.
3. Supply values in the following fields:
 - **Test Name.** Enter a name for the credential test.
 - **Credential Type.** Select the type of credential that can be used with this test. Possible types are *SNMP, Database, SOAP/XML, LDAP/AD, Basic/Snippet, SSH/Key, and PowerShell.*
 - **Execution Environment.** Select the execution environment to which you want to align the credential test. An execution environment contains the supporting modules, code, scripts, directories, and files (packaged in ScienceLogic Libraries) for the credential test. An execution environment includes its own installation directories, doesn't share libraries with other environments, and allows granular control of dependencies, versions, and permissions. The default execution environment is "EM7 Credential Tests". For more information, see the **ScienceLogic Libraries** manual.
 - **Steps.** Enter the JSON structure that defines how each step in the credential test will be executed. The JSON structure must specify an array (square-bracket notation) of objects (curly-braces notation). Each object in the list defines a step to be executed by the credential test. The object for a step must include the following keys:
 - *name.* The name of the step. This text will be displayed in the **Step** column in the credential test results.
 - *description.* A description of the step. This text will be displayed in the **Description** column in the credential test results.
 - *pass_message.* The log message to display when the success criteria of the specified *function* are met. To use the output from a function in the log message, you can include substitutions in this field.
 - *fail_message.* The log message to display when the success criteria of the specified *function* are not met. To use the output from a function in the log message, you can include substitutions in this field.
 - *step_tip.* Information for the user to troubleshoot their credential if this step fails. This text will be displayed if the user hovers over the information icon () in the credential test results.
 - *function.* The name of the function that will be called by SL1 to execute the step. For a list of available functions, see the [Available Step Functions](#) section.

For step functions that accept additional arguments, add an additional key/value pair in the object for that step to specify additional arguments.

The `pass_message` and `fail_message` can include substitutions. Substitutions are specified using the following format:

```
% ([Return Value Name]) s
```

Where `[Return Value Name]` is the name of the return value you want to substitute in to the `pass_message` and `fail_message`. For example, the ping function returns the latency in milliseconds in the variable "result". Suppose your step uses the ping function with the following `pass_message`:

```
Latency is %(result) s ms
```


Suppose that when a user runs the credential test, the ping function returns "10" in the variable "result". The following log message is displayed to the user:

```
Latency is 10 ms
```

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

Editing a Credential Test

The **Credential Test Management** page allows you to edit an existing credential test. To do so:

1. Go to the **Credential Test Management** page (System > Customize > Credential Tests).
2. Find the credential test that you want to edit and click its wrench icon (). The **Edit_Credential_Test** modal page appears.
3. Edit one or more parameters for the credential test. For a description of each field, see the [Creating a Credential Test](#) section.
4. Click the **[Save]** button to save your changes to the credential test.

Deleting Credential Tests

The **Credential Test Management** page allows you to delete one or more credential tests from SL1. To do so:

1. Go to the **Credential Test Management** page (System > Customize > Credential Tests).
2. Select the checkbox for each credential test you want to delete.
3. Click the **Select Actions** menu (in the lower right) and select *DELETE Credential Test*, then click the **[Go]** button.
4. In the pop-up window that appears, click **[OK]**. The selected credential tests will be deleted.

Available Step Functions

ping

This function executes the following ping command using the provided IP address or hostname:

```
sudo /bin/ping -c1 [IP address/Hostname]
```

- **Success/Failure Criteria.** Successful if a response is received.
- **Arguments.** None.
- **Return values on success:**
 - *success.* Returns True.
 - *result.* Returns the response time, in ms.
- **Return values on failure:**
 - *success.* Returns False.

nmap_udp

This function executes the following nmap command using the provided IP address or hostname and the port in the provided credential:

```
sudo /usr/bin/nmap -sU -p [port][IP address/Hostname]
```

- **Success/Failure Criteria.** Successful if the NMAP command returns "open" or "open|filtered" as the state of the port.
- **Arguments.** None.
- **Return values on success:**
 - *port.* Returns the port number from the credential.
 - *success.* Returns True.
 - *result.* Returns the state of the port from the NMAP output.
- **Return values on failure:**
 - *port.* Returns the port number from the credential or "Undefined" if no port is specified in the credential.
 - *success.* Returns False.
 - *result.* Returns the state of the port from the NMAP output.

nmap_tcp

This function executes the following nmap command using the provided IP address or hostname and the port in the provided credential:

```
sudo /usr/bin/nmap -P0 -p [port][IP address/Hostname]
```

- **Success/Failure Criteria.** Successful if the NMAP command returns "open" or "open | filtered" as the state of the port.
- **Arguments.** None.
- **Return values on success:**
 - *port.* Returns the port number from the credential.
 - *success.* Returns True.
 - *result.* Returns the state of the port from the NMAP output.
- **Return values on failure:**
 - *port.* Returns the port number from the credential or "Undefined" if no port is specified in the credential.
 - *success.* Returns False.
 - *result.* Returns the state of the port from the NMAP output.

nslookup_forward

This function executes the nslookup command-line utility using the provided hostname.

- **Success/Failure Criteria.** Successful if the forward lookup returns one or more results.
- **Arguments.** None.
- **Return values on success:**
 - *success.* Returns True.
 - *result.* Returns a string in the following format: Forward returned [*number*] result[s]
- **Return values on failure:**
 - *success.* Returns False.
 - *result.* Returns "Forward Failed".

nslookup

If the user provides a hostname, this function:

1. Executes the nslookup command-line utility using the provided hostname.
2. Executes the nslookup command-line utility using the IP address returned by the first nslookup command.

If the user provides an IP address, this function:

1. Executes the nslookup command-line utility using the provided IP address.
2. Executes the nslookup command-line utility using the hostname returned by the first nslookup command.

- **Success/Failure Criteria.** Successful if both the forward and reverse lookups return one or more results.
- **Arguments.** None.
- **Return values on success:**
 - *success.* Returns True.
 - *result.* Returns a string in the following format: *[direction]* returned *[number]* result[s], *[direction]* returned *[number]* result[s]
- **Return values on failure:**
 - *success.* Returns False.
 - *result.* Returns a string in one of the following formats:
 - *[direction]* failed, *[direction]* returned *[number]* result[s]
 - *[direction]* returned *[number]* result[s], *[direction]* failed
 - *[direction]* failed, *[direction]* failed

dynapp_execute

This function performs collection of a specified Dynamic Application using the credential and IP/hostname provided by the user.

- **Success/Failure Criteria.** Successful if Dynamic Application collection is successful.
- **Arguments.** One of the following arguments is required. If both are specified, only *app_id* is used:
 - *app_id.* Integer. The ID of the Dynamic Application to execute.
 - *app_guid.* String. The GUID of the Dynamic Application to execute.
- **Return values on success:**
 - *success.* Returns True.
- **Return values on failure:**
 - *success.* Returns False.

snmp_getnext

This function executes an SNMP getnext request on .1.3.6.1 using the credential and IP/hostname provided by the user. This function works only with SNMP credentials.

- **Success/Failure Criteria.** Successful if a value is returned by the getnext request.
- **Arguments.** None.
- **Return values on success:**

- *success*. Returns True.
- *result*. Returns the value returned by the request (typically the System Name).
- **Return values on failure:**
 - *success*. Returns False.

ssh_request

This function attempts to make an SSH connection using the following values:

- **The IP address/hostname from the provided Credential.** The host to use for the SSH connection.
- **The SSH Key from the provided SSH/Key Credential.** The private key to use for the SSH connection.
- **The Username from the provided Credential.** The username for the SSH connection.
- **The Password from the provided Credential.** The password for the SSH connection.
- **The Port from the provided Credential.** The port for the SSH connection. If no port is supplied, port 22 is used.
- **The command argument supplied to the function.** The command that is executed using the SSH connection. If the command argument is not supplied, no command is executed.

If the connection is successful and the command argument is supplied to the function, the function executes the command specified in the command argument.

- **Success/Failure Criteria.** If a command is not specified in the arguments, successful if an SSH connection is established. If a command is specified in the arguments, successful if the an SSH connection is established and the command returns an exit code of 0.
- **Arguments.** The following argument is optional:
 - *command*. String. SSH command to execute.
- **Return values on success:**
 - *success*. Returns True.
- **Return values on failure:**
 - *success*. Returns False.
 - *result*. Returns an error message.

db_query

This function attempts to make a database connection and execute a query using the credential and IP/hostname provided by the user. This function works only with Database credentials.

- **Success/Failure Criteria.** Successful if the database query returns rows.
- **Arguments.** The following argument is optional:

- *query*. String. Database query to execute. If no query is supplied, "SELECT 1;" is executed.

- **Return values on success:**

- *success*. Returns True.

- **Return values on failure:**

- *success*. Returns False.
- *result*. Returns an error message.

curl

This function executes a cURL request using the credential and IP/hostname provided by the user. Optionally, this function can perform an expression match on the returned content. This function works only with SOAP/XML credentials.

- **Success/Failure Criteria.** If match text is not specified in the arguments, successful if the cURL request returns an HTTP status code that does not begin with a 4 or 5. If match text is specified in the arguments, successful if the cURL request returns an HTTP status code that does not begin with a 4 or 5 and the supplied expression match is included in the response.

- **Arguments.** The following argument is optional:

- *match_text*. String. Text to match to the response.

- **Return values on success:**

- *success*. Returns True.
- *result*. Returns one of the following:
 - If no match text is specified, the string "HTTP [Status Code]" is returned.
 - If match text is specified, the string "Match text found" is returned.

- **Return values on failure:**

- *success*. Returns False.
- *result*. Returns one of the following:
 - If no match text is specified and the HTTP request returned a 400-series or 500-series status code, the string "HTTP [Status Code]" is returned.
 - If match text is specified and the HTTP request was successful, the string "Match text not found" is returned.
 - If an error is encountered executing the cURL request, an error message is returned.

aws_connect

Using the boto3 library, this function creates an IAM client object using the following values:

- **Username from the provided credential.** Used as the AWS Access Key ID.
- **Password from the provided credential.** Used as the AWS Secret Access Key.
- **%1 value from the provided SOAP/XML Credential.** Used as the AWS region. If the region is not supplied in the credential, "us-east-1" is used.

After creating the object, the function calls the `get_user()` request using the object.

- **Success/Failure Criteria.** Successful if the `get_user()` request is successful.
- **Arguments.** None
- **Return values on success:**
 - `success.` Returns True.
- **Return values on failure:**
 - `success.` Returns False.
 - `result.` Returns an error message.

aws_service_scan

Using the boto3 library, this function creates an AWS session object using the following values:

- **Username from the provided credential.** Used as the AWS Access Key ID.
- **Password from the provided credential.** Used as the AWS Secret Access Key.
- **%1 value from the provided SOAP/XML Credential.** Used as the AWS region. If the region is not supplied in the credential, "us-east-1" is used.

After creating the object, the function iterates through the list of services specified in the `expected_services` argument. For each `expected_services` argument, the function attempts to connect to the service using the AWS session object.

- **Success/Failure Criteria.** Successful if the connection to every service in the `expected_services` argument was successful.
- **Arguments.** The following argument is required:
 - `expected_services.` Specify a list of service names. The service names must match the possible service names returned by the `get_available_resources()` function for an AWS session object.
- **Return values on success:**
 - `success.` Returns True.
- **Return values on failure:**
 - `success.` Returns False.
 - `result.` Returns one of the following:

- If a client error occurs creating the AWS session object, returns an error message.
- If the AWS session object was created successfully, returns the following string: "cannot access the following services: *[comma-separated list of failed services]*"

nmap_aws

This function performs a port scan of port 443 on the URL of a specific AWS service and region. The function uses the following values to build the URL:

- **%1 value from the provided SOAP/XML Credential.** Used as the AWS region. If the region is not supplied in the credential, "us-east-1" is used.
- **The service argument supplied to the function.** Used as the AWS service. If the service argument is not supplied, "ec2" is used.

For services that do not use regions, this function executes the following nmap command:

```
sudo /usr/bin/nmap -P0 -p 443 [service].[region].amazonaws.com
```

For services that use regions, this function executes the following nmap command:

```
sudo /usr/bin/nmap -P0 -p 443 [service].amazonaws.com
```

- **Success/Failure Criteria.** Successful if the NMAP command returns "open" or "open | filtered" as the state of the port.
- **Arguments.** The following argument is optional:
 - *service.* The service to use in the URL. If the service argument is not supplied, "ec2" is used.
- **Return values on success:**
 - *port.* Returns "443".
 - *success.* Returns True.
 - *result.* Returns the state of the port from the NMAP output.
- **Return values on failure:**
 - *port.* Returns "443".
 - *success.* Returns False.
 - *result.* Returns the state of the port from the NMAP output.

nslookup_aws

This function executes the nslookup command-line utility URL of a specific AWS service and region. The function uses the following values to build the URL:

- **%1 value form the provided SOAP/XML Credential.** Used as the AWS region. If the region is not supplied in the credential, "us-east-1" is used.
- **The service argument supplied to the function.** Used as the AWS service. If the service argument is not supplied, "ec2" is used.

For services that do not use regions, the URL is in the following format:

```
[service].[region].amazonaws.com
```

For services that use regions, the URL is in the following format:

```
[service].amazonaws.com
```

- **Success/Failure Criteria.** Successful if the forward lookup returns one or more results.
- **Arguments.** The following argument is optional:
 - *service.* The service to use in the URL. If the service argument is not supplied, "ec2" is used.
- **Return values on success:**
 - *success.* Returns True.
 - *result.* Returns a string in the following format: Forward returned *[number]* result[s]
- **Return values on failure:**
 - *success.* Returns False.
 - *result.* Returns "Forward Failed".

ping_aws

This function executes a ping command to the URL of a specific AWS service and region. The function uses the following values to build the URL:

- **%1 value form the provided SOAP/XML Credential.** Used as the AWS region. If the region is not supplied in the credential, "us-east-1" is used.
- **The service argument supplied to the function.** Used as the AWS service. If the service argument is not supplied, "ec2" is used.

For services that do not use regions, the ping command is in the following format:

```
sudo /bin/ping -c1 [service].[region].amazonaws.com
```

For services that use regions, the ping command is in the following format:

```
sudo /bin/ping -c1 [service].amazonaws.com
```

- **Success/Failure Criteria.** Successful if a response is received.
- **Arguments.** The following argument is optional:

- *service*. The service to use in the URL. If the service argument is not supplied, "ec2" is used.

- ***Return values on success:***

- *success*. Returns True.
- *result*. Returns the response time, in ms.

- ***Return values on failure:***



- *success*. Returns False.

Using External Credential Services

Overview

This chapter provides instructions on how to use external credential services 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:

CyberArk	102
--------------------------------	-----

CyberArk

NOTE: CyberArk is not provided with SL1. You must purchase and manage the CyberArk Vault independently.

SL1 can source credential data from external CyberArk vaults using a credential gateway service. This service allows you to download the CyberArk Credential Provider—sometimes referred to as the CyberArk Credential Agent—from your instance of CyberArk and install the agent on SL1 Data Collectors. This agent syncs with credentials on the CyberArk vault server and uses those credentials to collect data from SL1-monitored devices when necessary. When an SL1 system has been set up with the CyberArk Agent and the credential gateway service, you can enter vault tags in any secret/encrypted credential field in SL1 instead of entering the password or secret value itself.

NOTE: CyberArk credentials are stored in the memory of the CyberArk Agent, not in SL1 Data Collectors.

CAUTION: The CyberArk credential gateway service integration is incompatible with SL1's Concurrent PowerShell feature. If you are using the CyberArk credential integration, you must have Concurrent PowerShell disabled. To disable Concurrent PowerShell, go to the **Behavior Settings** page (System > Settings > Behavior), ensure that the **Enable Concurrent PowerShell Collection** checkbox is not selected, and click **Save**.

Installing the CyberArk Credential Agent on Every Collector

Install the CyberArk Credential Provider (Agent) by following the instructions outlined in CyberArk's knowledge base article: [Install the Credential Provider](#).

NOTE: The CyberArk Agent must be installed on your SL1 Data Collectors. These appliances retrieve credential data from the CyberArk vaults.

Reinstalling the CyberArk Credential Agent on a Rebuilt Data Collector

In cases where you need to rebuild your Data Collector, you will first need to remove the previous Data Collector from the CyberArk Credential Agent. After rebuilding the Data Collector, you will then need to reinstall the agent on that rebuilt appliance.

To reinstall the agent, follow the steps outlined in CyberArk's knowledge base article: [Install the Credential Provider](#).

Defining a Custom CyberArk Delimiter

You can define a custom CyberArk delimiter when sourcing credential data from external CyberArk vaults. If you want to use a custom CyberArk delimiter, you must do so before starting the Credential Gateway Service:

1. Use the SSH protocol to connect to the Data Collector.
2. Enter the following commands

```
sudo vi /opt/em7/lib/python3/sl_credential_gateway/sl_credential_gateway.sh
```

3. Uncomment the line `#export CYBERARK_DELIMITER='!'` and change the `""!` character to a different delimiter character.
4. Save and exit the `sl_credential_gateway.sh` file.

Starting the Credential Gateway Service on Every Data Collector

The Credential Gateway Service is the method SL1 uses to interact with the CyberArk Credential Provider (Agent). You must complete the following steps for every Data Collector you are utilizing:

1. Use the SSH protocol to connect to the Data Collector.
2. Enter the following commands:

```
cd /opt/em7/lib/python3/sl_credential_gateway
```

```
sudo bash sl_credential_gateway.sh
```

3. A message appears indicating that the "Credential Gateway Service environment has been created and the service is running."

NOTE: If you already have a running Credential Gateway Service and then rerun this script, the service will stop and start a new one. Future development includes integrating this service into systemd instead of using this script so all systemd service functionality is supported.

Troubleshooting Credential Gateway Service Issues

If you are encountering errors while running the `sl_credential_gateway.sh` script, perform the following actions:

- Check that the directory `/opt/em7/lib/python3/sl_credential_gateway/venv/` exists. If it does, try to activate it with the command `source /opt/em7/lib/python3/sl_credential_gateway/venv/bin/activate` and see if it gives you any errors. You can also delete the `venv` directory and re-run `sl_credential_gateway.sh`.
- Check whether the gunicorn process runs with the command `pgrep -x "gunicorn"`. You should see two process IDs. If you see anything other than two IDs, kill them with the command `ps -ef | grep gunicorn | grep -v grep | awk '{print $2}' | xargs kill` and then re-run `sl_credential_gateway.sh`.

NOTE: If gunicorn is not running, check to see if `/var/run/em7/cgs.sock` exists. If it does, delete it and re-run `sl_credential_gateway`.

- Check `/var/log/sl1/cgs-error.log` for additional error information and search for troubleshooting guidance.

Setting Up Vault Tags in Your Credential

When an SL1 system has been set up with the CyberArk Credential Provider (Agent) and the SL1 Credential Gateway service, you can enter vault tags in any secret/encrypted credential field instead of entering the password/secret itself. Any time you use the credential in discovery or collections, SL1 will detect the vault tag, use the provided query string to retrieve a secret from the CyberArk credential provider, and then use it in this field.

NOTE: You can set up vault tags only in the classic SL1 user interface.

1. Go to the **Credential Management** page (System > Manage > Credentials).
2. Click the **[Actions]** button and then select the type of credential you want to create. The **Credential Editor** modal appears.
3. If there's a secret/encrypted credential field, enter the appropriate vault tag instead of entering the password/secret itself.

NOTE: The vault tag entered should be in the following format: `<vault id>:<appid>;<safe>;<object>`

For example, a query string for an SNMP Credential might appear as the following:

```
vault:123:appid=testappid;safe=Test;object=snmp_v2.
```

For CyberArk, the query string contains the following parameters:

- **appid.** The ID of the application registered in the CyberArk vault for SL1.
- **safe.** The ID of the safe in which the object resides.
- **object.** The ID of the object that holds the secret/encrypted data.

NOTE: These parameters must be lowercase in the vault tag.

NOTE: The vault ID value indicates which CyberArk vault SL1 should check first for the credential information. For SL1 12.1.0, any numeric value can be entered for the vault ID.

Example: Vault Tags

In this example, any time the *EM7 Default V2 Tag* credential is used in the **Profile Name** field, SL1 will retrieve the secret from the object *snmp_v2* contained in the **Test** safe as application *testappid*, from the CyberArk credential provider. Then, this secret will be used for the **SNMP Community (Read-Only)** field value.

Edit Credential X

Name *
EM7 Default V2 Tag

All Organizations ● Select the organizations the credential belongs to * Timeout (ms)
1500

SNMP Version	SNMP Retries	Port *
SNMP V2	1	161

SNMP Community
vault:123.appid=testappid.safe=Test:object=snmp_v2 SNMP Community
Read-Only Read/Write

Save & Test

Credential Tester

Select Credential test ▼

Select Collector ▼

IP or Hostname to test *

Test Credential



Save & Close

Overview of Discovery

Overview

This chapter provides an overview of the device discovery process 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>What is Discovery?</i>	107
<i>What Happens During Discovery?</i>	108
<i>What is a Dynamic Application?</i>	110
<i>Before You Run Discovery</i>	111
<i>System Settings that Affect Discovery</i>	112
<i>Device Settings that Affect Auto-Discovery and Re-Discovery</i>	116
<i>How File Systems are Hidden During Discovery</i>	117

What is 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.

SL1 also uses discovery to update existing information about a device and to add to existing information about a device. This type of discovery is called auto-discovery. For each existing device, SL1 automatically runs auto-discovery every night, to keep device data up-to-date.

You can manually trigger discovery at any time and update the data for one device or multiple devices.

What Happens During Discovery?

Discovery is executed:

- During the initial discovery of network devices and applications. This is called **initial discovery**.
- Automatically runs once a day to update the data on each discovered device. This is called **auto-discovery**, or sometimes referred to as "nightly discovery".
- On demand, when a user manually asks SL1 to rediscover a group of IP addresses or a list of fully-qualified domain names, rediscover a single device, or rediscover all devices to find those that should align with a selected Dynamic Application. This is called **re-discovery**.

Discovery uses the following processes:

- **Discovery: Auto (*discover_iprange.py*)**. This process examines one or more IP addresses or fully-qualified domain names and determines which IP addresses or fully-qualified domain names are aligned with a device. For each device, this process retrieves basic information about the device.
- **Discovery: Detail (*discover_detail.py*)**. This process is triggered by the **Discovery: Auto** process and retrieves details about each discovered device.
- **Discovery: Dynamic App (*discover_app.py*)**. This process is triggered manually from the **Dynamic Applications Manager** page and checks all existing devices against the selected Dynamic Application.
- **Discovery: Nightly Update (*discover_update.py*)**. SL1 runs this process each night for already discovered devices. This process updates the information collected during initial discovery.

During discovery of an IP address range (the Discovery: Auto process), SL1 does the following:

1. Performs a DNS lookup to determine the IP address for each fully-qualified domain name supplied in the discovery session. These IP addresses are added to the list of IP addresses supplied in the discovery session.
2. Pings each IP address to determine which are in use.
3. Based on the settings in the discovery session, runs nmap on each IP address to determine which are in use and which ports are open.
4. Searches DNS records to determine the fully qualified domain name at each IP address in range.
5. Tries each selected credential on each IP address to determine if each device is manageable (supports SNMP) or will be a device of type "pingable".

6. For devices that support SNMP, retrieves a system description, SysObject ID, system uptime, system contact, system name, and system location.
7. Assigns a Device Class to each device.

NOTE: If SL1 cannot determine the appropriate Device Class, it will assign the device to the Generic SNMP Device Class.

8. For devices that can be managed (supports SNMP or if non-SNMP discovery is enabled), assigns a device ID, device name, a primary IP address for use in SL1, and a primary credential.

For each device discovered by the Discovery: Auto process, the Discovery: Detail process does the following:

1. If interface discovery is enabled, finds all network interfaces for devices that support SNMP.
2. If the discovery scan level is set to *Initial Population of Apps* or higher, SL1 checks each discovered device (both those that support SNMP and those that don't) against the list of already-defined Dynamic Applications. SL1 searches each discovered device to find "discovery objects" and aligns devices with the appropriate Dynamic Application(s).
3. If the discovery scan level is set to *Discover SSL Certificates* or higher, checks for SSL certificates on port 443 (HTTPS).
4. If the discovery scan level is set to *Discover Open Ports*, *Advanced Port Discovery*, or *Deep Discovery*, SL1 performs a port scan using the settings appropriate for the scan level to determine the open ports for the device.

Immediately after the initial discovery session is completed, SL1 will use the aligned Dynamic Applications to collect additional data from devices. For more information about collection processes, see the manual ***Monitoring Device Infrastructure Health***.

What Happens During Discovery when the SL1 Agent is Installed?

If a device is monitored using the agent and is discovered as a SNMP or pingable device using the Discovery tool, the following default data collection methods, data display settings, and monitoring policies are applied during discovery:

- The method SL1 uses to monitor availability of the device is determined by the first method of discovery:
 - If the agent is installed and creates a device record before the device is discovered as an SNMP or pingable device, availability is based on whether the agent is reporting data to SL1.
 - If the device is discovered as an SNMP or pingable device before the agent is installed, availability is based on the method used to discover the device (SNMP, ICMP, or TCP).
- For Linux devices, the **[TCP/UDP Ports]** tab in the **Device Reports** panel will display open ports detected by the agent and open ports detected by the discovery process (using the NMAP command).
- The **[Processes]** tab in the **Device Reports** panel will display running processes detected by the agent and processes collected using SNMP.

- Port monitoring policies specify whether the policy will be executed using the agent or by a Data Collector using NMAP.
- Process monitoring policies are always executed using the agent.
- Data precedence settings specify whether Dynamic Applications that use data collected by the agent or Dynamic Applications that poll devices for data are used to represent CPU and memory utilization for devices.

For more information about configuring data precedence settings on the agent, or about monitoring ports, processes, and device availability with the agent, see the *Monitoring with the SL1 agent* manual.

What is a Dynamic Application?

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

How Does SL1 Align Dynamic Applications During Discovery?

Most Dynamic Applications include a discovery object. A discovery object enables SL1 to determine which devices to align with a Dynamic Application.

During discovery, SL1:

1. Searches the list of Dynamic Applications.
2. If a Dynamic Application includes a discovery object, SL1 adds that Dynamic Application to the list of Dynamic Applications to try to align during discovery.

3. For each Dynamic Application that includes a discovery object, SL1 checks the current discovery session for an appropriate credential. For example, for each database Dynamic Application, SL1 would look for one or more database credentials that have been selected for the discovery session.
4. For each discovered device, both those that support SNMP and those that don't, discovery tries to determine which Dynamic Applications to align. For each discovered device, SL1 tries to align each Dynamic Application in the list of Dynamic Applications to try during discovery. For each Dynamic Application in the list, SL1 tries to connect to each device with each of the appropriate credentials (until SL1 finds a working credential) and then tries to find the discovery object. If SL1 is able to connect to a device with one of the credentials and can then retrieve the discovery object, SL1 will align the Dynamic Application with the device.

NOTE: SL1 also includes more sophisticated logic that allows you to define multiple discovery objects, validate the value of the discovery object, and to align the Dynamic Application if a discovery object is not available. However, the most common use of a discovery object is as described above (discovery object exists).

5. If discovery aligns a Dynamic Application with a device, immediately after discovery completes SL1 will start the first collection from that device using the aligned Dynamic Application. This step is not performed for Dynamic Applications that meet all of the following three criteria:
 - Has a collection frequency of 1 minute, 2 minutes, 3 minutes or 5 minutes.
 - Does not have component mapping enabled (does not discover component devices).
 - Is aligned with a component device.

NOTE: During discovery, SL1 tries each SNMP credential specified in the discovery session on each discovered device, to determine if SL1 can collect SNMP details from the device. Later in the discovery session, during alignment of Dynamic Applications, discovery again tries each SNMP credential specified in the discovery session. If one of the SNMP credentials times out three times **without any response**, discovery will stop trying to use that SNMP credential to align SNMP Dynamic Applications. Note that "no response" means that a device did not respond at all. Note that if a device reports that "no OID was found" or "the end of the OID tree was reached", these are considered a legitimate response and would not cause SL1 to abandon the credential.

For details on Dynamic Applications, see the appropriate manual for each type of Dynamic Application.

Before You Run Discovery

To make your initial discovery session as productive as possible, you might want to perform these configuration tasks before running discovery:

- Determine the SNMP credentials for the devices and applications in your network. Define correlating credentials in SL1, to allow discovery to retrieve as much information as possible.

- If you want SL1 to immediately start collecting data from devices using Dynamic Applications, you should make sure that each of those Dynamic Applications includes a discovery object.
- If you want SL1 to immediately start collecting data from devices using Dynamic Applications, you should also define any additional credentials required for those Dynamic Applications. For example, if you want SL1 to immediately start monitoring all MySQL databases in your network, you should define credentials that allow SL1 to communicate with each MySQL database in your network. During discovery, SL1 will determine which devices can be monitored with a Dynamic Application for MySQL. After discovery, SL1 will use the database credential to collect data about each MySQL database in your network.
- Defining the global Behavior Settings that affect discovery and auto-discovery. For a description of the Behavior Settings that affect discovery, see the section [System Settings that Affect Discovery](#). For a detailed description of all the Behavior Settings, see the manual [System Administration](#).
- Define one or more organizations. During discovery, all discovered devices will be placed into a single organization. If you do not define an organization, SL1 will place all devices in the **System** organization. However, you can later assign one or more devices to another organization after discovery.
- If you want to perform bulk configuration of discovered devices by using device groups, device templates, or both, you should define device groups and device templates before performing discovery. However, you can apply device groups, device templates or both to one or more devices after initial discovery. For details on device groups and device templates, see the manual [Device Groups and Device Templates](#).

System Settings that Affect Discovery

Some of the parameters in the **Behavior Settings** page affect discovery functionality (discovery, auto-discovery, and re-discovery) in SL1.

NOTE: You can define global parameters for auto-discovery in this page, but you can always override those parameters on a per-device basis by editing the device settings in a device's **Device Properties** page. For details on configuring a device, see the manual [Device Management](#).

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

1. Go to the **Behavior Settings** page (System > Settings > Behavior).
2. In the **Behavior Settings** page, edit the values in one or more of the following fields:
 - **Ping & Poll Timeout (Msec.)**. This field specifies the number of milliseconds both the discovery tool and availability checks will wait for a response after pinging a device. After the specified number of milliseconds have elapsed without a response, the poll will timeout. The choices are from 100 to 5000 milliseconds.
 - **SNMP Poll Timeout (Msec.)**. This field specifies the number of milliseconds the discovery tool will wait for a response after sending an SNMP query to a device. After the specified number of milliseconds have elapsed without a response, the SNMP poll will timeout. The choices are from 100 to 5000 milliseconds.

- **SNMP Failure Retries.** This field specifies the number of times the discovery tool will try to communicate with a device after a timeout or failure. After that number of times has been met, the discovery tool will not retry unless the user manually restarts the discovery process. The choices are 0–6.
- **DHCP Community Strings.** SNMP "read only" community string, to use during discovery. This is required only if DHCP servers and devices use a different SNMP community string than other devices in the network. If the community string specified in the **Discovery Control Panel** (System > Manage > Classic Discovery) page does not work for DHCP devices, SL1 will automatically use the community string specified in this field.
- **NFS Detection Disable.** If selected, this checkbox prevents SL1 from monitoring and reporting on NFS "shared" hard drives. SL1 will monitor and report only on local hard drives.
- **Port Polling Type.** Specifies how SL1 should poll devices to discover open ports. The choices are:
 - *Half Open.* Uses a faster TCP/IP connection method and does not appear on device's logs.
 - *Full Connect.* Uses the standard TCP/IP connection to detect open ports.
- **Initial Discovery Scan Level.** Specifies the data to be gathered during the initial discovery session. You can override this setting for a single discovery session in the **Discovery Session Editor** modal page. The options are:
 - *0. Model Device Only.* Discovery tool will discover if device is up and running and if so, collect the make and model of the device. SL1 will then generate a device ID for the device, so it can be managed by SL1.
 - *1. Initial Population of Apps.* Discovery tool will search for Dynamic Applications to associate with the device. Discovery tool will attempt to collect data for the aligned Dynamic Applications. Discovery will also perform *0. Model Device Only* discovery.
 - *2. Discover SSL Certificates.* Discovery tool will search for SSL certificates and retrieve SSL data. Discovery tool will also perform *1. Initial Population of Apps* and *0. Model Device Only*.
 - *3. Discover Open Ports.* Discovery tool will search for open ports. Discovery tool will also perform *2. Discover SSL Certificates*, *1. Initial Population of Apps*, and *0. Model Device Only*. If your system includes a firewall and you select *3. Discover Open Ports*, discovery might be blocked and/or might be taxing to your network.
 - *4. Advanced Port Discovery.* Discovery tool will search for open ports, using a faster TCP/IP connection method. Discovery tool will also perform *2. Discover SSL Certificates*, *1. Initial Population of Apps*, and *0. Model Device Only*. If your system includes a firewall and you select *4. Advanced Port Discovery*, some auto-discovered devices might remain in a pending state (purple icon) for some time after discovery. These devices will achieve a healthy status, but this might take several hours.
 - *5. Deep Discovery.* Discovery tool will use nmap to retrieve operating system name and version. Discovery will also scan for services running on each open port and can use this information to match devices to device classes. Discovery tool will search for open ports, using a faster TCP/IP connection method. Discovery tool will also perform *2. Discover SSL Certificates*, *1. Initial Population of Apps*, and *0. Model Device Only*. For devices that don't

support SNMP, option 5. *Deep Discovery* allows you to discover devices that don't support SNMP and then align those devices with a device class other than "pingable".

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

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

TIP: ScienceLogic recommends that you delete all unused PowerPacks from your SL1 system to improve the performance of the nightly auto-discovery process.

- **Discovery Scan Throttle.** Specifies the amount of time a discovery process should pause between each IP address or hostname in a discovery session. (You specify the list of IP addresses or hostnames for a discovery session in the **IP Address/Hostname Discovery List** field in the **Discovery Session Editor** page.) Pausing discovery processes between IP addresses or hostnames spreads the amount of network traffic generated by discovery over a longer period of time. The choices are:
 - *Disabled.* Discovery processes will not pause.
 - *1000 Msec to 10000 Msec.* A discovery process will pause for a random amount of time between half the selected value and the selected value.
- **Port Scan All IPs.** Specifies whether SL1 should scan all IP addresses on a device for open ports. You can override this setting for a single discovery session in the **Discovery Session Editor** modal page. The choices are:
 - *0. Disabled.* SL1 will scan only the primary IP address (the one used to communicate with SL1) for open ports.
 - *1. Enabled.* SL1 will scan all discovered IP addresses for open ports.
- **Port Scan Timeout.** Length of time, in milliseconds, after which SL1 should stop trying to scan an IP address for open ports and begin scanning the next IP address (if applicable). You can override this setting for a single discovery session in the **Discovery Session Editor** modal page. Choices are between 60,000 and 1,800,000 milliseconds.
- **Restart Windows Services (Agent required).** Specifies whether SL1 should automatically restart failed Windows services that have been defined on the device with a startup type of "automatic". To use this feature, the managed device must be running the agent SNMP Informant, WMI Edition. For assistance or information on purchasing and installing this agent, please contact ScienceLogic. Users must also supply a value in the **SNMP Write** field in the **Device Properties** page for the device. The choices are:

- *0. Disabled.* SL1 will not automatically restart failed Windows services that have been defined on the device with a startup type of "automatic".
 - *1. Enabled.* SL1 will automatically restart failed Windows services that have been defined on the device with a startup type of "automatic".
- **Hostname Precedence.** Specifies which name SL1 will use for each discovered device. Choices are:
 - *SNMP System Name.* Use the device name specified in the device's SNMP System MIB. If *SNMP System Name* is selected and SL1 cannot find an SNMP name for the device, SL1 will assign the name returned by the DNS Reverse Lookup. If SL1 cannot find a DNS Reverse Lookup name for the device, SL1 will use the device's Admin Primary IP address as the device name in SL1.
 - *DNS Reverse Lookup.* Use the device name specified in the device's reverse-lookup record.
 - **Event Interface Name Format.** Specifies the format of the network interface name that you want to appear in events. If you selected *Interface Alias* for the deprecated **Interface Name Precedence** field in a previous release of SL1, the format for existing interfaces is set to {alias}. If you selected *Interface Name* for the deprecated **Interface Name Precedence** field in a previous release of SL1, the format for existing interfaces is set to {name}. The default format is {name}. You can use a combination of string text and the following tokens to define the interface name format for events, such as string_{name}, string_{alias}, {name}{alias}, or {ifdesc}:
 - {alias}
 - {name}
 - {state}
 - {ifdescr}
 - {if_id}
 - {did}
 - {ifindex}
 - {ifphysaddress}
 - {iftype}
 - {ifspeed}
 - {ifhighspeed}
 - {ifoperstatus}
 - {ifadminstatus}
 - **DNS Hostnames.** If SL1 will use the DNS Reverse Lookup name as the device name (see the description of the field **Hostname Precedence**), this field specifies whether SL1 will use the fully-qualified domain name or only the hostname for each discovered device. Choices are:

- *Strip Device Name (Hostname)*. SL1 will use only the device name as the DNS hostname for each device.
- *Use Full Domain Name (FQDN)*. SL1 will use the fully-qualified domain name as the device name for each device.

3. Click **[Save]**.

Device Settings that Affect Auto-Discovery and Re-Discovery

For each discovered device, the following settings in the **Device Properties** page (Devices > Device Manager > wrench icon) affect how nightly auto-discovery and rediscovery behaves for that device. These settings override any global settings defined in the **Behavior Settings** page (System > Settings > Behavior):

- **SNMP Read /Write**. The first drop-down lets you select an SNMP Read credential for read access to SNMP information on the device. The second drop-down let you select a n SNMP Write credential for read-and-write access to SNMP information on the device.
- **Availability Port**. Specifies the protocol and specific port SL1 should monitor to determine if the device is available. The list of ports will contain all the ports discovered by SL1. The data collected from this port will be used in device availability reports.
- **Run Availability Policy** (🔧). When you select this icon, SL1 immediately checks the availability of the device, using the port and protocol specified in the **Availability Port** fields. SL1 displays a **Session Logs** modal page that displays a detailed description of each step of the availability policy. This information is helpful when troubleshooting availability problems with a device.
- **Latency Port**. Specifies the protocol and specific port SL1 should monitor to determine latency for the device. The list of ports will contain all the ports discovered by SL1. The data collected from this port will be used in device latency reports.
- **Avail + Latency Alert**. Specifies how SL1 should respond when the device fails an availability check, when the device fails a latency check, and when the device fails both. These options allow you to create separate events when SNMP fails on a device and when a device is not up and running.
- **User Maintenance**. Specifies whether the device will be put into "user maintenance" mode. By default, when a device is in "user maintenance", SL1 will not generate events about the device.
- **User Maintenance Collection**. Specifies whether SL1 should poll the device during the "user maintenance". During normal operation, SL1 polls each device as specified by each device's policies and aligned Dynamic Applications.
- **Collection**. Specifies (among other things) if the device will be automatically updated each night with SL1's auto-discovery tool. To edit this field, select one of the following from the drop-down list:
 - *Enabled*. Device will be polled by the auto-discovery tool.
 - *Disabled*. Device will not be polled by the auto-discovery tool.
- **Collector Group**. Specifies which collector group will gather data from the device. You can select from a list of available collector groups.
- **Coll. Type**. Specifies how SL1 should perform auto-discovery. The choices are:

- *Standard*. SL1 will perform auto-discovery of each device based on the device's IP address. This method is appropriate for devices using standard DNS.
- *DHCP*. SL1 will perform auto-discovery of each device based on the device's MAC address. This method is appropriate for devices using DHCP.
- **Critical Ping**. Frequency with which SL1 should ping the device in addition to the five minute availability poll. If the device does not respond, SL1 creates an event.
- **Event Mask**. Events that occur on a single device within the selected time-interval are grouped together. This allows related events that occur in quick succession on a single device to be rolled-up and posted together, under one event description.

How File Systems are Hidden During Discovery

When you hide a file system:

- SL1 stops collecting information about the file system.
- SL1 does not generate events about the file system.
- SL1 does not monitor the file system for thresholds (defined in the **Device Thresholds** and **Global Threshold Settings** pages).
- SL1 does not include the file system in the **Device Summary** page.
- SL1 does not include the file system in file system reports in the **Device Performance** page.

The following rules are applied during discovery to automatically hide file systems:

- If the **NFS Detection Disable** checkbox is selected in the **Behavior Settings** page (System > Settings > Behavior), NFS file systems are automatically hidden during discovery.
- File systems of type "iso9660" are automatically hidden during discovery.
- File systems for which the storage size is not reported or the storage size is less than 1024 MB are automatically hidden during discovery.
- File systems of type "Other" are automatically hidden during discovery.

NOTE: If the type for a discovered file system changes, the auto-hide rules are re-applied to that file system. For example, suppose a Windows drive letter is initially discovered as a removable disk and is auto-hidden. If that drive-letter is later re-used for a fixed drive, this change will cause the file system to be automatically un-hidden.

NOTE: Hidden file systems are not monitored, only discovered. If the file system is monitored by a Dynamic Application, alerts or events can still be generated for a hidden file system. This only applies to internal collection alerts. For more information, see Knowledge Base article: <https://support.sciencelogic.com/s/article/4025>

To manually hide one or more file systems:

1. Go to the **Device Hardware** page (Devices > Hardware).
2. Filter the list to display only **Comp Type** of "file system".
3. Select the checkbox for one or more file systems you would like to hide.
4. From the **Select Actions** field (in the lower right), select *Hide File Systems*.
5. Click the **[Go]** button.
6. Each selected file system will be hidden in SL1 .

To manually unhide one or more file systems:


1. Go to the **Device Hardware** page (Devices > Hardware).
2. Filter the list to display only **Comp Type** of "file system".
3. Select the checkbox for one or more file systems you would like to unhide.
4. From the **Select Actions** field (in the lower right), select *Unhide File Systems*.
5. Click the **[Go]** button.
6. SL1 will resume collection for each selected file system and will include each selected file system in the **Device Summary** and **Device Performance** pages.

Chapter



5

Discovering Devices

Overview

This chapter describes how to use discovery in SL1 to find devices on your network. You can use the **[Add Devices]** button on the **Devices** () page or the **Discovery Sessions** page (Devices > Discovery Sessions) to start a discovery process, or you can run a "Classic Discovery" from the **Discovery Control Panel** page (System > Manage > Classic Discovery).

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>Prerequisites for Discovering Devices</i>	119
<i>Adding Devices Using Guided Discovery</i>	120
<i>Adding Devices Using Unguided Discovery</i>	123
<i>Working with Discovery Sessions</i>	129
<i>Managing Classic Discovery Sessions</i>	129

Prerequisites for Discovering Devices


To discover all of the devices on your network:

1. **Make a note of the range of IP addresses used on your network.** If your device does not have an IP address, make a note of the name of the root device. If you need help, ask your network administrator.

2. **An Organization must exist in SL1 for the new devices.** If you need to create an Organization, go to the **Organizations** page (Registry > Accounts > Organizations).
3. **A Collector Group must exist in SL1 that can reach the target device using a valid network path for the needed protocol.** For example, UDP 161 for SNMP and general ICMP traffic for Ping. If you are not sure of the Collector Group to use, consult an SL1 Architecture diagram or ask your SL1 System Administrator. You can access collector information on the **Collector Group Management** page (System > Settings > Collector Groups).
4. **You must create or use an existing credential.** You can access credential information on the **Credentials** page (Manage > Credentials). Because credential problems are the most common cause for discovery failure, you can test any credential that you create on the **Credential Test Management** page (System > Customize > Credential Tests).
5. **If you want to use a device template with a discovery session, you must use an existing template in SL1.** You can access device templates on the **Configuration Templates** page (Devices > Templates).
6. **The Grant All user needs to be used to access new discovery workflow,** as the SYS_SETTINGS_LICENSES_PAGE and SYS_SETTINGS_CUGS_PAGE access keys are needed to get collector or collector group information. For more information, see the **Access Keys** page (System > Manage > Access Keys).

NOTE: If you want to discover one of the third-party products that are available as an option when using guided discovery, you must have the corresponding PowerPack installed on your SL1 system. For example, if you want to discover an Amazon Web Services account, you must have the *Amazon Web Services* PowerPack installed.


Adding Devices Using Guided Discovery

On the **Devices** page () or the **Discovery Sessions** page (Devices > Discovery Sessions), you can add or "discover" new devices for monitoring in SL1. You add devices by creating a **discovery session**, which searches for devices on the network you specify.

The **guided discovery** process lets you select a discovery type specific to the type of devices you want to monitor, in addition to traditional SNMP discovery. The guided discovery wizard provides a filtered list of relevant credentials, the ability to create new credentials, and a reduced set of application-specific fields to help you efficiently discover the devices you need.

The following procedure uses Amazon Web Services as an example of the discovery type. Some steps and fields will vary depending on the discovery type.

To run a guided discovery:

1. On the **Devices** page () or the **Discovery Sessions** page (Devices > Discovery Sessions), click the **[Add Devices]** button. The **Select** page appears.
2. Select the type of devices you want to discover. You can choose one of the following device types:
 - Alibaba Cloud
 - AWS
 - AWS EC2
 - AWS IAM

- AWS Assume Role
- Azure
- Citrix
- Google Cloud
- IBM
- Ping
- SNMP
- VMware
- Windows
 - Windows WMI
 - Windows PowerShell
 - Windows Agent

TIP: Additional information about the requirements for device discovery appears in the **General Information** pane to the right.

NOTE: If you want to discover one of the third-party products that are available as an option when using guided discovery, you must have the corresponding PowerPack installed on your SL1 system. For example, if you want to discover an Amazon Web Services account, you must have the "Amazon Web Services" PowerPack installed.

WARNING: When executing AWS guided workflow discoveries, executing the same workflow or workflows with similar settings can result in asset duplication. For IAM guided workflows, this will result in completely duplicated account device component trees. For other AWS workflows, this might result in duplicated virtual devices that represent the AWS organization.

NOTE: If you want to do a more general discovery, you can select one of the options in the **Other ways to add devices** pane, such as **Unguided Network Discovery**. For more information, see [Adding Devices Using Unguided Discovery](#).

3. Click **[Select]**. The **Credential Selection** page appears, which is the first step in the guided discovery session.

NOTE: The contents of this page might vary depending on the discovery type you selected at the start of the Guided Discovery.

4. If the credential you want to use for discovery is already listed on the **Credential Selection** page and does not need to be edited, proceed to step 5. Otherwise, you can optionally do one of the following:
 - If the credential you need is not in the list, click the **[Create New]** button to open the **Create Credential** window, where you can specify the name and organization for the credential, the third-party username and password, and other data such as Cloud Type and Proxy information. You can also test the credential before you save using the **Credential Tester** panel. Click **[Save & Close]** to save the credential and return to the **Credential Selection** page of the guided discovery session. For more information on creating new credentials or testing credentials, see the section on [Defining Credentials](#) or [Using the Credential Tester Panel](#).
 - To edit a credential on the **Credential Selection** page, click the name of the credential you would like to edit from the **Name** column and edit that credential as needed. You can also test the credential before you save using the **Credential Tester** panel. Click the **[Save & Close]** button on the **Edit Credential** window to save your updates.

NOTE: During the guided discovery process, you cannot click **[Next]** until the required fields are filled on the page.

5. On the **Credential Selection** page of the guided discovery process, select a credential to allow SL1 to access the device that will act as the root device for the system being discovered, and then click **[Next]**. The **Discovery Session Name** page appears, which is the second step in the guided discovery session.

NOTE: The contents of this page might vary depending on the discovery type you selected at the start of the Guided Discovery.

6. Complete the following fields, as applicable:
 - **Discovery Session Name.** Type a name for the discovery session.
 - **Root Device Name.** Type the name of the root device for the application you want to monitor.
 - **Scope of Discovery.** If you are discovering Google Cloud devices, select whether you are discovering them at the *Organization Level* or the *Project Level*. For more information, see the **Monitoring Google Cloud** manual, the section on Monitoring Google Cloud.
 - **Select the organization to add discovered devices to.** Select the name of the organization to which you want to add the discovered device.
 - **Collector Group Name.** Select an existing collector group to communicate with the discovered device. This field is required.

NOTE: When assigning devices to a collector group, SL1's multi-tenancy rules will validate that the collector group you select belongs to the organization you selected in the previous field. If you attempt to run a discovery session where the devices, collector group, and credentials do not all belong to the same organization, you will receive an error message and will not be able to save or execute the discovery session.

- **IP.** Type the IP addresses for SL1 to scan during discovery.

NOTE: Some applications, such as VMware, require that you enter an IP address for the root device as well.

7. Click **[Next]**. SL1 creates a root device with the appropriate Device Class assigned to it and aligns the relevant Dynamic Applications. The **Device Discovery Completed** page appears, which is the third and final step of the guided discovery session. As SL1 discovers your devices, system messages relating to the discovery appear on the page under the heading "Discovery Logs".

NOTE: If SL1 cannot determine the appropriate Device Class, it will assign the device to the Generic SNMP Device Class.


8. When discovery has finished, click **[Close]**.

NOTE: The results of a guided discovery do not display on the **Discovery Sessions** page (Devices > Discovery Sessions). However, you can retrieve details of saved Guided Discovery Sessions with the `guidedDiscoverySessions` GraphQL query. Details for discovery sessions that create a virtual root device are not currently displayed in the user interface.

Adding Devices Using Unguided Discovery

Instead of running a Discovery for a specific discovery type, you can run an "unguided" discovery to find a range of devices using core credentials such as SNMP, Database, SOAP/XML, Basic/Snippet, SSH/Key, or PowerShell credentials.

To run an unguided discovery:

1. On the **Devices** page () or the **Discovery Sessions** page (Devices > Discovery Sessions), click the **[Add Devices]** button. The **Select** page appears.
2. Click the **[Unguided Network Discovery Workflow]** button. Additional information about the requirements for discovery appears in the **General Information** pane to the right.
3. Click **[Select]**. The **Add Devices** page appears.
4. Complete the following fields:

- **Discovery Session Name.** Type a unique name for this discovery session. This name is displayed in the list of discovery sessions on the **[Discovery Sessions]** tab.
 - **Description.** Type a short description of the discovery session. You can use the text in this description to search for the discovery session on the **[Discovery Sessions]** tab. Optional.
 - **Select the organization to add discovered devices to.** Select the name of the organization to which you want to add the discovered devices.
5. Click **[Next]**. The **Credentials** page of the **Add Devices** wizard appears.
 6. On the **Credentials** page, you can optionally do one of the following:
 - If the credential you need is not in the list, click the **[Create New]** button to open the **Create Credential** window, where you can specify the name and organization for the credential, the third-party username and password, and other data such as Cloud Type and Proxy information. You can also test the credential before you save using the **Credential Tester** panel. Click **[Save & Close]** to save the credential and return to the **Credential Selection** page of the guided discovery session. For more information on creating new credentials or testing credentials, see the section on [Defining Credentials](#) or [Using the Credential Tester Panel](#).
 - To edit a credential on the **Credential Selection** page, click the name of the credential you would like to edit from the **Name** column and edit that credential as needed. You can also test the credential before you save using the **Credential Tester** panel. Click the **[Save & Close]** button on the **Edit Credential** window to save your updates.
 7. On the **Credentials** page of the **Add Devices** wizard, select one or more credentials to allow SL1 to access a device's SNMP data and click **[Next]**. The **Discovery Session Details** page of the **Add Devices** wizard appears.
 8. Complete the following fields:
 - **List of IPs/Hostnames.** Provide a list of IP addresses, hostnames, or fully-qualified domain names for SL1 to scan during discovery. This field is required. In this field, you can enter a combination of one or more of the following:
 - One or more *single IPv4 addresses* separated by commas and a new line. Each IP address must be in standard IP notation and cannot exceed 15 characters. For example, "10.20.30.1, 10.20.30.2, 10.20."
 - One or more *ranges of IPv4 addresses* with "-" (dash) characters between the beginning of the range and the end of the range. Separate each range with a comma. For example, "10.20.30.1 – 10.20.30.254".
 - One or more IP address ranges in *IPv4 CIDR notation*. Separate each item in the list with a comma. For example, "192.168.168.0/24".
 - One or more *ranges of IPv6 addresses* with "-" (dash) characters between the beginning of the range and the end of the range. Separate each range with a comma. For example, "2001:DB8:0:0:0:0:0:0-2001:DB8:0:0:0:0:0:0003".
 - One or more IP address ranges in *IPv6 CIDR notation*. Separate each item in the list with a comma. For example, "2001:DB8:0:0:0:0:0:0/117".
 - One or more hostnames (fully-qualified domain names). Separate each item in the list with a comma.

TIP: You can also click the **[Upload File]** button to upload a comma-separated list of IPs.

- **Which collector will monitor these devices?** Select an existing collector group to monitor the discovered devices. Required.

NOTE: When assigning devices to a collector group, SL1's multi-tenancy rules will validate that the collector group you select belongs to the organization you selected in the previous field. If you attempt to run a discovery session where the devices, collector group, and credentials do not all belong to the same organization, you will receive an error message and will not be able to save or execute the discovery session.

- **Run after save.** Select this option to run this discovery session as soon as you click **[Save and Close]**.
- **Advanced options.** Click the down arrow icon (▼) to access additional discovery options.

In the **Advanced options** section, complete the following fields as needed:

- **Initial Scan Level.** For this discovery session only, specifies the data to be gathered during the initial discovery session. The options are:
 - *System Default (recommended).* Use the value defined in the **Behavior Settings** page (System > Settings > Behavior) in the classic user interface of SL1.
 - *1. Model Device Only.* Discovery will discover if the device is up and running and if so, collect the make and model of the device. SL1 will then generate a device ID for the device so it can be managed by SL1.
 - *2. Initial Population of Apps.* Discovery will search for Dynamic Applications to associate with the device. The discovery tool will attempt to collect data for the aligned Dynamic Applications. Discovery will later retrieve full sets of data from each Dynamic Application. Discovery will also perform *1. Model Device Only* discovery.
 - *3. Discover SSL Certificates.* Discovery will search for SSL certificates and retrieve SSL data. Discovery will also perform *2. Initial Population of Apps* and *1. Model Device Only*.
 - *4. Discover Open Ports.* Discovery will search for open ports. Discovery will also perform *3. Discover SSL Certificates*, *2. Initial Population of Apps*, and *1. Model Device Only*.

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

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

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

- 6. *Deep Discovery*. Discovery will use nmap to retrieve the operating system name and version. Discovery will also scan for services running on each open port and can use this information to match devices to device classes. Discovery will search for open ports, using a faster TCP/IP connection method. Discovery will also perform 3. *Discover SSL Certificates*, 2. *Initial Population of Apps*, and 1. *Model Device Only*.

NOTE: For devices that don't support SNMP, option 6. *Deep Discovery* allows you to discover devices that don't support SNMP and then align those devices with a device class other than "pingable". Note that option 6. *Deep Discovery* is compute-intensive.

NOTE: If SL1 cannot determine the appropriate Device Class, it will assign the device to the Generic SNMP Device Class.

- **Scan Throttle**. Specifies the amount of time a discovery process should pause between each specified IP address (specified in the **IP Address/Hostname Discovery List** field). Pausing discovery processes between IP addresses spreads the amount of network traffic generated by discovery over a longer period of time. The choices are:
 - *System Default (recommended)*. Use the value defined in the **Behavior Settings** page (System > Settings > Behavior) in the classic user interface for SL1.
 - *Disabled*. Discovery processes will not pause.
 - *1000 Msec to 10000 Msec*. A discovery process will pause for a random amount of time between half the selected value and the selected value.
- **Port Scan All IPs**. For the initial discovery session only, specifies whether SL1 should scan all IP addresses on a device for open ports. The choices are:
 - *System Default (recommended)*. Use the value defined in the **Behavior Settings** page (System > Settings > Behavior) in the classic user interface for SL1.
 - *Enabled*. SL1 will scan all discovered IP addresses for open ports.
 - *Disabled*. SL1 will scan only the primary IP address (the one used to communicate with SL1) for open ports.
- **Port Scan Timeout**. For the initial discovery session only, specifies the length of time, in milliseconds, after which SL1 should stop trying to scan an IP address for open ports and begin scanning the next IP address (if applicable). Choices are:

- *System Default (recommended)*. Use the value defined in the **Behavior Settings** page (System > Settings > Behavior).
 - Choices between 60 to 1,800 seconds.
- **Scan Ports**. Specify a list of ports to scan, separated by colons (:). The default is 21:22:25:80:136.
- **Interface Inventory Timeout (ms)**. Specifies the maximum amount of time that the discovery processes will spend polling a device for the list of interfaces. After the specified time, SL1 will stop polling the device, will not model the device, and will continue with discovery. The default value is 600,000 ms (10 minutes).
 - During the execution of this discovery session, SL1 uses the value in this field first. If you delete the default values and do not specify another value in this field, SL1 uses the value in the **Global Threshold Settings** page (System > Settings > Thresholds).
 - If you specify a value in this field and do not apply a device template to this discovery session, the **Interface Inventory Timeout** setting in the **Device Thresholds** page (Registry > Devices > Device Manager > wrench icon > Thresholds) is set to this value for each discovered device. If there is no device template applied to the discovery session and no value is supplied in this field, SL1 uses the value in the **Global Threshold Settings** page (System > Settings > Thresholds).
- **Maximum Allowed Interfaces**. Specifies the maximum number of interfaces per devices. If a device exceeds this number of interfaces, SL1 will stop scanning the device, will not model the device, and will continue with discovery. The default value is 10,000.
 - During the execution of this discovery session, SL1 uses the value in this field first. If you delete the default values and do not specify another value in this field, SL1 uses the value in the **Global Threshold Settings** page.
 - If you specify a value in this field and do not apply a device template to this discovery session, the **Maximum Allowed Interfaces** setting in the **Device Thresholds** page is set to this value for each discovered device. If there is no device template applied to the discovery session and no value is supplied in this field, SL1 uses the value in the **Global Threshold Settings** page.
- **Bypass Interface Inventory**. Specifies whether or not the discovery session should discover network interfaces.
 - *Selected*. SL1 will not attempt to discover interfaces for each device in the discovery session. For each discovered device, the **Bypass Interface Inventory** checkbox on the **Device Investigator [Settings]** tab will be selected.
 - *Not Selected*. SL1 will attempt to discover network interfaces, using the **Interface Inventory Timeout** value and **Maximum Allowed Interfaces** value.
- **Discover Non-SNMP**. Specifies whether or not SL1 should discover devices that don't respond to SNMP requests.
 - *Selected*. SL1 will discover devices that don't respond to the SNMP credentials selected in the **SNMP Credentials** field. These devices will be discovered as "pingable" devices.

- *Not Selected.* SL1 will not discover devices that don't respond to the SNMP credentials selected in the **SNMP Credentials** fields.

NOTE: You must either select a credential for the discovery session or select the **Discover Non-SNMP** option. SL1 will prevent you from proceeding with discovery if you have not met those conditions.

- **Model Devices.** Determines whether or not the devices that are discovered with this discovery session can be managed through SL1. Choices are:
 - *Enabled.* When a device is modeled, SL1 creates a device ID for the device; you can then access the device through the **Device Manager** page and manage the device in SL1.
 - *Disabled.* If a device is not modeled, you cannot access the device through the **Device Manager** page, and you cannot manage the device in SL1. However, each discovered device will still appear in the Discovery Session logs. For each discovered device, the discovery logs will display the IP address and device class for the device. This option is useful when performing an initial discovery of your network, to determine which devices you want to monitor and manage with SL1. For the amount of time specified in the **Device Model Cache TTL (h)** field, a user can manually model the device from the **Discovery Session** window.
 - **Enable DHCP.** Specifies whether or not the specified range of IPs and hostnames use DHCP.
 - *Selected.* SL1 will perform a DNS lookup for the device during discovery and each time SL1 retrieves information from the device.
 - *Not Selected.* SL1 will perform normal discovery.
 - **Device Model Cache TTL (h).** Amount of time, in hours, that SL1 stores information about devices that are discovered but not modeled, either because the **Model Devices** option is not enabled or because SL1 cannot determine whether a duplicate device already exists. The cached data can be used to manually model the device from the **Discovery Session** window.
 - **Log All.** Specifies whether or not the discovery session should use verbose logging. When you select verbose logging, SL1 logs details about each IP address or hostname specified in the **IP Address/Hostname Discovery List** field, even if the results are "No device found at this address."
 - *Selected.* This discovery session will use verbose logging.
 - *Not Selected.* This discovery session will not use verbose logging.
 - **Apply Device Template.** As SL1 discovers a device in the IP discovery list, that device is configured with the selected device template. You can select from a list of all device templates in SL1. For more information on device templates, see the manual on **Device Groups and Device Templates**.
9. Click **[Save and Close]** to save the discovery session. The **Discovery Sessions** page (Devices > Discovery Sessions) displays the new discovery session.
 10. If you selected the **Run after save** option on this page, the discovery session runs, and the **Discovery Logs** page displays any relevant log messages. If the discovery session locates and adds any devices, the **Discovery Logs** page includes a link to the **Device Investigator** page for the discovered device.

Working with Discovery Sessions

The **Discovery Sessions** page (Devices > Discovery Sessions) displays a list of all the existing **discovery sessions**, which are previous attempts to add devices using discovery.

On this page you can click the **[Actions]** button (☰) for a session and select one of the following actions:

- *Edit*. Run the **Add Device** wizard again so you can make changes to the selected discovery session.
- *Delete*. Delete the selected discovery session. You do not get a confirmation window after you click *Delete*; the session is immediately deleted.
- *Start*. Run the selected discovery session again. The **Discovery Logs** page appears when discovery completes.
- *Show Logs*. The **Discovery Logs** page for the selected discovery session displays data about the most recent run of a discovery session.

TIP: You can adjust the size of the rows and the size of the row text on this inventory page. For more information, see the section on "Adjusting the Row Density" in the *Introduction to SL1* manual.

From the **Discovery Sessions** page, you can also add devices using guided or unguided discovery by clicking the **[Add Devices]** button. For instructions on using guided discovery, see the section on [Adding Devices Using Guided Discovery](#). For instructions on using unguided discovery, see the section on [Adding Devices Using Unguided Discovery](#).

Managing Classic Discovery Sessions

The following sections describe how to view, create, and manage classic discovery sessions in SL1.

Viewing Information about a Classic Discovery Session

The **Session Register** pane in the **Discovery Control Panel** (System > Manage > Classic Discovery) page displays information about all discovery sessions defined in SL1.

TIP: To sort the list of discovery sessions, click on a column heading. The list will be sorted by the column value, in ascending order. To sort by descending order, click the column heading again. The **Last Edit** column sorts by descending order on the first click; to sort by ascending order, click the column heading again.

For each session, the **Session Register** displays:

- **Session Name.** Name of the discovery session. This field is optional.
- **IP/Hostname List.** The range of IP addresses and/or hostnames for SL1 to scan during discovery. This field can contain a combination of one or more of the following:
 - One or more *single IPv4 addresses* separated by commas. Each IP address must be in standard IP notation and cannot exceed 15 characters. For example, "10.20.30.1, 10.20.30.2, 10.20.30.3".
 - One or more *ranges of IPv4 addresses* with "-" (dash) characters between the beginning of the range and the end of the range. Separate each range with a comma. For example, "10.20.30.1 – 10.20.30.254".
 - One or more IP address ranges in *IPv4 CIDR notation*. Separate each item in the list with a comma. For example, "192.168.168.0/24".
 - One or more *ranges of IPv6 addresses* with "-" (dash) characters between the beginning of the range and the end of the range. Separate each range with a comma. For example, "2001:DB8:0:0:0:0:0:0-2001:DB8:0:0:0:0:0:0003".
 - One or more IP address ranges in *IPv6 CIDR notation*. Separate each item in the list with a comma. For example, "2001:DB8:0:0:0:0:0:0/117".
 - One or more fully-qualified domain names or hostnames.

NOTE: The following types of notation are **not supported**: IPv4 netmask with comma notation (192.168.168.0,24); a list of single IPv6 addresses, separated by comma.

- **Collector.** Data Collector used for the discovery session.
- **Organization.** The organization to which devices discovered during the discovery session will be assigned.
- **Pings.** This field specifies whether or not SL1 should discover devices that don't respond to the selected SNMP credentials. The possible values are:
 - Yes. SL1 will discover devices that don't respond to the SNMP credentials selected in the **SNMP Credentials** field for the discovery session. These devices will be discovered as "pingable" devices.
 - No. SL1 will not discover devices that don't respond to the SNMP credentials selected in the **SNMP Credentials** field for this discovery session.
- **Rediscovery.** Specifies whether or not SL1 is scheduled to run this discovery session, and if so, the frequency and time specified in the schedule.
- **User Edit.** Name of user who created or last edited the discovery session.
- **Last Edit.** Date and time discovery session was created or last edited.

To filter the list of discovery sessions in the **Session Register**, use the search fields at the top of each column. The search fields are find-as-you-type filters; as you type, the page is filtered to match the text in the search field, including partial matches. Text matches are not case-sensitive. Additionally, you can use the following special characters in each filter:

- , (comma). Specifies an "or" operation. For example:
"dell, micro" would match all values that contain the string "dell" OR the string "micro".
- & (ampersand). Specifies an "and" operation. For example:
"dell & micro" would match all values that contain the string "dell" AND the string "micro".
- ! (exclamation mark). Specifies a "not" operation. For example:
"!dell" would match all values that do not contain the string "dell".
- ^ (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.

Running a Classic Discovery Session

To perform a discovery session for one IP address, multiple IP addresses, or a range of IP addresses on the **Classic Discovery** page:

NOTE: To discover all the devices in your network, you must first know the range of IP addresses used in your network. If you need help, ask your network administrator.

1. Go to the **Discovery Control Panel** page (System > Manage > Classic Discovery).
2. In the **Discovery Control Panel**, click **[Create]**. The **Discovery Session Editor** page appears:
3. Supply values in the following fields:
 - **Name.** Type a name for the discovery session. This name is displayed in the list of discovery sessions in the **Discovery Control Panel** page.
 - **Description.** Optionally, type a description of the discovery session.
 - **IP Address/Hostname Discovery List.** Provide a list of IP addresses or fully-qualified domain names for SL1 to scan during discovery. In this field, you can enter a combination of one or more of the following:

NOTE: Instead of manually entering a list of IP addresses and hostnames, you can upload a file that contains the list of IP addresses and hostnames. See the description of the **Upload File** field.

- One or more *single IPv4 addresses* separated by commas. Each IP address must be in standard IP notation and cannot exceed 15 characters. For example, "10.20.30.1, 10.20.30.2, 10.20.30.3".
- One or more *ranges of IPv4 addresses* with "-" (dash) characters between the beginning of the range and the end of the range. Separate each range with a comma. For example, "10.20.30.1 – 10.20.30.254".
- One or more IP address ranges in *IPv4 CIDR notation*. Separate each item in the list with a comma. For example, "192.168.168.0/24".
- One or more *ranges of IPv6 addresses* with "-" (dash) characters between the beginning of the range and the end of the range. Separate each range with a comma. For example, "2001:DB8:0:0:0:0:0:0-2001:DB8:0:0:0:0:0:0003".
- One or more IP address ranges in *IPv6 CIDR notation*. Separate each item in the list with a comma. For example, "2001:DB8:0:0:0:0:0:0/117".
- One or more hostnames (fully-qualified domain names). Separate each item in the list with a comma.

CAUTION: If you enter both the hostname and IP address of the same devices, SL1 will discover two duplicate devices.

NOTE: The following types of notation are **not supported**: IPv4 netmask with comma notation (e.g., 192.168.168.0,24); a list of single IPv6 addresses, separated by comma.

NOTE: SL1 will display an error if your discovery session exceeds the maximum size for optimum performance. SL1 will display a warning message if your discovery session includes 100 or more IP addresses. The warning message will tell you that discovery with more than 100 IP addresses might "take a long time to discover".

- **Upload File.** Instead of manually entering a list of IP addresses and hostnames in the **IP Address/Hostname Discovery List** field, you can upload a file that contains a list of IP addresses and hostnames. The IP addresses and hostnames in the file must be in a format that is allowed for the **IP Address/Hostname Discovery List** field. Each address or range of addresses in the file must be separated by a newline character instead of a comma. You can browse to the file and then select it. After uploading the file, the **IP Address/Hostname Discovery List** field will display the IP addresses and hostnames from the file.
- **SNMP Credentials.** A community string that allows SL1 to access a device's SNMP data. SNMP credentials are defined in the **Credential Management** page (System > Manage > Credentials). If you want to retrieve SNMP data from one or more devices, you must select one or more working SNMP credentials in this field. You can select multiple credentials from this field. SL1 will try each selected credential when discovering devices and retrieving device data.
- **Other Credentials.** A username and password pair (among other fields) that allows SL1 to access a device's database data, SOAP data, XML data, WMI data, WBEM data, or data that is monitored with a Snippet Dynamic Application. These credentials are defined in the **Credential Management** page (System > Manage > Credentials). You can select multiple credentials from this field. SL1 will try each selected credential when searching for Dynamic Applications to align with each discovered device.

NOTE: You can use the field at the top of the **SNMP Credentials** field and the **Other Credentials** field to filter the list of credentials. If you enter an alpha-numeric string in the field, the **SNMP Credentials** field or the **Other Credentials** field will include only credentials that match the string.

NOTE: Your organization membership(s) might affect the list of credentials you can see in the **SNMP Credentials** field and the **Other Credentials** field.

- **Initial Scan Level.** For this discovery session only, specifies the data to be gathered during the initial discovery session. The options are:
 - *System Default (recommended).* Use the value defined in the **Behavior Settings** page (System > Settings > Behavior).

- *0. Model Device Only.* Discovery will discover if the device is up and running and if so, collect the make and model of the device. SL1 will then generate a device ID for the device so it can be managed by SL1.
- *1. Initial Population of Apps.* Discovery will search for Dynamic Applications to associate with the device. The discovery tool will attempt to collect data for the aligned Dynamic Applications. Discovery will later retrieve full sets of data from each Dynamic Application. Discovery will also perform *0. Model Device Only* discovery.
- *2. Discover SSL Certificates.* Discovery will search for SSL certificates and retrieve SSL data. Discovery will also perform *1. Initial Population of Apps* and *0. Model Device Only*.
- *3. Discover Open Ports.* Discovery will search for open ports. Discovery will also perform *2. Discover SSL Certificates*, *1. Initial Population of Apps*, and *0. Model Device Only*.

NOTE: If your system includes a firewall and you select *3. Discover Open Ports*, discovery might be blocked and/or might be taxing to your network.

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

NOTE: If your system includes a firewall and you select *4. Advanced Port Discovery*, some devices might remain in a pending state (purple icon) for some time after discovery. These devices will achieve a healthy status, but this might take several hours.

- *5. Deep Discovery.* Discovery will use nmap to retrieve the operating system name and version. Discovery will also scan for services running on each open port and can use this information to match devices to device classes. Discovery will search for open ports, using a faster TCP/IP connection method. Discovery will also perform *2. Discover SSL Certificates*, *1. Initial Population of Apps*, and *0. Model Device Only*.

NOTE: For devices that don't support SNMP, option *5. Deep Discovery* allows you to discover devices that don't support SNMP and then align those devices with a device class other than "pingable". Note that option *5. Deep Discovery* is compute-intensive.

NOTE: If SL1 cannot determine the appropriate Device Class, it will assign the device to the Generic SNMP Device Class.

- **Scan Throttle.** Specifies the amount of time a discovery process should pause between each specified IP address (specified in the **IP Address/Hostname Discovery List** field). Pausing discovery processes between IP addresses spreads the amount of network traffic generated by discovery over a longer period of time. The choices are:
 - *System Default (recommended).* Use the value defined in the **Behavior Settings** page (System > Settings > Behavior).
 - *Disabled.* Discovery processes will not pause.
 - *1000 Msec to 10000 Msec.* A discovery process will pause for a random amount of time between half the selected value and the selected value.
- **Port Scan All IPs.** For the initial discovery session only, specifies whether SL1 should scan all IP addresses on a device for open ports. The choices are:
 - *System Default (recommended).* Use the value defined in the **Behavior Settings** page (System > Settings > Behavior).
 - *0. Disabled.* SL1 will scan only the primary IP address (the one used to communicate with SL1) for open ports.
 - *1. Enabled.* SL1 will scan all discovered IP addresses for open ports.
- **Port Scan Timeout.** For the initial discovery session only, specifies the length of time, in milliseconds, after which SL1 should stop trying to scan an IP address for open ports and begin scanning the next IP address (if applicable). Choices are:
 - *System Default (recommended).* Use the value defined in the **Behavior Settings** page (System > Settings > Behavior).
 - Choices between 60,000 to 1,800,000 milliseconds.
- **Detection Method & Port.** During discovery, SL1 will scan the list of ports selected in this field to determine if the range of devices is up and running and which ports are open on each discovered device. If a device does not respond to SNMP or ICMP, SL1 uses an open port to collect availability data for that device. If you are not sure which ports are used by the range of devices, select the entry *Default Method*. SL1 will check ICMP (ping), FTP, SSH, Telnet, SMTP, and HTTP ports.

NOTE: You can use the field at the top of the **Detection Method & Port** field to filter the list of ports. If you enter an alpha-numeric string in the field, the **Detection Method & Port** field will include only ports that match the string.

- **Interface Inventory Timeout (ms).** Specifies the maximum amount of time that the discovery processes will spend polling a device for the list of interfaces. After the specified time, SL1 will stop polling the device, will not model the device, and will continue with discovery. The default value is 600,000 ms (10 minutes).

- During the execution of this discovery session, SL1 uses the value in this field first. If you delete the default values and do not specify another value in this field, SL1 uses the value in the **Global Threshold Settings** page (System > Settings > Thresholds).
- If you specify a value in this field and do not apply a device template to this discovery session, the **Interface Inventory Timeout** setting in the **Device Thresholds** page (Registry > Devices > Device Manager > wrench icon > Thresholds) is set to this value for each discovered device. If there is no device template applied to the discovery session and no value is supplied in this field, SL1 uses the value in the **Global Threshold Settings** page (System > Settings > Thresholds).
- **Maximum Allowed Interfaces.** Specifies the maximum number of interfaces per devices. If a device exceeds this number of interfaces, SL1 will stop scanning the device, will not model the device, and will continue with discovery. The default value is 10,000.
 - During the execution of this discovery session, SL1 uses the value in this field first. If you delete the default values and do not specify another value in this field, SL1 uses the value in the **Global Threshold Settings** page.
 - If you specify a value in this field and do not apply a device template to this discovery session, the **Maximum Allowed Interfaces** setting in the **Device Thresholds** page is set to this value for each discovered device. If there is no device template applied to the discovery session and no value is supplied in this field, SL1 uses the value in the **Global Threshold Settings** page.
- **Bypass Interface Inventory.** Specifies whether or not the discovery session should discover network interfaces.
 - *Selected.* SL1 will not attempt to discover interfaces for each device in the discovery session. For each discovered device, the **Bypass Interface Inventory** checkbox in the **Device Properties** page will be selected.
 - *Not Selected.* SL1 will attempt to discover network interfaces, using the **Interface Inventory Timeout** value and **Maximum Allowed Interfaces** value.

NOTE: If a device has already been discovered and then is rediscovered through the **Discovery Session Editor** page, the **Bypass Interface Inventory** checkbox in the **Device Properties** page will retain its previous value, regardless of what is selected in the **Discovery Session Editor** page.

- **Discover Non-SNMP Devices.** Specifies whether or not SL1 should discover devices that don't respond to SNMP requests.
 - *Selected.* SL1 will discover devices that don't respond to the SNMP credentials selected in the **SNMP Credentials** field. These devices will be discovered as "pingable" devices.
 - *Not Selected.* SL1 will not discover devices that don't respond to the SNMP credentials selected in the **SNMP Credentials** fields.

- **Model Devices.** Determines whether or not the devices that are discovered with this discovery session can be managed through SL1. Choices are:
 - *Enabled.* When a device is modeled, SL1 creates a device ID for the device; you can then access the device through the **Device Manager** page and manage the device in SL1.
 - *Disabled.* If a device is not modeled, you cannot access the device through the **Device Manager** page, and you cannot manage the device in SL1. However, each discovered device will still appear in the Discovery Session logs. For each discovered device, the discovery logs will display the IP address and device class for the device. This option is useful when performing an initial discovery of your network, to determine which devices you want to monitor and manage with SL1. For the amount of time specified in the **Device Model Cache TTL (h)** field, a user can manually model the device from the **Discovery Session** window.
- **DHCP.** Specifies whether or not the specified range of IPs and hostnames use DHCP.
 - *Selected.* SL1 will perform a DNS lookup for the device during discovery and each time SL1 retrieves information from the device.
 - *Not Selected.* SL1 will perform normal discovery.
- **Device Model Cache TTL (h).** Amount of time, in hours, that SL1 stores information about devices that are discovered but not modeled, either because the **Model Devices** option is not enabled or because SL1 cannot determine whether a duplicate device already exists. The cached data can be used to manually model the device from the **Discovery Session** window.
- **Collection Server PID.** This field contains a list of all Data Collectors on the network. Select the Data Collector that is local or closet to the devices to be discovered.
 - For SL1 appliances, only the name of the appliance will appear in this field.


NOTE: After initial discovery, each device will use the collector group that contains this Data Collector for collection and rediscovery.

- **Organization.** This field contains a list of all organizations defined in SL1. Devices discovered during the discovery session will be assigned to the selected organization.

NOTE: Make sure you have the desired organization created and selected before running the discovery process. This field assigns all devices and networks in the specified IP range to a single organization. However, you can later assign individual devices and networks to different organizations.


- **Add Devices to Device Group(s).** When SL1 discovers a device in the IP discovery list, that device is added to each selected device group. You can select one or more device groups from a list of device groups in SL1 that have "Discovery" selected in the **Visibility** field. For more information on device groups, see the manual on **Device Groups and Device Templates**.

NOTE: You can use the field at the top of the **Add Devices to Device Group(s)** field to filter the list of device groups. If you enter an alpha-numeric string in the field, the **Add Devices to Device Group(s)** field will include only device groups that match the string.


- **Apply Device Template.** As SL1 discovers a device in the IP discovery list, that device is configured with the selected device template. You can select from a list of all device templates in SL1. For more information on device templates, see the manual on **Device Groups and Device Templates**.
 - **Log All.** Specifies whether or not the discovery session should use verbose logging. When you select verbose logging, SL1 logs details about each IP address or hostname specified in the **IP Address/Hostname Discovery List** field, even if the results are "No device found at this address."
 - *Selected.* This discovery session will use verbose logging.
 - *Not Selected.* This discovery session will not use verbose logging.
4. Click the **[Save]** button to **save the discovery session**. Close the **Discovery Session Editor** page.
 5. In the **Discovery Control Panel** page, click the **[Reset]** button. The new discovery session will appear in the **Session Register** pane.
 6. To launch the new discovery session, click its **Queue this Session** icon ().
 7. If no other discovery sessions are currently running, the session will be executed immediately. If another discovery session is currently running, your discovery session will be queued for execution.

Viewing Information about Classic Discovery

To view information about a discovery session that has already run:


1. Go to the **Discovery Control Panel** page (System > Manage > Classic Discovery).
2. In the **Discovery Control Panel** page, in the Session Register, find the discovery session you are interested in. Click its magnifying glass icon (.
3. The **Discovery Session** modal page appears. This page provides details on the discovery session.
4. The **Discovery Session** page includes an entry for each action performed during the discovery session. Each entry in the **Discovery Session** page contains:
 - **Date Time.** Date and time the action was executed.
 - **Discovery Log Message.** When applicable, starts with the IP address of the discovered device. Also includes a description of the action that was executed by the discovery session.

NOTE: If you did not select **Auto-Update** in the **Device Properties** page for one or more devices, when the discovery process tries to discover one of those devices, the log will display the message "Auto-Update is disabled and prevents discovery from updating any device properties."

- **Class Type.** The device class for each discovered device. SL1 will determine the device class for each device, even if a device will not be modeled by SL1.
 - **Checkbox.** If a device was discovered but not modeled, you can select this checkbox and click the **[Model]** button to model the device. If this device is a potential duplicate, the **Discovery Session** page displays the message "Not modeled, potential duplicate device". If you are certain that the device is not a duplicate, you can choose to model the device.
5. To save the log to the local computer, go to the **Discovery Control Panel** page (System > Manage > Classic Discovery), and click the **Export** icon () for the session.



Creating a New Classic Discovery Session with the "Save As" Button

You can edit an existing discovery session, make one or more changes, and then save the edited discovery session as a new session. The previous session still exists, unedited. To do this:

1. In the **Discovery Control Panel** page (System > Manage > Classic Discovery), in the **Session Registry** pane, find the discovery session you want to edit. Click its wrench icon (). The **Discovery Session Editor** page appears.
2. On the **Discovery Session Editor** page, you can edit one or more values listed in the [Running a Discovery Session](#) section.
3. Click the **[Save As]** button to save the discovery session as a new session. The new session appears in the **Discovery Control Panel** page (System > Manage > Classic Discovery).

Editing a Classic Discovery Session

You can edit the parameters of a discovery session in the **Discovery Control Panel** page. To do this:

1. Go to the **Discovery Control Panel** page (System > Manage > Classic Discovery).
2. Find the discovery session you want to edit. Click its wrench icon (.
3. The **Discovery Session Editor** page appears with the values from the previous discovery session. You can edit any of the fields described in the section [Running a Discovery Session](#).
4. Click the **[Save]** button to save your changes. To save the edited discovery session **as a new session** (the previous session will still exist), click the **[Save As]** button. Close the **Discovery Session Editor** page.
5. To manually run discovery using the edited session, find the edited session in the **Session Register** pane and click its **Queue this Session** icon (). If no other discovery sessions are currently running, the session will be executed immediately. If another discovery session is currently running, your discovery session will be queued for execution.

Scheduling a Classic Discovery Session

You can schedule one-time and recurring re-execution of a selected discovery session. You can use the **Schedule Manager** page (System > Manage > Classic Discovery > calendar icon) to:

- Specify regularly recurring discovery of a specific IP range. This is helpful when you want to monitor an IP range where new devices are frequently added.

- Specify one-time re-discovery of a specific IP range. This is helpful when you are aware of hardware or software changes within that IP range that you want SL1 to monitor.


NOTE: Scheduled re-execution of a discovery session is slightly different than SL1's automatic, nightly rediscovery. Nightly rediscovery is applied only to already discovered devices and uses the policies and configuration applied to each device. Re-execution of a discovery session can discover new devices within an IP range and uses only the settings specified in the **Discovery Control Panel** 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 (System > Manage > Classic Discovery > calendar icon) displays the following information about each scheduled or recurring discovery session:

- **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 discovery session, click its wrench icon () and update the discovery session as needed on the **Schedule Editor** modal page. (For more information, see the section [Defining a Scheduled or Recurring Discovery Session](#).)

Defining a Scheduled or Recurring Discovery Session

You can schedule a discovery session in SL1 from the **Schedule Manager** page. SL1 will automatically run the discovery session at the scheduled time.

To define a scheduled or recurring discovery session:

1. Go to the **Schedule Manager** page (System > Manage > Classic Discovery > calendar icon).
2. Click **[Create]**. The **Schedule Editor** modal page appears.
3. On the **Schedule Editor** modal page, make entries in the following fields:

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.
- **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.

- **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 Date** field.
 - If you specify "10 Days", then the scheduled process recurs every 10 days from the date listed in the **Start Date** field.
 - If you specify "2 Weeks", then the scheduled process recurs every two weeks, on the same day of the week as the **Start Date**.
 - If you specify "3 Months" the ticket recurs every three months, on the same day of the month as the **Start Date**.
- **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 enter 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 Discovery Sessions

You can enable or disable one or more scheduled or recurring discovery sessions from the **Schedule Manager** page (System > Manage > Classic Discovery > calendar icon). To do this:

1. Go to the **Schedule Manager** page (System > Manage > Classic Discovery > calendar icon).

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 Discovery Sessions

You can delete one or more scheduled or recurring discovery sessions from the **Schedule Manager** page (System > Manage > Discovery > calendar icon). To do this:

1. Go to the **Schedule Manager** page (System > Manage > Classic Discovery > calendar icon).
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.

Manually Re-Running Classic Discovery for a Dynamic Application

From the **Dynamic Applications Manager** page, you can manually run the Dynamic Application alignment portion of discovery for all existing devices in the system. That is, you can ask SL1 to check each Dynamic Application and each existing device and align each device with each appropriate Dynamic Application.

For each Dynamic Application you select for re-discovery, SL1 tries to connect to each existing device using the credentials already aligned with each device. If SL1 is able to connect to a device with one of the credentials and can then retrieve the discovery object associated with the Dynamic Application, SL1 will align the Dynamic Application with the device.

To manually run discovery for a *single Dynamic Application*:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Dynamic Applications).
2. In the **Dynamic Applications Manager** page, find the Dynamic Application you want to use for network-wide discovery. Click its lightning bolt icon (⚡).
3. If no other discovery tasks are currently running, SL1 will immediately perform discovery of the entire network, using the selected Dynamic Application.
4. If other discovery tasks are currently running, SL1 will add the request to the discovery queue.

To manually run discovery for *multiple Dynamic Applications*:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Dynamic Applications).
2. In the **Dynamic Applications Manager** page, find the Dynamic Applications you want to use for network-wide discovery. Select the checkbox for each Dynamic Application.
3. In the **Select Actions** menu (in the lower right of the page), select **DISCOVER Applications**.
4. Click the **[Go]** button.
5. If no other discovery tasks are currently running, SL1 will immediately perform discovery of the entire network, using the selected Dynamic Application.
6. If other discovery tasks are currently running, SL1 will add the request to the discovery queue.

Manually Re-Running Classic Discovery for a Device



You can manually re-discover a device, using the settings and configuration in the **Device Properties** page for the device.

Remember that the credentials and settings defined in the **Device Properties** page override:

- Settings in the **Behavior Settings** page (System > Settings > Behavior).
- Credentials and settings in the **Discovery Control Panel** page (System > Manage > Classic Discovery) from the initial discovery of the device.

SL1 will update the device with the data from the discovery session. The discovery session does not change, overwrite, or affect the existing, historical data associated with the device.

To re-discover a device:

1. Go to the **Device Manager** page (Devices > Device Manager).
2. In the **Device Manager** page, find the device you want to re-discover. Click its wrench icon (.
3. In the **Device Properties** page, click the **Rediscover** icon (.
4. If no other discovery tasks are currently running, SL1 will immediately perform discovery of the selected device.
5. If other discovery tasks are currently running, SL1 will add the request to the discovery queue.

Chapter



6

Collection

Overview

This chapter describes collection and collection processes 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:

What is Collection?	146
Collection Processes	147

What is Collection?

Collection is the tool that retrieves policy-based information and Dynamic Application-based information from a device. After a device is discovered, you can define monitoring policies for that device in SL1. For example, if you define a policy to monitor a system process, the collection tool retrieves that information.

SL1 uses the following methods for collection:

- Dynamic Applications use collection processes to collect data.
- Monitoring Policies for devices also trigger collection. These policies include:
 - Domain Name policies
 - Email Round-Trip policies

- SOAP/XML Transaction policies
- System Process policies
- TCP/IP Port policies
- Web Content policies
- Windows Services policies
- SL1 automatically collects the following about each managed device:
 - Device availability and device latency
 - Network topology
 - File system information, if available
 - A list of open ports
 - Bandwidth usage
- The SL1 agent automatically collects the following about each device on which it is installed:
 - Device availability
 - Device performance and configuration metrics
 - A list of open ports
 - Log information
 - System processes

Collection Processes

Unlike discovery, collection tasks run at scheduled intervals throughout the day. Collection tasks collect the types of data described below. The interval specified is the default interval and can be modified.

- Device availability and device latency (based on the port through which SL1 communicates), every five minutes.
- CDP relationships between devices, every two hours.
- LLDP relationships between devices, every two hours.
- Critical device availability (if enabled, based on ping to specified port), every 5 seconds, 30 seconds, 60 seconds or 120 seconds (defined by user).
- Critical port availability (if enabled, based on ping to specified port), every 5 seconds, 30 seconds, 60 seconds or 120 seconds (defined by user).
- DNS availability based on DNS-monitoring policies, every five minutes.

- Data specified in Dynamic Applications. Collection tasks retrieve data from each aligned device, at the frequency specified in the Dynamic Application.
- Email round-trip statistics based on Email-monitoring policies, every five minutes.
- File system information, every five minutes.
- File system inventory, every two hours.
- Bandwidth usage on managed interfaces, every minute, 5 minutes, 10 minutes, 15 minutes, 30 minutes, 60 minutes, or 120 minutes (defined by user).
- Layer-3 relationships between devices, every two hours.
- List of all discovered system processes on all discovered devices, every two hours.
- Availability of system processes based on process-monitoring policies, every five minutes.
- List of all discovered Windows services on all discovered devices, every two hours.
- Availability of Windows services based on service-monitoring policies, every five minutes.
- SNMP details for each discovered device, every five minutes.
- Availability of ports based on port-monitoring policies, every five minutes.
- Layer-2 relationships between devices, every hour.
- Virtual machine relationships between devices, every hour.
- Availability of web content based on web content-monitoring policies, every five minutes.
- Web-transaction statistics based on a SOAP/XML-monitoring policy, every five minutes.
- If the SL1 agent is installed, SL1 collects a list of all processes running on a device, every five minutes.



For details on collection processes, go to the **[Processes]** tab of the **Device Investigator** or the **Process Manager** page (System > Settings > Admin Processes) and look for processes with names that start with "Data Collection".

Duplication Protection During Discovery

Overview

This chapter describes the duplication protection features that SL1 employs during discovery.

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>Duplicate IP Addresses and Duplicate MAC Addresses During Discovery</i>	149
<i>Duplicate MAC Addresses for Component Devices</i>	151
<i>Managing MAC Vendor Records</i>	152

Duplicate IP Addresses and Duplicate MAC Addresses During Discovery

NOTE: SL1 applies duplication protection to only those devices discovered by IP address. Therefore, the content in this section does not apply to devices discovered by hostname.

NOTE: Component devices are discovered using Dynamic Applications instead of using a discovery session. The description in this section does not apply to component devices. For details on how SL1 handles duplicate MAC addresses for component devices, see the section [Duplicate MAC Addresses for Component Devices](#).

During discovery, SL1 compares the IP addresses and MAC addresses of a newly discovered device with the IP addresses and MAC addresses of devices that have already been discovered to ensure that devices are not duplicated and IP conflicts do not occur. There are three possible outcomes of this comparison:

- The newly discovered device is considered a duplicate of an existing device and the information collected during discovery is used to update that existing device record.
- The newly discovered device is not a duplicate of an existing device and the information collected during discovery is used to create a new device record.
- The newly discovered device could be considered a duplicate of an existing device. SL1 does not automatically use the information collected during discovery to either create or update a device record. **The Discovery Session Logs** page includes an option for a user to manually create a new device record using the information collected during discovery.

NOTE: If SL1 discovers an existing MAC address that is not is part of a range of MAC addresses that are marked as "virtual", it will be considered a duplicate of an existing device, regardless of its collector group.

Each managed device in SL1 can have three types of IP addresses:

- **Admin Primary.** This is the IP address that SL1 uses to communicate with a device. This IP address is always a primary address and cannot be demoted to a secondary address. Within a single Collector Group, devices cannot have duplicate **Admin Primary** IP addresses.
- **Primary.** One or more IP addresses that SL1 uses to match incoming log messages (traps and syslog messages) with a device. You can define a primary IP address in the **Device Properties** page for a device.
- **Secondary.** SL1 gathers information about this IP address, but does not use this IP address to match incoming messages (traps and syslog messages) with a device. You can define a secondary IP address in the **Device Properties** page for a device.

If any of the following conditions are true, the newly discovered device is considered a duplicate of an existing device and the information collected during discovery is used to update that existing device record:

- The primary IP address of the newly discovered device is not unique to the Collector Group that discovered the device.
- The primary IP address of the newly discovered device is not unique within the system and the primary IP address is in a public address space.
- A secondary IP address of the newly discovered device is not unique within the system and that secondary IP address is in a public address space.

If any of the following conditions are true, the newly discovered device could be considered a duplicate of an existing device. SL1 does not automatically use the information collected during discovery to either create or update a device record. The **Discovery Session Logs** page includes an option for a user to manually create a new device record using the information collected during discovery:

- The primary IP address of the newly discovered device is unique within the system; the secondary IP addresses associated with the newly discovered device are either unique within the system or are in a private address space; and the MAC addresses associated with the device match the MAC addresses associated with an existing device.
- The primary IP address of the newly discovered device is not unique within the system; the primary IP address is in a private address space; the secondary IP addresses associated with the newly discovered device are either unique within the system or are in a private address space; and the MAC addresses associated with the device match the MAC addresses associated with an existing device.
- The primary IP address of the newly discovered device is not unique within the system; the primary IP address is in a private address space; and no secondary IP address information has been discovered for the device.
- The primary IP address of the newly discovered device is not unique within the system; the primary IP address is in a private address space; the secondary IP addresses associated with the newly discovered device are either unique within the system or are in a private address space; and no MAC address information has been discovered for the device.

NOTE: For more information about configuring virtual MAC addresses, see the section [Managing MAC Vendor Records](#).

Duplicate MAC Addresses for Component Devices

SL1 handles duplicate MACs for component devices differently than duplicate MACs for physical devices. When a component device is assigned a MAC address, SL1 does not enforce uniqueness and will allow a component device to be created with the same MAC address as existing physical devices and/or existing component devices.

Unlike how SL1 discovers physical devices, SL1 uses Dynamic Applications to retrieve data from a management device and "discover" each entity managed by that management device as a component 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, physical devices are identified by IP address and MAC address. In SL1, component devices are identified by a device name, a unique identifier, and a device class. A Dynamic Application that creates component devices can assign a MAC address to each component device, but is not required to.

- 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 example, in a Cisco UCS system, SL1 might discover a physical server that hosts the UCS manager. This physical server is the **root device**. SL1 might discover a chassis on the root device. The chassis is a **component device**. The chassis is a child device to the physical server. SL1 might also discover a blade as a component device that is part of the chassis. The blade is a child device to the chassis. The chassis is the **parent device**.

SL1 does not automatically combine new component devices with any existing device record using the MAC address of the new component device. A component device can be combined with an existing device record under the following conditions:

- Dynamic Applications that create component devices can assign a globally unique identifier (GUID) to each component device. When SL1 performs collection for a Dynamic Application, and the Dynamic Application includes a collection object with a GUID component identifier, SL1 compares the collected values for that collection object with all GUID values for all component devices discovered in the system. If a newly collected value matches a GUID value for an existing component device, the device from which SL1 collected the new value will become the parent of the existing component device. The existing component device will no longer be associated with its previous parent device. No new component device will be created.
- You can merge a physical device and a component device. You can do this in the **[Actions]** menu in the **Device Properties** page (Devices > Device Manager > wrench icon) for either the physical device or the component device. When you merge a physical device and a component device, the device record for the component device is no longer displayed in the user interface; the device record for the physical device is displayed in user interface pages that previously displayed the component device. For example, the physical device is displayed in lieu of the component device in the **Device Components** page and the **Component Map** page. All existing and future data for both devices will be associated with the physical device. You can unmerge a component device from a physical device in the **[Actions]** menu in the **Device Properties** page for the physical device (Devices > Device Manager > wrench icon).

Managing MAC Vendor Records

The **MAC Vendor Records** page (System > Customize > MAC Vendors) allows you to view and edit the list of MAC Vendor Records in SL1. MAC Vendor Records include vendor information about each MAC address prefix. A MAC address prefix is the first three groups of hexadecimal digits in a MAC address. The MAC Address prefix uniquely identifies the vendor of the network interface. Some vendors use multiple MAC Address prefixes, but each vendor's MAC Address prefixes are unique to that vendor and are not used by other vendors.

Viewing the List of MAC Vendor Records

The **MAC Vendor Records** page (System > Customize > MAC Vendors) displays information about each MAC Vendor Record in SL1. The **MAC Vendor Records** page displays the following for each MAC Vendor Record:

- **MAC Hex.** The MAC address prefix for the vendor record.
- **Vendor.** The name of the vendor of the network interfaces that use the MAC address prefix.
- **Vendor Notes.** Additional information about the vendor.
- **Virtual.** Indicates whether the vendor associated with this MAC address prefix allows the same MAC address to be re-used on multiple devices:
 - Yes. The vendor allows the same MAC address to be re-used on multiple devices. If a new interface is discovered during nightly discovery and that interface has a MAC address with this prefix that is already associated with an interface record in the system, SL1 will create a new interface record for the newly discovered interface.
 - No. The vendor does not allow the same MAC address to be re-used on multiple devices. If a new interface is discovered during nightly discovery and that interface has a MAC address with this prefix

that is already associated with an interface record in the system, SL1 will NOT create a new interface record for the newly discovered interface.

Filtering the List of MAC Vendor Records

The **MAC Vendor Records** page includes four filters, in the top row in the list of MAC Vendor Records. You can specify one or more parameters to filter the display of MAC Vendor Records. Only MAC Vendor Records that meet all the filter criteria will be displayed in the **MAC Vendor Records** page.

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

- For each filter, you must enter text to match against. SL1 will search for MAC Vendor Records that match the text, including partial matches. Text matches are not case sensitive. You can use the following special characters in each filter:
 - , (comma). Specifies an "or" operation. For example:
"dell, micro" would match all values that contain the string "dell" OR the string "micro".
 - ! (exclamation mark). Specifies a "not" operation. For example:
"!dell" would match all values that do not contain the string "dell".
- **MAC Hex**. You can enter text to match, including special characters, and the **MAC Vendor Records** page will display only MAC Vendor Records that have a matching prefix.
- **Vendor**. You can enter text to match, including special characters, and the **MAC Vendor Records** page will display only MAC Vendor Records that have a matching vendor name.
- **Vendor Notes**. You can enter text to match, including special characters, and the **MAC Vendor Records** page will display only MAC Vendor Records that have a matching vendor note.
- **Virtual**. You can enter text to match, including special characters, and the **MAC Vendor Records** page will display only MAC Vendor Records that have a matching virtual setting.

Editing the Virtual Setting for MAC Vendor Records

The **MAC Vendor Records** page contains a drop-down field in the lower right called **Select Action**. This field allows you to apply an action to multiple MAC Vendor Records simultaneously.

To apply an action to multiple MAC Vendor Records:

1. Go to the **MAC Vendor Records** page (System > Customize > MAC Vendors).
2. In the **MAC Vendor Records** page, select the checkbox for each MAC Vendor Record you want to apply the action to. To select all checkboxes for all MAC Vendor Records, select the red checkbox at the top of the page.
3. In the **Select Action** drop-down list, select one of the following actions:
 - **Set Virtual Flag To: Yes**. Sets the virtual setting for the Mac Vendor Record to "yes".
 - **Set Virtual Flag To: No**. Sets the virtual setting for the Mac Vendor Record to "no".

4. Click the **[Go]** button to apply the selected action to the selected MAC Vendor Records.

Chapter

8

Troubleshooting Discovery

Overview

If discovery is not working as you expected, you can try these troubleshooting steps to try to fix any problems. If problems persist, please contact ScienceLogic Customer Care.

To perform the troubleshooting steps in this manual, you must be allowed root-level access to SL1 appliances from a shell session.

This chapter covers the following topics:

<i>Checking Network Security</i>	155
<i>Debugging the Discovery Processes</i>	156
<i>Checking Communication between Data Collectors and the Database Server</i>	157

Checking Network Security

Your network security and network configuration can prevent SL1 from communicating with each device in your network. To ensure that discovery can access each device in your network, check the following:

- To discover a device as a "pingable" device, SL1 must be able to either:
 - Ping the device (access through ICMP).
 - Access at least one of the ports selected in the **Discovery Control Panel** page.
- To discover a device as "manageable" (that is, the device supports SNMP), SL1 must be able to access the UDP port defined in the working SNMP credential for that device.
- On each DNS server(s) for your network, SL1 must access UDP port 53.

- If there are firewalls between the Data Collectors and devices to be monitored, make sure that the firewalls are configured to allow SL1 access to those devices.

Debugging the Discovery Processes


When you debug a process, you tell SL1 to use verbose logging for that process. You can then view the log file to view detailed log files. If discovery is not performing as you would expect, you might find it helpful to debug one or more of the discovery processes.

In general:

- If SL1 is not discovering one or more devices that you know exist, debug the process **Discovery: Auto**.
- If SL1 is discovering devices but not retrieving the appropriate information, debug the process **Discovery: Detail**.
- If SL1 is not aligning Dynamic Applications with devices during discovery, debug the process **Discovery: Dynamic App**.

WARNING: ScienceLogic recommends that you enable the debug option only while troubleshooting a problem and that you immediately turn off debugging when you have completed troubleshooting. Don't leave the debug option enabled during normal operation of SL1. When you turn on debugging, SL1 will run significantly more slowly.

To enable the debug option for a discovery process:

1. Go to the **Process Manager** page (System > Settings > Admin Processes).
2. In the **Process Manager** page, find the process you want to debug and click its wrench icon (.
3. The **Process Editor** modal page appears.
4. Edit the following field:
 - **Debug Mode.** Enables or disables debugging information for a process. Select *Enabled*.
5. Click the **[Reset]** button.
6. Go to the **Discovery Control Panel** page (System > Manage > Classic Discovery) and *run the discovery session again*.
7. Log in to the console of the SL1 appliance where the process is running. Alternately, you can use SSH to open a shell session on the SL1 appliance. In most cases, you will log in as "root".
 - If you are using a SL1 system, log in to the current SL1 appliance.
 - If you are using a distributed SL1 system, log in to the Data Collector associated with the discovery session.

NOTE: For details on enabling and using SSH with SL1, information about root access, and instructions on making root access secure, see the manual **Security**.

TIP: To view a list of IP addresses for all SL1 appliances in your network, go to the **Appliance Manager** page (System > Settings > Appliances).

8. Navigate to the directory `/data/logs`. View the file `silol.log`. The most recent entries will be posted at the end of the file.
9. After you have finished troubleshooting the process, remember to disable debugging.

Checking Communication between Data Collectors and the Database Server

For distributed SL1 systems, discovery can fail if the Data Collectors and the Database Server cannot communicate with each other.

To check communication between the Database Server and a Data Collector:

1. Log in to the console of the SL1 appliance where the process is running. Alternately, you can use SSH to open a shell session on the SL1 appliance. Log in as **em7admin**.
 - If you are using a SL1 system, log in to the current SL1 appliance.
 - If you are using a distributed SL1 system, log in to the Database Server.

TIP: To view a list of IP addresses for all SL1 appliances in your network, go to the **Appliance Manager** page (System > Settings > Appliances).

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

2. From the command line, enter the following:

```
silomysql -P 7707 -h <IP address of the Data Collection Server  
associated with the discovery session> -u root -p
```

You will be prompted to enter the MySQL root password.

3. If you can successfully execute this command from the Database Server, this means that the Database Server is successfully communicating with the Data Collector.
4. If you cannot successfully execute this command from the Database Server:

- Go to the **Appliance Manager** page (System > Settings > Appliances) and ensure that the settings for the Database Server and the Data Collector server are correct.
- Ensure that a network firewall is not preventing the Database Server and the Data Collector from communicating with each other.
- Log in to the console of the Data Collector. Alternately, you can use SSH to open a shell session on the Data Collector. Log in as **em7admin**. To ensure that MySQL is running on the Data Collector, enter the following at the command line:

```
mysqladmin -u root -p status
```

You will be prompted for the password.

- If the database is running, the command will return statistics about the database.
- If the database is not running, the command will return an error message. To restart the database, enter the following at the command line:

```
sudo service em7_db_start
```

© 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